Tutorial

Hybrid Oscillator-Qubit Quantum Processors: Instruction Set Architectures, Abstract Machine Models, and Applications

Yuan Liu,[1,2,3,*] Shraddha Singh,[4,5,*] Kevin C. Smith,[4,5,6,*] Eleanor Crane,[1] John M. Martyn,[1] Alec Eickbusch,[4,5,†] Alexander Schuckert,[7] Richard D. Li,[4,5] Jasmine Sinanan-Singh,[1] Micheline B. Soley,[8] Takahiro Tsunoda,[4,5] Isaac L. Chuang,[1,9] Nathan Wiebe,[10,11,12] and Steven M. Girvin[4,5,‡]

[1] *Department of Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*
[2] *Department of Electrical and Computer Engineering, North Carolina State University, Raleigh, North Carolina 27606, USA*
[3] *Department of Computer Science, North Carolina State University, Raleigh, North Carolina 27606, USA*
[4] *Yale Quantum Institute, Yale University, New Haven, Connecticut 06511, USA*
[5] *Departments of Applied Physics and Physics, Yale University, New Haven, Connecticut 06520, USA*
[6] *Brookhaven National Laboratory, Upton, New York 11973, USA*
[7] *Joint Quantum Institute and QuICS, National Institute of Standards and Technology (NIST) and University of Maryland, College Park, Maryland 20742, USA*
[8] *Department of Chemistry, University of Wisconsin-Madison, Madison, Wisconsin 53706, USA*
[9] *Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*
[10] *Department of Computer Science, University of Toronto, Toronto, ON M5S1J7, Canada*
[11] *Canadian Institute for Advanced Research, Toronto, ON Canada*
[12] *Pacific Northwest National Laboratory, Richland, Washington, USA*

(Received 28 August 2024; revised 15 August 2025; published 7 January 2026)

This tutorial offers a pedagogical guide to hybrid quantum processors that integrate discrete-variable (DV) qubits and continuous-variable (CV) oscillators. Aimed at computer scientists, engineers, and physicists, it provides an overview of the experimental, algorithmic, and architectural aspects of this novel and rapidly developing hardware model. Experimental realizations of this model include superconducting, trapped-ion, and neutral-atom platforms. By combining DV and CV components, hybrid oscillator-qubit processors enable a powerful new paradigm that offers complementary strengths for quantum control, error correction, computation, and simulation. Working toward the goal of a full-stack system connecting applications to CV-DV hardware, we define and formulate abstract machine models and instruction set architectures. These essential abstractions enable codesign of hardware and software, and resource estimation for exploring the potential of current and future hardware for computational and simulation tasks. Using these abstractions, we present both new and existing examples that illustrate the benefits of hybrid CV-DV processors relative to traditional DV-only hardware in computation as well as quantum simulation of physical models. Examples include algorithms for transferring states between DV and CV systems, performing the quantum Fourier transform, and simulation of lattice gauge theories. Relative to qubit-only hardware, the bosonic degrees of freedom natively available in hybrid architectures can substantially reduce the circuit complexity of simulations for physical models containing bosons. A key technique is the extension of quantum signal processing ideas to CV-DV systems. This work is intended to serve as a timely and comprehensive guide to this relatively unexplored yet promising approach to quantum computation and to provide a road map to guide future development.

DOI: 10.1103/4rf7-9tfx

[*] These authors contributed equally to this work (ordered alphabetically).
[†] Present address: Google Quantum AI, Santa Barbara, California 93111, USA.
[‡] Contact author: steven.girvin@yale.edu

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

CONTENTS

I. INTRODUCTION 3
A. Outline of this paper 4
B. Reader's guide 4
II. HYBRID CV-DV QUANTUM COMPUTING ARCHITECTURE 6
A. Physical model of hybrid CV-DV quantum computation 6
B. Abstract machine models (AMMs) 7
C. Instruction set architectures 12
D. Quantum error correction 14
E. Programming languages and compilation 14
F. Algorithms and applications 15
III. STATES AND OPERATORS 17
A. Qubit states and operations 17
B. The Hilbert space of oscillators 18
C. Representation and bases 18
D. Truncation of the Hilbert space 24
IV. INSTRUCTION SETS FOR OSCILLATOR-QUBIT SYSTEMS 25
A. Gaussian oscillator control 27
B. Gaussian operations 29
C. Gaussian operations are analogous to Clifford gates 36
D. Hybrid non-Gaussian operations 39
E. Instruction sets for universal control of hybrid oscillator-qubit systems 42
F. Control flow and benchmarking 45
V. QUANTUM ERROR CORRECTION WITH OSCILLATORS 50
A. Dissipation and decoherence 50
B. Overview of bosonic QEC 52
C. Logical ISA for bosonic QEC 55
VI. COMPILATION METHODS 64
A. Introduction and overview of compilation schemes 64
B. Representations of hybrid CV-DV algorithms 65
C. Exact analytical compilation of entangling gates 66
D. Systematic compilation of single-qubit-oscillator unitaries 74
E. Approximate numerical compilation 82
VII. ALGORITHMS AND APPLICATIONS 84
A. Arbitrary qubit-oscillator state transfer 85
B. Quantum Fourier transform 89
C. Quantum simulation 91
D. Quantum random walks 100
VIII. ADVANTAGES OVER QUBIT-ONLY ARCHITECTURES 102
A. Mapping between oscillators and qubits 102
B. Compilation of bosonic operators into and from standard qubit gate sets 104
C. Efficiency of bosonic instruction set relative to qubit-only simulations 109
IX. CONCLUSIONS AND OUTLOOK 110
A. Challenges and opportunities for AMMs and ISAs 110
B. Open questions and directions for future work 112
ACKNOWLEDGMENTS 114
DATA AVAILABILITY 115
APPENDIX A: PHYSICAL IMPLEMENTATIONS 115
1. Superconducting and cavity QED 115
2. Trapped-ion quantum processors 119
3. Neutral-atom quantum processors 120
APPENDIX B: BLOCH-MESSIAH DECOMPOSITION 122
1. SUM gate 122
2. Two-mode squeezing 123
APPENDIX C: OSCILLATOR NOISE MODELS 123
1. Kraus operator representation of an error channel 123
2. Amplitude damping 124
3. Dephasing 124
4. Heating 124
APPENDIX D: UNIVERSALITY OF ISAs AND CROSS-COMPILATION 125
1. Compilation between the Fock-space ISA and the phase-space ISA 125
2. Compilation between sideband ISA and phase-space ISA 127
3. Compilation between the sideband ISA and the Fock-space ISA 127
4. Universality of multi-qubit–oscillator systems 128
5. Oscillator-only universality: Controllability and the cubic instruction set 129

APPENDIX E: MEASUREMENTS AND TOMOGRAPHY 131
1. Hybrid qubit-cavity measurements 131
2. Dispersive readout of the qubit 131
3. Cavity photon-number measurement 132
4. Cavity photon-number "superparity" measurement 133
5. Homodyne detection 134
6. Wigner and characteristic function tomography 135
7. The Husimi Q-function 137
8. Calibrating controlled displacements 137
APPENDIX F: COMPARISON BETWEEN ANALYTICAL AND NUMERICAL COMPILATION 137
1. Trotterization 138
2. QSP 138
APPENDIX G: BRIEF REVIEW OF QSP AND QSVT FORMALISMS FOR QUBITS 139
1. Intuition from composite pulses 139
2. Qubit QSP 140
APPENDIX H: DETAILS ON QFT IMPLEMENTATION 141
1. Details on QFT method 1: Single-variable QSP 141
2. Details on QFT method 2: Non-Abelian QSP 143
REFERENCES 144

I. INTRODUCTION

The ability to harness quantum resources in physical systems has revolutionary potential for computation, information processing, and communication [1,2]. Most work to date has largely focused on leveraging two-state, discrete-variable (DV) systems acting as qubits for quantum computation. These include nuclear [3,4] and electron spins [5], pairs of levels in neutral atoms [6–9], ions [10–20] and color centers [21], as well as synthetic atoms such as quantum dots [22] and superconducting qubits [23–29], and various possible hybrid systems [30–32]. Separately, quantum resources that fundamentally differ from two-level spin systems have been proposed for use, including bosons [33,34] and fermions [35–39]. Of particular importance is the fact that a single bosonic mode, also known as a harmonic oscillator, has a countable infinity of states and can therefore be characterized by a continuous variable (CV), in stark contrast to both qubits and fermions.

In this tutorial, we focus our discussion on *hybrid quantum CV-DV systems* composed of both oscillators and qubits, a paradigm that offers great potential for quantum computation as well as quantum simulation of physical models of interacting bosons (e.g., the boson Hubbard model [40]) or bosons coupled to spins or fermions (e.g., lattice gauge theories [41] and polaritonic chemistry [42]). Thus far, a systematic discussion of such systems has not been established in the literature. This is largely due to several theoretical challenges. First, the physics of bosonic modes is very different from qubits, and the creation and annihilation operators (or equivalently, the position and momentum operators) used to describe the Hamiltonian of oscillators satisfy completely different commutation relations than the familiar Pauli matrices in the qubit case. This can make it challenging to treat oscillators and qubits in the same mathematical framework for quantum computation, while hinting that bosonic modes can be a resource of considerable computational power. Second, going beyond the physical systems to a more abstract computer science point of view, answers to crucial questions have been lacking. How should such hybrid quantum CV-DV processors be programmed? How can their computational power be reasoned about? How can the required resources for such processors be estimated?

Although the first set of challenges is more fundamental, the key to overcoming the computer science challenges highlighted above is to construct *abstract machine models* (AMMs) and corresponding *universal instruction set architectures* (ISAs) for such hybrid processors. An ISA [43] contains an inventory of the discrete set of fundamental operations and measurements that are possible in the hardware. As is well known on the classical side, AMMs such as Turing machines, random access machines, finite automata, and even lambda calculus [44] facilitate formal reasoning about the computational power of classical computers (e.g., the computability of specified functions such as the halting function). For qubit-only quantum processors, AMMs and ISAs have been useful for quantifying the quantum advantage of qubit-based quantum computers [45,46]. Similarly, the establishment of AMMs and ISAs for hybrid CV-DV quantum processors will be crucial to quantifying their computational power and performing resource estimation, as well as for codesign of algorithms, software, and hardware [47,48] to demonstrate the potential quantum advantage of not only quantum versus classical computers but also hybrid oscillator-qubit versus qubit-only quantum computers.

The goal of the present work is to provide a compelling tutorial overview of hybrid CV-DV quantum computation through a complete, bottom-up picture spanning seven main layers, from physical mechanisms through models, architectures, programming languages and compilation, and algorithms and applications:

(a) *Physical mechanisms*: foundational physical primitives and formalism for hybrid quantum CV-DV systems and their physical hardware realizations.

(b) *Abstract machine models*: theoretical abstractions of hybrid quantum CV-DV hardware that allow exploration of design decisions, bridging between computational complexity and architecture design.
(c) *Instruction set architectures*: hardware-software interfaces that allow translation of algorithms into well-defined primitive and abstract operations.
(d) *Programming languages and compilation*: languages for algorithmic design and techniques to map these instructions onto circuits executable within CV-DV ISAs, including:

   (i) analytic techniques for compiling multi-qubit and multi-oscillator entangling gates
   (ii) extension of quantum signal processing to hybrid CV-DV systems, and
   (iii) methods for realizing bosonic quantum error correction (QEC) through compilation

(e) *Algorithms and applications*: computational problems (described in a high-level programming language) naturally compilable into this model, their resource requirements, and debugging procedures, including such examples as:

   (i) quantum state transfer between CV and DV systems, with application to performing the quantum Fourier transform
   (ii) efficient techniques for state and process tomography in hybrid CV-DV systems
   (iii) quantum simulation of physical models involving both oscillators and qubits

A. Outline of this paper

The remainder of this paper is structured as follows. Section II presents the hybrid CV-DV architecture, discussing in detail the seven layers of abstraction. We then dive deeper into the theoretical aspects and experimental prospects of each architecture layer. Beginning with the physical layer, in Sec. III we provide a pedagogical introduction to the representation of bosonic states, intended for nonphysicists who are familiar with qubits but not oscillators. In Sec. IV, we present the hierarchy of bosonic and hybrid gates beginning with simple Gaussian operations on bosonic modes (analogous to Clifford operations on qubits), followed by non-Gaussian hybrid oscillator-qubit gates, and ending with descriptions of various universal instruction sets. In Sec. V, we give an overview of QEC in hybrid CV-DV processors. In Sec. VI, we present exact and approximate analytical and numerical compilation and circuit synthesis techniques, including the extension of the quantum signal processing framework to hybrid CV-DV systems. In Sec. VII, we leverage the compilation techniques developed in the preceding section for small-scale applications including useful computational subroutines and quantum simulation of important physical models. In Sec. VIII, we discuss the advantages of oscillator-qubit hardware over qubit-only hardware and, *en route*, introduce novel compilation methods to simulate native oscillator gates in the latter setting. Finally, Sec. IX presents conclusions, outlooks, and open research problems. This is followed by a series of technical appendixes on physical implementations and microscopic Hamiltonians, realistic noise models, measurement and tomography, circuit synthesis techniques, universality and cross-compilation of different ISAs, and other details relevant to the main text.

About vocabulary and notation. Throughout the text, we use the terms bosonic mode, qumode, oscillator, resonator, and cavity mode interchangeably. To indicate the identity and Pauli matrices, we use both the physics notation $\sigma_0, \sigma_x, \sigma_y, \sigma_z$ and the quantum computer science notation I, X, Y, Z. It is also important to note that we use the Nielsen and Chuang [1] convention that $|0\rangle$ represents the excited state of the qubit and $|1\rangle$ represents the ground state. This convention is carefully defined in Sec. III A. For clarity, we often use $|\cdot\rangle_{\text{osc}}$ to distinguish the oscillator subspace from the qubit subspace.

B. Reader's guide

Given the length of this paper and the interdisciplinary nature of the material covered, it is useful to augment the above outline with a "reader's guide." In Fig. 1 we provide a simple map for readers interested in navigating to particular topics. In the following, we give recommendations for how readers with different backgrounds can engage with the material.

1. Beginners

For nonexpert readers, following the architectural overview in Sec. II, we begin in Sec. III with an introduction to quantum harmonic oscillators and the language from the field of quantum optics used to describe them. It is essential to understand this foundational material. For a first reading, the subsequent sections are best covered in the order presented.

2. CV or quantum optics background

For readers with a strong background in quantum optics but less background in quantum computing, we suggest using this tutorial as follows. We recommend that these readers familiarize themselves with the discussion on instruction set architectures and abstract machine models in Sec. II, as these ideas are likely to be new. We recommend skimming through the majority of Sec. III, but paying attention to the truncation discussion in Sec. III D and the mapping between oscillators and qubits in Sec. VIII A. Unlike traditional quantum optics, circuit QED (cQED, the microwave quantum optics of superconducting circuits) and trapped-ion systems offer

powerful gate-based control of qubits and oscillators and novel tomographic techniques. We recommend closely reading the discussion in Sec. IV to understand how the specific features of the programming model we consider work. We further recommend paying close attention to Sec. V C, which discusses compiling for bosonic QEC schemes that may be of interest. The following Sec. VI will also be of great use, but we recommend paying particular attention to the methods of Sec. VI D, which give systematic compilation techniques for gates including Baker-Campbell-Hausdorff (BCH) and Trotter-based approaches and quantum signal processing based strategies. We then recommend reading the discussion in Sec. VII C, which covers the application to Hamiltonian simulation and is likely to be of great interest to the reader.

3. Computer science background

For computer scientists familiar with qubits, but unfamiliar with quantum oscillators, we recommend paying close attention to Sec. III to understand the basic formalism of bosonic quantum computation as well as methods to encode qubits in such systems. As a basic motivation for the gate complexity advantages of oscillators over qubits, we recommend reading Sec. VIII, which shows the overheads encountered when simulating oscillators with qubits. The discussion in Sec. IV is also vital to understanding the gates used in our instruction sets, and special attention should be paid to the boxes given within that describe the fundamental operations included within our various ISAs. The abstract machine models and ISAs discussed in Sec. IV E will be of great interest to computer scientists approaching this from an architecture perspective, although the discussion in Sec. IV F can be skimmed over. The discussion in Secs. V C and VI is very important for the computer science reader, as it shows strategies of interest for those with algorithms and error-correction backgrounds. The discussion of the algorithms in Sec. VII provides several novel state-transfer protocols, and a novel approach for performing the quantum Fourier transform (QFT) using properties of quantum harmonic oscillators.

4. The expert reader

For readers who are familiar with the concepts of bosonic quantum computation as well as error correction on these platforms, much of the early material can be used as a reference. For the benefit of these readers, we summarize in the following our key results, which are novel either because they are original or because they have recently appeared in narrow venues and are synthesized in this paper into a broader context for a wider audience.

The discussion on control flow and benchmarking in Sec. IV F contains several important results, including an elaboration and formal error-bound analysis on Wigner process tomography on hybrid oscillator-qubit devices.

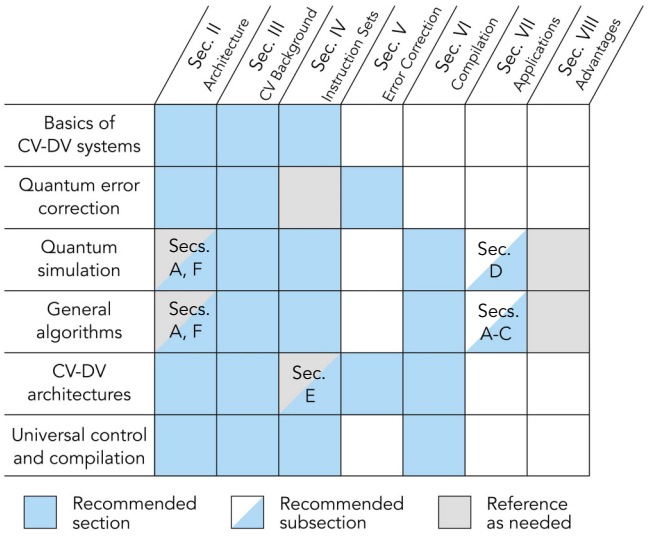


FIG. 1. A reader's guide for this work based on particular subtopics. We note that the indicated recommended sections and subsections are nonexhaustive, but are intended to give a rough outline of possible trajectories in reading this work. The introductory (Sec. I) and concluding sections (Sec. IX) are recommended for all readers.

The highly developed experimental state of the art for hybrid CV-DV Wigner function and characteristic function tomography in ion-trap and circuit QED (cQED) systems is not widely known outside these communities. Appendix E is devoted to general tomographic methods.

Section VI contains a host of new and recently reported results. Specifically, the discussions on the synthesis of unitaries in Secs. VI C and VI D provide several useful ways to perform such syntheses that may be of great interest to an expert reader. In particular, the extension of quantum signal processing (QSP) approaches to hybrid CV-DV systems presented therein is new and demonstrates a connection between block-encodings and qubit-controlled oscillator gates. In traditional qubit-only QSP [49], one deals with sequences of qubit rotations through specified angles to achieve the desired transformation of a block-encoded operator. We have extended this to the case in which the rotation angles are noncommuting quantum operators associated with the position and momentum of an oscillator. This section should be reviewed carefully even if methods such as group-commutator-based approaches to unitary synthesis are already well known to the reader. Specific pedagogical and graphically illustrated examples of the application of hybrid CV-DV QSP are also presented, which we hope will be beneficial to theorists and experimentalists alike.

Section VI C presents new results for exact synthesis (i.e., without Trotter error) for a variety of long-distance qubit-qubit and oscillator-oscillator entangling gates for our hybrid hardware layout. The application of these ideas

to quantum simulation, state transfer, and the QFT given in Sec. VII is also of great interest. Section VII C, for example, outlines new compilation techniques that take advantage of bosonic modes natively available in the hardware for efficient simulation of strongly correlated many-body systems and lattice gauge theories containing bosons. For describing oscillators with qubit-only hardware, Sec. VIII contains novel, scalable compilation techniques and a costing of the gate count. Most of the remaining sections not mentioned above can be skimmed or used as a reference for an expert reader, but the aforementioned results are novel and may be of great interest.

In addition to formal circuit synthesis techniques, we place into context tools that have somewhat narrowly arisen in computer science and quantum computer science and should be of broader interest to quantum information physicists. Examples include STRAWBERRY FIELDS [50], a compiler for photonic circuits, and BOSEHEDRAL [51], a compiler for beam-splitter networks relevant to boson sampling circuits. Another is BOSONIC QISKIT [52–55], an extension of QISKIT that provides a quantum intermediate representation for hybrid CV-DV systems with standardized naming and classification schemes for the bosonic gate sets used in the ISAs we present in this work. Finally, we mention GENESIS [56] and the SYMBOLIC HAMILTONIAN COMPILER [57], which are compilers designed to support Hamiltonian simulation on hybrid CV-DV quantum computing systems.

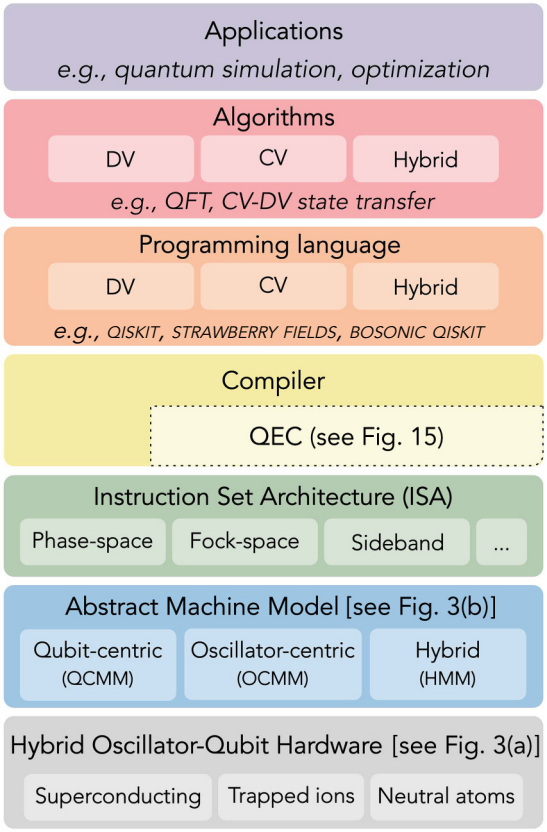


FIG. 2. The seven layers of hybrid oscillator-qubit quantum computer architecture. Each horizontal layer represents an abstraction of the one below it, while inset boxes provide options and examples for each layer. From top to bottom, a user first decides on an application to which they can employ discrete-variable (DV), continuous-variable (CV), or hybrid CV-DV algorithms. A program is written using an appropriate programming model for the user-level abstract machine model (AMM) (see Sec. II B 4). The compiler takes high-level program instructions and returns a sequence of logical gates expressed in a low-level language, taking into account the chosen quantum error correction (QEC) options (see Fig. 15). The logical gates are then transpiled into a physical gate sequence that is interpreted by the instruction set architecture. Depending on the nature of the instructions, they are implemented using a qubit-centric, oscillator-centric, or hybrid machine model. Finally, the instruction sequence is translated into a pulse sequence executed on the physical hardware (see Fig. 3).

II. HYBRID CV-DV QUANTUM COMPUTING ARCHITECTURE

This section presents an overview of hybrid CV-DV architectures, beginning with a perspective on the physical model of hybrid CV-DV quantum computation in Sec. II A, followed by principled definitions of three main abstract machine models in Sec. II B, an introduction to instruction set architectures in Sec. II C, the concept of compilation and error correction in Secs. II D and II E, and potential applications and opportunities to be expected for hybrid CV-DV quantum computation in Sec. II F.

A. Physical model of hybrid CV-DV quantum computation

Modern computing systems are assembled from multiple abstractions, and hybrid CV-DV quantum processors naturally share this approach, as illustrated in Fig. 2. This abstraction stack draws from pure DV quantum processors [45], but also possesses important differences. Although other concepts for hybrid CV-DV system architectures are known [58], and while current hybrid CV-DV quantum systems do not yet have mature realizations of all the layers depicted, a natural starting point is how such machines are physically realized.

Tremendous theoretical and experimental progress has been made in the past two decades on quantum control [59–62] and QEC [63–69] based on small versions of a hybrid device architecture in which quantum information can be stored and manipulated in both superconducting qubits and photonic states of microwave oscillators (bosonic modes) within the "circuit QED" paradigm, notionally illustrated in Fig. 3(a) and outlined in Appendix A 1. Circuit QED is the microwave analog of

optical cavity quantum electrodynamics with natural atoms replaced by synthetic atoms (superconducting qubits) the macroscopic size of which permits extremely strong coupling between the atom and the microwave optical field in a two-dimensional (2D) or three-dimensional (3D) superconducting resonator, yielding cooperativities many orders of magnitude larger than are possible in traditional cavity QED [28,29,58,77,78]. This strong coupling permits high-fidelity quantum nondemolition (QND) measurement of the qubit state using the dispersive coupling to a readout resonator. The same coupling permits QND measurement of the photon number (and binary functions of the photon number such as the parity) in a storage resonator. As we will discuss in detail, the strong coupling also opens up remarkable opportunities for quantum control and the creation of exotic hybrid quantum states of oscillators and qubits. For example, it is relatively easy to displace the oscillator conditioned on the state of the qubit (as is done in trapped-ion systems) and to rotate the qubit on its Bloch sphere conditioned on the photon number in the resonator (which is not readily achievable in trapped ions). Such gates will play important roles in our discussion. Another feature of cQED, shared to an extent with cold atoms in optical lattices, is the ability to engineer parametric processes which give experimenters control of the phase and amplitude of the hopping of quanta between lattice sites; physically, such quanta may be photons in cQED and atoms in optical lattices. This turns out to be important for simulating physical models in which neutral bosons are coupled to either static gauge fields as if they were charged particles in a magnetic field, or to dynamical gauge fields in simulations of lattice gauge theories.

There have likewise been rapid advances in trapped-ion systems [16,20,72,73,75,76,79–97] that combine DV ion qubits with CV motional degrees of freedom [notionally illustrated in Fig. 3(b) and outlined in Appendix A 2]. These advances include the realization of scalable 2D ion arrays [90,93–95], Penning microtraps [91,92], the integration of optical tweezers and dynamic potentials [96,97], and the implementation of strong higher-order nonlinear interactions, including trisqueezing and quadsqueezing [72,84], opening new opportunities for non-Gaussian resource generation. These technological advancements have enabled cutting-edge experimental demonstrations including the analog simulation of chemistry [86,89] and other fermionic models [18], the observation of finite-energy transitions [98] and emergent hydrodynamics [99], and the realization of non-Abelian topological order [20]. Trapped-ion systems tend to be used in a qubit-centric paradigm with the bosonic motional modes used to realize all-to-all connectivity [QCMM in Fig. 3(b)]. However, there are notable exceptions focusing on the creation and measurement of novel bosonic (phonon) operations and states [84,100–105], the preparation of GKP code words for error correction [106–109], bosonic simulations of quantum chemistry [86,110,111], and the use of bosonic modes for machine-learning tasks [112].

Neutral atoms in optical tweezers [illustrated in Fig. 3(a) and outlined in Appendix A 3] have traditionally been qubit-centric in their approach, viewing the mechanical oscillations of the atoms within the tweezer trap as a source of decoherence rather than as a resource. This is beginning to change as experiments have started to explore hybrid qubit-boson operations [7,113,114]. In particular, Ref. [115] has recently realized mode lifetimes exceeding those of the qubits, and further demonstrated hyperentanglement of two tweezers in both the qubit and motional degrees of freedom. In addition, bosonic neutral atoms in conjunction with precision placement of atoms using optical tweezers have been used to realize the bosonic Hong-Ou-Mandel effect [116] and, more recently, boson sampling [117]. Theory proposals have explored realizing bosonic quantum error-correcting codes [118] and exploiting the intrinsic nonlinearity of the potential [119].

At a high level, this tutorial envisions a system composed of oscillators (see Sec. III) and qubits coupled via the hybrid CV-DV gates introduced in Sec. IV, while abstracting away physical implementation details wherever possible. We assume that only the qubit subsystems are directly measurable, and that beam splitters enable interactions between oscillators. Throughout, we primarily focus on architectures where beam-splitter operations serve as the main mechanism for building connectivity. However, we note that in some platforms such as neutral-atom arrays, direct two-qubit gates (e.g., Rydberg interactions) may be more natural for connecting nearby qubits [see Fig. 3(a)]. All coupling operations are constrained by a specified interaction graph, and algorithms must be designed to respect this hardware connectivity.

To ground the discussion, we will often specialize to the simple yet experimentally realistic hardware layout shown in Fig. 3(a) and featuring near-neighbor connectivity. In this setup, each qubit is tightly integrated with a nearby oscillator, enabling fast hybrid operations. The qubits are conceptually arranged in a 2D array, with one oscillator paired to each qubit. While this layout is fixed in circuit-QED architectures, ion-trap and optical-tweezer platforms allow for dynamic reconfiguration, as qubits can be physically moved to reconfigure the geometry on the fly.

B. Abstract machine models (AMMs)

An abstract machine model (AMM) plays an important role between the hardware architecture and an instruction set architecture. For classical computers, AMMs provide valuable theoretical models that bridge between computational complexity and practical hardware architecture design, allowing exploration of design decisions and analyses of their impact before physical implementation. For example, several AMMs have played important roles

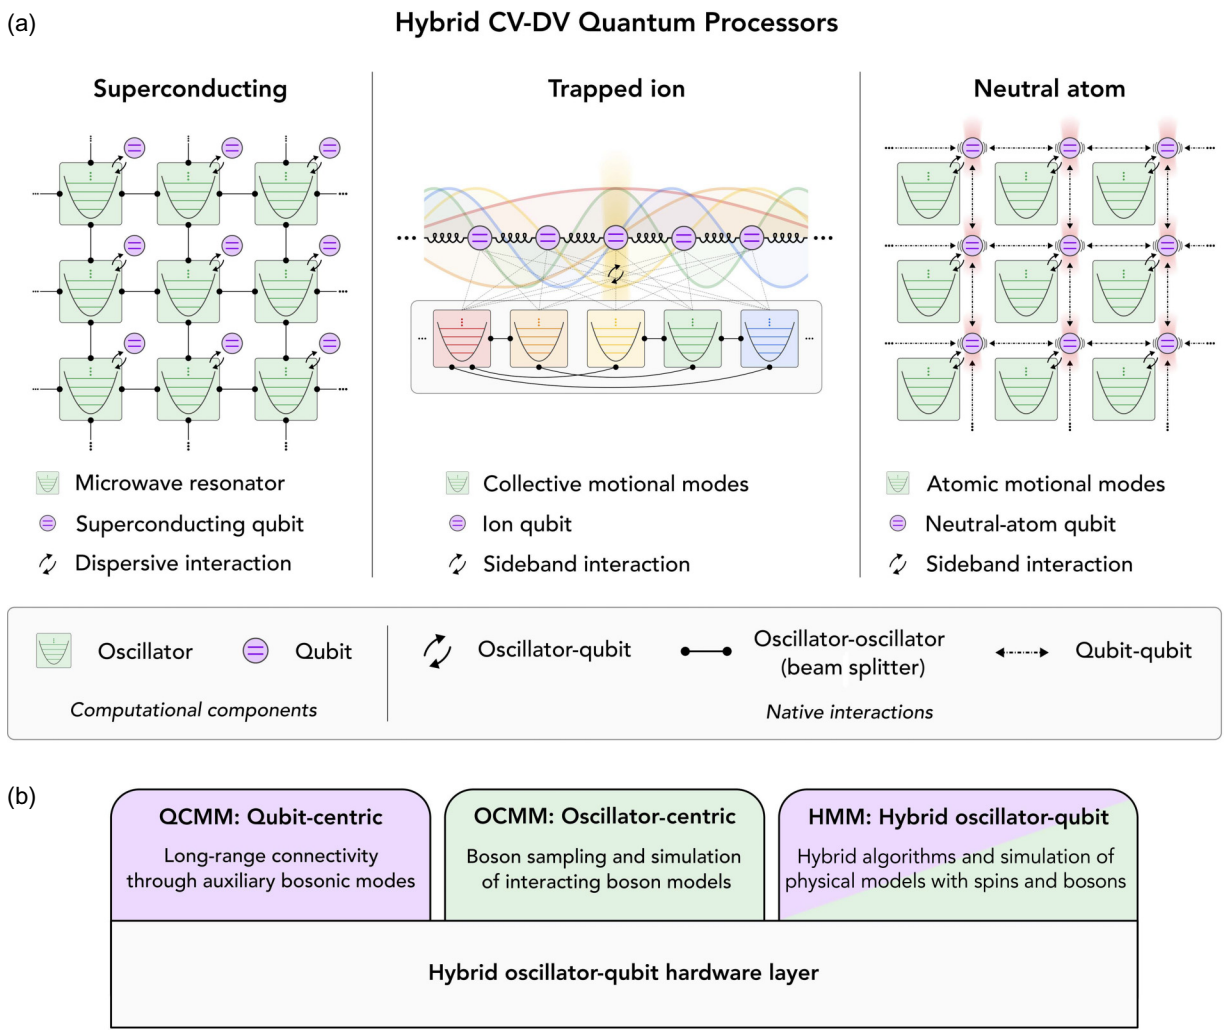


FIG. 3. A schematic illustration of hybrid CV-DV hardware and its abstraction. (a) *Left panel:* A hybrid CV-DV quantum processor composed of superconducting microwave resonators dispersively coupled to individual superconducting qubits. Adjacent microwave resonators are coupled via microwave-controlled beam splitters that permit fast, high-fidelity SWAP gates (∼100 ns, 99.92% fidelity in the single-photon subspace [70]) that enable remote communication and entanglement of both oscillators and qubits. As each qubit is coupled to only a single bosonic mode, control and crosstalk issues are minimized. (Reference [71] considers a similar architecture but allows for native two-qubit gates between neighboring qubits.) This superconducting architecture will be our primary focus throughout this work, though much of our discussion is extendable to hybrid CV-DV trapped-ion and neutral-atom platforms. *Center panel:* In contrast, an ion-trap implementation of a hybrid CV-DV processor has ionic-spin qubits each coupled to all $3N$ collective motional modes of the ion array (as indicated by the dashed gray lines). Native beam splitters between selected motional modes can be realized through, e.g., laser-based methods [72,73], laser-free methods such as modulation of trap electrode potentials [72,74,75], or via mutual coupling to an ionic qubit [76]. *Right panel:* A hybrid CV-DV quantum processor using neutral atoms in an optical-tweezer array. Here, DV components correspond to atomic qubits, while CV degrees of freedom correspond to the motional modes of individual atoms. Similar to the superconducting architecture, each qubit is coupled to a single motional mode or oscillator, minimizing crosstalk. In place of beam splitters, however, connectivity between oscillator-qubit "units" is realized using qubit-qubit Rydberg interactions. (b) As discussed in the text, the flexibility of the hybrid oscillator-qubit hardware layer enables three distinct abstract machine models visible to the next layer of the stack: the qubit-centric machine model (QCMM) presents a standard qubit architecture to the compiler layer, abstracting away the bosonic modes which play a secondary but powerful role as the communication fabric for remote entanglement and gates; the oscillator-centric machine model (OCMM) presents the bosonic modes to the compiler layer with the quantum information encoded in the oscillators (possibly via bosonic error-correction codes), abstracting away the qubits which play a secondary role as nonlinear elements needed for universal control of the bosonic modes; finally, the hybrid machine model (HMM) exposes both CV and DV degrees of freedom on equal footing to the transpiler layer, enabling applications such as quantum simulation of physical models containing both spins (and/or fermions) and oscillators. This taxonomy and an alternative taxonomy are discussed in Sec. II B.

in classical computer architectures, providing a mechanism to understand and optimize the balance between computation, communication, and storage resources allocated in hardware to best suit algorithmic and application needs. These seminal AMMs include [43] the random access machine (RAM), which models the cost of accessing memories with different latencies and throughputs; the message-passing interface (MPI), which models the trade-off between communication and local computation when no global shared memory is provided; and the JAVA virtual machine (JVM), which abstracts away the underlying processor and operating system, providing a model for "write once, run anywhere" software development. Other important early AMMs [120] include finite state machines, Turing machines, and even cellular automata. All these models have provided invaluable mechanisms for analytical reasoning about computational complexity given constraints on memory, concurrency, connectivity, and other fundamental resources such as energy, space, and time.

Compared with the sophistication of AMMs for classical computation, AMMs for quantum computation are still in their infancy, but they certainly aspire to satisfy many of the same goals. Quantum AMMs seek to allow analytical study of computational complexity given constrained resources, just as classical AMMs do. But these must extend to include quantum specific resources, such as entanglement and coherence, as is apparent from considering uniquely quantum aspects of communication, storage, and computation:

(a) *Communication* can be performed using qubits instead of just with bits. In certain hardware schemes such as trapped-ion systems, physically moving qubits is very slow [121,122] compared with gate speeds. But if entangled quantum states are available, then they can be employed to teleport quantum states over long distances at the speed of classical communication; quantum teleportation can thus serve as a kind of "quantum wire" [123], with its own latency and throughput trade-offs.

(b) *Storage* of quantum states can also be much more challenging than storing bits. Good quantum memories are not yet available and may require costly encoding in error-correction codes and constant refreshing [124]. Quantum memories also come with many different access models, including quantum random-access memories (QRAMs) with parallel or serial access [124–133].

(c) *Computation* is often expressed in the language of gates and circuits, and while quantum computation may use this language much like classical computation, alternative and distinctly nonclassical models are also available. This notably includes *measurement-based* quantum computation (MBQC), in which computation proceeds through a sequence of measurements alone, with no gates being applied [134]. Remarkably, quantum states can embody the ability to perform gates of complexity beyond what is possible with native hardware [135]; this is the idea behind the magic states used with gate teleportation to enable universal fault-tolerant (FT) quantum computation with hardware that performs only Clifford gates [136].

Quantum AMMs have arisen to capture many of these uniquely quantum aspects, but these models largely build on qubits (DV quantum systems) as the primitives and eschew a distinct role for bosonic (CV) quantum hardware. Naively, a quantum computation consists of only three steps: state preparation of the input quantum registers, application of a unitary transformation representing the algorithm, and readout (measurement) of the output registers. This native model is the dominant one for nearly all modern quantum programming models such as QISKIT [137]. However, with the rise of bosonic quantum hardware, programming models are beginning to emerge for such systems, including BOSONIC QISKIT [52] and STRAWBERRY FIELDS [50].

Most interesting is not just an AMM that captures classical computation, or CV or DV quantum computation, but instead, all of these combined. A compelling example of the potential for combining qubit-based quantum computation with classical computation is the "local alternating quantum-classical computation" (LAQCC) model [138–140] of tightly coupled quantum and classical computation and communication with frequent measurements and feedforward. In other words, unitary quantum computation is followed by further unitary operations conditioned on classically processed measurement results from the first stage, and so on. This is a substantially more powerful computational model than quantum alone, or classical alone, and strongly indicates the desirability of including classical computation aspects in quantum AMMs. We note that the communication via quantum "wires," the gate teleportation, and MBQC mentioned above all require the mixed quantum-classical operations and feedforward of the LAQCC model.

In discussing quantum AMMs, a key high-level question arises: What is the appropriate level of abstraction for a quantum AMM, given the infancy of quantum hardware? Specifically, two contrasting viewpoints arise naturally, as suggested in Fig. 2. One is the viewpoint from the perspective of a user who wishes to compile efficient applications utilizing a given computational model. A *user-visible quantum AMM* (and its associated high-level instruction sets) for CV-DV hybrid computation should abstract away all details of the hardware and its connectivity and simply present a set of abstract qubits and abstract oscillator objects. Such an AMM should hide details of the fact that

the abstract qubit might be a physical qubit or a logical qubit comprising many physical qubits or a logical qubit encoded in a physical oscillator.

In contrast, consider the viewpoint from the perspective of the transpiler, which must deal with the task of mapping applications into primitive operations on physical systems. A *transpiler-visible quantum AMM* (and its associated low-level ISAs) is an abstraction of the physical layer (at the level shown schematically in Fig. 3). Transpiler-visible AMMs can be very helpful for resource estimation. The ISAs that we primarily focus on in this paper sit between the transpiler and the physical hardware layer and allow the transpiler to convert high-level instructions into low-level instructions that are consistent with the connectivity and error properties of the physical hardware. Given this, and also since resource estimations are a goal of the present work, we largely focus here on transpiler-visible quantum AMMs, but we do provide some perspective on user-visible quantum AMMs as well.

Specifically, we elaborate below on three types of transpiler-visible quantum AMMs, conceptualizing qubit-centric (QCMM), oscillator-centric (OCMM), and hybrid quantum CV-DV (HMM) models.

These three make distinct use of the relative strengths of qubits and oscillators to serve different architectural purposes, in particular for quantum computation, internal communication, and storage. We follow this with brief remarks on the higher level of abstraction in user-visible AMMs.

1. Qubit-centric abstract machine model (QCMM)

The qubit-centric model, QCMM in Table I, utilizes qubits for computation, but takes advantage of bosonic subsystems for key tasks they may perform better than qubits.

In particular, bosonic superconducting resonators are known to offer long coherence times which exceed transmon-qubit lifetimes by an order of magnitude or more. As will be discussed in Sec. V, bosonic quantum error correction can be employed to lengthen coherence times of states encoded into a resonator, with a hardware cost overhead that is significantly lower than that required for qubits. This is because the codes can focus on correcting photon loss, the single dominant error mechanism faced by resonators; in contrast, qubits generically suffer from more error mechanisms, including both bit-flip and phase-flip errors, and in addition, multiple qubits are required to achieve the larger Hilbert-space

TABLE I. Three types of transpiler-visible AMMs, with examples for each, specifying the roles of the physical qubits and oscillators, and the nature of instruction sets employed. The QCMM gives an abstraction for settings where the bosonic modes and qubits are both used to harness physical qubits. The OCMM provides a model for bosonic experiments where the qubit is used to facilitate state preparation, Hamiltonian evolution, and measurement in the bosonic modes. The HMM describes models which seek to utilize both continuous-variable and discrete operations for computation. Applications utilizing the HMM have been under explored to date, but we argue in Sec. VII C that there may be circumstances where the devices modeled by the HMM permit simulations that would otherwise be out of reach with present-day technology. Complementing these three transpiler-visible AMMs, a different but related set of user-visible AMMs could also be formulated; such AMMs may be appropriate for describing logical qubits that are built on top of the qubits and qumodes exposed by the above models, as described in the text.

AMM type	Subtype	Physical qubit role	Physical oscillator role	Physical qubit instructions	Physical bosonic instructions
QCMM (qubit-centric)	Qubits + Bosonic memory	Computation	Memory (error corrected)	Universal	Single-oscillator error correction
	Bosonic bus	Computation	Communication	Universal	Entanglement swapping
	Bosonic sensing	Computation and I/O	Sensing	Universal	Displacement, Jaynes-Cummings
OCMM (oscillator-centric)	Boson sampling	I/O	Computation	Boson state preparation and measurement	Displacement, multimode squeezing, multimode beam splitter
	Bosonic system simulation	I/O	Simulation	Boson state preparation, evolution, and measurement	Universal
HMM (hybrid CV-DV)	Hybrid	Computation simulation of spins or fermions	Computation simulation	Universal	Universal

dimensions needed for redundant encoding. Hardware-efficient bosonic memories may thus be employed in aid of qubit-based quantum computation. In this "qubits + bosonic memory" model, a universal set of gates is chosen for qubits but a limited set of gates is chosen for the oscillators, enabling error correction for their use as quantum memories. Qubit states can be written into and read out from the bosonic memories, but direct logic gates on the quantum memories are not allowed. Different QEC codes could be employed for different resonators, enabling realization of a hierarchy of quantum memories [141]; e.g., a short code that is easy to decode could provide a short-term cache memory, while a large code (with longer read-write times for the memory) could serve as a longer-term static memory.

Bosonic systems also offer the potential for long-range interconnection between qubit subsystems. Electrical transmission lines, coplanar waveguides, optical cavities, and optical fibers are all bosonic systems. When patches of qubits are coupled into such bosonic systems, these resonators and transmission lines can provide a kind of "bus" interconnection between the patches [142]. Such bosonic interconnects have been envisioned as being the glue enabling modular quantum computers [143], which may be many quantum processing units tied together within a single chip, or superconducting processors connected between different dilution refrigerators, or trapped-ion processors linked across a network spanning a laboratory or even a metropolitan area. As we discuss in Sec. VI C, in this "bosonic bus" model, the gates needed for qubit entanglement swapping used to realize quantum communication include qubit-oscillator entanglement (as in trapped-ion systems) and oscillator mode SWAPs (for remote communication). Other qubit operations can also be implemented using the bosonic modes to implement gates on the qubits. This reinforces the fact that in such models, the qumodes need not be solely used for memory or encoding qubits within such AMMs.

Another unique feature of bosonic systems is their utility for sensing physical phenomena. Magnetic and electric fields are all bosonic, and if quantum processing is to be applied to information arising from such fields, then the qubits must couple to them. Such a coupling between qubits and bosonic systems is precisely the native physical Jaynes-Cummings interaction presented in Appendix A and underlies the oscillator and hybrid qubit-oscillator gates defined in Sec. IV. The oscillator displacement gate embodies a prototypical signal to be sensed. In the "bosonic sensing" model, a bosonic subsystem is explicitly introduced for transducing physical fields into qubit systems used for further quantum processing. This model is exemplified by the quantum logic spectroscopy performed using, e.g., trapped ions [144] and digital homodyne or heterodyne [145] and digital photon-number detection [146–148] of microwave (and axion dark matter [149]) signals in resonators. As another example, by using the ISAs we present, it is straightforward to construct large two- and four-legged Schrödinger cat states that are sensor states capable of detecting displacements of an oscillator by distances much less than the zero-point uncertainty in the position through the rapid oscillations of the photon-number parity as a function of the displacement [26,150]. Such small displacements can also be detected by using entanglement-enhanced interferometry [151–153] based on two-mode squeezing operations.

2. Oscillator-centric abstract machine model (OCMM)

The oscillator-centric model, OCMM in Table I, utilizes bosonic resonators for computation, and employs qubits only sparsely in support of the computation. This is natural to do when the computation closely matches the dynamics of coupled resonators. For example, this is the case for the well-known boson sampling problem [154], which samples the output of a large system of coupled beam splitters and phase shifters, when fed (ideally) with single-photon inputs. In principle, sampling of the output distribution is very difficult to compute classically [154], but in realistic (optical) implementations, the input is a squeezed state of light instead of a single photon, and imperfect squeezing and photon losses [147,155–160] can simplify classical simulation of the system [161–163]. Nevertheless, it is known that coupled oscillator systems can be very hard to classically simulate [164], so their computational potential is of great interest. In the microwave domain, auxiliary qubits allow high-fidelity preparation of the requisite input single-photon states, and with processing of the output states for sampling. In addition, auxiliary qubits can be used to induce interactions among the oscillators for quantum simulations of, e.g., the Bose-Hubbard model (see Sec. VII C 5). We note that the BOSEHEDRAL compiler [51] has recently been developed to simplify the task of efficiently compiling beam-splitter networks for the Gaussian boson sampling problem that has applications in solving graph-theoretic problems.

3. Hybrid abstract machine model (HMM)

The hybrid quantum CV-DV model, HMM in Table I, employs qubits as well as bosonic resonators for computation. A full complement of qubit gates and bosonic gates is also utilized, to enable arbitrary transforms to be performed on both systems. This hybrid model is the main focus of this paper. After a brief discussion of user-visible AMMs in the next subsection, we continue in the rest of this section by describing several instruction set architectures natural to the HMM.

4. User-visible AMMs as an alternative taxonomy

The above discussion presents one possible taxonomy of transpiler-visible AMMs which explicitly makes reference to the relative roles of the physical components (qubits and bosonic modes). An alternative AMM taxonomy at a slightly higher user-visible level of abstraction could be based solely on the (effective) quantum degrees of freedom (defined by their Hilbert spaces) exposed to the software stack. Thus, for example, the QCMM would only provide access to two-level systems without distinguishing between physical qubits (i.e., transmons) or logical qubits comprising either multiple physical qubits or logical bosonic qubit encodings in oscillators. In this model, the only operations visible to the upper levels of the stack might be logical qubit rotations and two-qubit gates (independent of the physical hardware involved). An example of how bosonic systems can model DV quantum computation is given by PERCEVAL, a quantum programming framework for DV quantum computation with photonic systems [165]. The QCMM might additionally expose error models and connectivity for these objects, which would, of course, depend on their underlying physical hardware. The oscillator-centric machine model, or OCMM, exposes all of the natively available Gaussian operations as well as the non-Gaussian oscillator gates discussed in Sec. IV. The non-Gaussian oscillator gates require the use of auxiliary physical qubits, but the OCMM would abstract away the details of the qubit portion of the associated hybrid gates. Finally, the CV-DV HMM would expose to the user both the CV and DV degrees of freedom along with one of the hybrid ISAs presented in Sec. IV, but might allow the user to, e.g., assume all-to-all connectivity or abstract away the physical qubits comprising logical qubits if error correction is invoked.

C. Instruction set architectures

Next above the physical level of the architecture stack of Fig. 2 is the instruction set architecture (ISA). Across the entire hybrid quantum computing stack, the ISA plays a pivotal role as the hardware-software interface that translates quantum algorithms or circuits into a well-defined discrete (but continuously parametrized) instruction set. An ISA builds upon a choice of AMM and specifies a universal instruction set. There may be a high-level ISA suitable for programming, as well as a low-level ISA suitable as the target for compilation and hardware implementation. While the mathematically detailed construction of a hybrid CV-DV ISA is presented in Sec. IV, here we begin with a broader perspective about the intent of such ISAs, the design choices involved, and differences compared with classical ISAs.

1. Goals for a good ISA

Traditionally, the ISA plays a vital role in classical computer design by defining how software controls the hardware. The ISA can be considered to provide two viewpoints: one to a programmer interacting with the hardware at a low level, and in this sense the ISA serves as a programmer's manual to the hardware. The ISA also provides a crucial viewpoint to compilers, which serve to translate high-level languages to hardware operations via the ISA. An ISA provides such viewpoints by specifying what instructions can be performed by an information processing unit, what data types are provided, the memory and memory hierarchy provided for information storage, and how the processor accesses memory and with what costs.

A well-designed ISA offers many important features. For example, a portable ISA may allow software to be written once and compiled to work on many different hardware platforms implementing the same ISA. Good ISA designs seek to also be complete, in terms of providing a sufficient set of instructions that capture capabilities provided by hardware implementations. Instructions should ideally also be efficient, enabling fast execution of desirable operations while also capturing some sense of regularity and consistency such that high-level languages can be optimally compiled into ISA operations. At the same time, a well-designed ISA should remain sufficiently simple that it is reasonably easy to implement with hardware, and understandable to programmers who must implement compilers.

Designing an ISA thus entails a multi-faceted balancing act, trading off between simplicity versus complexity, and optimality versus universality. We seek to design a thoughtful ISA for hybrid CV-DV quantum computers, and ideally this should inherit lessons learned from ISAs designed for classical computers. Two well-known examples are the reduced instruction set computer (RISC) and the complex instruction set (CISC) ISAs, exemplified by the Arm and Intel processors, respectively. These two very different examples provide a crude map for locating quantum ISAs in perspective. Because quantum hardware is far less mature, it is evident that the ISAs we present here share similarities with the RISC model. At the same time, the quantum ISAs lack features that are common in almost all rudimentary classical computer ISAs.

2. ISA design assumptions

The quantum ISAs presented here assume an architecture that tightly integrates qubits and oscillators, where each qubit is located near another oscillator [as notionally illustrated in Fig. 3(a)] with fast hybrid instructions available between the qubit and oscillator. Qubits may be directly linked to each other, but we assume here that they are not. In the pedagogically oriented model employed in this tutorial, each oscillator is linked to its paired qubit and connected to neighboring oscillators via microwave-activated couplers. Moreover, depending on whether QEC is employed or not, the architecture model may capture

either noisy intermediate-scale quantum (NISQ) (no QEC) or fault-tolerant (with QEC) applications.

Building on the physical model for hybrid CV-DV quantum processors, our execution model assumes that the qubits and the qumodes are tightly coupled, meaning that the operations are synchronized. This means that we do not need to consider message passing to execute quantum code in this setting. We assume that (most) gates may be executed in parallel on different qumodes and qubits. We further assume that measurement results can be used to conditionally control subsequent quantum operations. This is needed for quantum error correction, linear optical quantum computing, as well as measurement-based quantum computing, and more broadly for the general LAQCC model mentioned earlier.

Finally, the programming model that we take assumes that the sequence of raw gates is represented using a gate description language that describes the precise sequences of gates applied to the qumodes and qubits in the system. This language is interpreted by a compiler and translated into physical gates that the quantum computer then executes. There is no assumption that the code provided is stored within the quantum computer; rather, we assume that it is held by the external classical computer.

Note that, at present, memory architecture for quantum computers is often not discussed (though see Refs. [141, 166,167] on this topic), nor are specialized processing elements such as arithmetic logic units (ALUs).

As a result, the role of data structures is thus far limited in quantum computation, particularly when compared to classical computation where the power of these ideas often stems from the assumption that cheap memory access is available. The closest quantum analog is perhaps the QRAM structure, which provides a low-depth circuit implementation for a lookup table in a quantum computer [124–126,128–131,168]. However, even here the assumption of constant cost unit memory access may be questionable. For these reasons, we do not discuss the implementation of such computational elements in hybrid oscillator-qubit hardware; however, bosonic operations may aid in the replacement of some costly arithmetic functions and we thus note the possibility that bosonic hardware may play a role in future quantum architectures with ALUs and potentially even QRAM structures [132].

The choice of instruction set can be made using two different approaches, as is elaborated further upon in Sec. IV. In the top-down approach, the algorithms or user-visible AMMs choose a convenient option from the various instruction sets. For example, a general application needs to be compiled into the native ISA on the hybrid processor. In the bottom-up approach, we choose a hardware-convenient instruction set and then decide the algorithms which can be compiled efficiently with this instruction set. Information about the physical layers of the oscillator-qubit device enables the estimation of resource requirements for execution within the ISA (possibly with QEC) for a particular application or algorithm of interest. Section IV thus prepares the way for compilation methods for logical operations of an error-corrected encoding in an oscillator (see Sec. V), or quantum circuits used in various CV-DV applications (see Secs. VI and VII).

3. Hybrid CV-DV ISA versus qubit and classical ISAs

Hybrid oscillator-qubit mode hardware systems are not of wide familiarity to the quantum computer science community, but such systems are rapidly becoming experimentally accessible and deserve increased attention given their potential advantages.

As these advantages are now starting to be realized experimentally, the time is ripe to develop and analyze formal AMMs and ISAs for existing and future hardware, but there are challenges to be overcome. We are in the very earliest stages of learning how to program such hardware, and ISAs are needed to give developers (from both computer science and quantum physics backgrounds) an understanding of what is possible with current and future hardware for creating fault-tolerant circuits, modules, and processors, and for compiling efficient algorithms. Also, while the present-day development of digital-analog hybrid computing in classical computer architectures [169,170] shares some common themes with hybrid CV-DV quantum systems, such classical hybrids are yet fundamentally different from the quantum case.

One feature of quantum ISAs that is important to understand is that even for hardware based solely on DV objects (i.e., qubits), quantum machines share some features with analog computers (but unlike analog computers, do permit error correction) [171]. For example, the family of superposition states of a single qubit is continuous, much like the continuous states of an analog computer. Any single-qubit superposition state can be obtained from a fixed starting state ($|0\rangle$, say) using only a discrete set of two gates describing rotations about the x and y axes of the single-qubit Bloch sphere,

$$R_x(\theta) = e^{-i\frac{\theta}{2}X}, \tag{1}$$

$$R_y(\theta) = e^{-i\frac{\theta}{2}Y}, \tag{2}$$

using notation standard in the field [1], with X and Y being the usual Pauli matrices. We see however that, unlike the discrete gate sets used in classical computation, quantum gates are generally parametrized by one or more continuous variables (rotation angles). To develop and formally reason about circuit synthesis and fault tolerance, these continuous parameters may themselves be discretized (e.g., into Clifford operations plus one fixed non-Clifford gate [1]). Hybrid architectures are "even more analog" in the sense that bosonic modes are CV degrees of

freedom, as described above. Techniques for general circuit synthesis and fault-tolerance analysis in CV systems are only beginning to be developed, our ability to formally reason about the properties of such circuits [60,172–177] is still quite limited, and the complexity classification of bosonic states and operations is only beginning to be explored [178,179].

Two principal hybrid CV-DV quantum computation instruction set architectures which go beyond qubit-only and classical ISAs are the Fock-space ISA and the phase-space ISA, as introduced in Sec. IV. The Fock-space instruction set is based on primitive gates conditioned on the number of excitations held in the oscillator. This is in contrast with the phase-space instruction set, in which qubit operations are conditioned on the position or momentum of the oscillator. A discussion of cross-compilation among the various universal instruction sets is given in Appendix D.

D. Quantum error correction

Above the ISA layer of Fig. 2, as an optional part of the compilation process we have an optional step of QEC. Unlike modern classical computing systems, quantum computers remain highly prone to errors at their physical layers. For example, a typical CMOS logic gate may fail with probability perhaps much less than one in 10^{20} operations, whereas current two-qubit quantum gates currently fail one in 10^5 operations at best. Unfortunately, many quantum applications call for 10^{15} to 10^{20} nontrivial quantum operations. It is widely believed that such levels of high-fidelity operation may be achieved by encoding a "logical" qubit using multiple physical qubits using quantum error correction and by employing methods of fault-tolerant quantum computation (FTQC) [1]. Some overhead of physical gates and qubits is needed to achieve such FTQC, and in principle the required overhead is only polynomial in the scale of the ideal circuit being simulated. However, the constant prefactors and other scaling terms can be prohibitive in view of the physically realistic numbers of qubits and control circuits accessible today. Thus, the employment of such QEC is not always desired, so we envision the architecture providing a flag to determine whether or not QEC is used before the next layer. Section V is dedicated to the discussion of QEC, and provides a definition of a "logical" level of instructions which arises when QEC is included.

For hybrid processors, bosonic quantum error-correcting codes can be used to encode (DV) logical qubits into physical oscillators to realize DV-only logical machines using hybrid oscillator-qubit physical resources. These logical DV machines will have a particular layout and connectivity between logical qubits depending on specific physical layer hardware connectivity as well as the QEC codes used. They will also have particular (residual, uncorrected) errors that may vary with the particular gate or measurement being performed—abstract machine model information that can be passed up the stack to help with noise-aware compilation. For bosonic Hamiltonian simulation applications, oscillator-to-oscillator QEC encodings may be possible (Sec. V B 4) but would be challenging. Alternatively, we can omit the QEC stage to retain the full CV characteristics of the modes that are essential to bosonic simulations. In this case, physical layer oscillator-qubit instruction sets are directly used as "logical"-level instruction sets with connectivity and layout determined by the hardware implementation.

For hybrid qubit-boson Hamiltonian simulation, QEC can be employed in both the Fock-space ISA and the phase-space ISA in a semi-fault-tolerant way by using the higher levels of an auxiliary transmon-qubit mode while keeping the main cavity mode intact. For example, in Ref. [180], it has been shown that the $|g\rangle$-$|f\rangle$ (ground and second excited state) manifold of a transmon can be used to perform SNAP gates (see Box IV.10) fault tolerantly on a rotationally symmetric code. Similarly, controlled-displacement gates could be performed using the $|g\rangle$-$|f\rangle$ manifold, using the $|e\rangle$ state to flag decay errors. We discuss the details of the error-correction layer in Sec. V.

In the NISQ era, we skip QEC and directly transition from the compiler to the hardware stack. In this scenario, the compiler outputs gate sequences composed of physical gates that can be processed directly by the control pulse generator.

E. Programming languages and compilation

The next two integral layers of the architecture of Fig. 2 are the compilation and programming language layers. Compilation involves the mapping of programming languages into the instruction set. The programming language takes as input the algorithmic description from the layer above and outputs to the compiler layer a high-level description of the underlying quantum algorithms [181–185] in the associated software development kit (SDK) such as QISKIT [137] or BOSONIC QISKIT [52,53] with proper circuit language descriptions and syntax.

Any programming language specialized to express algorithms that we wish to deploy on bosonic hardware needs to have appropriate abstractions to reflect the architecture of the system that we wish to program at a level of abstraction appropriate for the problem. For any programming language that is well suited for our architecture, we need to have appropriate abstractions for registers of oscillators (or qumodes) while also retaining appropriate abstractions for registers and operations on ordinary qubits, as well as for hybrid CV-DV operations.

As an example, BOSONIC QISKIT [52,53] achieves this by building on top of QISKIT [137], which is a leading low-level programming language embedded in PYTHON

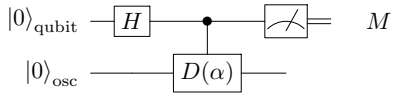
that can be used to program qubit-based quantum computers. Specifically, BOSONIC QISKIT inherits the great strength of QISKIT, which is the ease with which it can be used to provide a low-level description of a quantum circuit, and merges this with new functionality that allows it to describe bosonic modes and operations. BOSONIC QISKIT provides a new class for describing oscillators called a qumode register. This class provides information about the state of a qumode within a specified mode occupation number cutoff $2^n - 1$ for integer n. This choice is made to allow easy simulation within existing QISKIT simulators which work with qubit register Hilbert-space dimensions that are a power of 2. A collection of associated qumodes is brought together into a concept called a qumode register. Qumodes, qubits, and classical memory can all be wrapped together in a single class known as a `CVCircuit`. This class can be initialized with the following instructions for the case of a single qumode (qmr) with a cutoff of $2^6 - 1$, one qubit (qbr), and a classical memory register (cr) [52,53]:

```
% Initialize qumode register
qmr = c2qa.QumodeRegister(num_qumodes = 1,
    num_qubits_per_qumode = 6)
% Initialize qubit register
qbr = qiskit.QuantumRegister(1)
% Initialize classical register
cr = qiskit.ClassicalRegister(1)
% Initialize hybrid CV-DV quantum circuit
circuit = c2qa.CVCircuit(qmr, qbr, cr)
```

We can then apply a Hadamard gate to the qubit and a controlled displacement for $\alpha = 1$ to the qumode followed by a measurement of the qubit register that stores the result in the classical register via the following code:

```
% Assign gate parameters for CV gates
alpha = 1
% Append a DV gate
circuit.h(qbr[0])
% Append a hybrid CV-DV gate
circuit.cv_c_d(alpha, qmr[0], qbr[0])
% Measure CV and DV components
circuit.measure(qbr[0], cr[0])
```

This code generates the sequence of bosonic and qubit gates illustrated in the following quantum circuit diagram:



This provides an example of the syntax and circuit description for a simple hybrid algorithm expressed through BOSONIC QISKIT [52,53].

Section VI describes how the compiler takes as input the high-level language description and connectivity graph of the hardware (and potentially also the heterogeneous hardware noise model) and outputs a low-level language description of the program as logical gate sequences. This can be a challenging task. Translation from high- to low-level languages often proceeds hierarchically. More specifically, the quantum program description in a high-level language can first be translated into the so-called quantum intermediate representation (QIR), which is a representation of the original program independent of different source programming languages and is also agnostic to hardware instruction details. The QIR serves as a middle ground to connect high-level programming languages to instruction sets to enable the modular design of quantum compilers. The QIR can be further compiled to the elementary logical gate set, given the connectivity graph of the device.

For DV-only applications (including bosonically encoded logical DVs), a standard DV QIR may be used [186–190], while for Hamiltonian simulation of bosonic systems, an additional "bosonic QIR" can be helpful. On a CV mode, the most general unitary operation is a phase-space volume-preserving diffeomorphism generated by a control Hamiltonian that is a polynomial in the oscillator positions and momenta multiplied by Pauli operators (or the identity) (see Sec. VI B). Such polynomials naturally serve as a QIR for bosonic system gates. As an example, such a bosonic QIR might contain terms such as the conditional beam-splitter operation (see Table IV)

$$e^{i\theta Z(a^\dagger b + b^\dagger a)} = e^{i2\theta Z(\hat{x}_a\hat{x}_b + \hat{p}_a\hat{p}_b)}, \tag{3}$$

which has linear couplings between the position x and momentum p of two different oscillators a and b multiplied by a single Pauli operator Z. As noted in Sec. VII C 7, this gate is particularly useful for $\mathbb{Z}_2$ lattice gauge theory simulation (see Ref. [41] for more details). More generally, higher-order polynomials can be used as well.

The programming languages and compilation stack layers process the instruction set gates to enact circuits for the algorithms layer.

F. Algorithms and applications

At the top of the stack of Fig. 2 sit *applications* and *algorithms*.

Of course, hybrid CV-DV quantum computers should be able to realize general quantum algorithms specific to DV-only hardware, which is a large topic of its own right. We instead focus here on just the special quantum algorithms and applications particular to the hybrid architecture. Specifically, what advantages might be found for applications utilizing hybrid CV-DV hardware, instruction sets, and compilers? Section VII details how one likely advantage of hardware natively containing bosonic modes is that each such mode (formally) has a countably infinite Hilbert space dimension while having a relatively simple error model [191], because there are very few "moving parts" relative to a collection of qubits with a large composite Hilbert-space dimension. This hardware efficiency is one of the key reasons that the superconducting cavity architecture was the first one to achieve memory QEC operating above the break-even point [192] using bosonic codes [63–65,67–69,107,193].

A second potential advantage of hardware natively containing bosonic modes is in the construction of programmable quantum simulators for physical models containing bosons (e.g., the bosonic Hubbard model from condensed-matter physics and lattice gauge theories from particle physics). Qubit-only systems are naturally efficient for the simulation of quantum magnets (especially spin-1/2 magnets). Qubits can also simulate fermions using the Jordan-Wigner transformation to insert the appropriate fermion minus signs [35,194–198] but this is expensive since it involves very high-weight operators, especially in greater than one spatial dimension. It is less well known that there are also significant difficulties in efficiently representing bosons using qubits [199,200]. As we discuss, this is partially because each bosonic mode can contain many bosons.

However, if we invoke a cutoff $N_{\max}$ on the number of bosons per mode, we can represent a Hilbert space of the same dimension using only $\sim \log_2[N_{\max}+1]$ qubits (per mode), so this is not a severe problem if the maximum boson number is below this cutoff, both in the initial state and throughout the simulation. A more serious difficulty lies in the matrix elements of the field operators a and $a^\dagger$ which contain factors related to the square root of the boson number. These field operators are natively available in hardware containing bosonic modes but are quite difficult to synthesize in qubit representations of bosonic modes because a very large number of quantum gates are needed to accurately synthesize these square-root factors [201,202] (see Secs. VII C and VIII B). The practical benefits of the ability to efficiently represent bosons have recently been experimentally demonstrated by the construction of a highly hardware-efficient quantum simulation of the Franck-Condon vibrational spectra of small molecules using Gaussian boson sampling with optical modes [160,203,204] and with microwave modes [147,205] representing the mechanical vibrational modes. These simple devices with natively available bosonic modes have accurately carried out simulations that would be impossible on any currently existing superconducting qubit-only quantum computer [147]. A related experiment has recently simulated dissipative molecular quantum dynamics near a conical intersection [206].

A third potential advantage of the hybrid hardware we discuss is that state and process tomography for CV systems (see Appendix E) can be straightforwardly performed through simple protocols for measurement of either the Wigner function or the characteristic function [60,64,65,108,207–212]. State and process tomography makes it possible to understand and calibrate the error model for the particular hardware. This in turn allows one to develop a corresponding AMM containing a description of the appropriate error models associated with the gates, measurements, and time evolution of the hardware within the framework of the ISA. This information is useful as a basis for integration into either high-level or intermediate representation languages such as QISKIT [137], Q$^{\#}$ [213], QIR [188,189], and Multi-Level Intermediate Representation Compiler [112]. And as mentioned earlier, BOSONIC QISKIT [52] is an extension of the QISKIT language that can represent gates, measurements, and error models for hybrid oscillator-qubit processors, and is available to the community [53–55] as a codesign tool.

The above advantages are largely directly applicable to trapped-ion systems [76,86], neutral atoms [115,118,214], and superconducting qubit or microwave resonator systems [28,29,215,216]. In this work, we primarily specialize in the superconducting platform but also comment on key differences in trapped-ion and neutral-atom platforms where applicable. Recent introductory reviews of superconducting cQED concepts can be found in Refs. [28,29,216,217]. Useful works on continuous-variable quantum information processing include Refs. [218–224]. The power of oscillators for information processing tasks such as phase estimation is surveyed in Ref. [223]. Readers interested in background information on the microscopic Hamiltonians for hardware implementations of superconducting circuits, trapped-ion systems, and neutral atoms are referred to Appendix A. See Ref. [225] for a textbook on the principles of superconducting quantum computers.

The hybrid CV-DV systems notionally illustrated in Fig. 3 enjoy many of the useful features of linear optics quantum computation (LOQC) [226–232], but have key advantages over LOQC because of the ability to use the qubits as auxiliary controllers to create and manipulate complex and highly non-Gaussian photon states in the resonators (e.g., bosonic error-correction code words) and to deterministically perform nontrivial gate operations on the bosonic modes [233–236] without relying on the measurement- and fusion-based protocols to supply the required nonlinearity, as is done in LOQC. An additional advantage in the superconducting case [Fig. 3(a)] is that

each qubit is coupled to a single bosonic mode (to minimize crosstalk) and the beam splitters connecting the resonator modes are real-time controllable (microwave pulse activated) [70,237], allowing for rapid and high-fidelity routing of the bosonic quantum information throughout the (2D or potentially 3D) hardware fabric. With proper scheduling to avoid collisions, this routing can be efficiently parallelized.

Finally, we note that a critically important advantage of the superconducting architecture is that, unlike in LOQC, photon-number measurements [147,148,238,239] (and photon-number parity measurements [240]) are very nearly QND (no photons are absorbed by the detector). This means that, e.g., in the dual-rail architecture (originally developed for LOQC [227,241]) where the logical qubit states involve one photon shared between two resonators, the dominant error (photon loss) can be converted with very high efficiency to an erasure error via QND measurement of the joint photon number in the dual-rail qubit [233,234,236,242] which jumps from one to zero when a photon is lost, or from one to two in the less likely event that a photon is gained (see also the qubit version of the dual-rail in Refs. [243,244]). Because the location of erasure errors is flagged, such errors have lower entropy and are vastly easier to correct [245–248].

Having established the details of the hybrid CV-DV architecture, we are now ready to deep dive into each layer of the stack in the next sections.

III. STATES AND OPERATORS

The physical hardware is the lowest layer in the CV-DV architecture stack (Fig. 2), and in that context, this section describes the basics of a CV-only architecture, i.e., oscillator states and operations. This is intended to provide a pedagogical introduction to the nontrivial aspects of a continuous-variable space from a mathematical perspective, assuming prior familiarity with standard qubit-based quantum computation. The notation introduced in this section is used throughout this tutorial.

After summarizing standard qubit states and gates in Sec. III A, we define in Sec. III B the Hilbert space of a multi-qubit and multi-oscillator quantum system that our bosonic quantum processor operates on.

Section III C then introduces the representation of bosonic states in two bases, the Fock basis and phase-space basis. While the former can be directly translated into the language of qudits, the latter does not have any one-to-one mapping with a DV architecture. Hence, in that section, while introducing the two bases, we attempt to give a formal understanding of the CV space to computer scientists and DV algorithm experts, while using the opportunity to standardize our notation. Section III C also introduces the concept of phase space and stellar-holomorphic representation. Next, in Sec. III D, the issue of truncating the infinite-dimensional Hilbert space is discussed briefly. The various representations discussed here lead to different ISAs for hybrid CV-DV processors as discussed later in Sec. IV.

A. Qubit states and operations

It is useful to review the standard representation that we employ for single- and multi-qubit states before discussing the harmonic oscillator. To be unambiguous, we employ the following conventions here [249]:

$$|0\rangle = |\uparrow\rangle = |e\rangle = \begin{pmatrix}1\\0\end{pmatrix}, \tag{4}$$

$$|1\rangle = |\downarrow\rangle = |g\rangle = \begin{pmatrix}0\\1\end{pmatrix}, \tag{5}$$

$$\sigma_0 = \hat{I} = \begin{pmatrix}1 & 0\\0 & 1\end{pmatrix}, \tag{6}$$

$$\sigma_x = \sigma_1 = X = \begin{pmatrix}0 & 1\\1 & 0\end{pmatrix}, \tag{7}$$

$$\sigma_y = \sigma_2 = Y = \begin{pmatrix}0 & -i\\+i & 0\end{pmatrix}, \tag{8}$$

$$\sigma_z = \sigma_3 = Z = |e\rangle\langle e| - |g\rangle\langle g| = \begin{pmatrix}+1 & 0\\0 & -1\end{pmatrix}, \tag{9}$$

$$\sigma^+ = \frac{X+iY}{2} = |e\rangle\langle g| = \begin{pmatrix}0 & 1\\0 & 0\end{pmatrix}, \tag{10}$$

$$\sigma^- = \frac{X-iY}{2} = |g\rangle\langle e| = \begin{pmatrix}0 & 0\\1 & 0\end{pmatrix}. \tag{11}$$

Consistent with physics conventions, σ_z is still the energy operator, σ^+ raises the energy, and σ^- lowers the energy. Single-qubit gates are simply rotations

$$U = e^{-i\frac{\theta}{2}\vec{h}\cdot\vec{\sigma}}, \tag{12}$$

parametrized by a rotation angle θ and an axis on the Bloch sphere represented by the unit vector $\hat{h} = (h_x, h_y, h_z)$. Here, the "spin" vector is

$$\vec{\sigma} = (\sigma_x, \sigma_y, \sigma_z). \tag{13}$$

We will have occasion to make frequent use of rotations about an axis in the equatorial plane lying at an angle φ away from the x axis, and so we define the following notation for this purpose:

$$\sigma_\varphi = \sigma_x \cos\varphi + \sigma_y \sin\varphi. \tag{14}$$

The basis for multi-qubit states that we use is also the standard one. For example, for two qubits we have four states

that are direct products of two single-qubit states:

$$|\psi_0\rangle = |00\rangle = \begin{pmatrix}1\\0\end{pmatrix} \otimes \begin{pmatrix}1\\0\end{pmatrix} = \begin{pmatrix}1\\0\\0\\0\end{pmatrix}, \tag{15}$$

$$|\psi_1\rangle = |01\rangle = \begin{pmatrix}1\\0\end{pmatrix} \otimes \begin{pmatrix}0\\1\end{pmatrix} = \begin{pmatrix}0\\1\\0\\0\end{pmatrix}, \tag{16}$$

$$|\psi_2\rangle = |10\rangle = \begin{pmatrix}0\\1\end{pmatrix} \otimes \begin{pmatrix}1\\0\end{pmatrix} = \begin{pmatrix}0\\0\\1\\0\end{pmatrix}, \tag{17}$$

$$|\psi_3\rangle = |11\rangle = \begin{pmatrix}0\\1\end{pmatrix} \otimes \begin{pmatrix}0\\1\end{pmatrix} = \begin{pmatrix}0\\0\\0\\1\end{pmatrix}., \tag{18}$$

In this standard basis, multi-qubit operators are given by (sums of) direct products of individual qubit operators. For example, the joint parity operator for two qubits is

$$Z \otimes Z = \begin{pmatrix}+1 & 0\\0 & -1\end{pmatrix} \otimes \begin{pmatrix}+1 & 0\\0 & -1\end{pmatrix} \tag{19}$$

$$= \begin{pmatrix}+1 & 0 & 0 & 0\\0 & -1 & 0 & 0\\0 & 0 & -1 & 0\\0 & 0 & 0 & +1\end{pmatrix}. \tag{20}$$

Further discussion on qubit-only gates can be found in Table VIII of Appendix A.

B. The Hilbert space of oscillators

A bosonic mode is typically hosted in a superconducting microwave resonator, the mechanical oscillation of trapped ions in radio-frequency potentials, or neutral atoms in optical potentials, as illustrated in Fig. 3. Here, we represent a bosonic mode of frequency $\omega_R = 2\pi f$ as a simple harmonic oscillator with Hamiltonian (in units with Planck's constant $\hbar = 1$)

$$H_0 = \omega_R \left[\hat{n} + \frac{1}{2}\right], \tag{21}$$

where $\hat{n}$ is the number operator denoting the number of photons (or phonons, i.e., bosons) in the resonator mode. The corresponding energy eigenstate $|m\rangle$ is called a number state or Fock state as

$$\hat{n}|m\rangle = m|m\rangle; \quad m = 0, 1, 2, 3, \ldots \tag{22}$$

The number operator is thus

$$\hat{n} = \sum_{m=0}^{\infty} m\hat{P}_m, \tag{23}$$

where

$$\hat{P}_m \equiv |m\rangle\langle m| \tag{24}$$

is the projector onto Fock state $|m\rangle$. The Fock state with zero photons is referred to as the vacuum state.

A key feature of the harmonic oscillator is that the energy-level spacing is constant, which means that, like its classical counterpart, the quantum harmonic oscillator is isochronous—its period of oscillation is independent of the amplitude of the oscillation. This makes the transformation to the interaction representation very simple and allows us to eliminate the rapid time dependence generated by H_0 and study only the slower evolution under additional control terms (see Appendix A). That is, in essence, we can move to a rotating frame in which H_0 is effectively zero. We therefore use the interaction representation almost exclusively throughout this work.

C. Representation and bases

The fundamental difference between the state space that is often used in quantum computing and the one that we discuss here is that the oscillator state space is no longer conveniently represented in terms of two-level systems, i.e., qubits. Instead, quantum information is stored and manipulated in the joint Hilbert space of both harmonic oscillators and qubits. Similar in spirit to the nomenclature of a "qubit," the vector space of the oscillator is referred to as a "qumode" [250]. We discuss a few representations of the qumode state and operators in this section.

1. Fock basis

The simplest basis for representing the qumode state is just the eigenstates of the number operator $\hat{n}$, the Fock basis $|m\rangle$, defined in Eq. (22). Two other physically important operators are the boson creation operator $a^\dagger$ and its adjoint, the destruction operator a, which obey

$$a^\dagger|m\rangle = \sqrt{m+1}|m+1\rangle, \tag{25}$$

$$a|m+1\rangle = \sqrt{m+1}|m\rangle, \tag{26}$$

$$\left[a, a^\dagger\right] = 1, \tag{27}$$

$$\hat{n} = a^\dagger a. \tag{28}$$

$a^\dagger$ and a are also known as raising and lowering operators or ladder operators since they move the excitation number

up and down the ladder of Fock states. Fock state $|m\rangle$ can be created from the vacuum state $|0\rangle$ by the application of m raising operators:

$$|m\rangle = \frac{1}{\sqrt{m!}}(a^\dagger)^m|0\rangle. \tag{29}$$

The eigenstates $|\alpha\rangle$ of the lowering (boson destruction) operator are known as coherent states and are a continuous family of states labeled by a complex parameter α:

$$|\alpha\rangle = e^{\alpha a^\dagger - \alpha^* a}|0\rangle = e^{-\frac{|\alpha|^2}{2}} e^{\alpha a^\dagger}|0\rangle \tag{30}$$

$$= e^{-\frac{|\alpha|^2}{2}} \sum_{n=0}^{\infty} \frac{\alpha^n}{\sqrt{n!}}|n\rangle. \tag{31}$$

This family of eigenstates of the lowering operator obeys

$$a|\alpha\rangle = \alpha|\alpha\rangle, \tag{32}$$

with (complex) eigenvalue α. The operators $a, a^\dagger$ are non-Hermitian and are defective in the Jordan sense. That is, $a^\dagger$ has no right eigenstates and a has no left eigenstates. It seems strange that there exist quantum states from which one can remove a photon and still be in the same state. This is a "Hilbert hotel" phenomenon associated with the countable infinity of Fock states.

The fact that coherent states form a continuum means that they are uncountable and must necessarily be overcomplete since the Hilbert-space dimension is countably infinite. Indeed, it follows from Eq. (31) that the resolution of the identity is

$$\hat{I} = \frac{1}{\pi}\int d^2\alpha |\alpha\rangle\langle\alpha| = \sum_{n=0}^{\infty} |n\rangle\langle n|. \tag{33}$$

Coherent states are thus our first hint that, unlike DV systems (qubits), oscillators have CV features.

It can be argued [251] that coherent states are the least "quantum" (since, as we show in Box IV.1, they are simple classical displacements of the Gaussian vacuum) and Fock states are the most "quantum" (since they are highly non-Gaussian). Gaussianity is explained in the next section.

2. Position-momentum basis

The fact that oscillators admit a continuous-variable description becomes more clear when we consider the position and momentum operators for oscillators, as we now discuss. In quantum optics, the position coordinate of an electromagnetic oscillator mode is generally taken to be the electric field of the mode at some particular spatial location, and the conjugate momentum is then related to the magnetic field at that point. As we see below, the ground state wave function in the position basis is a Gaussian so the electric field has a Gaussian distribution even in the absence of excitations. These ground state fluctuations in the electric field are referred to as *vacuum noise* (even though the underlying state is pure) [252]. The wave function is also a Gaussian in the momentum representation as the Fourier transform serves to convert between the position and momentum representations, and Gaussians are mapped to Gaussians under Fourier transforms.

We can form dimensionless Hermitian position ($\hat{x}$) and momentum ($\hat{p}$) operators from the ladder operators via

$$\hat{x} = \lambda_x(a + a^\dagger) \tag{34}$$

$$\hat{p} = -i\lambda_p(a - a^\dagger), \tag{35}$$

where λ_x and λ_p are real constants. One conventional choice is $\lambda_x = \lambda_p = \frac{1}{\sqrt{2}}$, which yields the usual commutator (again with $\hbar = 1$),

$$[\hat{x}, \hat{p}] = +i. \tag{36}$$

We refer to this case as employing "standard units." We, however, also use a different convention $\lambda_x = \lambda_p = \frac{1}{2}$, which we refer to as employing "Wigner units" and which yields

$$[\hat{x}, \hat{p}] = +\frac{i}{2}. \tag{37}$$

We alert the reader as to which unit choice is being made as needed.

The choice of Wigner units has the advantage that we may write the ladder operators in the simple form

$$a = \hat{x} + i\hat{p}, \tag{38}$$

$$a^\dagger = \hat{x} - i\hat{p}, \tag{39}$$

and more importantly, we have the convenient feature that for coherent states the mean position and momentum obey

$$\langle\alpha|\hat{x}|\alpha\rangle = \mathrm{Re}\,\alpha, \tag{40}$$

$$\langle\alpha|\hat{p}|\alpha\rangle = \mathrm{Im}\,\alpha, \tag{41}$$

and the corresponding variances are given by

$$\sigma_x^2 = \langle\alpha|\hat{x}^2|\alpha\rangle - \langle\alpha|\hat{x}|\alpha\rangle^2 = \frac{1}{4}, \tag{42}$$

$$\sigma_p^2 = \langle\alpha|\hat{p}^2|\alpha\rangle - \langle\alpha|\hat{p}|\alpha\rangle^2 = \frac{1}{4}. \tag{43}$$

With this choice of dimensionless "Wigner" units, we can now write the Hamiltonian in Eq. (21) in the form familiar

from a mass-and-spring harmonic oscillator,

$$H_0 = \omega_{\mathrm{R}}[\hat{x}^2 + \hat{p}^2]. \tag{44}$$

The operator a is defective, having only a single right eigenvector, namely the vacuum state, which has an eigenvalue of 0:

$$a|0\rangle = 0|0\rangle. \tag{45}$$

The Hermitian position operator is the sum of the creation and annihilation operators and has a continuous spectrum of non-normalizable eigenvectors labeled by the position at which they are infinitely sharply peaked:

$$\hat{x}|x\rangle = x|x\rangle. \tag{46}$$

In terms of these position eigenstates, the resolution of the identity is

$$I = \int_{-\infty}^{+\infty} \mathrm{d}x\, |x\rangle\langle x|. \tag{47}$$

From the relation $I^2 = I$ it follows that the orthonormality condition on the position eigenvectors is

$$\langle y|x\rangle = \delta(x - y), \tag{48}$$

where δ is the Dirac delta function, illustrating that the position eigenstates are non-normalizable and, strictly speaking, can only be defined in terms of a distribution, i.e., sequences of narrower and narrower normalizable wave packets, e.g., Gaussians.

In the so-called position representation, the state vector of the oscillator,

$$|\psi\rangle = \int_{-\infty}^{+\infty} \mathrm{d}x\, \psi(x)|x\rangle, \tag{49}$$

is represented by a (possibly complex) wave function

$$\psi(x) = \langle x|\psi\rangle, \tag{50}$$

the argument of which is the position of the oscillator. In this position representation, the position operator $\hat{x}$ is represented by multiplication by the number x:

$$\hat{x}|\psi\rangle \to x\psi(x). \tag{51}$$

From the commutation relation in Eq. (37), we see that the momentum operator acting on the state can be written as a derivative acting on the wave function

$$\hat{p}|\psi\rangle \to -\frac{i}{2}\frac{d}{dx}\psi(x). \tag{52}$$

In this representation, Eq. (32) becomes the differential equation

$$\left(x + \frac{1}{2}\frac{d}{dx}\right)\psi_\alpha(x) = \alpha\psi_\alpha(x), \tag{53}$$

which has a normalized solution for the coherent state wave function

$$\psi_\alpha(x) = \left(\frac{2}{\pi}\right)^{\frac{1}{4}} e^{-(x-\alpha)^2}. \tag{54}$$

We see that for real $\alpha = x_0$, the center of the Gaussian is shifted from the origin to position x_0, while for imaginary $\alpha = iK$ we have

$$\psi_{iK}(x) = \left(\frac{2}{\pi}\right)^{\frac{1}{4}} e^{-(x-iK)^2} \tag{55}$$

$$= \left(\frac{2}{\pi}\right)^{\frac{1}{4}} e^{K^2} e^{2iKx} e^{-x^2}, \tag{56}$$

which, because of the factor of $1/2$ in Eq. (52), indeed corresponds to a momentum boost of K.

Returning to Eq. (30) we see that the operator that takes the vacuum state to coherent state $|\alpha\rangle$ is a unitary translation operator in phase space (a concept defined in the next section):

$$D(\alpha) \equiv e^{\alpha a^\dagger - \alpha^* a} = e^{-\frac{|\alpha|^2}{2}} e^{+\alpha a^\dagger} e^{-\alpha^* a} \tag{57}$$

$$= e^{2i\{[-\mathrm{Re}\,\alpha]\hat{p} + [\mathrm{Im}\,\alpha]\hat{x}\}}. \tag{58}$$

We see from this that when using Wigner units, it is $2\hat{p}$ that is the generator of displacements and $2\hat{x}$ that is the generator of momentum boosts.

3. Introduction to phase space

Phase space is a concept borrowed from classical mechanics: a single point in this space corresponds to a two-component position or momentum vector $\vec{R} = (x, p)$. While the state of a classical particle moving in one spatial dimension is fully specified by its position and momentum (position times velocity), in quantum mechanics, this point becomes a "fuzzy blob" because the operators $\hat{x}$ and $\hat{p}$ do not commute. The Heisenberg uncertainty principle guarantees a lower bound on the product of the variances of $\hat{x}$ and $\hat{p}$. Since coherent states are Gaussians, it is straightforward to show that the variance of each is 1/4 in our dimensionless "Wigner" units.

As noted above, CV systems have the interesting property that their states can be represented in the (countably) infinite discrete Fock basis, or in terms of smooth continuous wave functions $\psi(x)$, or continuous density "matrices" $\rho(x,x')$ in the position (or the momentum) basis. To develop intuition about these continuous representations, it is helpful to begin with a study of the classical oscillator, the state of which is not defined in Hilbert space, but rather by its position $\vec{R} = (x,p)$ in the 2D "phase space" of position and momentum. Statistical mixtures (ensembles) of classical states are then described by a probability distribution $P(x,p)$. Quantum mechanically, the numbers (x,p) are replaced by noncommutative operators $(\hat{x},\hat{p})$ that satisfy the canonical commutations in Eqs. (36) and (37). Mixtures of quantum oscillator states are described by the Wigner function (see Appendix E 6), $W(x,p)$, a *quasiprobability* distribution which is real but can take negative values [210,253]. There are two other commonly used representations of CV states, the characteristic function and the Husimi Q-function in Appendixes E 6 and E 7, respectively.

Universal quantum control requires the ability to execute an arbitrary unitary transformation on the Hilbert space of the combined qubit-oscillator system. It is useful to think of the problem of unitary synthesis as one of Hamiltonian synthesis: how we synthesize a Hermitian operator (Hamiltonian) H such that it generates the desired unitary $U = e^{-iHt}$ under time evolution, or more generally for a time-dependent Hamiltonian

$$\frac{dU}{dt} = -iH(t)U(t). \tag{59}$$

It is useful to begin with the classical dynamics of a single oscillator. The intuition gained from this will prove very useful in designing quantum Hamiltonians to achieve desired unitary transformations in Hilbert space. In the classical limit, we evaluate the quantum equations of motion only to zeroth order in Planck's constant [254]. Rather than Hilbert space, we describe the state of the oscillator in terms of a point $\vec{R} = (x,p)$ in the 2D phase space, where the numbers (rather than operators) x and p are, respectively, the position and the corresponding canonical momentum. The Hamiltonian $H(x,p)$ determines how the system moves through phase space according to Hamilton's equations of motion [255]:

$$\vec{\Lambda}(\vec{R}) = \frac{d\vec{R}}{dt} = \left(\frac{\partial H(x,p)}{\partial p}, -\frac{\partial H(x,p)}{\partial x}\right). \tag{60}$$

Time evolution maps the phase space onto itself, and it turns out to be useful to think of this mapping as described by the velocity field $\vec{\Lambda}(\vec{R})$ of a fluid filling all of the phase space—the position $\vec{R}$ of each fluid element evolves according to the Hamilton equation of motion. An interesting feature of this fluid is that it is guaranteed to be incompressible because the flow field is automatically divergenceless (Liouville's theorem):

$$\vec{\nabla}\cdot\vec{\Lambda}(\vec{R}) = \frac{\partial}{\partial x}\frac{\partial H(x,p)}{\partial p} - \frac{\partial}{\partial p}\frac{\partial H(x,p)}{\partial x} = 0. \tag{61}$$

This means that we can think of the mapping of phase space onto itself under Hamiltonian evolution as a volume-preserving diffeomorphism.

As a first illustrative example, consider the simple classical Hamiltonian

$$H = gxp, \tag{62}$$

where g is a real constant. Because the Hamiltonian is quadratic, the equation of motion for the divergenceless flow field $\vec{\Lambda}$,

$$\vec{\Lambda}(x,p) = \left(\frac{dx}{dt}, \frac{dp}{dt}\right) = g\,(x,-p)\,, \tag{63}$$

is linear and thus readily solved:

$$x(t) = x(0)e^{gt}, \tag{64}$$

$$p(t) = p(0)e^{-gt}, \tag{65}$$

where $x(0), p(0)$ are the initial position and momentum at time $t = 0$. From this we see that H is the "squeezing" Hamiltonian because under its action, a circular region of the phase-space "fluid" evolves into a squeezed ellipse (of precisely the same area) as illustrated in Fig. 4. This is a simple example of a Gaussian operation, a concept that is discussed in more detail for the quantum case in Sec. IV A.

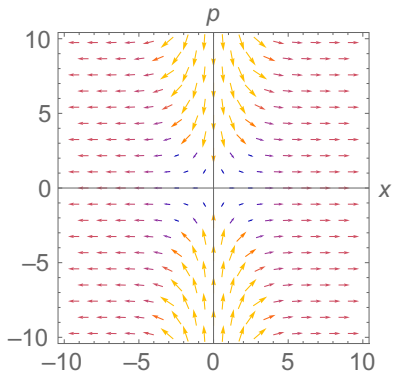


FIG. 4. (a) The phase-space flow field $\vec{\Lambda}$ under the squeezing Hamiltonian in Eq. (62). (b) The time evolution of an initial circular region into a squeezed ellipse (of the same total area) under the phase-space flow field for times $gt = 0.0$, 0.2, 0.4, and 0.8 (blue, orange, green, and red).

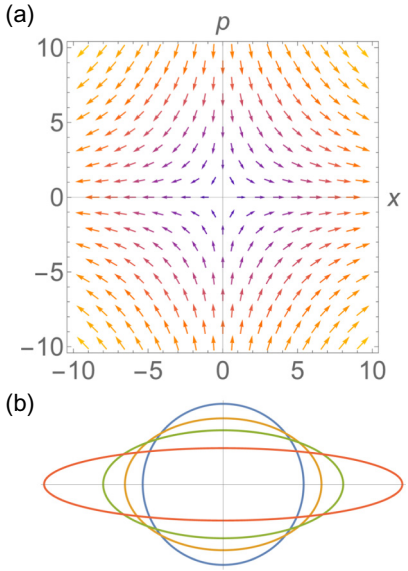


FIG. 5. The flow field for the Hamiltonian given in Eq. (66). Note that the flow field for large $|x|$ and small $|p|$ is much more uniform (less squeezing) than that of the squeezing Hamiltonian shown in Fig. 4.

As a second example, we consider the nonquadratic Hamiltonian

$$H = g \tanh\left(\frac{x}{2}\right) p. \tag{66}$$

The phase-space flow field for this Hamiltonian is illustrated in Fig. 5. Note that, because the hyperbolic tangent is a constant for a large argument, the flow field is very uniform for large $|x|$, being well approximated by

$$\vec{\Lambda}(x,p) \sim \mathrm{sgn}(x)\,(g,0). \tag{67}$$

Thus, application of this Hamiltonian for time t leads to a simple displacement (for large $|x|$) rather than squeezing, with the sign of the displacement matching the sign of x

$$x \to x + gt\,\mathrm{sgn}(x). \tag{68}$$

This particular flow will prove useful for the "two-legged cat amplifier" operation described in Sec. VI D 3. An elaboration of this Hamiltonian,

$$H = g\left[-\tanh\left(\frac{x+p}{2}\right)(x-p) + \tanh\left(\frac{x-p}{2}\right)(x+p)\right], \tag{69}$$

produces a flow field with fourfold rotational symmetry illustrated in Fig. 6. As described in Sec. V C, this proves useful for QEC using the bosonic four-leg cat code.

As we discuss in more detail when we present compilation techniques in Sec. VI, these examples illustrate the fact that the classical intuition associated with visualizing

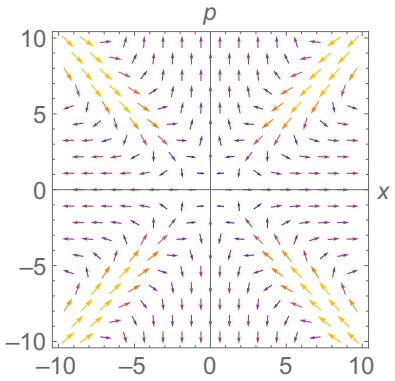


FIG. 6. A flow field exhibiting the fourfold rotational symmetry of the Hamiltonian given in Eq. (69).

flow fields can be extremely helpful to keep in mind when designing quantum unitary transformations that perform various useful tasks. We see that the nonlinearities needed for oscillator Hamiltonians can be implemented by either non-Gaussian oscillator controls or hybrid oscillator-qubit gates as defined in Tables III and IV. We also define a minimum universal gate set in Table V.

4. Stellar-holomorphic representation

A less well-known representation of CV states is via analytical functions in the complex plane (i.e., holomorphic functions), the so-called stellar representation [178, 256–258], which has also been used in condensed-matter physics [259] to describe quantum Hall states. In the stellar representation, the bosonic creation and annihilation operators are mapped to a complex variable and its derivative operator. This representation provides a visualization of CV state dynamics in terms of the movement of the zeros of analytical functions. Additionally, it quantifies the non-Gaussian character of CV states (a useful resource for universal CV computation; see Sec. IV for more details) in terms of the number of zeros of the analytical functions. We provide an introduction to the stellar representation in this section and toward the end also highlight open questions.

The basic idea is to use the following mapping from bosonic operators to a complex variable z and the corresponding derivative,

$$a^\dagger \to z, \quad a \to \partial_z, \tag{70}$$

such that the canonical commutation relationship $[a, a^\dagger] = [\partial_z, z] = 1$ is preserved. Then, any CV state in the Fock representation $\sum_{n=0}^{\infty} \alpha_n/\sqrt{n!}(a^\dagger)^n\,|0\rangle$ (with α_n being a complex number) can be mapped to an analytical function of z,

$$\sum_{n=0}^{\infty} \frac{\alpha_n}{\sqrt{n!}}(a^\dagger)^n\,|0\rangle \to f(z)\,|0\rangle\,, \tag{71}$$

where

$$f(z) = \sum_{n=0}^{\infty} \frac{\alpha_n}{\sqrt{n!}} z^n, \tag{72}$$

and the vacuum state is replaced by the integration measure

$$|0\rangle \to \frac{1}{\sqrt{\pi}} e^{-\frac{1}{2}|z|^2}. \tag{73}$$

By convention, the derivative $a \to \partial_z$ does not act on the integration measure. Thus, for example, we have the translation

$$\langle 0|a^n (a^\dagger)^n|0\rangle \to \frac{1}{\pi} \int d^2 z e^{-|z|^2} \left[\frac{\partial^n}{\partial z^n} z^n\right] = n! \tag{74}$$

$$\to \frac{1}{\pi} \int d^2 z e^{-|z|^2} \left[(z^*)^n z^n\right]. \tag{75}$$

Here, the first line follows from the picture that we are calculating the vacuum expectation value of the operator $a^n (a^\dagger)^n$ using the translation in Eq. (70), while the second line follows from integration by parts or from the picture that we are computing the norm of the state $|\psi\rangle = (a^\dagger)^n|0\rangle$ using the translation $a^\dagger \to z$.

Moreover, the analytical function $f(z)$ admits the following factorization [260]:

$$f(z) = e^{-\frac{1}{2}Az^2 + Bz + C} g(z), \quad \text{where,}$$
$$g(z) = \prod_{j=1}^{S} (z - \lambda_j), \tag{76}$$

for $A, B, C \in \mathbb{C}, \lambda_j \in \mathbb{C}$. The factorization in Eq. (76) naturally decomposes $f(z)$ into a Gaussian part $(e^{-\frac{1}{2}Az^2 + Bz + C})$ and a non-Gaussian (polynomial) part $g(z)$, which is an entire function and hence uniquely determined [261] by its zeros $\{\lambda_j;\ j = 1, \ldots S\}$. The number of zeros, S, is called the stellar rank and is a measure of the non-Gaussianity of the state [262] (and if $S = 0$, we take $g(z) = 1$). Here, the parameters A, B, and C, respectively, play the role of squeezing, displacing, and adding a constant phase (and normalization) to the CV states. In particular, for $A = |A|e^{i\phi}$, the magnitude $0 \le |A| < 1$ controls the squeezing of the CV state, while the phase ϕ controls the direction of squeezing. Similarly, $|B|$ controls the displacement, and $\arg(B)$ determines the direction of the displacement. For $\phi = 0$ (real A), $B = 0$ we have for the purely Gaussian case $(S = 0, g(z) = 1)$,

$$\langle(\hat{x})^2\rangle = \frac{1}{4}\left(\frac{1+A}{1-A}\right), \tag{77}$$

$$\langle(\hat{p})^2\rangle = \frac{1}{4}\left(\frac{1-A}{1+A}\right), \tag{78}$$

which demonstrates the squeezing relative to the vacuum variances of $\frac{1}{4}$ for each that were given in Eqs. (42) and (43).

Interestingly, it has been established that the difficulty of preparing a CV state can be quantified by the stellar rank in the above representation, while the dynamics of a qumode can be visualized by the movement (shifting, merging, and splitting) of the roots in the complex plane [258]. Therefore, a time-dependent CV state is associated with a corresponding time-dependent analytical function $f(z, t)$.

An arbitrary unitary operation on a qumode can always be written as a function of $a^\dagger$ and a. Therefore, in the stellar representation, a CV unitary operation may be written as

$$U(z, \partial_z) = e^{-iH(z, \partial_z)}, \tag{79}$$

such that the dynamics of a CV state is governed by the following differential equation on the corresponding analytical function $f(z, t)$:

$$i\partial_t f(z, t) = H(z, \partial_z) f(z, t). \tag{80}$$

This agrees with our discussion about visualizing CV state dynamics in Sec. III C 3 as flows in phase space.

As a simple example to better understand the CV dynamics in stellar representation, let us examine how the SNAP gate (described in Sec. IV) moves the roots of the stellar function. By definition, the SNAP gate unitary $e^{i\sum_n \theta_n |n\rangle\langle n|}$ imparts a different phase θ_n to each Fock state $|n\rangle$. In the stellar picture, this means that the initial stellar function $\sum_{n=0}^{\infty} a_n z^n \propto \prod_j (z - \lambda_j)$ is changed to $\sum_{n=0}^{\infty} a_n e^{i\theta_n} z^n \propto \prod_k (z - \lambda'_k)$. In the special case of $\theta_n = n\theta$ for a constant θ, the resulting stellar function is

$$\sum_n a_n e^{in\theta} z^n = \sum_n a_n (e^{i\theta} z)^n \propto \prod_j (e^{i\theta} z - \lambda_j). \tag{81}$$

This means that the new roots λ'_k are related to the old roots λ_j via

$$\lambda'_k = e^{-i\theta} \lambda_k, \quad \forall\, k \tag{82}$$

and the number of roots does not change. We note that this special SNAP gate is nothing but the free evolution of the oscillator, which induces a global phase transformation on all the roots, i.e., a global rotation in the phase space. It

remains an open question as to how the roots of the stellar function move subject to the action of an arbitrary SNAP gate or other arbitrary unitary.

More concretely, it would be interesting and useful to develop theories and protocols to: (i) shift an individual root of the stellar function while keeping all other roots unchanged; (ii) delete an individual root; and (iii) create an additional root, by leveraging the instruction set architecture in Sec. IV. This is an open problem at present, but developments along this line will provide insights that would allow us to better compile CV algorithms as we discuss more in Sec. VI. For example, a universal instruction set based on the stellar representation, i.e., "stellar ISA," may be constructed composed of operations that can manipulate individual roots of the stellar function, in a similar spirit to the ISAs that we introduce in Sec. II. Furthermore, when coupling to a qubit, the joint oscillator-qubit system requires two separate stellar functions to describe the oscillator state entangled with each of the two-qubit bases. Finally, we note that the stellar representation has been generalized to the multimode setting in Ref. [257]. For a more detailed discussion of non-Gaussian quantum states and stellar representations, we refer the reader to the tutorial given in Ref. [262] and to a recent paper on "photon catalysis" for multimode non-Gaussian state preparation [263].

D. Truncation of the Hilbert space

The representation discussed in the previous section concerns an infinite Hilbert space which is not feasible in practice. Here, we discuss possible truncations that can be applied to the oscillator space in realistic settings while keeping the advantages of a bosonic system. The Hilbert space of an ideal harmonic oscillator has infinite dimension, which may lead to technical difficulties in theoretical treatments. Of course, for physically realizable states of an oscillator in real hardware, we can never approach infinite energy, and the inevitable effects of even small anharmonicities increase as the number of bosons increases. For theoretical discussion of an abstract machine model in terms of an ISA, truncation is necessary to avoid any notion of infinity as well as the possible inception of anharmonicity in physical hardware. This former issue is particularly relevant as it prevents us from proposing algorithms that would use systems with infinite energy to instantaneously solve problems and also to ensure that standard algorithmic analysis techniques such as remainder estimates from Taylor's theorem can be rigorously used to bound approximation errors.

Truncations may be introduced in various forms. In the phase-space picture, truncation can be applied by a Gaussian damping envelope centered at the origin of the phase space. It may also be implemented by limiting the number of roots of the stellar function (or equivalently, the degree of the polynomial $g(z)$) throughout the computation. Within the Fock basis, truncation can be defined by omitting all Fock levels with energy higher than a certain threshold. Such a sharp energy cutoff in the Fock basis can also be replaced by a smoother damping function. A rigorous discussion of truncation errors is presented in Ref. [264].

In the following discussion, for simplicity, we use a sharp cutoff in Fock levels to truncate the Hilbert space. This is achieved by defining a maximum allowed boson number, $N_{\max}$, to restrict the oscillator Hilbert-space dimensionality to $d = N_{\max} + 1$. Physically, this places an upper bound on the energy. For example, for $N_{\max} = 3$ we have for the matrix representation of the destruction operator,

$$a = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & \sqrt{2} & 0 \\ 0 & 0 & 0 & \sqrt{3} \\ 0 & 0 & 0 & 0 \end{pmatrix} \tag{83}$$

and for the creation operator,

$$a^\dagger = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & \sqrt{2} & 0 & 0 \\ 0 & 0 & \sqrt{3} & 0 \end{pmatrix}. \tag{84}$$

Note that these truncated representations correctly reproduce

$$a^\dagger a = \hat{n} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \end{pmatrix}, \tag{85}$$

but incorrectly yield

$$aa^\dagger = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 0 & 0 & 3 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \neq \hat{n} + \hat{I}. \tag{86}$$

Thus the commutator is also not the identity

$$[a, a^\dagger] = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -3 \end{pmatrix} \neq \hat{I}. \tag{87}$$

It is therefore best if, before the numerical evaluation of any expression, one moves all the creation operators to the left of all the destruction operators (putting them in "normal order") taking into account the correct commutation relations. In any case, $N_{\max}$ must be set large enough that

the quantum amplitudes in the largest boson number states are extremely small throughout the entire calculation.

For low-energy states, it has been shown [265] that a value of $N_{\max}$ that is logarithmic in the truncation error desired will often suffice for the simulation. However, for dynamical simulation with a large number of bosons, cruder but more widely applicable bounds can be used. Specifically, if we assume that we have a state $|\psi(0)\rangle$ such that

$$\langle\psi(0)|\, H\, |\psi(0)\rangle = E, \tag{88}$$

it is easy to see that for any time-independent Hamiltonian H this mean energy is a constant of motion, meaning that the time-evolved state satisfies $\langle\psi(t)|\, H\, |\psi(t)\rangle = E$ for all $t \geq 0$. We then have, if $H = \omega a^\dagger a + V$ for positive semidefinite V, that

$$\begin{aligned} E &= \omega\, \langle\psi(t)|\, a^\dagger a\, |\psi(t)\rangle + \langle\psi(t)|\, V |\psi(t)\rangle \\ &\geq \omega\, \langle\psi(t)|\, a^\dagger a\, |\psi(t)\rangle := \omega \mathbb{E}(N(t)), \end{aligned} \tag{89}$$

where $\mathbb{E}$ denotes an expectation value over the state, and $N(t)$ is the time-dependent number operator. Markov's inequality states that the probability that a non-negative random variable achieves a value of $\aleph$ times its mean value is at most $1/\aleph$. From this, it follows that for any constant $\aleph \geq 1$,

$$\sum_{j \geq \aleph E/\omega} |\langle\psi(t)|j\rangle|^2 \leq \frac{1}{\aleph}, \quad \forall\, t \geq 0. \tag{90}$$

Thus, if we wish for the total probability of observing a boson number greater than the cutoff $N_{\max} = \aleph E/\omega$ to be at most δ_{cut} for all times, then it suffices to choose

$$N_{\max} \geq \frac{E}{\delta_{\text{cut}}\omega}. \tag{91}$$

Tighter bounds are possible given more information, but the simplicity and general applicability of this bound make it a useful if crude estimate of the cutoff for most cases with quadratic Hamiltonians in the field operators.

Much later in this tutorial, in Sec. VIII A, various mappings between the oscillator representation and the qubit representation are discussed, to bridge the gap between these CV and the familiar DV systems. The nontriviality of the CV space is an exemplary resource that is hard to beat via a discrete architecture and hence becomes the central focus of this work in later sections. To highlight this nontriviality on the unique features of CV architecture, Sec. VIII B represents a quantitative compilation of simple bosonic operators into DV gate sets.

IV. INSTRUCTION SETS FOR OSCILLATOR-QUBIT SYSTEMS

With the basics of oscillators and qubits laid out in the previous section, we are ready to present the important elementary gates that are natively available in current experiments on superconducting and/or ion-trap hybrid qubit-bosonic systems. Given an understanding of the natively available gates, we can select subsets of the gates to define instruction set architectures and learn how to compile these to synthesize arbitrary unitaries required to execute QEC, computation, and simulation circuits. We relegate the discussion of the microscopic Hamiltonians and physical implementation details behind the natively available gates to Appendix A.

We organize our discussion of the natively available and more complex synthesized gates according to the complexity of their corresponding Hamiltonians. The coarsest classification of gates and their corresponding Hamiltonians is whether they are qubit-only (Table VIII), oscillator-only (Table III), or hybrid entangling gates (Table IV). Tables III and IV list commonly available oscillator and qubit-oscillator hybrid gates arranged, respectively, and will be frequently referred to throughout this section. For completeness, we also list the standard qubit gates in Table VIII of Appendix A.

To guide the reader familiar with qubits and Bloch spheres, we have introduced the notion of phase space for continuous-variable systems and provided intuition for CV unitaries in terms of visualization of Hamiltonian flows in phase space in Sec. III C 3. In Sec. IV A, we leverage this intuition to explain the parallels between rotations on the DV Bloch sphere and simple displacements in the CV phase space generated by Hamiltonians that are linear in $\hat{x}, \hat{p}$. As mentioned previously, we approach general unitary synthesis as a Hamiltonian synthesis problem by writing the unitary in the form of Eq. (59) or more simply, for the case of a time-independent Hamiltonian, in the form

$$U(t) = e^{-iHt}, \tag{92}$$

where the Hermitian operator H is the Hamiltonian associated with the gate. In general, this is not simply the physical Hamiltonian but rather, an effective Hamiltonian synthesized in the hardware layer through the application of various microwave pulses. We refer the interested reader to the following experimental papers which discuss these issues [60,64,65,106–108,180,266–269] and to Appendix A. Within the oscillator-only gate family, we can express the Hamiltonian $H(\hat{x}, \hat{p})$ as a polynomial in the position ($\hat{x}$) and momentum ($\hat{p}$) operators of the oscillator mode(s) and classify gates according to the degree of the polynomial [218].

Section IV B presents standard Gaussian gate operations, the corresponding Hamiltonians of which are at most

TABLE II. Qubit Clifford circuits versus Gaussian bosonic circuits [218,219]. TMS stands for two-mode squeezing. The Weyl-Heisenberg group is the group of (noncommuting) translations in phase space effected (in standard dimensionless units) via $X(s) = e^{-is\hat{p}}$ which translates by s in position and $Z(t) = e^{it\hat{x}}$ which produces translations by t in momentum. Under non-Gaussian measurements, $|n\rangle\langle n|$ is the projector onto the Fock state with n photons, and $a^\dagger a$ stands for photon-number-resolving detection. Note that in general, there is no unique one-to-one mapping between the qubit gates in the second column and the bosonic gates in the third column. Note the difference between the Fourier transform F and the Hadamard operation: $F^\dagger \neq F$ while $H^\dagger = H$. As a result, the Fourier transform maps $\hat{x}$ into $\hat{p}$ and $\hat{p}$ into $-\hat{x}$, while the Hadamard maps X into Z and Z back into X (with no minus sign).

Aspect	Qubits	Bosonic modes
Standard basis	Z, $\{0,1\}$	$\hat{x}$, $\|x\rangle$
Conjugate basis	X, $\|0\rangle \pm \|1\rangle$	$\hat{p}$, $\|p\rangle = \int \mathrm{d}x\, e^{ixp}\|x\rangle$
Basis transformation	Hadamard	Fourier transform (i.e., $e^{i(\pi/2)\hat{n}}$)
Group	Pauli	Weyl-Heisenberg (e.g., $e^{i\alpha\hat{x}}$, $e^{i\beta\hat{p}}$)
	Clifford	Gaussian unitaries (e.g., $e^{i\theta\hat{n}}$, $e^{i\lambda\hat{x}^2}$, $e^{i\hat{x}_1\hat{p}_2}$)
Channel	Pauli or Clifford	Displacement or Gaussian
Measurement	Pauli-basis measurement of X and Z requires Clifford operations	Homodyne measurement of $\hat{x}$ and $\hat{p}$ requires Gaussian operations
	Non-Clifford ($\vec{b}\cdot\vec{\sigma}$)	Non-Gaussian ($a^\dagger a$, $\|n\rangle\langle n|$, aa, $\hat{x}^2$, $\hat{p}^3$, $\cos k\hat{x}$)
Representation	State tomography	Wigner- (or characteristic) function tomography
Benchmarking	Clifford operations are 2-design	Gaussian operations are not 2-design
Example entangled state	Bell: $\|00\rangle + \|11\rangle$	EPR (TMS, SUM): $\sum_n \|n\rangle_1\|n\rangle_2$, $\int \mathrm{d}x\, \|x\rangle_1\|x\rangle_2$
Universal computation	A non-Clifford gate, e.g., $T = e^{i(\pi/8)Z}$, augments the Clifford group with magic (non-Cliffordness)	A non-Gaussian unitary, cubic gate: $e^{i\chi\hat{x}^3}$ augments Gaussian operations with Wigner negativity

quadratic in $\hat{x}$ and $\hat{p}$. In Sec. IV C, we also discuss parallels (summarized in Table II) between the group of Clifford gates for qubits and Gaussian gates for oscillators. Like their discrete-variable cousins, Gaussian gates are nonuniversal, and Gaussian circuits can be efficiently simulated classically (assuming that the initial state is Gaussian). The efficiency for qubit Clifford circuit simulation comes from being able to easily track the transformation of the stabilizers [270,271]. For Gaussian circuits, one simply follows the evolution of the means and the covariance matrix for the oscillator position and momentum [272,273].

Universal oscillator control requires being able to generate polynomials of arbitrary order using combinations of gates [33,218], as discussed in Appendix D 5. In Sec. IV D, we explain that this can be achieved given access to a single non-Gaussian gate having a Hamiltonian of degree 3 or higher. Note that this definition assumes no decoherence; in a realistic setting, decoherence needs to be accounted for, leading to a trade-off between using more gates for higher unitary precision and the corresponding increase in infidelity due to decoherence increasing with circuit depth. Section IV D also introduces useful hybrid oscillator-qubit gates.

Importantly, we show in Sec. IV E that hybrid oscillator-qubit gates with Hamiltonians of the form

$$H = \vec{h}(\hat{x},\hat{p})\cdot\vec{\sigma} \tag{93}$$

can be compiled to generate universal unitaries for the full hybrid system even when the oscillator part $\vec{h}(\hat{x},\hat{p})$ is only of maximum degree 1. When $\vec{h}$ is homogeneous of degree 0 we have the ordinary single-qubit rotation gate in Eq. (12). When $\vec{h}$ is homogeneous of degree 1, we see that we can interpret the unitary as a rotation of the qubit that depends linearly on the position and momentum of the oscillator. We also see that we can alternatively view the same unitary as a displacement of the oscillator in phase space, conditioned on the state of the qubit. By combining these viewpoints we hope to give the reader intuition about how such hybrid gates open up new possibilities for QSP [49,274–278] beyond those available in qubit-only and oscillator-only systems.

We note that by including measurements in a circuit, the concept of universal control can be extended to quantum channels that describe both unitary as well as nonunitary transformations of oscillators. Nonunitary or irreversible maps can be achieved by entangling the oscillator with an auxiliary qubit followed by measurement or reset of the qubit. In Ref. [279] the authors have shown that efficient universal channel construction is possible using a single auxiliary qubit with QND readout and adaptive control. It is admissible to replace readout and adaptive control with auxiliary-qubit-assisted gates followed by a reset of the qubits (see, e.g., Fig. 7). The same technique can be used to efficiently construct channels for oscillator systems.

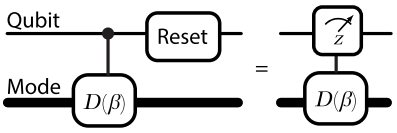


FIG. 7. Conditional displacement (see Table IV) followed by a reset of the qubit is equivalent to the measurement of the qubit followed by feedback displacement conditioned on the measurement result. In the first circuit, the information about which state was produced ends up in the environment (due to the reset). In the second circuit, the experimentalist may have access to the information from the measurement result. The equivalence is strictly true only if the measurement result is not available for subsequent use in the circuit.

A. Gaussian oscillator control

Building upon the visualization of classical flows in phase space in Sec. III C 3, we begin our discussion of Gaussian control with its extension to the quantum case, which involves unitary transformations in Hilbert space. Section IV A 1 gives a general discussion about Gaussian oscillator control using quadratic Hamiltonians, which induce linear (symplectic) transformations on the phase space; Sec. IV A 2 discusses the composition of such CV transformations.

1. Basics of Gaussian control

Gaussian states are defined to be those the wave functions $\Psi(x_1,\ldots,x_N)$ of which are Gaussians. The ground state of a harmonic oscillator is a prototypical Gaussian state. Gaussian unitaries are defined to be those generated by Hamiltonians that are at most quadratic in the position and momentum operators $\hat{x}_i$ and $\hat{p}_i$ for the various modes (or equivalently, at most quadratic in the creation and annihilation operators a_i and $a_i^\dagger$) such as the quantum version of the classical squeezing Hamiltonian given in Eq. (62). We return to this specific example shortly and see that all Gaussian unitaries applied to Gaussian states preserve the Gaussian nature of such states. But first, we begin with some general remarks on the quantum version of phase space where the position and momentum operators do not necessarily commute.

Given N oscillators, let us define the quantum operator phase-space coordinate vector of length $2N$:

$$\hat{\mathbf{R}} = \left(\hat{x}_1,\hat{p}_1,\hat{x}_2,\hat{p}_2,\ldots,\hat{x}_N,\hat{p}_N\right). \tag{94}$$

The most general quadratic Hamiltonian H_G can be written as

$$H_{\mathrm{G}} = \sum_{i=1}^{2N}\lambda_i\hat{R}_i + \frac{1}{2}\sum_{i,j}^{2N}\hat{R}_i\Lambda_{ij}\hat{R}_j, \tag{95}$$

where the vector $\boldsymbol{\lambda}$ is real and the matrix $\boldsymbol{\Lambda}$ is real and symmetric.

It is useful to define the antisymmetric matrix

$$\boldsymbol{\Omega} = \bigoplus_N \begin{pmatrix} 0 & +1 \\ -1 & 0 \end{pmatrix}, \tag{96}$$

which allows us to represent the commutator of different elements of $\hat{\mathbf{R}}$ as

$$[\hat{R}_j,\hat{R}_k] = +i\Omega_{jk}. \tag{97}$$

All Gaussian unitaries take the form

$$U_G(t) = \exp\left(-iH_{\mathrm{G}}t\right) \tag{98}$$

and (since the Heisenberg equation of motion is linear) produce a linear transformation of the form

$$\hat{\mathbf{R}}(t) = U_G^\dagger(t)\hat{\mathbf{R}}(0)U_G(t) \tag{99}$$

$$= \boldsymbol{\Delta}(t) + \mathbf{M}(t)\hat{\mathbf{R}}(0), \tag{100}$$

where $\boldsymbol{\Delta}$ is a real-valued phase-space displacement vector and $\mathbf{M}(t)$ is a real matrix. It is important to note that Eq. (99) describes element-wise conjugation of $\hat{\mathbf{R}}(0)$ by the unitary. It does not represent matrix multiplication of matrices with compatible dimensions. By contrast, the second term in Eq. (100) does represent ordinary multiplication of the column vector $\hat{\mathbf{R}}(0)$ by the matrix $\mathbf{M}(t)$.

In order to preserve the commutation relation in Eq. (97), $\mathbf{M}(t)$ must obey

$$\mathbf{M}^{\mathrm{T}}(t)\boldsymbol{\Omega}\mathbf{M}(t) = \boldsymbol{\Omega}, \tag{101}$$

meaning that $\mathbf{M}(t)\in \mathrm{Sp}(2N,\mathbb{R})$ must be a real symplectic matrix.

Equation (99) is solved using the Heisenberg equation of motion,

$$\frac{d\hat{\boldsymbol{R}}(t)}{dt} = +i[H_{\mathrm{G}},\hat{\boldsymbol{R}}(t)]. \tag{102}$$

Two limiting cases are particularly simple:

Case I: $H_{\mathrm{G}} = \boldsymbol{\lambda}\hat{\boldsymbol{R}}$

$$\frac{d\hat{\boldsymbol{R}}(t)}{dt} = \boldsymbol{\Omega}\boldsymbol{\lambda}, \tag{103}$$

$$\hat{\boldsymbol{R}}(t) = \hat{\boldsymbol{R}}(0) + \boldsymbol{\Delta}(t), \tag{104}$$

$$\boldsymbol{\Delta}(t) = \boldsymbol{\Omega}\boldsymbol{\lambda}t. \tag{105}$$

Case II: $H_{\rm G} = \frac{1}{2}\hat{\boldsymbol{R}}\boldsymbol{\Lambda}\hat{\boldsymbol{R}}$

$$\frac{d\hat{\boldsymbol{R}}(t)}{dt} = \boldsymbol{\Omega}\boldsymbol{\Lambda}\hat{\boldsymbol{R}}(t), \tag{106}$$

$$\hat{\boldsymbol{R}}(t) = \boldsymbol{M}(t)\hat{\boldsymbol{R}}(0), \tag{107}$$

$$\boldsymbol{M}(t) = \exp\left(\boldsymbol{\Omega}\boldsymbol{\Lambda}t\right). \tag{108}$$

It is straightforward to show that this exponential form of $\boldsymbol{M}(t)$ is indeed symplectic and obeys Eq. (101). The general case of the Hamiltonian in Eq. (95) can be reduced to case II by simply completing the square, yielding

$$\hat{\boldsymbol{R}}(t) = -\boldsymbol{\Lambda}^{-1}\boldsymbol{\lambda} + \exp\left(\boldsymbol{\Omega}\boldsymbol{\Lambda}t\right)\left(\hat{\boldsymbol{R}}(0) + \boldsymbol{\Lambda}^{-1}\boldsymbol{\lambda}\right). \tag{109}$$

It is sometimes convenient to represent the symplectic transformations not in terms of the Hermitian position and momentum operators, $\hat{x}$ and $\hat{p}$, but rather, in terms of the complex mode operators, a and $a^\dagger$. In this basis, the symplectic transformation matrices can be complex. We note that the displacement vector $\boldsymbol{\Delta}$ does not affect the commutation relations. Further details can be found in Refs. [222,280].

As an example of case I for a single oscillator, the simplest nontrivial control Hamiltonian is the linear Hamiltonian

$$H_G = \lambda_p\hat{p} + \lambda_x\hat{x}, \tag{110}$$

for which the Heisenberg equations of motion are (in standard units)

$$\frac{d}{dt}\hat{x} = +\lambda_p, \tag{111}$$

$$\frac{d}{dt}\hat{p} = -\lambda_x, \tag{112}$$

which yields the displacement canonical transformation with

$$\boldsymbol{\Delta}(t) = \boldsymbol{\Omega}\boldsymbol{\lambda}t = (+\lambda_p, -\lambda_x)t, \tag{113}$$

$$\mathbf{M}(t) = \hat{I}, \tag{114}$$

in Eq. (100), as illustrated in Fig. 8.

As a second simple example, this time for case II, let us examine (in standard units) the quantum version of the squeezing Hamiltonian in Eq. (62),

$$H_G = \frac{g}{2}[\hat{x}\hat{p} + \hat{p}\hat{x}], \tag{115}$$

which must be symmetrized for it to be Hermitian and corresponds to the case

$$\boldsymbol{\lambda} = \mathbf{0}, \tag{116}$$

$$\boldsymbol{\Lambda} = g\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}. \tag{117}$$

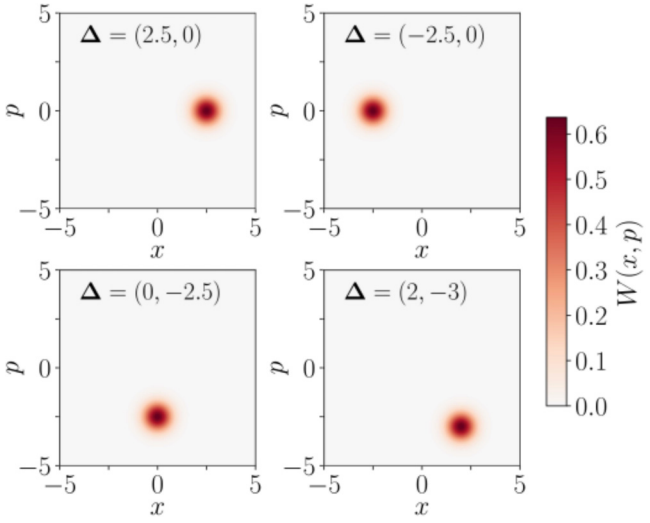


FIG. 8. The Wigner quasiprobability distribution after evolution under the linear displacement Hamiltonian in Eq. (110) for various values of $\boldsymbol{\Delta} = (+\lambda_p t, -\lambda_x t)$.

The solution of the Heisenberg equation of motion is identical to the classical case discussed above:

$$\boldsymbol{\Delta} = \mathbf{0}, \tag{118}$$

$$\mathbf{M} = \exp\left(\boldsymbol{\Omega}\boldsymbol{\Lambda}t\right) = \begin{pmatrix} e^{gt} & 0 \\ 0 & e^{-gt} \end{pmatrix}, \tag{119}$$

$$\hat{x}(t) = e^{+gt}\hat{x}(0), \tag{120}$$

$$\hat{p}(t) = e^{-gt}\hat{p}(0), \tag{121}$$

and preserves the (standard units) commutation relation

$$[\hat{x}(t), \hat{p}(t)] = [\hat{x}(0), \hat{p}(0)] = +i. \tag{122}$$

Reverting to the Schrödinger picture and time evolving the vacuum state under the squeezing Hamiltonian yields the so-called "squeezed vacuum,"

$$|\psi(t)\rangle = e^{-iHt}|0\rangle. \tag{123}$$

The Wigner (quasiprobability) function (defined in Appendix E 6) of this state is shown in Fig. 9 for the same parameter values as in Fig. 4. Note the close similarity of the shape of the Wigner function to the shape of the circular region of classical phase space that has been squeezed into an ellipse in Fig. 4(b). In the classical case, the initial circular region has been chosen arbitrarily. For the Wigner function, the area of the initial circular region and the subsequent squeezed regions is fixed by the size of the (squeezed) vacuum fluctuations given in Eqs. (155) and (156).

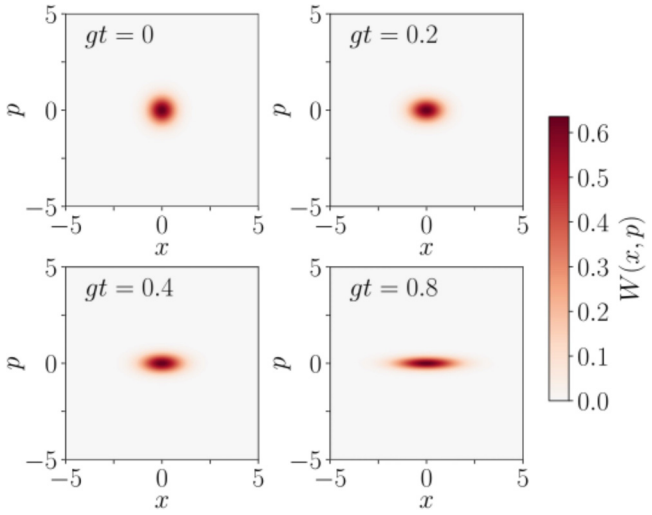


FIG. 9. The Wigner quasiprobability distribution of an initial vacuum state evolved under the squeezing Hamiltonian in Eq. (117) for the same values of gt used in Fig. 4(b).

2. Composition of two Gaussian transformations

A subtlety arises when evaluating Eq. (100) for the case of applying two successive Gaussian transformations, U_1 and U_2, that individually obey

$$\hat{\mathbf{R}}_1 = U_1^\dagger \hat{\mathbf{R}}_0 U_1 = \boldsymbol{\Delta}_1 + \mathbf{M}_1 \hat{\mathbf{R}}_0, \tag{124}$$

$$\hat{\mathbf{R}}_2 = U_2^\dagger \hat{\mathbf{R}}_0 U_2 = \boldsymbol{\Delta}_2 + \mathbf{M}_2 \hat{\mathbf{R}}_0. \tag{125}$$

To illustrate this point, for simplicity, we specialize to the case $\boldsymbol{\Delta}_1 = \boldsymbol{\Delta}_2 = \mathbf{0}$ and consider the sequential transformation

$$\hat{\mathbf{R}}_3 = U_1^\dagger U_2^\dagger \hat{\mathbf{R}}_0 U_2 U_1. \tag{126}$$

In the Schrödinger picture, we can understand this sequential transformation as an initial active transformation of the state $|\Psi\rangle$ to a new state $U_1|\Psi\rangle$, followed by a subsequent active transformation resulting in the final state $U_2(U_1|\Psi\rangle)$. In the Heisenberg picture, however, one can understand $U_1^\dagger U_2^\dagger \hat{\mathbf{R}}_0 U_2 U_1$ as carrying out a sequence of passive transformations to change the basis. Naively, Eq. (126) appears to tell us to carry out the passive transformations in reverse order: U_2 first and then U_1 so that the net result would be $\hat{\mathbf{R}}_3 = \mathbf{M}_1\mathbf{M}_2\hat{\mathbf{R}}_0$. However, this is incorrect, as can be seen by rewriting Eq. (126) in the form recommended by Schwinger [281]:

$$\hat{\mathbf{R}}_3 = [U_1^\dagger U_2^\dagger U_1][U_1^\dagger \hat{\mathbf{R}}_0 U_1][U_1^\dagger U_2 U_1]. \tag{127}$$

Here, the interpretation is that one first applies the passive transformation U_1 followed by the passive transformation U_2 (in the new basis generated by the first passive transformation). Thus the correct final expression is

$$\hat{\mathbf{R}}_3 = \mathbf{M}_2\mathbf{M}_1\hat{\mathbf{R}}_0, \tag{128}$$

showing that the symplectic transformations of the operators are applied in the same order as the (active) unitary transformations U_2U_1 are applied to the state.

B. Gaussian operations

A generating set of Gaussian gates [282] is given by displacement (Sec. IV B 1), phase-space rotation which turns out to be related to the Fourier transform (Sec. IV B 2), single-mode squeezing (Sec. IV B 3), and beam splitting (Sec. IV B 4) or two-mode squeezing (Sec. IV B 5); see Table III. As noted in the caption of Table IX, all of the Gaussian operations on oscillators that we discuss here are natively available in the superconducting platform. Compilation of these primitive gates into arbitrary Gaussian unitaries can be efficiently performed using the Bloch-Messiah decomposition of the corresponding symplectic transformation matrices (see Appendix B) and other standard techniques [222,300,301]. An over-complete list of primitive operations that can be used to generate arbitrary multimode Gaussian operations is the following:

(a) displacement
(b) Fourier transform
(c) single-mode squeezing
(d) beam splitter
(e) two-mode squeezing

Displacements are used to generate the elements of the Heisenberg-Weyl group, H(n), while the next three are used to generate unitaries corresponding to the $\mathbf{M} \in \mathrm{Sp}(2n, \mathbb{R})$ in Eq. (101) generated by the algebra of the Lie brackets of the homogeneous quadratic polynomials. In this section, we elaborate on these operations and the decomposition of a few other well-known Gaussian operations mentioned in Table III in terms of the primitives. For convenience, we switch here from using position and momentum operators to creation and annihilation operators.

1. Displacement gate

The phase-space translation operation described above and listed in Box IV.1 is a very important primitive for any bosonic mode instruction set. The displacement in phase space is parametrized by the complex number α and the gate has the following properties:

$$D^\dagger(\alpha) = D(-\alpha), \tag{129}$$

$$D^\dagger(\alpha) a D(\alpha) = a + \alpha. \tag{130}$$

From this and Eq. (38) it follows that $D(+\alpha)$ is a translation in phase space by an amount (in Wigner units)

TABLE III. Common oscillator-only gates. This table is not comprehensive; however, it contains some of the most common gates on bosonic oscillators. While for simplicity we list these as oscillator-only gates, the experimental realization of many of them requires the use of an auxiliary nonlinear system (e.g., a qubit) for their execution.

		Gate name	Parameters	Definition	Reference theory or experiment
Oscillator-only gates	Gaussian (Sec. IV B)	Displacement (Box IV.1)	$\alpha \in \mathbb{C}$	$D(\alpha) = \exp\left[\alpha a^\dagger - \alpha^* a\right]$	[283]/[29]
		Phase-space rotation (Box IV.2)	$\theta \in [0, 2\pi)$	$R(\theta) = \exp\left[-i\theta a^\dagger a\right]$	[222]/
		Fourier transform	$\theta = \pi/2$	$F = \exp\left[-i\frac{\pi}{2} a^\dagger a\right]$	
		(Single-mode) squeezing (Box IV.3)	$\zeta \in \mathbb{C}$	$S(\zeta) = \exp\left[\frac{1}{2}(\zeta^* a^2 - \zeta a^{\dagger 2})\right]$	[283]/[284]
		beam splitter (Box IV.4)	$\theta \in [0, 4\pi), \varphi \in [0, \pi)$	$\mathrm{BS}(\theta, \varphi) = \exp\left[-i\frac{\theta}{2}\left(e^{i\varphi} a^\dagger b + e^{-i\varphi} a b^\dagger\right)\right]$	[283,285]/[70, 237]
		SWAP	$\theta = \pi, \varphi = \pi/2$	$\mathrm{SWAP} = \exp\left[\frac{\pi}{2}(a^\dagger b - a b^\dagger)\right]$	
		Two-mode squeezing (Box IV.5)	$r \in \mathbb{R}, \varphi \in [0, \pi)$	$\mathrm{TMS}(r, \varphi) = \exp\left[r\left(e^{i\varphi} a^\dagger b^\dagger - e^{-i\varphi} a b\right)\right]$	[283]/[286]
		Two-mode summing (Box IV.6)	$\lambda \in \mathbb{R}$	$\mathrm{SUM}(\lambda) = \exp\left[\frac{1}{2}\lambda\left(a + a^\dagger\right)\left(b^\dagger - b\right)\right]$	[219]/[287,288]
	Non-Gaussian (Sec. IV D)	Generalized single-mode squeezing[a]	$z \in \mathbb{C}, N \geq 3$	$U_N(z) = \exp\left[z a^{\dagger N} - z^* a^N\right]$	[289]/[290,291]
		Generalized k-mode squeezing	$z \in \mathbb{C}, k \geq 3\ \vec{n} = \{n_1, \ldots, n_k\}$	$U_{\vec{n}}(z) = \exp\left[z \prod_{p=1}^k a_p^{\dagger n_p} - z^* \prod_{p=1}^k a_p^{n_p}\right]$	[292]/[290,293]
		Cubic phase	$r \in \mathbb{R}$	$C(r) = \exp\left[-ir\hat{x}^3\right]$	[294,295]/[267, 290,291]
		Self-Kerr	$\theta \in \mathbb{R}$	$K(\theta) = \exp\left[-i\frac{\theta}{2} a^{\dagger 2} a^2\right]$	[217,283]/[296]
		SNAP[b] (Box IV.10)	$\vec{\varphi} = \{\varphi_n\}, \varphi_n \in [0, 2\pi)$	$\mathrm{SNAP}(\vec{\varphi}) = \sum_n e^{-i\varphi_n} \lvert n\rangle \langle n\rvert$	[59]/[266]
		eSWAP (Box IV.11)	$\theta \in [0, 4\pi)$	$\exp\left[-i\frac{\theta}{2}\,(\mathrm{SWAP})\right]$	[297]/[298]

[a]Generalized squeezing is non-Gaussian for $N \geq 3$. U_3 is often called *trisqueezing*, and U_4 can be called *quadsqueezing*.

[b]SNAP stands for *S*elective *N*umber-dependent *A*rbitrary *P*hase gate. The SNAP gate is implemented with a sequence of SQR gates (defined in Box IV.9) using an auxiliary qubit, here not included for clarity as the oscillator and qubit are left unentangled if the qubit starts in the computational basis. However, the sign of the phase imparted is dependent upon the qubit state, and as such the SNAP gate is more completely expressed as $\mathrm{SNAP}(\vec{\varphi}) = \sum_n e^{-i\varphi_n \sigma_z} \lvert n\rangle \langle n\rvert$. The simple phase-space rotation gate $R(\theta)$ can be performed "virtually" in software by simply resetting the phase of subsequent drives on the oscillator in analogy with the virtual Z gate used in qubit systems [299].

$\Delta x = \mathrm{Re}\,\alpha$ and $\Delta p = \mathrm{Im}\,\alpha$. From Eq. (130) it follows that the coherent state $|\alpha\rangle$ is simply the displaced vacuum state

$$|\alpha\rangle = D(\alpha)|0\rangle, \tag{131}$$

and more generally, that displacement maps coherent states to coherent states

$$D(\alpha)|\beta\rangle = |\alpha+\beta\rangle. \tag{132}$$

The normal-ordered form of the displacement operator in Eq. (134) is preferred for numerical calculations where the Hilbert space of the oscillator is truncated to the lowest $d = N_{\max}+1$ states.

Box IV.1. Displacement Gate

$$D(\alpha) \equiv e^{\alpha a^\dagger - \alpha^* a} \tag{133}$$

$$= e^{-\frac{\alpha^*\alpha}{2}} e^{\alpha a^\dagger} e^{-\alpha^* a}, \tag{134}$$

Composition of displacement gates:

$$D(\beta)D(\alpha) = D(\alpha+\beta)e^{+\frac{1}{2}[\beta\alpha^* - \beta^*\alpha]}, \tag{135}$$

$$D(\alpha)D(\beta) = D(\alpha+\beta)e^{-\frac{1}{2}[\beta\alpha^* - \beta^*\alpha]}, \tag{136}$$

$$= D(\beta)D(\alpha)e^{-[\beta\alpha^* - \beta^*\alpha]}, \tag{137}$$

$$= D(\beta)D(\alpha)e^{-2iA(\alpha,\beta)}, \tag{138}$$

where $A(\alpha,\beta)$ is the oriented area of the parallelogram formed by the displacements α and β.

In terms of the phase-space coordinates, the multimode displacement gate is given by Eq. (98) with only linear terms in the Hamiltonian, and the resulting displacement is given by Eq. (105).

Mathematically, the composition (multiplication) of two translations gives another translation, and hence these operations form a group known as the Weyl-Heisenberg group. Because the generators $\hat{x}$ and $\hat{p}$ (or equivalently, a and $a^\dagger$) do not commute, the group is non-Abelian. In analyzing the structure of the group of translations, it is helpful to use the following identity:

$$e^{B+C} = e^{-\frac{1}{2}[B,C]} e^B e^C, \tag{139}$$

which is valid provided that both B and C commute with their commutator $[B,C]$. From this, the displacement composition rules shown in Eqs. (135)–(138) follow, with $A(\alpha,\beta)$ being the oriented area of the parallelogram in the complex plane defined by α and β, as shown in Fig. 10.

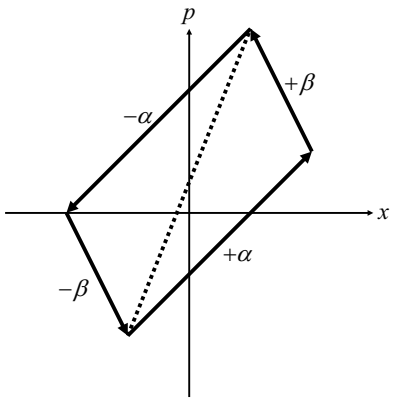


FIG. 10. Traversing a right-handed closed path of area A in phase space yields a geometric phase factor $e^{2iA} = e^{\alpha^*\beta - \alpha\beta^*}$ (in the dimensionless Wigner units that we are using).

In the Wigner units that we have chosen, if this area is an integer multiple of π (2π in standard dimensionless units), the two displacements commute. Using Eq. (129) and the composition rules, we see that if we take the system along a right-handed loop around the parallelogram, the system returns to its original state up to a geometric phase

$$D^\dagger(\beta)D^\dagger(\alpha)D(\beta)D(\alpha) = e^{+2iA(\alpha,\beta)}\hat{I}. \tag{140}$$

This geometric phase plays a crucial general role in the control of bosonic modes and a particular role in the Gottesman-Kitaev-Preskill (GKP) quantum error-correction code [282,302]. Separately, it can be used to compile entangling gates between physical qubits, as we explain in Sec. VI C 2.

2. Phase-space rotation gate

Another simple gate for a cavity is the phase-space rotation gate given in Box IV.2.

The phase-space rotation gate has the following properties:

$$R(\theta)|\alpha\rangle = |e^{-i\theta}\alpha\rangle, \tag{141}$$

$$R^\dagger(\theta) = R(-\theta), \tag{142}$$

$$R^\dagger(\theta)aR(\theta) = e^{-i\theta}a, \tag{143}$$

$$R^\dagger(\theta)a^\dagger R(\theta) = e^{+i\theta}a^\dagger, \tag{144}$$

$$R^\dagger(\theta)\hat{x}R(\theta) = \cos\theta\hat{x} + \sin\theta\hat{p}, \tag{145}$$

$$R^\dagger(\theta)\hat{p}R(\theta) = \cos\theta\hat{p} - \sin\theta\hat{x}, \tag{146}$$

corresponding to a clockwise rotation of the state in phase space by angle θ via the symplectic transformation

$$\begin{pmatrix}\hat{x}(\theta)\\ \hat{p}(\theta)\end{pmatrix} = \begin{pmatrix}\cos\theta & +\sin\theta\\ -\sin\theta & \cos\theta\end{pmatrix}\begin{pmatrix}\hat{x}\\ \hat{p}\end{pmatrix}. \tag{147}$$

The rotation gate $R(\pi/2)$ executes a quantum Fourier transform from the position basis to the momentum basis

by simply rotating the state in phase space [223]. If we view the oscillator motion in the nonrotating laboratory frame, we see that the fact that the potential energy of an oscillator is quadratic in the position coordinate means that if one simply waits for a quarter cycle of the oscillatory motion, the position is precisely converted into (negative) momentum. This is true quantum mechanically as well as classically. For example, a position-squeezed state (described in Sec. IV B 3) that has highly uncertain momentum is converted upon rotation in phase space by $\pi/2$ into a state with definite momentum and highly uncertain position. See Sec. VII B for applications of the quantum Fourier transform gate.

Box IV.2. Phase-Space Rotation Gate

$$R(\theta) = e^{-i\theta\hat{n}}. \tag{148}$$

Quantum Fourier transform gate:

$$F = R\left(\frac{\pi}{2}\right) = e^{-i\frac{\pi}{2}\hat{n}}. \tag{149}$$

$$F^\dagger \hat{x} F = +\hat{p} \tag{150}$$

$$F^\dagger \hat{p} F = -\hat{x}. \tag{151}$$

Note that free evolution in the laboratory frame under Hamiltonian H_0 in Eq. (21) for a time t yields a clockwise rotation by $\theta = \omega_R t$. Consequently, it is not necessary to apply any perturbation or drive to the cavity to carry out the gate. Instead, one simply needs to shift the time reference of the control system. This shift of the time reference is equivalent to a shift in the phase of the carrier wave at the resonator frequency used to effect displacement gates. Thus one simply has to keep track at the classical level (i.e., "in software") that any subsequent drive amplitude applied to the resonator is phase shifted by θ and any instruction set can readily capture this type of *in silico* gate (also known as a "virtual" gate).

3. Single-mode squeezing gate

Squeezing is a fundamental quantum resource for CV systems [218]. The single-mode squeezing gate in Box IV.3 can be used to reduce the quantum fluctuations in the oscillator position by carrying out the transformation

$$\hat{x}_r = e^{-r}\hat{x}, \tag{152}$$

$$\hat{p}_r = e^{+r}\hat{p}. \tag{153}$$

It is clear that this is a canonical (i.e., symplectic) transformation because it antisqueezes the momentum by just the right amount to preserve the commutation relation (here in Wigner units),

$$[\hat{x}_r, \hat{p}_r] = [\hat{x}, \hat{p}] = +\frac{i}{2}. \tag{154}$$

As a result, squeezed states saturate the Heisenberg inequality by minimizing the product of the variances of position and momentum,

$$\langle 0|\hat{x}_r^2|0\rangle - \langle 0|\hat{x}_r|0\rangle^2 = \frac{1}{4}e^{-2r}, \tag{155}$$

$$\langle 0|\hat{p}_r^2|0\rangle - \langle 0|\hat{p}_r|0\rangle^2 = \frac{1}{4}e^{+2r}, \tag{156}$$

just as coherent states do in Eqs. (42) and (43).

Unlike the displacement transformation, which is generated by terms linear in a and $a^\dagger$, the squeezing transformation is generated by quadratic terms

$$s(r) = e^{\frac{1}{2}r(a^2 - a^{\dagger 2})}, \tag{157}$$

$$= e^{ir(\hat{x}\hat{p} + \hat{p}\hat{x})}. \tag{158}$$

Note that the latter form matches the time evolution using the Hamiltonian in Eq. (115) with $r = -2gt$. These quadratic terms ("two-photon driving") can be generated by the nonlinear quantum optics technique of three- or four-wave mixing, or through dissipation engineering [60, 268,269,284,303,304]. The current state of the art in superconducting circuits yields squeezing up to factors $e^{2r} \sim 10$ dB [60], where 1 dB means a factor of $10^{1/10}$ and n dB represents a factor of $10^{n/10}$. In superconducting circuits, the maximum squeezing is limited in part by the self-Kerr effect (weak anharmonicity induced in the resonator by the presence of the qubit), an effect we have neglected here.

Box IV.3. Single-Mode Squeezing Gate

$$S(r,\theta) = e^{+i\theta\hat{n}} s(r) e^{-i\theta\hat{n}} \tag{159}$$

$$= e^{\frac{1}{2}(\zeta^* a^2 - \zeta a^{\dagger 2})}, \tag{160}$$

where $\zeta = re^{i2\theta}$ and $s(r)$ is defined in Eq. (158). This operation performs the transformation

$$(+\cos\theta\hat{x} + \sin\theta\hat{p}) \to e^{-r}(+\cos\theta\hat{x} + \sin\theta\hat{p}), \tag{161}$$

$$(-\sin\theta\hat{x} + \cos\theta\hat{p}) \to e^{+r}(-\sin\theta\hat{x} + \cos\theta\hat{p}). \tag{162}$$

The symplectic transformation matrix (in the mode basis) induced by $s(r)$ is given in Eq. (167).

To verify that the unitary in Eq. (158) is correct, let

$$a_r = s^\dagger(r)as(r), \tag{163}$$

$$a_r^\dagger = s^\dagger(r)a^\dagger s(r), \tag{164}$$

which obey

$$\frac{d}{dr}\begin{pmatrix} a_r \\ a_r^\dagger \end{pmatrix} = -\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}\begin{pmatrix} a_r \\ a_r^\dagger \end{pmatrix}, \tag{165}$$

and have the solution

$$\begin{pmatrix} a_r \\ a_r^\dagger \end{pmatrix} = \mathbf{M}_s(r)\begin{pmatrix} a \\ a^\dagger \end{pmatrix}, \tag{166}$$

where the symplectic mode transformation matrix is given by

$$\mathbf{M}_s(r) = \begin{pmatrix} \cosh r & -\sinh r \\ -\sinh r & \cosh r \end{pmatrix}. \tag{167}$$

These results in turn yield Eqs. (152) and (153), showing that the symplectic transformation matrix in the position or momentum basis is diagonal. The unitary transformation in Eq. (158) squeezes the position fluctuations of the oscillator. As shown in Box IV.3, this transformation can be generalized to squeeze any quadrature (i.e., any linear combination of $\hat{x}$ and $\hat{p}$) by simply applying a rotation transformation. Note that rotating the squeezing axis by $\theta = \pi$ returns us to the original squeezed state with $\theta = 0$ as it should.

4. Two-mode beam-splitter gates

Recent experimental progress using parametric modulation of the coupling between two bosonic modes [70,237] has allowed the high-fidelity realization of Hamiltonians which are bilinear in operators for the two modes. The first of these is the beam-splitter Hamiltonian

$$\hat{V}_{\rm BS} = \frac{1}{2}(\chi^* ab^\dagger + \chi a^\dagger b). \tag{168}$$

Evolution under this Hamiltonian for time t with $\chi t = \theta e^{i\varphi}$ yields the beam-splitter unitary gate given in Box IV.4. The corresponding 4×4 symplectic mode operator transformation matrix is given in Eq. (B6), found in Appendix B.

Box IV.4. Beam-Splitter Gate

$$\mathrm{BS}(\theta,\varphi) = e^{-i\frac{\theta}{2}\left[e^{i\varphi}a^\dagger b + e^{-i\varphi}ab^\dagger\right]}. \tag{169}$$

In the interaction picture, the beam-splitter gate transforms the two bosonic mode operators as follows:

$$a_{\theta,\varphi} = \mathrm{BS}^\dagger(\theta,\varphi)a\,\mathrm{BS}(\theta,\varphi) = \cos\frac{\theta}{2}a - i\sin\frac{\theta}{2}e^{+i\varphi}b, \tag{170}$$

$$b_{\theta,\varphi} = \mathrm{BS}^\dagger(\theta,\varphi)b\,\mathrm{BS}(\theta,\varphi) = \cos\frac{\theta}{2}b - i\sin\frac{\theta}{2}e^{-i\varphi}a. \tag{171}$$

A key feature of the beam-splitter gate is that it preserves the total excitation number in the two cavities

$$[\mathrm{BS}(\theta,\varphi), \hat{n}_a + \hat{n}_b] = 0. \tag{172}$$

This fact plays an important role in the Schwinger-boson representation of spins discussed in Sec. VII C 6.

Note from Eqs. (170) and (171) that for $\theta = \pi$ and $\varphi = 0$,

$$a_{\pi,0} = -ib, \tag{173}$$

$$b_{\pi,0} = -ia, \tag{174}$$

the beam splitter swaps the two mode operators. Reverting to the Schrödinger picture, we find that this gate swaps the states of the two cavities,

$$\begin{aligned}\mathrm{BS}(\pi,0)|\Psi_a,\Psi_b\rangle &= e^{-i\frac{\pi}{2}[\hat{n}_a+\hat{n}_b]}\mathrm{SWAP}|\Psi_a,\Psi_b\rangle \\ &= e^{-i\frac{\pi}{2}[\hat{n}_a+\hat{n}_b]}|\Psi_b,\Psi_a\rangle,\end{aligned} \tag{175}$$

where the exponential of the total number operator in front accounts for the $(-i)$ factor in Eqs. (173) and (174). For convenience of notation, we have assumed here that the two cavities are in a product state, but by the linearity of

quantum mechanics, the gates perform equally well on all states (including states of indefinite total photon number) since any two-mode state can be written as a linear superposition of product states. A key and powerful feature of the bosonic SWAP gate is that it is agnostic as to the particular state of the cavities (e.g., the particular bosonic code being utilized to hold the quantum information). This is highly useful for routing and transporting bosonic qubits across the hardware fabric, a possibility that is used for compilation in Sec. VI C.

The reader should not confuse this gate with the so-called iSWAP gate, which is an entangling gate for two-level qubits. Importantly, one might think that $\mathrm{BS}(\pi/2,0)$ is related to the entangling gate $\sqrt{\mathrm{iSWAP}}$, but it is not. $\mathrm{BS}(\pi/2,0)$ is the 50:50 beam splitter familiar from optics. This gate has been used experimentally to demonstrate the Hong-Ou-Mandel effect for microwave photons [305,306], in which two identical photons incident in each input port of a 50:50 beam splitter always exit from the same port (or in this case end up in the same resonator):

$$\mathrm{BS}(\pi/2,0)|1\rangle|1\rangle = -\frac{1}{\sqrt{2}}\left[|0\rangle|2\rangle + |2\rangle|0\rangle\right]. \tag{176}$$

If one encodes quantum information (i.e., a qubit) in the $n=0,1$ Fock states of each cavity,

$$|\Psi\rangle = \alpha|0\rangle + \beta|1\rangle, \tag{177}$$

we see that the Hong-Ou-Mandel effect can take the system out of that code space by populating the $n=2$ Fock states (while still preserving the total photon number in the two modes). Thus, this cannot be related to $\sqrt{i\mathrm{SWAP}}$ despite the fact that applying $\mathrm{BS}(\pi/2,0)$ twice does indeed yield $\mathrm{BS}(\pi,0)$, which, up to phase factors, is SWAP.

5. Two-mode squeezing gate

Another useful bilinear interaction is the two-mode squeezing (TMS) interaction [307]

$$\hat{V}_{\mathrm{TMS}} = i g_{\mathrm{TMS}}\left(e^{i\varphi}a^\dagger b^\dagger - e^{-i\varphi}ab\right). \tag{178}$$

Time evolution under this interaction yields the two-mode squeezing gate in Box IV.5. From the Heisenberg equations of motion,

$$\frac{d}{dt}\begin{pmatrix} a(t) \\ b^\dagger(t)\end{pmatrix} = g_{\mathrm{TMS}}\begin{pmatrix} 0 & e^{i\varphi} \\ e^{-i\varphi} & 0\end{pmatrix}\begin{pmatrix} a(t) \\ b^\dagger(t)\end{pmatrix}, \tag{179}$$

we see that the mode operators transform as

$$\begin{pmatrix} a(t) \\ b^\dagger(t)\end{pmatrix} = \left\{\cosh r\hat{I} + \sinh r\begin{pmatrix} 0 & e^{i\varphi} \\ e^{-i\varphi} & 0\end{pmatrix}\right\}\begin{pmatrix} a(0) \\ b^\dagger(0)\end{pmatrix}, \tag{180}$$

where $r = g_{\mathrm{TMS}}t$. From this result, we can obtain the symplectic transformation matrix $\mathbf{TMS}(r,\varphi)$ (here expressed in the basis of mode creation and annihilation operators rather than positions and momenta) corresponding to the TMS unitary,

$$\begin{pmatrix} a(t) \\ a^\dagger(t) \\ b(t) \\ b^\dagger(t)\end{pmatrix} = \mathbf{TMS}(r,\varphi)\begin{pmatrix} a \\ a^\dagger \\ b \\ b^\dagger\end{pmatrix}, \tag{181}$$

where

$$\mathbf{TMS}(r,\varphi) = \begin{pmatrix} \cosh r & 0 & 0 & e^{i\varphi}\sinh r \\ 0 & \cosh r & e^{-i\varphi}\sinh r & 0 \\ 0 & e^{i\varphi}\sinh r & \cosh r & 0 \\ e^{-i\varphi}\sinh r & 0 & 0 & \cosh r\end{pmatrix}. \tag{182}$$

Box IV.5. Two-Mode Squeezing Gate

$$\mathrm{TMS}(r,\varphi) = e^{-i\hat{V}_{\mathrm{TMS}}t} = e^{r\left(e^{i\varphi}a^\dagger b^\dagger - e^{-i\varphi}ab\right)}, \tag{183}$$

where $r = g_{\mathrm{TMS}}t$. The corresponding symplectic mode transformation matrix $\mathbf{TMS}(r,\varphi)$ is given in Eq. (182).

The two-mode squeezed vacuum state [283],

$$\mathrm{TMS}(r,\varphi)|0,0\rangle = \frac{1}{\cosh r}\sum_{m=0}^{\infty}\left[e^{i\varphi}\tanh r\right]^m|m,m\rangle, \tag{184}$$

is an entangled state of two modes characterized by having equal photon numbers in both modes. For $\varphi = 0$ one can think of this entangled state as the CV analog of the DV Bell state

$$|\mathrm{Bell}\rangle = \frac{1}{\sqrt{2}}\left[|00\rangle + |11\rangle\right], \tag{185}$$

which is characterized by an equal excitation number in both qubits.

The interferometer circuit in Fig. 11 realizes the Bloch-Messiah decomposition [222,300,301] of the two-mode squeezing gate using a pair of beam splitters and single-mode squeezers as follows:

$$\mathrm{TMS}(r,\pi/2) = \mathrm{BS}(\pi/2,0)[s(r)\otimes s(r)]\mathrm{BS}(\pi/2,\pi). \tag{186}$$

Here, the gate symbols represent the full unitary operators acting on the Hilbert space, not the symplectic matrix

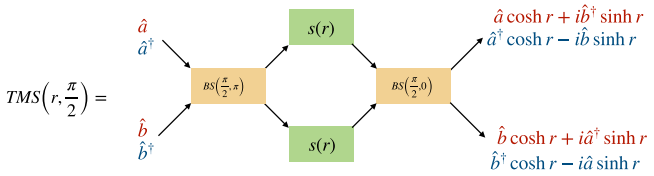


FIG. 11. The Bloch-Messiah decomposition of the two-mode squeezing TMS $(r,\pi/2)$ gate using photon-number-preserving beam-splitter gates and single-mode squeezing gates derived in Appendix B. The transformed mode operators on the right are given by the equations in Box IV.5. Here, a and b denote different oscillators. The effect of this operation is described by how the annihilation and creation operators are transformed. The red (blue) operators on the left are transformed to red (blue) operators on the right for each oscillator. Note that the terms in Eq. (186) are applied right to left, while the circuit diagram should be read left to right.

representation of the gates. The direct product $s(r)\otimes s(r)$ represents single-mode squeezing applied to each arm of the interferometer in Fig. 11. The derivation of this circuit is given in Appendix B, as is the symplectic transformation of the quadrature coordinates.

In the microwave domain, two-mode squeezing has been performed with parametric pumping of a nonlinear Josephson mixer to realize $\hat{V}_{\rm TMS}$ directly. To the best of our knowledge, the largest TMS coupling rate between two superconducting microwave oscillators achieved to date is $g_{\rm TMS}/2\pi = 12.6$ MHz using a pumped Josephson ring modulator [288]. Two-mode squeezing has also been demonstrated between phonons in surface-acoustic-wave (SAW) resonators [308]. We also note that two-mode squeezing of itinerant microwave photons is commonly used for quantum limited parametric amplification [286, 309–311]. An entanglement-enhanced interferometer can be constructed from a pair of two-mode squeezers and used to detect small displacements of the electromagnetic field [151–153] below the vacuum fluctuation limit.

6. *Two-mode* SUM *gate*

Let $|x_a\rangle|x_b\rangle$ be position eigenstates for two oscillators, a and b. As shown in Box IV.6, the two-mode SUM gate displaces one oscillator by an amount determined by the position of the other oscillator. The SUM gate and any other gate generated by bilinears in the oscillator operators can be created by simultaneous use of the beam splitter and two-mode squeezing interactions:

$$\hat{V} = \left[\chi ab^\dagger + \chi^* a^\dagger b\right] + i\left[\zeta^* a^\dagger b^\dagger - \zeta ab\right]. \tag{187}$$

For example, we see from Eq. (192) that the required Hamiltonian has the form

$$\begin{aligned}\hat{V} &= -ig[a+a^\dagger][b-b^\dagger]\\ &= +ig[ab^\dagger - a^\dagger b] + ig[a^\dagger b^\dagger - ab],\end{aligned} \tag{188}$$

for which the beam splitter and two-mode squeezing rates are identical. Examples of other gates that can be realized in this way include

$$S_{xx} = e^{i2\lambda \hat{x}_a \hat{x}_b}, \tag{189}$$

$$S_{pp} = e^{i2\lambda \hat{p}_a \hat{p}_b}, \tag{190}$$

$$S_{xp-px} = e^{i2\lambda[\hat{x}_a\hat{p}_b - \hat{x}_b\hat{p}_a]}. \tag{191}$$

Box IV.6. Two-Mode SUM Gate

$$\text{SUM}(\lambda) = e^{-i2\lambda \hat{x}_a \hat{p}_b} = e^{\frac{\lambda}{2}(a+a^\dagger)(b^\dagger - b)} \tag{192}$$

$$\text{SUM}(\lambda)|x_a\rangle|x_b\rangle = |x_a\rangle|x_b + \lambda x_a\rangle \tag{193}$$

where λ is an arbitrary real scale factor. The factor of 2 comes from the choice of Wigner units in which the generator of displacements is $2\hat{p}$, as shown in Eq. (58).

In the momentum representation, the action of the SUM gate is

$$\text{SUM}(\lambda)|p_a\rangle|p_b\rangle = |p_a - \lambda p_b\rangle|p_b\rangle \tag{194}$$

The SUM gate can be used to create Bell-like entangled states of the following form in the position representation:

$$\begin{aligned}&\text{SUM}(\lambda=1)\int_{-\infty}^{+\infty} dx\, \Psi(x)\,[|x\rangle\otimes S(r,0)|0\rangle]\\ &\approx \int_{-\infty}^{+\infty} dx\, \Psi(x)\,[|x\rangle\otimes|x\rangle],\end{aligned} \tag{195}$$

where we have assumed $r \gg 1$ so that the squeezed vacuum state $S(r,0)|0\rangle$ is approximately a position eigenstate centered at position 0.

The Bloch-Messiah decomposition for the SUM gate [312] has the advantage that two-mode squeezing can be replaced by simpler single-mode squeezing,

$$\begin{aligned}\text{SUM}(\lambda) &= \text{BS}(\pi+2\theta, -\pi/2)[s(r)\otimes s(-r)]\\ &\quad\times \text{BS}(2\theta, -\pi/2),\end{aligned} \tag{196}$$

$$\sinh r = \frac{\lambda}{2}, \tag{197}$$

$$\cos(2\theta) = \tanh(r), \tag{198}$$

$$\sin(2\theta) = -\text{sech}(r), \tag{199}$$

where we use the tensor-product ordering convention that $B\otimes A$ means that A is applied to the upper arm of the interferometer and B is applied to the lower arm. Thus, $s(-r)$ is

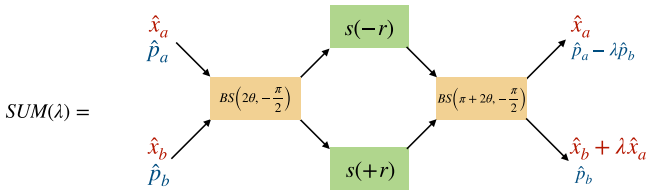


FIG. 12. The Bloch-Messiah decomposition of SUM(λ) gate [312] using photon-number-preserving beam-splitter gates and single-mode squeezing gates. The transformation from the initial mode quadrature operators on the left to the final quadratures on the right is given by the equations in Box IV.6. Here, a and b denote different oscillators. The effect of this operation is described by how the position and momentum operators are transformed. The red (blue) operators on the left are transformed into red (blue) operators on the right for each oscillator.

applied to the upper arm and $s(+r)$ is applied to the lower arm of the interferometer in Fig. 12. The derivation of this decomposition can be found in Appendix B.

C. Gaussian operations are analogous to Clifford gates

In this section, we discuss the analogies summarized in Table II between Gaussian operations on oscillators and Clifford operations on qubits. The analogies are close but not perfect, as we shall see in this pedagogical discussion, which combines some previously known facts and new insights.

Before considering CV gates, let us briefly review the classification of qubit gates according to the Clifford hierarchy, which is defined as follows:

$$C_1 P C_1^\dagger \in C_1, \tag{200}$$

$$C_n P C_n^\dagger \in C_{n-1}, \quad \text{for } n > 1, \tag{201}$$

$$C_n \subset C_{n+1}, \quad \forall n, \tag{202}$$

where C_1 is the set of *Pauli gates* $\{P\}$. The second level of the Clifford hierarchy C_2 is the *Clifford group* the elements of which conjugate a Pauli gate (i.e., a string of Pauli operators) into another Pauli gate; i.e., the Clifford group is the *normalizer* of the Pauli group. If a gate does not belong to C_2, it falls under the class of non-Clifford gates.

The Clifford group for an arbitrary number of qubits, N, can be generated by just three types of gates: H (the Hadamard gate), $S = Z^{1/2}$, and CNOT. Even though Clifford circuits can generate large superposition states such as maximally entangled Bell and GHZ states (using the CNOT gate), the Gottesman-Knill theorem [270,271] tells us that they are easy to simulate classically. To see this, consider an N-qubit starting state $|000\ldots000\rangle$. This state is "stabilized" (i.e., is a $+1$ eigenstate of) the set of N single-qubit Pauli Z operators, $\{Z_N, Z_{N-1}, \ldots, Z_1, Z_0\}$, and is uniquely defined by this list of stabilizers. Under the action of an arbitrary Clifford circuit, this set of stabilizers is mapped (in the Heisenberg picture) by conjugation to new Pauli strings (generally of weight higher than 1 under the action of CNOT gates that create entanglement) and the list of transformed stabilizers continues to uniquely define the quantum state at the output of the circuit. Consider, for example, the Bell-state generation circuit

$$\text{CNOT}_{10}\,[\text{H} \otimes I]\,|00\rangle = \frac{1}{\sqrt{2}}\,[|00\rangle + |11\rangle], \tag{203}$$

where the notation CNOT_{10} is used to indicate that qubit 1 is the control and qubit 0 is the target (and the qubits are numbered right to left starting with ordinal number zero). This circuit transforms the initial stabilizer set $\{Z_1, Z_0\}$ to the set of weight-2 stabilizers $\{X_1X_0, Z_1Z_0\}$ that uniquely defines the Bell state.

The efficiency advantage of the stabilizer formalism can be seen by noting that if we had a set of $N = 2M$ qubits and a Clifford circuit that produced M randomly selected Bell pairs, the quantum state would be described by a large superposition of 2^N quantum amplitudes, whereas the set of stabilizers would still only be of size N. Simple classical algorithms exist [271] to update the list of stabilizers of the state according to the Clifford transformations and thus Clifford circuits can be efficiently simulated classically and therefore do not represent the full power of quantum computation. We define a "stabilizer state" as any state that can be produced from the all-zero state using a Clifford circuit. Equivalently, a stabilizer state for N qubits is a $+1$ eigenstate of N independent generators (Pauli strings) of the stabilizer group.

An analogous classification of CV gates [218,219] can be carried out following the structure of the Clifford hierarchy [313,314]. Table II summarizes the overall analogy that we discuss here. The analog of qubit Pauli operations are oscillator displacements (in position and momentum). The analog of the Pauli group for oscillators is the Heisenberg-Weyl group H(1), which is a continuous Lie group composed of displacements in phase space. We have already seen that the product of two displacements is (up to a global phase) another displacement (see Box IV.1) and that every displacement has an inverse. The algebra $\mathfrak{h}(1)$ that generates this group is spanned by the canonical position and momentum operators [see Eq. (94)] and the identity operator, $\{\hat{\boldsymbol{R}}, \hat{I}\}$, satisfying the required commutation relations.

By analogy with the DV case, the CV analog of the Clifford group C_2 is associated with those CV operations that preserve (up to overall phases) the group of displacements under conjugation. The only such operations are the set of Gaussian operations defined in Eq. (98). As shown in Eq. (100), Gaussian operations induce a linear transformation (rotation, translation, and symplectic rescaling) on the phase-space coordinates. It follows from this that a phase-space displacement is mapped into a different displacement

under conjugation by any Gaussian operation

$$U_{\rm G}^{\dagger} e^{i\boldsymbol{\lambda}\hat{\boldsymbol{R}}} U_{\rm G} = e^{i\theta} e^{i\boldsymbol{\lambda}'\hat{\boldsymbol{R}}}, \tag{204}$$

$$\boldsymbol{\lambda}' = \boldsymbol{\lambda}\boldsymbol{M}, \tag{205}$$

$$\theta = \boldsymbol{\lambda}\boldsymbol{\Delta}, \tag{206}$$

where $\boldsymbol{M}$ and $\boldsymbol{\Delta}$ are defined in Eq. (100).

From the above definition of a stabilizer state in the DV setting, one might imagine that the CV analog of a DV stabilizer state is a Gaussian state—any state created from the CV vacuum state (Fock state $|0\rangle$) via Gaussian operations. We will see shortly that this definition has some problems, but for the moment we pursue it further to see where it leads. The vacuum (ground state, or a coherent state with displacement $\alpha = 0$) wave function of a harmonic oscillator is, of course, a Gaussian function of the coordinate, as shown in Eq. (54). We have seen that under the influence of any single- or multimode Hamiltonian that is quadratic in position and momentum (or equivalently, quadratic in the creation and annihilation operators), the Heisenberg equations of motion for the operators are linear. This immediately implies that Gaussian states evolve into other Gaussians. Time evolution under the application of quadratic control Hamiltonians can displace the position, boost the momentum, and/or squeeze the Gaussian, but the wave function always remains a Gaussian. This means that it is easy to classically compute and represent the time evolution in terms of a few parameters representing the mean values and covariances of the positions and momenta, $\vec{R}$, or equivalently, of the creation and annihilation operators, $\langle a\rangle$, $\langle a^{\dagger}a\rangle$, $\langle aa\rangle$, $\langle a^{\dagger}b^{\dagger}\rangle$, etc. We do not need to keep track of higher-order correlators, because for Gaussian states, Wick's theorem [315] guarantees that they are simple products of lower-order correlators. For example (assuming for simplicity that $\langle a\rangle = \langle b\rangle = 0$), Wick's theorem tells us that

$$\langle a^{\dagger}ab^{\dagger}b\rangle = \langle a^{\dagger}a\rangle\langle b^{\dagger}b\rangle + \langle a^{\dagger}b^{\dagger}\rangle\langle ab\rangle + \langle a^{\dagger}b\rangle\langle ab^{\dagger}\rangle. \tag{207}$$

In condensed-matter physics and quantum field theory, terms such as $\langle aa\rangle$ and $\langle ab\rangle$ are known as anomalous correlators. In quantum optics, they are described as representing one- and two-mode squeezing, respectively.

We can draw an analogy [219] between the simplicity of computing time evolution under quadratic bosonic Hamiltonians (Gaussian operations) and the Gottesman-Knill theorem [270,271] that a quantum computer based on qubits and using only Clifford group operations is easy to simulate classically. The fact that DV Clifford circuits can be efficiently simulated classically is related not just to the ability to do stabilizer updates but also to the fact that the full nonclassical correlations inside Bell states cannot be revealed without making non-Clifford (e.g., T-gate 45°) rotations on the Bloch sphere (or rotating the measurement axis by 45°) to violate the Bell inequalities. Such non-Clifford rotations are sometimes described as introducing "magic" (or "nonstabilizerness") to a state [316–320]. Similarly, CV transformations by Gaussian operations only update the mean and variance of Gaussian states (analogous to stabilizer updates) and fail to introduce any negativity in the Wigner function (defined in Appendix E 6). If we start with a state with a nonnegative Wigner function, then that property is preserved under all Gaussian operations. Hence the expectation value of bosonic observables can be readily obtained by classical importance sampling of the wave function, without suffering from any sign problems [272,273]. This is not generically the case for non-Gaussian states and thus Wigner negativity for bosonic systems is akin to "magic" in qubit states. Reference [321] highlights the similarity between Gaussian CV states and stabilizer DV states.

Even though Gaussian operations on Gaussian states cannot produce Wigner negativity, they *can* produce entanglement [218], just as Clifford operations on qubits can. Box IV.5 shows that two-mode squeezing is a Gaussian operation that produces entanglement (most readily apparent in the Fock basis) and Box IV.6 shows that the SUM gate is a Gaussian operation that produces entanglement (most readily apparent in the position basis).

Despite the rough analogy described above between DV stabilizer states and Clifford gates on the one hand and CV Gaussian states and gates, there are important differences between them. For the DV case, Clifford gates are (for the case of a single qubit) a subgroup of the compact group SU(2), whereas for the CV case, Gaussian operations on a single oscillator mode belong to the noncompact group SU(1, 1). The reason for this difference originates with the Bogoliubov transformation of the mode operators effected by conjugation with the squeezing transformation [see Eqs. (167), (181), and (182) and Boxes IV.3 and IV.5]. We have seen already that Gaussian operations can be represented by symplectic transformations of the N coordinates and momenta so that SU(1, 1) is isomorphic to $\mathrm{Sp}(2N, \mathbb{R})$.

At this point we must note one glaring problem with taking the CV analog of stabilizer states to be Gaussian states—namely, Gaussian states are not eigenstates of translation operators (which are the analog of Pauli operators).

Working in standard dimensionless units, consider a displacement operation by a distance Δ in position applied to a position eigenstate $|x\rangle$

$$D_x(\Delta)|x\rangle = |x + \Delta\rangle. \tag{208}$$

The (unnormalizable) eigenvectors of this displacement operation are highly non-Gaussian GKP-like (see

Sec. V C) states [280,282]

$$|\psi_{k,x_0}\rangle = \sum_{j=-\infty}^{+\infty} e^{+ikj\Delta}|x_0 - j\Delta\rangle, \quad \text{obeying} \tag{209}$$

$$D_x(\Delta)|\psi_{k,x_0}\rangle = e^{ik\Delta}|\psi_{k,x_0}\rangle, \tag{210}$$

where the fiducial position x_0 is arbitrary.

From the above, we see that we need to rethink our definition of stabilizer states. One reasonable possibility is to restrict our attention to the two GKP code states

$$|0\rangle_{\mathrm{L}} = \sum_{m=-\infty}^{+\infty} |2m\sqrt{\pi}\rangle_x, \tag{211}$$

$$|1\rangle_{\mathrm{L}} = \sum_{m=-\infty}^{+\infty} |(2m+1)\sqrt{\pi}\rangle_x, \tag{212}$$

and to define the Pauli group to be the subgroup of the Gaussian operations which act like DV Pauli operators within this restricted code space. These are the phase-space translations

$$X_{\mathrm{L}} = D_x(\sqrt{\pi}) = e^{-i\sqrt{\pi}\hat{p}}, \tag{213}$$

$$Z_{\mathrm{L}} = D_p(\sqrt{\pi}) = e^{+i\sqrt{\pi}\hat{x}}, \tag{214}$$

$$Y_{\mathrm{L}} = +iX_{\mathrm{L}}Z_{\mathrm{L}}. \tag{215}$$

It is straightforward to show from the commutation properties of these translations given in Eq. (140) that they precisely obey the Pauli algebra

$$X_{\mathrm{L}}Y_{\mathrm{L}} = iZ_{\mathrm{L}}, \quad \text{et cyc.}, \tag{216}$$

and show that

$$X_{\mathrm{L}}|0\rangle_{\mathrm{L}} = |1\rangle_{\mathrm{L}}, \tag{217}$$

$$X_{\mathrm{L}}|1\rangle_{\mathrm{L}} = |0\rangle_{\mathrm{L}}, \tag{218}$$

$$Z_{\mathrm{L}}|0\rangle_{\mathrm{L}} = +|0\rangle_{\mathrm{L}}, \tag{219}$$

$$Z_{\mathrm{L}}|1\rangle_{\mathrm{L}} = -|1\rangle_{\mathrm{L}}, \tag{220}$$

as expected. Furthermore, both GKP code words are stabilizer states that are $+1$ eigenstates of the two stabilizers

$$S_x = X_{\mathrm{L}}^2, \tag{221}$$

$$S_p = Z_{\mathrm{L}}^2. \tag{222}$$

As a result, the Pauli operators effectively square to the identity as they do in the usual DV case. We see that the GKP code effectively compactifies the phase-space plane into a torus of surface area 4π that contains a Hilbert space of dimension 2. One might ask how it is that only two stabilizers are needed to reduce the infinite oscillator Hilbert-space dimension down to two. This is because, unlike ordinary DV Pauli string stabilizers that only take on two values, $\pm$, S_x and S_p are unitaries with a continuum of eigenvalues on the unit circle in the complex plane, and continuous displacements of the code words in phase space cause continuous changes in the phase of the stabilizer eigenvalues. Thus, specifying the phase of the stabilizer requires an infinite number of bits of information, and it is this fact that causes the infinite reduction in Hilbert-space dimension.

We do not delve into the more general Gaussian operations that are analogous to Clifford group operations on GKP bosonic qubits other than to note that the CNOT gate is the following Gaussian operation on two oscillators (denoted a and b):

$$\mathrm{CNOT} = e^{-i\hat{p}_b\hat{x}_a}, \tag{223}$$

which is nothing but the SUM gate defined in Box IV.6. Mode b is displaced by an odd number of half lattice constants (i.e., and X_{L} operation is performed) if and only if mode a is in state $|1_{\mathrm{L}}\rangle$. Finally, we mention that experimentally realistic, finite-energy (normalizable) approximations to the GKP code words are discussed in Sec. V C.

Note that in Table II we compare qubit Pauli measurements to homodyne measurements of the oscillator phase-space coordinates $\hat{x}$ and $\hat{p}$. The latter are the generators of displacements, but they are not the displacements themselves (which are the objects that are analogous to the Pauli operators). While homodyne measurements are considered to be Gaussian measurements (typically performed by amplifiers involving one- or two-mode squeezing or using beam splitters and local oscillators as described in Appendix E 5), measurement of displacement operators requires non-Gaussian operations (see Appendix E 6).

Another point related to measurement is the following. One important task for quantum computation is to benchmark the fidelity of different operations. Given the analogy between Clifford operations and Gaussian operations, it may be tempting to conclude that techniques such as randomized benchmarking [322] for Clifford gates can be directly applied to CV systems. However, this is not the case because Gaussian operations do not form a unitary t-design. We will come back to this in more detail in Sec. IV F 2.

The fact that Gaussian states are not eigenstates of displacements also leads to a major point of difference found in the context of QEC. For the case of qubits, even though our adversary, the "noise demon," may have the power to do non-Clifford operations, we can (under certain conditions) defeat it and carry out QEC using only Clifford circuits and measurements [171]. The analogous statement for continuous-variable systems is not true. There is an important no-go theorem stating that Gaussian resources

(i.e., a Gaussian channel comprising Gaussian operations plus homodyne measurement) are insufficient to do CV QEC [323]. This is discussed further in Sec. V B 4.

In addition to Table II, which summarizes the overall analogies we have discussed here, the reader is directed to Table III, which summarizes experimentally relevant Gaussian and non-Gaussian gates for oscillators.

D. Hybrid non-Gaussian operations

We present a discussion of the gate operations (listed in Table IV) that are possible (under certain physics approximations) for the hybrid oscillator-qubit systems, from a computer science perspective, leaving details of practical implementation in different hybrid CV-DV platforms to Appendix A. The reader is also directed to important prior work in Ref. [325] describing control techniques to create non-Gaussian states in a hybrid system comprising a nanomechanical resonator and a charge-based qubit.

In the context of universal quantum channel construction, combining the qubit-assisted universal oscillator control, discussed in this section, with qubit readout or reset can thus yield universal channel control for oscillators using single auxiliary qubits. There have been several experimental demonstrations of channel constructions using oscillator-qubit coupling. Dissipative cooling of an oscillator to a squeezed vacuum state has been achieved in trapped ions [268,269] and in microwave resonators [284]. Dissipation to a superposition of squeezed states, demonstrated in trapped ions as well as superconducting circuits [64,65,326], has been used in the context of QEC. Error-correction maps are irreversible quantum channels and thus such universal channel construction methods can be used to remove the excess entropy or excitations of a bosonic code. We refer the readers to Sec. V C for details regarding these examples.

In the case of hybrid oscillator-qubit gates, throughout this work, we interchangeably use the prefixes "conditional" and "controlled" to indicate an oscillator gate that depends on the state of the qubit (or vice versa). We emphasize that this use of "controlled" differs from the conventional usage for multi-qubit gates, where a controlled gate corresponds to an operation that is carried out if and only if the control qubit is in a particular state. In contrast, here these prefixes indicate a gate that more generally carries out a different operation for each possible state of the control system.

1. Conditional displacement

Qubit-controlled cavity displacement operations (see Box IV.7) have been discussed in the context of superconducting platforms in [60,327]. In trapped ions and neutral atoms, these gates are realized natively using sideband interactions in Refs. [106,107,115,328].

TABLE IV. Common hybrid oscillator-qubit gates. This table is not comprehensive; however, it contains some of the most common gates for hybrid oscillator-qubit systems.

		Gate name	Parameters	Definition
Hybrid oscillator-qubit gates	Single-oscillator (Sec. IV D)	Conditional rotation (Box IV.8)	$\theta \in [0, 4\pi)$	$\mathrm{CR}(\theta) = \exp\left[-i\frac{\theta}{2}\sigma_z a^\dagger a\right]$
		Conditional parity (Box IV.8)	$\theta = \pi$	$\mathrm{CP} = \exp\left[-i\frac{\pi}{2}\sigma_z a^\dagger a\right]$
		SQR[a] [324] (Box IV.9)	$\theta_n \in [0, 4\pi)\ \varphi_n \in [0, 2\pi)$	$\mathrm{SQR}(\vec{\theta}, \vec{\varphi}) = \sum_n R_{\varphi n}(\theta_n) \otimes \lvert n\rangle\langle n\rvert$
		Jaynes-Cummings[b]	$\theta, \varphi \in [0, 2\pi)$	$\mathrm{JC}(\theta, \varphi) = \exp\left[-i\theta\left(e^{i\varphi}\sigma_- a^\dagger + e^{-i\varphi}\sigma_+ a\right)\right]$
		Anti-Jaynes-Cummings[c]	$\theta, \varphi \in [0, 2\pi)$	$\mathrm{AJC}(\theta, \varphi) = \exp\left[-i\theta\left(e^{i\varphi}\sigma_+ a^\dagger + e^{-i\varphi}\sigma_- a\right)\right]$
		Conditional displacement [60,64] (Box IV.7)	$\alpha \in \mathbb{C}$	$\mathrm{CD}(\alpha) = \exp\left[\sigma_z\left(\alpha a^\dagger - \alpha^* a\right)\right]$
		Rabi interaction	$\theta \in \mathbb{R}$	$\mathrm{RB}(\theta) = \exp\left[-i\sigma_x(\theta a^\dagger + \theta^* a)\right]$
		Conditional squeezing	$\zeta \in \mathbb{C}$	$\mathrm{CS}(\zeta) = \exp\left[\frac{1}{2}\sigma_z(\zeta^* a^2 - \zeta a^{\dagger 2})\right]$
	Multi-oscillator (Sec. IV D)	Conditional beam splitter [237]	$\theta \in [0, 4\pi)\ \varphi \in [0, \pi)$	$\mathrm{CBS}(\theta, \varphi) = \exp\left[-i\frac{\theta}{2}\sigma_z(e^{i\varphi} a^\dagger b + e^{-i\varphi} a b^\dagger)\right]$
		Conditional two-mode squeezing	$\xi \in \mathbb{C}$	$\mathrm{CTMS}(\xi) = \exp\left[\sigma_z(\xi a^\dagger b^\dagger - \xi^* a b)\right]$
		Conditional SUM gate	$\lambda \in \mathbb{R}$	$\mathrm{CSUM}(\lambda) = \exp\left[\frac{\lambda}{2}\sigma_z(a^\dagger + a)(b^\dagger - b)\right]$

[a] SQR stands for number-Selective Qubit Rotation. The single-qubit rotation $R_\varphi(\theta)$ is defined in Table VIII of Appendix A.
[b] Also called "red sideband."
[c] Also called "blue sideband."

Box IV.7. Conditional Displacement

$$D_c(\alpha,\beta) = |0\rangle\langle 0| \otimes D(\alpha) + |1\rangle\langle 1| \otimes D(\beta), \quad (224)$$

where $D(\alpha)$ is the unconditional displacement gate given in Eq. (134). We also define the symmetric conditional displacement gate

$$\mathrm{CD}(\alpha) = D_c(+\alpha,-\alpha) \quad (225)$$

$$= \exp\left[\sigma_z \otimes \left(\alpha a^\dagger - \alpha^* a\right)\right]. \quad (226)$$

For β real, we have (in Wigner units)

$$\mathrm{CD}(\beta) = e^{-2i\beta\sigma_z\hat{p}}, \quad (227)$$

$$\mathrm{CD}(i\beta) = e^{+2i\beta\sigma_z\hat{x}}. \quad (228)$$

which are, respectively, conditional position displacement and momentum boost by β.

A more general form of the conditional displacement can be obtained by conjugating $\mathrm{CD}(\alpha)$ (or its asymmetric variant $D_c(\alpha,\beta)$) with auxiliary single-qubit rotations, yielding

$$\mathrm{CD}_{\hat{b}\cdot\vec{\sigma}}(\alpha) = e^{\hat{b}\cdot\vec{\sigma}\otimes[\alpha a^\dagger - \alpha^* a]} \quad (229)$$

in the former case, where $\hat{b}$ is a unit vector defining an axis on the auxiliary-qubit Bloch sphere.

We will also utilize still further generalizations of these gates to the multimode, multi-qubit case in which the displacement of mode j is controlled by qubit k. These are denoted $\mathrm{CD}^{(j,k)}(\alpha)$ [see Eq. (336)] and $D_c^{(j,k)}(\alpha,\beta)$ and are discussed in Sec. VI C.

These gates are used to displace oscillators in phase space along a vector defined by the complex number α or β, conditioned on the qubit state. For example, if $\alpha,\beta\in\mathbb{R}$, then the above-defined gate displaces the oscillator by an amount α (β) from the origin along $\hat{x}$ if the qubit is in state $|0\rangle$ ($|1\rangle$). Commonly, we realize conditional displacement gates with $\alpha=-\beta$ which yield the gate defined in Table IV. Throughout this paper, we most often use the term "conditional displacement" to refer to this latter symmetric case, and this is made explicit when referencing the more general form. Furthermore, we note the useful algebraic identity for composing two conditional displacements given in Sec. VI C 1 c. This section also provides some intuitive pictures to aid in understanding the application of these important gates.

Note that from a control perspective, the conditional displacement gate can also be viewed as an oscillator-controlled qubit rotation. For example, the left-hand side of the controlled-momentum boost

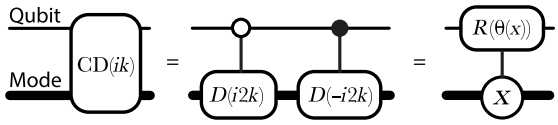


FIG. 13. Circuit representations of the two equivalent interpretations of the conditional displacement gate $\mathrm{CD}(i2k)$ in Eq. (230). On the left, a displacement is applied to the oscillator conditioned on the qubit state; this is equivalently viewed on the right as a qubit z rotation where the rotation angle depends on the position of the oscillator.

$$\mathrm{CD}(ik) = e^{i2k\hat{x}\sigma_z} = \int \mathrm{d}x|x\rangle\langle x|e^{-i\frac{\theta(x)}{2}\sigma_z}, \quad \theta(x) = -4kx \quad (230)$$

can be viewed as a momentum boost of the oscillator by $\pm k$ controlled by the state of the qubit and the right-hand side can be viewed as a rotation by angle $\theta(x)$ controlled by the oscillator position x. These two complementary views are represented in circuit language in Fig. 13.

2. Conditional rotation and parity gates

In addition to displacements of the bosonic mode conditioned on the state of an auxiliary qubit, we can have conditional rotations of the bosonic mode in phase space (see Box IV.8).

Box IV.8. Conditional Rotation and Parity Gates

This is the conditional version of the phase-space rotation gate defined in Box IV.2

$$\mathrm{CR}(\theta) = e^{-i\frac{\theta}{2}\sigma_z a^\dagger a} \quad (231)$$

This gate is natively available through the dispersive interaction between the bosonic mode and the auxiliary qubit [see Eq. (A24) of Appendix A]. The natively available dispersive interaction can be synthetically enhanced or compensated using a scheme utilizing the SNAP gate (defined in Box IV.10), as described in Sec. A 1 c b.

The conditional photon-number parity gate can be defined either as

$$\mathrm{CP} = \mathrm{CR}(\pi) \quad (232)$$

or, using a different phase convention, as

$$\mathrm{CP}' = e^{-i\pi\frac{I-\sigma_z}{2}a^\dagger a}. \quad (233)$$

Either way, the phase kickback from the controlled parity gate (see Fig. 43) can be used to make a QND

measurement of the photon-number parity [240] and even the full photon-number distribution via phase estimation or related methods [59,147,329].

A generalization of these gates, $\mathrm{CR}^{kj}(\theta)$, to the multi-qubit, multimode case in which oscillator k is rotated conditioned on the state of qubit j is defined in Sec. VI C [see Eq. (333)].

3. Photon-number-selective qubit rotation (SQR) gate

The conditional rotation gate in Box IV.8 can be interpreted as a Z rotation on the qubit by a rotation angle that is proportional to the boson number (operator). However, it would be even more powerful if a different arbitrary qubit rotation could be applied conditioned on each oscillator Fock state. This is accomplished with the photon-number-selective qubit rotation (SQR) gate (see Box IV.9). Combined with the displacement gate, the SQR gate suffices for universal hybrid control of the oscillator-qubit system (Table V).

Box IV.9. SQR Gate

The cavity-conditioned qubit rotation gate, also known as the photon-number-selective qubit rotation (SQR) gate, is defined as

$$\mathrm{SQR}(\vec{\theta},\vec{\varphi}) = \sum_{n=0}^{N_{\max}} R_{\varphi_n}(\theta_n)\otimes|n\rangle\langle n| \tag{234}$$

$$= e^{-i\sum_{n=0}^{N_{\max}}\frac{\theta_n}{2}\vec{b}_n\cdot\vec{\sigma}\otimes\hat{P}_n}, \tag{235}$$

where $R_{\varphi_n}(\theta_n)$ (defined in Table VIII in Appendix A) is a qubit rotation by angle θ_n about the axis $\vec{b}_n = (\cos\varphi_n, \sin\varphi_n, 0)$ (lying on the equator of the Bloch sphere), $\hat{P}_n$ is the Fock state projector defined in Eq. (24), and the total rotation angle of the qubit is given by θ_n for cavity Fock state n. Different qubit rotations conditioned on different photon numbers can be executed in parallel by supplying multiple simultaneous drive tones to achieve this gate (see Appendix A).

4. Photon-number-selective phase gate

In some cases, one may be interested in only controlling the oscillator while keeping the qubit state unchanged. This can be realized by a reduced version of the SQR gate, the photon-number-selective phase (SNAP) gate [266], which applies a different Berry phase φ_n to each photon Fock state $|n\rangle$, as described in Box IV.10. Together with (unconditional) cavity displacements, this is sufficient for universal control of the oscillator (Table V). An efficient numerical gate synthesis protocol has been given to optimize short SNAP gate sequences to implement any given unitary on the oscillator [177]. There has also been recent theoretical progress in analytically solving for efficient gate sequences [330]. The use of SNAP gates has also been proposed to enhance nonlinearities in optomechanical systems [331].

Box IV.10. SNAP Gate

The SNAP gate [59] is a powerful gate that imparts a different phase (chosen by the user) on each Fock state

$$\mathrm{SNAP}(\vec{\varphi}) = e^{-i\sum_n \varphi_n|n\rangle\langle n|} \tag{236}$$

$$\mathrm{SNAP}(\vec{\varphi})\sum_n \Psi_n|n\rangle = \sum_n e^{-i\varphi_n}\Psi_n|n\rangle, \tag{237}$$

where $\vec{\varphi} = (\varphi_0, \varphi_1, \varphi_2, \ldots, \varphi_{N_{\max}})$.

TABLE V. Common instruction sets for oscillator and hybrid oscillator-qubit architectures using gates defined in Tables III, IV, and VIII. Gaussian-only operations can achieve linear oscillator control. Adding non-Gaussian operations such as the general squeezing operation or SNAP gate can allow universal oscillator control. Moreover, to enable universal hybrid (oscillator-qubit) control, entangling gates (such as $\mathrm{CD}(\beta)$ or $\mathrm{SQR}(\vec{\theta},\vec{\varphi})$) between oscillators and qubits are needed. See Appendix D for proof of the universality of the hybrid instruction set and cross-compilation between different instruction sets. $U_N(z)$ ($N \geq 3$) is the generalized squeezing gate as defined in Table III. $R_\varphi(\theta)$ is a single-qubit rotation gate defined in Table VIII. The above constitute the minimal instruction sets, but we often have access to more gates such as dispersive interactions, phase rotations, etc.

	Instruction set name	Minimum gate set	Sections
Linear oscillator control	Gaussian	$\mathcal{G} = \{D(\alpha), S(\zeta), \mathrm{BS}(\theta,\varphi)$ or $\mathrm{TMS}(r,\varphi)\}$	Sec. IV B and Table III
Universal oscillator control	Cubic	$\mathcal{G} + U_3(z)$	
	Quartic	$\mathcal{G} + U_4(z)$	
	SNAP	$\{D(\alpha), \mathrm{SNAP}(\vec{\varphi}), \mathrm{BS}(\theta,\varphi)$ or $\mathrm{TMS}(r,\varphi)\}$	
Universal hybrid control	Phase-space	$\{\mathrm{CD}(\beta), R_\varphi(\theta), \mathrm{BS}(\theta,\varphi)\}$	Sec. IV D and Table IV
	Fock-space	$\{\mathrm{SQR}(\vec{\theta},\vec{\varphi}), D(\alpha), \mathrm{BS}(\theta,\varphi)\}$	
	Sideband	$\{R_\varphi(\theta), \mathrm{JC}(\theta), \mathrm{BS}(\theta,\varphi)\}$	

While frequently used as an oscillator-only gate, in practice this gate is achieved using an auxiliary qubit dispersively coupled to the bosonic mode and hence is more realistically written as a hybrid gate,

$$\text{SNAP}(\vec{\varphi}) = e^{-i\sum_n \sigma_z \varphi_n |n\rangle\langle n|}. \tag{238}$$

To understand how to synthesize the SNAP gate, we make use of the cavity-conditioned qubit rotation (SQR) gate in Box IV.9 and the identity [59]

$$\text{SNAP}(\vec{\varphi}) = \text{SQR}(-\vec{\pi}, \vec{\varphi})\text{SQR}(+\vec{\pi}, \vec{0}), \tag{239}$$

where $\vec{\pi} = (\pi)(1, 1, \ldots, 1)$. This identity follows from the fact that if we start the auxiliary qubit in state $|0\rangle$, this pair of gates causes the qubit trajectory to return to the initial state, but the enclosed area on the Bloch sphere imparts a Berry phase φ_n on the cavity if it is in Fock state n. For simplicity, we have written the SNAP gate as if it were a cavity-only non-Gaussian gate. However, as we have seen, the auxiliary qubit is used as a "catalyst" to effect the gate. If the qubit starts out in state $|1\rangle$, the sign of the Berry phase is reversed. Thus there is an implicit factor of σ_z associated with the factor φ_n in Eqs. (236) and (237) and this is actually a hybrid qubit-cavity gate.

We also direct the reader to interesting recent work [332] that uses QSP methods to synthesize SNAP gates and other non-Gaussian hybrid CV-DV operations.

5. *Exponential-*SWAP *gate*

Another important hybrid gate that has been realized [298] is the exponential-SWAP (eSWAP) gate, which can be used as an entangling gate for universal computation that is bosonic code agnostic [297,298]. As defined below, the eSWAP gate produces a linear combination of the original and the swapped bosonic states on two modes with the relative fraction defined by θ, conditioned on the qubit state.

Box IV.11. eSWAP Gate

$$\begin{aligned}\text{eSWAP}(\theta) &= e^{i\frac{\theta}{2}\,\text{SWAP}} \\ &= \cos\left(\frac{\theta}{2}\right) I + i \sin\left(\frac{\theta}{2}\right) \text{SWAP},\end{aligned} \tag{240}$$

where the SWAP gate is defined as a special case of the beam-splitter gate in Box IV.4 which swaps the states in two bosonic modes.

Similar to the SNAP gate, the eSWAP gate is often used as an oscillator-only gate, but in practice is achieved using an auxiliary qubit that is initialized to $|0\rangle$. Generalizing to an arbitrary initial ancillary state, the eSWAP gate is written as

$$\text{eSWAP}(\theta) = e^{i\frac{\theta}{2}\sigma_z \otimes \text{SWAP}}. \tag{241}$$

One way to look at the SWAP operation in a two-mode encoding (e.g., dual-rail encoding), is to view it as a logical Pauli-X operation which flips between $|0\rangle_\text{L} = |01\rangle$ and $|1\rangle_\text{L} = |10\rangle$. In this sense, when $\theta = \pi/2$, $\text{eSWAP}(\pi/2) = (I + i\text{SWAP})/\sqrt{2}$ can be used as the $\sqrt{i\text{SWAP}}$ gate to entangle between two modes. Remarkably, this remains true even if $|01\rangle$ refers to arbitrary logically encoded states (e.g., binomial code states) rather than Fock states in the two cavities. This is because both identity and SWAP are agnostic to the contents of the two cavities, a powerful feature of the SWAP and eSWAP gates.

E. Instruction sets for universal control of hybrid oscillator-qubit systems

Having laid out in the previous section a catalog of hybrid CV-DV gates, we turn next to assembling different subgroups of these gates into instruction set architectures. We will consider six distinct ISAs. The first two are oscillator-only ISAs. The Gaussian ISA is easy to realize experimentally since it is generated by polynomial Hamiltonians linear or quadratic in the oscillator position and momentum operators. It is, however, nonuniversal. Universal control can be obtained by adding any higher-order polynomial, the simplest of which yields the cubic ISA.

Turning to hybrid CV-DV ISAs that give universal control over both the oscillators and the qubits, we begin with the phase-space ISA based on oscillator displacements in phase space controlled by the state of an ancilla qubit. This ISA is useful in hardware that has moderate dispersive coupling between the ancilla and the oscillator. The Fock-space ISA focuses on SQR gates that rotate the ancilla conditioned on the number of quanta in the oscillator. This is most useful in hardware with stronger dispersive coupling between the oscillator and the ancilla. The sideband ISA focuses on gates generated by the Jaynes-Cummings (or anti-Jaynes-Cummings) Hamiltonian under which control pulses can add or subtract a photon from a resonator while simultaneously flipping the state of the ancilla qubit. Finally, we will return to the Gaussian ISA, which can be made universal if a supply of non-Gaussian resource states is available.

1. Oscillator-only ISAs

The Gaussian operations comprising displacements (in position and momentum), single-mode squeezing, and beam splitters together form a parametrized instruction set, known as the Gaussian instruction set. The availability of this instruction set allows one to construct arbitrary Hamiltonians quadratic in $\hat{a}$ and $\hat{a}^\dagger$. For convenience, we may also include two-mode squeezing (TMS), but it is redundant since, as illustrated in Fig. 11, TMS can be synthesized from single-mode squeezing using the Bloch-Messiah decomposition. The Gaussian instruction set contains only linear $(a^\dagger, a)$, quadratic $(a^\dagger a, a^\dagger a^\dagger, aa)$, and bilinear $(a^\dagger b, ab^\dagger)$ control Hamiltonians and hence, as discussed above, is analogous to a nonuniversal Clifford-only gate set. It maps Gaussian states to Gaussian states and thus cannot create Wigner negativity, but can create entanglement between modes (just as Clifford gates can create Bell states but not magic states for qubits). The universal control of oscillators requires a nonlinear resource beyond the quadratic Hamiltonians that generate the Gaussian instruction set. One way to achieve this is with a unitary realized via a cubic Hamiltonian (see Appendix D 5), using, e.g., specialized nonlinear elements such as the superconducting nonlinear asymmetric inductive element (SNAIL) [291,333,334]. In hybrid CV-DV processors, we have instruction sets that can achieve universality for oscillators via linear (or quadratic) coupling to DV ancillas—another nonlinear resource. In fact, these instruction sets allow for universal control over the full joint states of hybrid CV-DV processors.

2. Phase-space instruction set

We turn next to the hybrid system control. Hybrid cavity-qubit gates can be used to compile an effectively nonquadratic Hamiltonian using the BCH rules [335–341] and Trotter-Suzuki expansion (see Sec. VI). Remarkably, the coupling term between qubits and oscillators only needs to be linear in the oscillator mode operators (and is therefore easy to implement experimentally) to generate arbitrary nonlinear controls on the oscillators.

The phase-space instruction set is simple and experimentally relevant [60] for universal control of hybrid systems containing both oscillators and qubits. It comprises qubit-controlled oscillator displacements, qubit rotations, and beam splitters if there are multiple oscillators. We begin our analysis of the universality with the case of a single qubit coupled to a single resonator. The most general unitary for such a hybrid system is defined by the (effective) Hamiltonian

$$H = h_0(\hat{x},\hat{p}) + \vec{h}_1(\hat{x},\hat{p})\cdot\vec{\sigma}, \tag{242}$$

where $\vec{\sigma} = (X, Y, Z)$ is the vector of Pauli matrices. The first term h_0 above is a resonator-only Hamiltonian, while the second is a joint resonator-qubit Hamiltonian. As we discuss further below (see also Ref. [60]), h_0 is redundant since it can be generated from $\vec{h}_1$. We therefore focus solely on $\vec{h}_1$. Remarkably, any arbitrary polynomial Hamiltonian of the form given in Eq. (242) can be generated from the Lie brackets of a very small control set

$$G_{\text{phase-space}} = \{X, Z, \hat{x}Z, \hat{p}Z\}, \tag{243}$$

which can then be used for universal control. The first two terms generate universal control of the qubit. The second two terms individually generate controlled displacements in phase space defined in Box IV.7. As noted previously, controlled Gaussian operations are (roughly) analogous to controlled-Clifford operations (e.g., the Toffoli gate) which are non-Clifford and yield universality. Using rotations of the qubit, we can generate $\hat{x}X$ from $\hat{x}Z$ and $\hat{p}Y$ from $\hat{p}Z$

$$-i[\hat{x}X, \hat{p}Y] = Z(\hat{x}\hat{p} + \hat{p}\hat{x}). \tag{244}$$

Note that we obtain the anticommutator for the cavity operators instead of the commutator as before, thereby increasing the degree of the polynomial to yield a controlled-squeezing operation (which is discussed further in Sec. VI D 1 b). This Hamiltonian generates controlled-squeezing of position and antisqueezing of momentum, or vice versa. An alternative standard form of the (conditional) squeezing Hamiltonian can be obtained by rotating the axes of the phase-space displacements by $\pi/4$ to yield

$$-i\left[\frac{(\hat{x}+\hat{p})}{\sqrt{2}}X, \frac{(\hat{x}-\hat{p})}{\sqrt{2}}Y\right] = Z\left(\hat{x}^2 - \hat{p}^2\right). \tag{245}$$

Iteration of the above commutator procedure (with appropriate qubit rotations) on the control set in Eq. (243) allows us to use polynomials of degree m, n to achieve higher-order terms of degree $m + n$,

$$[\hat{x}^n X, \hat{p}^m Y] = i[\hat{x}^n\hat{p}^m + \hat{p}^m\hat{x}^n]Z, \tag{246}$$

and thus we can generate arbitrary polynomials for $\vec{h}_1(\hat{x},\hat{y})$ giving universal qubit-controlled operations on the oscillator. Examples of conditional Gaussians that can be achieved in superconducting cavity-qubit circuits are described in Refs. [60,233]. In trapped ions and neutral atoms, conditional displacements using simultaneous red and blue sideband interactions are native gates [16,106,107,115,328].

How do we create the oscillator-only Hamiltonian, $h_0(\hat{x},\hat{p})$? We can simply generate two polynomial terms of the form $A(\hat{x},\hat{p})Z$ and $B(\hat{x},\hat{p})Z$ such that their commutator

yields the identity operator on the qubit and

$$h_0(\hat{x},\hat{p}) = [A(\hat{x},\hat{p})Z, B(\hat{x},\hat{p})Z] = [A(\hat{x},\hat{p}), B(\hat{x},\hat{p})], \tag{247}$$

acting only on the oscillator.

Again, the use of beam-splitter gates allows us to extend this to the multi-oscillator, multi-qubit case. In summary, we see that the phase-space instruction set yields universal hybrid system control (see Appendix D). Example applications of this are presented in Sec. VI. Another extension of the phase-space ISA to multiple modes can also be achieved using displacements conditioned on a single-qubit ancilla state that is coupled to multiple oscillators [342].

3. Fock-space instruction set

The Fock-space instruction set comprises single-qubit rotations conditioned on the cavity photon number (SQR), simple oscillator displacements ($D(\alpha)$, *not* conditioned on the state of the qubit), and beam splitters. This should be contrasted with the phase-space instruction set in which the operations on the cavity are conditioned on the state of the qubit. The Fock-space instruction set is thus dual to the phase-space instruction set in the sense that it interchanges the control and target for conditional gates.

Unconditional displacements are relatively simple in both cQED and trapped-ion platforms. The SQR and closely related SNAP gates were developed for superconducting systems to take advantage of the natively available dispersive coupling [266,267] but have not been reported for the trapped-ion and neutral-atom platforms, where they would likely require synthesis using sideband gates [343,344]. Conversely, the spin-dependent forces used to create conditional displacements are natively available in ion traps but must be synthesized [60] from the dispersive coupling in superconducting systems, but as discussed in Sec. VI C, they offer greater control over the individual bosonic modes. Numerically [177] and analytically [330] optimized sequences of Fock-space instructions have been theoretically demonstrated to be highly expressive and efficient for state preparation.

4. Sideband instruction set

The sideband instruction set comprises qubit rotations, Jaynes-Cummings Hamiltonian sideband transitions (see Table IV) which convert a Fock state excitation between the qubit mode and the cavity mode, and beam splitters. The Jaynes-Cummings interaction is natively available (but "always on") in superconducting qubit systems [77,78,345,346]. Universal control of oscillator-qubit systems has been achieved using the sideband instruction set in trapped ions (see Appendix A 2 b), both approximately [61] and analytically [62], by strictly confining the control protocol within a certain energy budget. This is achieved by carefully choosing a subset of sideband rotations such that no leakage is possible beyond a given Fock level. Similar techniques of restricting the sideband rotations in the Jaynes-Cummings interaction to a subset have been employed to generate nontrivial qubit states and rotations without residual entanglement between the qubit and the oscillator [347,348]. Numerical synthesis methods have been presented in [349]. With the sideband instruction set, the Law-Eberly protocol can be readily employed to prepare Fock states [209]. Sideband operations have also been demonstrated in neutral atoms in the context of sideband cooling [114] and manipulation of the lowest two Fock levels [115], which could be straightforwardly generalized to control of the higher levels (see Appendix A 3).

5. Gaussian instruction set with non-Gaussian resource states

An alternative model for universal computation is to consider having only Gaussian operations plus a supply of fixed non-Gaussian resource states (which would of course have to be created using non-Gaussian gates such as those described above). The non-Gaussian "state factory" is a specialized system that handles the resource production (possibly in a nondeterministic way that utilizes distillation and "repeat-until-success" strategies). The algorithm is then run using only relatively simple Gaussian operations on these resource states. This CV model is analogous to universal DV computation with Clifford circuits using a T-state factory as the non-Clifford resource. This model will not be discussed further here, but the reader is directed to Refs. [222,294,350–354] for an entrée into the literature on this subject.

6. Choice of instruction set

The choice of the instruction set depends on the parameter details of the available hardware. It also depends on the nature of specific applications, because the underlying algorithm for different applications may favor certain gates over others. In Sec. VI, we use single-oscillator-qubit instruction sets to discuss examples of state transfer, channel construction, and unitary applications. We use the phase-space instruction set when discussing examples related to states with a concise description in terms of the phase-space variables $\hat{x}$ and $\hat{p}$, e.g., GKP logical code states [282]. We use the Fock-space instruction set for tasks that are better represented using the photon-number creation and annihilation operators $\hat{a}^\dagger$ and $\hat{a}$. Section VII C presents multi-oscillator-qubit instruction set applications in terms of the advantage of using hybrid hardware over qubit-only systems for Hamiltonian simulations. The SNAP gate from the SQR gate in the Fock-space instruction set is useful for gates that conserve photon number, as shown

TABLE VI. Common native and synthesized Hamiltonian terms in condensed-matter physics and quantum chemistry. "Higher-order" terms refer to those composed of >2 creation or annihilation operators multiplied together (i.e., polynomials of order 3 or higher in terms of $\hat{x}$ and $\hat{p}$). $f_i^\dagger$ and f_i are fermionic creation and annihilation operators, realized via fermion-to-qubit mappings such as Jordan-Wigner or Bravyi-Kitaev [355]. We note that there will, in general, be multiple avenues to synthesize e^{-iH_it}, and the aim here is to give a rough, nonexhaustive guide to the native gates and techniques needed. Here, "Approx." refers to the use of approximate compilation techniques, such as the Baker-Campbell-Hausdorff (BCH) and Trotter-Suzuki formulas (see Sec. VI D). See Table IV for definitions of common hybrid gates listed in the Hamiltonian synthesis column. See Ref. [41] for a thorough manual on how to simulate fermions, bosons, and gauge fields in CV-DV hardware.

Term H_i	Description	Synthesis of e^{-iH_it}	Example models
$a_i^\dagger a_i$	Mode energy, chemical potential	$R(\theta)$	Jaynes-Cummings [356], Bose-Hubbard [357]
$\hat{n}_i(\hat{n}_i-1)$	On-site interactions	SNAP($\vec{\varphi}$)	Bose-Hubbard [357]
$e^{i\phi}\hat{a}_i^\dagger\hat{a}_j$ + h.c.	Hopping	BS(θ,φ)	Bose-Hubbard [357], Su-Schrieffer-Heeger [358]
$Z_i(a+a^\dagger)$	Spin-dependent displacement	CD(α)	Rabi [359], Dicke [360], Hubbard-Holstein [361]
$Z_iZ_j(a+a^\dagger)$	Spin-spin-dependent displacement	CD(α), SWAP, CP (Sec. VI C 1)	Photosynthetic complex [362]
$Z_i(e^{i\varphi}a_i^\dagger a_j$ + h.c.)	Spin-dependent hopping	CBS(θ,φ)	Bosonic $\mathbb{Z}_2$ LGT [363]
$a_i^\dagger a_j a_k$, $a_i^\dagger a_j^\dagger a_k a_l$, etc.	Higher-order terms	$R(\theta)$, $D(\alpha)$, TMS(r,φ), BS(θ,φ) + Approx.	Bosonic SYK [364,365]
$S_i^zS_j^z$, $(\vec{S}_i\cdot\vec{S}_j)^k$, etc.	(Large)-spin interactions	$R(\theta)$, BS(θ,φ) + Approx.	AKLT [366]
$f_i^\dagger f_j$, $f_i^\dagger f_j^\dagger f_k f_l$, etc.	Fermionic matter	Multi-qubit gates (Sec. VI C 2)	Electronic structure [367], Hubbard-Holstein [361] Fermi-Hubbard [368]

in Table VI. From a computer architecture point of view, choosing different instruction sets is essentially using different hybrid operations in the bosonic QIR (Quantum Intermediate Representation; see Fig. 14) for a hybrid CV-DV quantum computer. See Sec. II for more architecture discussions.

F. Control flow and benchmarking

In this section, we discuss two important aspects of architectures: control flow (Sec. IV F 1) and benchmarking (Sec. IV F 2) of hybrid quantum processors. These aspects are vitally important to understanding the capabilities of hybrid qubit-bosonic hardware because they allow us

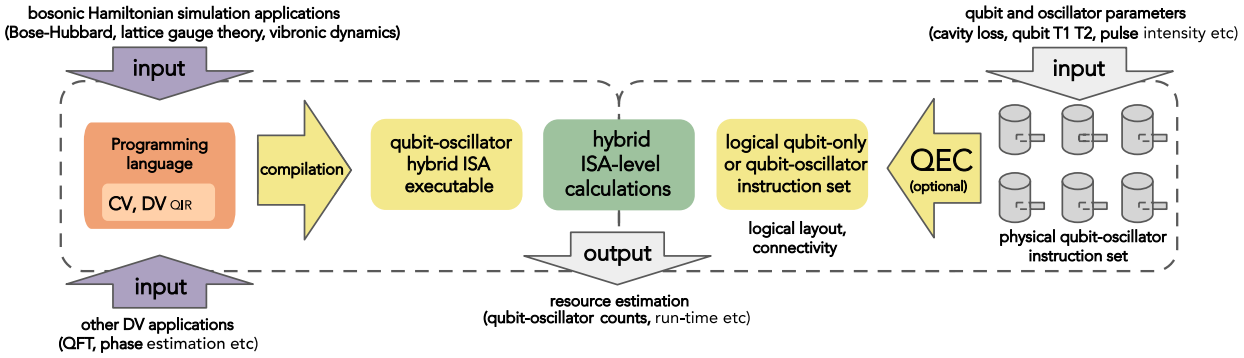


FIG. 14. The resource-estimation work flow for a hybrid oscillator-qubit quantum processor, inspired by Ref. [45]. From the left to the center: a native hybrid qubit-bosonic quantum application such as bosonic Hamiltonian simulation or a DV application such as phase estimation or quantum Fourier transform (QFT) is given as input and is described by proper programming languages including bosonic QIR (for bosonic Hamiltonian simulation) or standard DV QIR (for DV-only applications). The QIR program is then compiled into oscillator-qubit hybrid ISA executables which are agnostic to any physical device information. From the right to the center: the physical parameters of the hardware (such as cavity loss or qubit T_1 and T_2 coherence times) are taken as input which is used to estimate the cost of physical layer oscillator-qubit instruction sets (the gray cylinders are sketches of 3D post cavities with transmon qubits). Optional QEC is then applied to synthesize logical-level qubit-only or oscillator-qubit instruction sets from the physical ones. The two sides are combined in the middle to output the overall resource requirements for a given application.

to assess the performance of such devices and predict whether such hardware is adequately calibrated for a given computational task and how classical controls can be used to allow adaptive protocols to be run on such systems. We discuss both of these issues in turn below.

1. Control flow and feedback

In quantum computers, control is realized using some (classical or quantum) information from prior circuit execution to better guide the subsequent quantum computation. Quantum feedback control has been proposed and demonstrated experimentally [369,370], and the essence of it can be most simply understood by combining the control and the target systems as part of a large interacting quantum system upon which the usual analysis can be performed without resorting to control theory.

More commonly, control signals appear as classical information, which is obtained from the measurement of part of the quantum system from prior execution or during run-time. The measurement outcomes can either be used to design better measurement schemes in later parts of the circuits or as classical control signals that parametrize (i.e., modify) later unitary quantum gates. As the most evident example, classical controls have been extensively explored for the construction of variational quantum algorithms [371–373]. Another good use of classical control is to improve the robustness of the hybrid CV-DV system for autonomous QEC, as discussed in Sec. V. Error-mitigation techniques [374–378] in general and particularly on bosonic systems [379] are examples that demonstrate the power of classical feedback controls as well.

At a broader scale, control flow will play an increasingly important role as hybrid CV-DV quantum processors acquire more complicated designs in the near future, especially for aspects regarding modular architectures and distributed quantum computation [167,228,380,381], since a significant amount of information and signals need to be communicated across the entire computing hardware fabric. Specifics concerning what properties of bosonic quantum systems can be measured and how measurements are performed are detailed in Appendix E.

Finally, control flow can be used to overcome connectivity limitations in quantum hardware for which two-qubit gates can typically only be natively performed between physically adjacent qubits in the absence of a quantum bus that can transmit photons across the device. The *locally alternating quantum classical computation* (LAQCC) model [138,140,382,383] uses run-time measurements, classical computation, communication, and feedforward conditional quantum operations to enhance the power of purely unitary circuits for tasks such as constant-time quantum state preparation [20,383–386] and quantum communication and entanglement distribution [138,382] for, e.g., the realization of QRAM [124] and fault-tolerant quantum circuits with local operations [387].

2. Benchmarking and certification

Benchmarking and certification of hybrid CV-DV quantum processors are challenging due to the infinite-dimensional nature of the oscillator Hilbert space (see Ref. [388] and references therein for a discussion of these issues). Some protocols developed for benchmarking generic quantum systems [389,390] may be applied to CV systems, including device-independent protocols [391] for certifying entanglement between qubits and oscillators or between oscillators. Benchmarking protocols specifically designed for CV states have also been proposed by leveraging Gaussian measurements such as heterodyne detection [392] to directly measure the Husimi Q-function (Appendix E), which has been applied to verify Gaussian boson sampling experiments [393]. Such methods are suitable in the optical domain but do not take full advantage of the modern gate (see Sec. IV) and measurement (see Appendix E) capabilities now available in the microwave domain of cQED. In cQED [394] and in ion traps [395], multimode Wigner tomography (see discussion below and in Appendix E) based on joint photon-number parity measurements have been demonstrated.

A major open challenge that remains in the field is to determine whether there exists a protocol like randomized benchmarking for bosonic systems without Hilbert-space truncation. Such a protocol would be invaluable because randomized benchmarking (and related protocols) gives a way to efficiently determine average gate fidelity for gates implemented within a gate set.

Clifford randomized benchmarking works by drawing a random sequence of Clifford operations followed by the inverse of the sequence (which can be efficiently computed using the tableau algorithm of Gottesman [396]). The state is then measured to see if it successfully matches the starting state. The rate of decay of the success probability as a function of the length of the random Clifford circuit then yields the average gate fidelity for the Clifford operations. Similar extensions to non-Clifford gates such as interleaved randomized benchmarking also exist and can be used to benchmark an entire gate set [397,398].

There is a strong mathematical analogy between the Clifford operations and Gaussian operations in optics (see Sec. IV C) since both have positive Wigner representations. The Wigner transform of a bosonic operator G is

$$g(x,p) = \int \mathrm{d}y\, e^{ipy} \langle x - y/2|\, G\, |x + y/2\rangle \tag{248}$$

and can be classically simulated (when coupled with stabilizer or Gaussian measurements). Furthermore, Gaussian operations are also closed under multiplication, which suggests that the randomized benchmarking protocol could

work immediately out of the box. Unfortunately, however, this is not the case.

The major problem facing this idea is that Gaussian operations do not form a unitary t-design, which is to say that averaging over the set of Gaussian gates does not necessarily match all (t,t) moments (i.e., t moments in the real and imaginary components) of the marginal distributions of the phase-space quadratures. This inability to quickly scramble uniformly over the entire CV space means that the randomization process will not be able to give an estimate of the average gate fidelity over the set of all Gaussian operations in this context. This means that, at present, randomized benchmarking-like protocols do not exist in such settings, and finding a suitable alternative benchmarking procedure remains an open challenge in the field. We note that recent progress has been made toward this goal, however, with regularized-rigged designs (i.e., in a rigged Hilbert space) offering an approximate solution for CV systems [399].

A more ambitious goal is to benchmark a generic hybrid CV-DV quantum computation or simulation via *process tomography* [400] which may involve many non-Gaussian operations. One approach toward addressing this is to use the stellar hierarchy [258], mentioned in Sec. III C 4, to quantify the non-Gaussian character of CV states from zeros of the Husimi Q-function. This allows us to use heterodyne measurements to perform fidelity estimation and witness Wigner negativity (which is equivalent to non-Gaussianity for pure states) [401] for general CV states.

Benchmarking necessarily involves measurement. One interesting question is what measurement protocols beyond simple Gaussian measurements such as homodyne and heterodyne detection can be useful to improve the benchmarking efficiency, as simple non-Gaussian gates can be compiled easily from our ISA gate sets. In addition, progress on the complexity of classical shadow tomography of CV systems can shed light on benchmarking as well [399,402,403]. A rigorous analysis of measurement complexity, computation complexity, sampling complexity, and postprocessing complexity is needed to fully characterize each benchmarking protocol [390].

For completeness, let us consider how to characterize a state using Wigner function tomography (also described in Appendixes E and E 6). $W(\alpha)$, the Wigner function at an arbitrary point in phase space represented by the complex number α, can be found by measuring

$$W(\alpha) = \mathrm{Tr}[O(\alpha)\rho], \tag{249}$$

where ρ is a quantum state associated with the system the Wigner function of which we wish to identify, and we have defined [404]

$$O(\alpha) \equiv \frac{2}{\pi}\sum_{n=0}^{\infty}(-1)^n D(\alpha)\,|n\rangle\langle n|\,D^\dagger(\alpha). \tag{250}$$

For the case of a hybrid system, we need to characterize correlations between the qubits and the resonator modes. To achieve this, it is useful to consider a discrete operator basis for the qubit portion of the system. For example, using n-qubit Pauli operators P_i and the orthonormality of the Pauli basis, we can always write the joint CV-DV density matrix as

$$\rho = \frac{1}{2^n}\left(\sum_i \mathrm{Tr}_{\mathrm{qubit}}\left[\rho(I\otimes P_i)\right]\otimes P_i\right). \tag{251}$$

From this, we can then define the conditional Wigner function to be

$$W(\alpha|P_i) = \mathrm{Tr}\left[(O(\alpha)\otimes P_i)\rho\right], \tag{252}$$

the expectation values of which can be used to compute the expectation values of ρ. In Fig. 27 we present model examples of conditional Wigner functions associated with oscillator cat state production using a linear combination of unitaries (LCU) protocol. Experimental realization of such conditional Wigner functions can be found in Ref. [212].

A useful aim is to understand the complexity of learning $W(\alpha|P_i)$ within error ϵ and the probability of failure at most δ. Each experiment to learn a particular value of the conditional density matrix requires a preparation of a density operator ρ and then we need to measure $O(\alpha)\otimes P_i$ for each prepared state. For simplicity, let us assume that we are promised that $|\mathrm{Tr}[(D(\alpha)\,|n\rangle\langle n|\,D^\dagger(\alpha)\otimes P_i)\rho]| \le e^{\gamma_i-\lambda_i n}$ for constants γ_i and λ_i. Let us further define $P(n|i) = (1-e^{-\lambda_i})e^{-\lambda_i n}$ which corresponds to an exponentially decreasing probability distribution that we aim to introduce to approximately model the support of the distribution on large photon numbers. Under these circumstances, we can compute importance functions for each of the variables such that

$$\mathrm{Tr}\left([O(\alpha)\otimes P_i]\,\rho\right) = \frac{2}{\pi}\sum_n (-1)^n P(n|i) \times\left[\mathrm{Tr}\left(\left[D(\alpha)\,|n\rangle\langle n|\,D^\dagger(\alpha)\otimes P_i\right]\rho\right)/P(n|i)\right]. \tag{253}$$

We see that this takes the form of an expectation value of a random variable $(-1)^n(\mathrm{Tr}([D(\alpha)\,|n\rangle\langle n|\,D^\dagger(\alpha)\otimes P_i]\rho)/P(n|i))$. We can then estimate the sum by computing the mean of this weighted quantity. The variance in the mean is a simple function of the population variance,

which is given by

$$\begin{aligned}\sigma_i^2 &:= \frac{4}{\pi^2}\sum_n \left[\mathrm{Tr}([D(\alpha)\,|n\rangle\langle n|\,D^\dagger(\alpha)\otimes P_i]\rho)^2/P(n|i)\right] \\ &\quad - \mathrm{Tr}\left((O(\alpha)\otimes P_i)\rho\right)^2 \\ &\leq \frac{4}{\pi^2}\left(\sum_n e^{2\gamma_i-\lambda_i n}/(1-e^{-\lambda_i})\right) \\ &= \frac{4e^{2\gamma_i}}{\pi^2(1-e^{-\lambda_i})^2}. \end{aligned} \tag{254}$$

It follows that the number of samples needed to estimate the error within sample variance ϵ^2 is

$$N_{\mathrm{samp}} \leq \frac{\sigma_i^2}{\epsilon^2}. \tag{255}$$

From the Chernoff bound, it is straightforward to see that the scaling of the number of samples needed to learn just one point in the phase-space density within error ϵ and probability of failure at most δ under these assumptions scales as $O(\sigma_i^2\log(1/\delta)/\epsilon^2)$.

Unfortunately, as the space of the oscillator is formally unbounded unless cutoffs are imposed, we would need to repeat this an infinite number of times to learn the Wigner function even inside some compact region of phase space. Instead, we can assume that the underlying distribution is twice differentiable with the Hessian matrix of the Wigner function satisfying $\max_i \|H_W\| \leq \Gamma$. This assumption can be justified when the boson number is strictly bounded above by a constant because the Hermite polynomials used to represent a photon-number eigenstate are each twice differentiable and a finite sum of twice differentiable functions is differentiable. This also implies that the wave function vanishes superexponentially quickly outside a finite phase-space volume because the eigenfunctions decay as a Gaussian multiplied by a polynomial of bounded degree. Under these circumstances, we can bound the error in any phase-space integral over a finite region A_i of area $|\Delta_A|^2$ and further assume that each of the lengths of the regions is at most a constant multiple of all of the others for simplicity (otherwise, we would need to specifically maximize over the maximum lengths rather than stating the bound more simply in terms of the area). From the midpoint rule, if α is the midpoint of the region, then

$$\left|\int_{x\in A_i} W(x)\mathrm{d}x - W(\alpha)|\delta|^2\right| \in \mathcal{O}(|\Delta_A|^4\Gamma). \tag{256}$$

Thus, if we are promised that ρ only has support within a bounded region of phase space with area A, then there are at most $A/|\Delta_A|^2$ such regions that we need to evaluate. Thus, if we wish to obtain an error of ϵ in any phase-space volume, then the area of each of the small phase-space regions, A_i, is

$$|\Delta_A|^2 \in \mathcal{O}\left(\frac{\epsilon}{\Gamma A}\right). \tag{257}$$

Then in turn the number of samples is simply the number of points in the grid multiplied by the required samples per point. This gives us

$$N_{\mathrm{samp}} \in \widetilde{\mathcal{O}}\left(\frac{A^3\Gamma^2\max_i\sigma_i^2\log(1/\delta)}{\epsilon^4}\right). \tag{258}$$

Here, we say that $f \in \tilde{\mathcal{O}}(g)$ if f is asymptotically less than or equal to g up to polylogarithmic multiplicative factors. This process needs to be repeated for each of the P_i in the Pauli basis. So if we wish to compute each of the conditional Wigner functions within this fixed error budget, then from the union bound we need to perform each measurement within error $\epsilon/4^n$ to ensure that the total error is at most ϵ. Thus the total number of measurements needed is in

$$N_{\mathrm{samp}} \in \widetilde{\mathcal{O}}\left(4^n\frac{A^3\Gamma^2\max_i\sigma_i^2\log(1/\delta)}{\epsilon^4}\right). \tag{259}$$

This shows that the error in reconstructing a tomographically complete set of conditional Wigner functions for the state is comparable to the cost of a tomographic reconstruction of a state under similar assumptions if the variance is modest and the second derivative is also small. In contrast, even a small qubit system may be intractable to characterize if these continuity and variance assumptions are not met.

Next, let us consider specializing this problem to the problem of ancilla-assisted process tomography. Ancilla-assisted process tomography has long been used in qubit systems, and the idea behind it is to use the Choi-Jamiołkowski isomorphism [405,406] to construct a state that is equivalent to the quantum channel that we wish to characterize. We then can learn the channel by tomographically reconstructing the corresponding state. The Choi state is operationally constructed for a qudit system of dimension D_c in the following manner. First, an entangled state of the form

$$|\psi\rangle_c := \frac{1}{\sqrt{D_c}}\sum_{j=0}^{D_c-1}|j\rangle\,|j\rangle \tag{260}$$

is prepared. Next, we apply a quantum channel, Λ, that we wish to characterize to one of the two subspaces:

$$\frac{1}{D_c}\sum_{jk}\Lambda(|j\rangle\langle k|)\otimes|j\rangle\langle k|\,. \tag{261}$$

Each j and k on the rightmost qudit then flags the density matrix found by acting on a complete operator basis,

which, using the linearity property of quantum channels, maps is sufficient to reconstruct the action of the channel on any density matrix.

We aim to repeat the same logic here for hybrid qubit-oscillator systems. First, we need to construct an analog of the Choi state. The most natural input would be to consider the following unnormalizable state acting on two copies of the oscillator space and the qubit space:

$$|\psi_c\rangle = \sum_i \int_{-\infty}^{\infty} dq(|q\rangle\,|q\rangle)(|i\rangle\,|i\rangle), \tag{262}$$

where the integral is over the oscillator position q. Conceptually, this state is nearly identical to the previous proposal, but the fact that the state needs to be distributed over all values of q necessitates infinite squeezing, which is unphysical. The uniform superposition over the qubit states $|i\rangle$ can be efficiently prepared using Hadamard transformations, which are conceptually much less problematic. We can address this issue with the oscillator states by considering the use of states with finite squeezing and aim only to reconstruct the state over a region of fixed cutoff. Specifically, we replace the state $|\psi_c\rangle$ with the state

$$\int_{-\infty}^{\infty} \frac{e^{-x^2/2\zeta^2}}{\pi^{1/4}\sqrt{\zeta}}\,|x\rangle\,|0\rangle\,\mathrm{d}x \otimes \left(\frac{1}{\sqrt{2^n}}\sum_i |i\rangle\right)|0\rangle\,. \tag{263}$$

Next, by performing the two-mode SUM gate (Box IV.6) and n controlled-NOT operations on the qubits we can transform this to

$$|\psi\rangle_c' = \int_{-\infty}^{\infty} \frac{e^{-x^2/2\zeta^2}}{\pi^{1/4}\sqrt{\zeta}}\,|x\rangle\,|x\rangle\,\mathrm{d}x \otimes \left(\frac{1}{\sqrt{2^n}}\sum_i |i\rangle\,|i\rangle\right). \tag{264}$$

We can then follow the same reasoning as employed before: we apply a channel to the first qubit oscillator pair to produce a state of the form

$$\mathcal{C} := \frac{1}{2^n}\int_{-\infty}^{\infty} \mathrm{d}x\,\mathrm{d}y \frac{e^{(-x^2-y^2)/2\zeta^2}}{\sqrt{\pi}\zeta} \times \sum_{ij} \Lambda(|x\rangle\langle y| \otimes |i\rangle\langle j|) \otimes |x\rangle\langle y| \otimes |i\rangle\langle j|\,. \tag{265}$$

In the limit as $\zeta \to \infty$ we see that this gives the original definition of the Choi state; however, in this context, the resulting state is not only more physically realistic (states of infinite squeezing are experimentally impossible to prepare) but also is confined to a finite region following the analysis given previously to discuss the problem of quantum state tomography of the hybrid system.

Next, we have to argue about the truncation of the function. We have that if we wish to truncate the positions to $(x,y) \in \kappa[-\zeta,\zeta] \times [-\zeta,\zeta]$ for $\kappa \geq 1$, the maximum error can be computed by considering the volume of the integrated Gaussian that is neglected from the truncated integral. This integral is given by the error function and from the asymptotics of the error function, we have that the neglected portion of the integral is in $\mathcal{O}(e^{-\kappa^2\zeta^2})$ which, if we set this error to be ϵ_{trunc}, then it suffices to take $\kappa \in \mathcal{O}(\zeta^{-1}\sqrt{\log(1/\epsilon_{\text{trunc}})})$. Thus we can always construct a finite region in position such that the support of the state is well defined within that space and the remainder is small.

With such a Choi state we can then compute the channel fidelity with the ideal channel. This process, discussed in Ref. [407], involves the introduction of several operators that do not have direct physical interpretations that are needed to reshape the information in the Choi state (or Choi matrix) into an appropriate form. The first of these operators, F, is the Choi representation of the transpose map T. The projection onto the symmetric space is then given by $P_S = (1+F)/2$. Finally, let λ_1 be the maximum eigenvalue of $P_S(T\otimes I)(\hat{C})P_S$. If we assume that we discretize our state as $\hat{C}$ then the minimum gate fidelity is given by

$$\begin{aligned}\mathcal{F}_{\min} &= \lambda_1 - \|[\lambda_1 P_S - P_S(T\otimes I)(\hat{C})]P_S)\|_{S_1} \\ &= \lambda_1 - \|[\lambda_1 P_S\|_{S_1} + \|P_S(T\otimes I)(\hat{C})]P_S)\|_{S_1},\end{aligned} \tag{266}$$

where $\|X\|_{S_1} := \sup_{|\psi\rangle,|\phi\rangle} \langle\psi|\,\langle\phi|\,X\,|\psi\rangle\,|\phi\rangle$. The first term in the evaluation of $\mathcal{F}_{\min}$ can be computed as follows. If we let $[X]_W$ be the discretized Wigner representation of an operator X and $\hat{C}_W$ be our discretized Wigner representation of the Choi matrix, then (neglecting discretization error)

$$\begin{aligned}&\|P_S(T\otimes I)(\hat{C})P_S\|_{S_1} \\ &\quad = \sup_{|\psi\rangle,|\phi\rangle} \mathrm{Tr}([P_S(|\psi\rangle\langle\psi| \otimes |\phi\rangle\langle\phi|)P_S]_W[(T\otimes I)(\hat{C})]_W),\end{aligned} \tag{267}$$

where here the trace is taken to be over the discretized positions of the Wigner function as well as the qubit values. This shows that if we perform discretized Wigner function tomography to learn the transposed Choi state $(T\otimes I)(\hat{C})$, we can then find the channel fidelity by optimizing overall states. This optimization can be carried out using semidefinite programming in time that is polynomial in the dimension of the discretized operator.

The next step is to show how we can estimate the discrete Wigner for a hybrid qubit-oscillator state. Specifically, we perform the multi-dimensional Wigner transformation of the channel to learn the Wigner representation of the result. Specifically, the Wigner representation of the operator is (see equivalent definitions of the Wigner

function in Appendix E 6)

$$W(\mathbf{p},\mathbf{q}) = \int_{\mathbf{x}\in\mathbb{R}^N} d\mathbf{x}\, \langle \psi(\mathbf{q}+x/2)|\, e^{i\mathbf{p}\cdot\mathbf{x}}\, |\psi(\mathbf{q}-x/2)\rangle\,. \tag{268}$$

We see from the dot product in the exponential, that the integral computed for the multi-dimensional Wigner transform is simply the nested integral over each of the components individually. Thus, the generalization of the formula for the value of the conditional Wigner function for the multi-dimensional case is simply

$$W(\boldsymbol{\alpha}|P_i) = \mathrm{Tr}([O(\boldsymbol{\alpha})\otimes P_i]\rho), \tag{269}$$

where $O(\boldsymbol{\alpha})$ is the higher-dimensional generalization of the one-dimensional analog considered earlier in Eqs. (249)–(252):

$$O(\boldsymbol{\alpha}) = \bigotimes_{j=1}^{\dim(\boldsymbol{\alpha})} \left(\frac{2}{\pi}\sum_{n=0}^{\infty}(-1)^n D(\alpha_j)\,|n\rangle\langle n|\,D^\dagger(\alpha_j)\right), \tag{270}$$

where $\dim(\cdot)$ represents the dimension of a vector, and α_j is the jth component of $\boldsymbol{\alpha}$. As before, reconstructing the channel from the Choi state is then a simple matter of repeating the previous tomographic protocol.

By repeating this process, we can perform Wigner function tomography to learn the input and output pairs of Wigner functions that describe all gate operations, fiducial preparations, and measurements. Thus, up to unobservable gauges, we can perform quantum gate set tomography on such a system to characterize every possible operation in a hybrid bosonic qubit device. This shows that, while no randomized benchmarking-like protocol has yet been developed, protocols can be easily constructed that resemble gate set tomography [408,409]. A full characterization of the gauge orbits that are present in such a reconstruction is left for future work.

V. QUANTUM ERROR CORRECTION WITH OSCILLATORS

QEC is essential for mitigating noise, faulty gates, and decoherence in a functional quantum computer. In the hybrid CV-DV architecture (Fig. 2), QEC sits between the compiler and hardware layers, ensuring reliable logical qubits or oscillators and operations, in contrast to the noisy physical counterparts controlled via the low-level ISA. Various error-correction strategies exist, leveraging multi-level quantum systems such as oscillators to enable flexible encoding choices—qubits, qudits, or encoded oscillators—based on data requirements.

Encoding options differ based on machine model. For example, one approach that will be discussed in this section is to encode a two-level system in an oscillator, protecting an abstract qubit from errors locally (or "at the hardware level"). This model, though oscillator-based, corresponds to an abstract qubit model that provides the user access to effective DV degrees of freedom corresponding to an abstraction of the underlying CV-DV hardware. This is one example of how QEC can change abstractions. As another example, a qudit can be encoded in a single oscillator or an oscillator in multiple oscillators. If used for bosonic system simulations, this forms an abstract oscillator model. However, qudit encodings qualify as effective oscillators only when the encoded system has enough levels for efficient bosonic simulations; otherwise, they primarily reduce resource overhead compared to an abstract qubit model (forming, instead, an abstract qudit model).

The remainder of this section is organized as follows. Section V A gives an overview of the noise models for oscillators. As it is more widely known, we do not explicitly discuss noise models on qubits, and instead refer to Ref. [1]. Section V B gives a comprehensive overview of encoding options for each machine model. Next, Sec. V C explores ISA compatibility with different encodings using schemes from previous works. As the QCMM is the most developed error-correcting model using oscillators, widely classified through translation symmetric and rotational symmetric codes, we focus on compiling logical instruction sets for qubit encodings. This discussion highlights how ISAs and AMMs influence encoding choices and vice versa. Once a QEC code is selected, encoding circuits, syndrome measurements, and logical operations are compiled into the high-level ISA as primitive resources. In Fig. 15 we illustrate how a hybrid architecture integrates encoded logical layers for both DV and CV computation from the physical layer.

A. Dissipation and decoherence

Like qubits, resonators suffer amplitude damping. For a qubit, amplitude damping takes the excited state to the ground state at rate γ. For a collection of L qubits, the typical energy and typical energy loss rate scale linearly in L. For an oscillator, amplitude damping causes the average energy to relax at rate κ:

$$\frac{d}{dt}\langle\hat{n}\rangle = -\kappa\,\langle\hat{n}\rangle. \tag{271}$$

This equation demonstrates why it is not advantageous to encode many qubits in a single oscillator—the mean boson number and therefore the excitation loss rate in a typical state scales *exponentially* in the number of qubits L. Nevertheless, such replacements can have significant advantages for modest L because, as discussed in the next section, superconducting microwave resonators storing bosonic codes have very desirable error-correction capabilities.

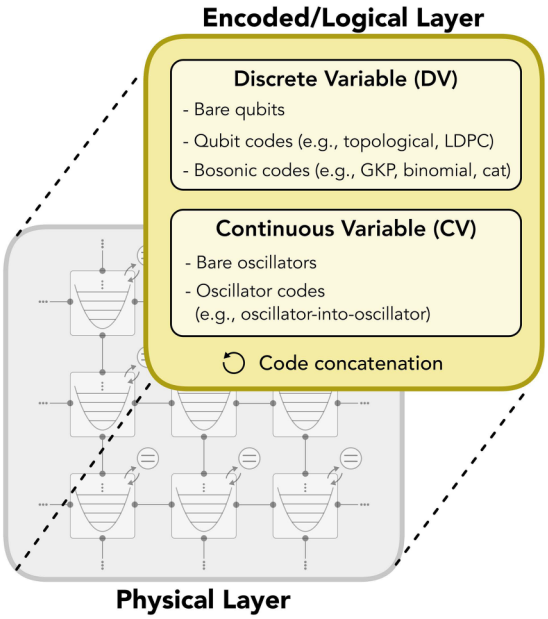


FIG. 15. The QEC stack. In place of computation using physical CV and DV elements, it is beneficial to instead compute using encoded CV and DV computational objects that enable quantum error detection and/or correction (also see Fig. 2). Possibilities for DV codes include both qubit codes (DV to DV) and bosonic codes (CV to DV), while CV codes include oscillator-to-oscillator codes. These codes can be recursively "stacked" using what is known as code concatenation.

Damping and decoherence of quantum systems coupled to a bath are typically described using the master equation for the density matrix derived by making the Born-Markov approximation on the assumption that the coupling to the bath is weak and the bath is memoryless. The Lindblad form of the master equation guarantees that the time evolution of the density matrix corresponds to a completely positive trace-preserving (CPTP) map

$$\frac{d}{dt}\rho = -i[H,\rho] + \sum_j \mathcal{D}(E_j)\rho, \tag{272}$$

where $\mathcal{D}(\mathcal{O})$ is a "superoperator" the action of which on the density matrix is given by

$$\mathcal{D}(\mathcal{O})\rho = \mathcal{O}\rho\mathcal{O}^\dagger - \frac{1}{2}\left[\mathcal{O}^\dagger\mathcal{O}\rho + \rho\mathcal{O}^\dagger\mathcal{O}\right] \tag{273}$$

and the E_j are "jump" operators acting on the Hilbert space of the oscillator and describing the effects of coupling to the bath.

For oscillator Hilbert-space dimension $d = N_{\max} + 1$, the density matrix ρ and the jump operators are $d \times d$ matrices. The density matrix is Hermitian, positive semidefinite, and has a unit trace, but the jump operators need not be Hermitian. In numerical solutions of the Lindblad master equation, it is sometimes convenient to "vectorize" the density matrix, turning it into a vector of length d^2. The advantage of this is that it allows the superoperators to be written as ordinary (very large, typically sparse) $d^2 \times d^2$ matrices and $\mathcal{D}(\mathcal{O})\rho$ becomes ordinary matrix multiplication into a vector.

Typical jump operators for a microwave resonator include amplitude damping

$$\sqrt{\kappa(\bar{n}+1)}a, \tag{274}$$

describing the loss of energy due to (linear) coupling to the bath, and heating,

$$\sqrt{\kappa\bar{n}}a^\dagger, \tag{275}$$

describing the gain of energy from the bath. Here, κ is the rate of energy (not amplitude) damping of the oscillator and $\bar{n}$ is a phenomenological parameter giving the steady-state mean excitation number in the oscillator caused by its interaction with the (possibly nonequilibrium) bath. If the above two jump operators are the only ones present in the dynamics, then the oscillator number distribution in the steady state becomes the Bose-Einstein distribution and we have

$$\bar{n} = \frac{1}{e^{\beta\hbar\omega} - 1}, \tag{276}$$

where $\omega = 2\pi f$ is the oscillator (angular) frequency and $\beta = 1/k_B T$ is the inverse temperature. Since $k_B/h \sim 21$ GHz/K, a typical microwave resonator with $f \sim 5$ GHz should have all its excitations frozen out at a dilution refrigerator temperature of 20 mK. This tends to be a relatively good approximation for resonators but not for qubits, as the latter suffer more strongly from nonequilibrium heating effects and can be in their first excited state with a probability as high as ∼1–10%.

Unlike superconducting qubits, superconducting resonators typically do not suffer much from intrinsic frequency fluctuations that lead to dephasing (without energy relaxation). But if such terms are present and if the frequency fluctuations $\delta\omega(t)$ are classical Gaussian white noise having a flat spectral density [410] κ_φ such that

$$\langle\langle \delta\omega(t)\delta\omega(t')\rangle\rangle = \kappa_\varphi \delta(t-t'), \tag{277}$$

where the double brackets refer to the ensemble average over the noise, then the corresponding jump operator would be

$$\sqrt{\kappa_\varphi}\hat{n}, \tag{278}$$

where $\hat{n} = a^\dagger a$ is the number operator. Under the influence of pure dephasing, $\hat{n}$ is a constant of the motion but the

amplitude autocorrelation function,

$$\langle a^\dagger(t)a(0)\rangle = \langle \hat{n}\rangle \langle\langle e^{-i\int_0^t d\tau\, \delta(\omega(\tau)}\rangle\rangle = \langle \hat{n}\rangle e^{-\frac{\kappa_\varphi}{2}t}, \quad (279)$$

decays exponentially in time.

In practice, the primary source of dephasing errors in cavities is an extrinsic effect, namely the dispersive coupling of the cavities to the transmon qubits used to control them. An unexpected change in the transmon state due to coupling to its environment produces a large sudden change in the resonator frequency. This dispersive coupling is discussed further below in Sec. IV D. For a more elaborate discussion on noise mechanisms (including amplitude damping, dephasing, and heating) on different experimental platforms, we refer the readers to Appendix C.

B. Overview of bosonic QEC

A correctable logical qubit requires a large state space to provide the redundancy and entanglement needed to hide and protect quantum information from decoherence processes. For physical-qubit-only models, the large state space of dimension $D = 2^n$ comes from using n physical qubits to form the logical qubit. If a single logical qubit is encoded, the logical code words span a 2D (entangled) subspace of the large state space. The formally infinite-dimensional Hilbert space of an oscillator offers a distinctly different model [62,302,411–416] in which logical code words are superpositions of various (low-lying) photon-number states within just a single mode or a small number of modes [280,417]. Physical errors may be corrected without losing information if they deform the code subspace without distorting the encoded information. The control overhead, i.e., the additional cost in time and circuit depth for creating and correcting logical qubits, together with errors in the imperfect control processes, are the primary limitations in achieving useful QEC (bosonic or qubit-based) in many experiments today. In addition to control hardware, auxiliary registers are required for logical state preparation, error-syndrome measurement, and logical state readout. Error correction of logical qubits encoded into bosonic modes has recently been achieved at and beyond the break-even point (for memory) where the quantum coherence of the error-corrected logical qubit exceeds the coherence of the best of the available uncorrected qubits in the system [63,65,68,418] so that the error-correction gain $G > 1$ exceeds unity.

Protection of quantum memories is an important goal but large-scale fault-tolerant (FT) computation will also require noise-resilient logical operations as well as minimal overhead. An FT architecture is one in which errors in all inherently noisy procedures, such as state preparation, gates, measurement, and the error-correction process itself, propagate only to a limited extent, such that they can be corrected in subsequent rounds of QEC. The road to fully FT computation with superconducting circuits is long, but there is steady progress in creating better bosonic circuit components and protocols such as FT detection of quantum errors [419], "error-transparent" logical gates [67,420,421], and "path-independent" gates [180,422,423].

The dominant noise source in microwave oscillators is photon loss, which provides a simple error model, particularly since, unlike the multi-qubit case, all the information and the errors occur in a single mode. Other noise such as heating, dephasing, and displacements may also be corrected with QEC. Noise is inherently platform dependent, and the reader is directed to Appendixes A and C for more details about platform-specific noise and noise models.

Discrete-variable QEC requires at a minimum $n = 5$ physical qubits to construct one logical qubit [424] that can correct a general single-qubit error. For the case of only amplitude damping noise, $n = 4$ qubits (corresponding to Hilbert-space dimension $D = 16$) suffice to correct the damping to lowest order [425]. Because there are many physical modes (qubits), the noise channel naturally has increased entropy associated with (missing) information about the physical location of the errors. By way of comparison, the smallest single-mode bosonic code that corrects amplitude damping [416] to lowest order uses only states with photon numbers ≤ 4 corresponding to a Hilbert-space dimension of only $D = 5$. When combined with necessary strategies to determine the n possible locations of various types of errors (e.g., the three possible Pauli errors on each qubit), reaching fault tolerance in qubit-based codes becomes extremely challenging. Algorithms such as Shor's factoring, which requires at least $\sim$4000 noiseless logical qubits to factor a 2048-bit integer, when considered practically, lead to resource estimations of 3 orders of magnitude more physical qubits [426]. One recent performance analysis of a repetition cat code QEC architecture [427] has estimated that the computation of a 256-bit elliptic curve logarithm could be done in about 9 h with approximately 125 000 highly noise-biased cat qubits.

Bosonic QEC has emerged as a promising strategy for detecting and correcting physical errors directly at the hardware level. By enabling local error correction, it reduces both the decoding complexity and resource overhead typically associated with FTQC [427,428]. This advantage arises from hardware-efficient syndrome extraction and correction protocols, made possible by the relatively simple error model of harmonic oscillators. Furthermore, the infinite-dimensional Hilbert space of bosonic systems provides a vast design space, allowing researchers to tailor encodings to the specific error mechanisms encountered in real-world hardware. Given that the choice of encoding fundamentally influences hardware control—and, by extension, the ISA—this section introduces an overview of the most popular bosonic encodings currently under investigation, with a particular emphasis

on those demonstrated in experimental laboratories. Further details about compiling these codes can be found in Sec. V C.

1. Single-mode encodings

The extension of QEC to bosonic modes was first carried out for amplitude damping (i.e., photon loss errors), using a multimode Fock basis scheme [429] and then a superposition of coherent states (cat states) scheme [430]. Following these, various codes have been developed that leverage more complex phase-space patterns to protect against other oscillator decoherence mechanisms, such as small displacements in GKP codes [64,65,107,282,418], and against arbitrary polynomials of fixed degree in creation or annihilation operators with binomial codes [67,68,416].

Cat codes and binomial codes are a subset of a code class with discrete rotation symmetry [415], whereas GKP codes have discrete translation symmetry. Typically, all these bosonic codes perform better with increasing average photon number (even though the photon loss rate scales linearly with photon number) because a larger code space provides room for more redundancy. This is analogous to qubit-based logical encodings where the error rate increases linearly with the number of physical qubits, but the logical error rate decreases (provided that the physical error rate is below the threshold) because of the increasing code distance. A study of the performance of cat, binomial, and GKP codes under the same photon loss error model [414] has found that GKP codes significantly outperform other codes for most loss rates, and approaches (within a constant factor) the quantum capacity of the Gaussian thermal loss channel [431].

2. Error-correction techniques and break-even performance

Generally, QEC requires auxiliary registers for encoding, error detection, and recovery from errors, but these auxiliary elements also suffer from decoherence and excess measurement backaction. For an entré into the literature on fault-tolerant control of oscillators by qubits, see Refs. [180,422,423,432,433]. One path forward is to use real-time measurement-based feedback [63,64,68,193], or semiautonomous [65,434] or fully autonomous [69,435–439] feedback on the QEC circuit.

A more complete unified approach is to consider the combined quantum channel of the encoding, error, and recovery steps and to design recovery protocols that minimize hardware and measurement overheads. For example, autonomous QEC does not rely on syndrome error measurements but rather, uses an engineered dissipation channel to stabilize some error-syndrome operators. Some protocols require only a single driven auxiliary nonlinear element to create the appropriate dissipative bath [413,440–442] or double-well Hamiltonian [443–448] to stabilize, e.g., two-legged cat states which strongly and autonomously suppress certain errors. References [107,326] have independently proposed an autonomous QEC protocol for GKP states, and Ref. [107] has demonstrated a dissipative pumping technique for GKP code space stabilization in trapped ions, leading to a $\sim 3\times$ longer lifetime of the encoded state than without stabilization. (Note that this is not, however, the error-correction gain G.) Reference [65] has employed a semiautonomous version of the protocol in Ref. [326] that has yielded an error-correction gain (relative to the longest-lived encoding [0, 1 Fock] of the cavity) of $G \sim 2.3$. As decay errors in the auxiliary qubit used to detect error syndromes in the logical GKP qubit still limit the performance of GKP encodings, using a biased-noise auxiliary qubit such as the Kerr-cat [443,445] or fluxonium qubit [449,450] in which decay errors are exponentially suppressed compared to dephasing errors [302] is helpful to further improve the error-correction gain G. Recently, this gain has been seen to be preserved in the first-error-corrected qudit encoding of the $d = 3, 4$ levels using the GKP encoding [66].

3. Multimode encodings

Encoding QEC codes across multiple oscillators offers additional advantages in addressing specific hardware implementation issues. Recently discovered encodings of logical qubits in two oscillators [280,451] have suggested ways around this problem. In Ref. [414], the authors encode a cat qubit in two oscillators where error syndromes can be detected with a circuit fault tolerant to auxiliary-qubit errors. Examples of four-dimensional (4D) symplectic lattices for GKP encoding in two oscillators have been studied in Ref. [280]. The authors show that the tesseract lattice, where all stabilizers are represented by orthogonal translation vectors in 4D phase space, can measure stabilizers using a "corner-circuit" which reduces the errors due to decay of the auxiliary qubit to essentially zero for an ideal GKP code. In practical (i.e., finite-energy) versions of this code, the auxiliary qubit contributes to logical error rates an amount which is inversely proportional to the finite-energy envelope size of the 4D GKP code.

Extensions of these encodings to more than two oscillators have been suggested for cat codes as well as GKP codes in Refs. [280,451–456] as a strategy for scaling up; however, a practical analysis with circuit-level noise is still pending for such encodings. To use these multimode encodings realistically, one must limit the support of the stabilizers to nearest-neighbor modes, as in surface codes, or at least make the code rate constant, as in good LDPC codes (which are challenging because the stabilizers are not geographically local).

Another prominent example of multimode encoding is the Schwinger-boson representation of higher-spin Hamiltonians for the SU(2) group (and generalizations to the SU(N) group [457,458]) described in Sec. VII C 6. Traditionally, such mappings have been important tools for formal analysis in low-energy [459,460] and high-energy physics. As we discuss in Sec. VII C, the inception of hybrid CV-DV quantum processors discussed in the present work can make good use of the Schwinger-boson representations for efficient quantum simulation of spin Hamiltonians on bosonic hardware. Here, we focus on QEC using the Schwinger-boson representation of a spin-1/2 (i.e., a qubit), known in the optics community as the dual-rail encoding [227,233,234,236,241].

The dual-rail encoding has several interesting features, particularly in the cQED architecture, where important gate, control, and measurement capabilities are more readily available than in traditional quantum optics. First, as discussed in Sec. VII C 6, arbitrary single-qubit rotations can be achieved using (microwave pulse activated) beamsplitter couplings between the two rails (described in Box IV.4). Second, for superconducting resonators the dominant intrinsic error is (slow) photon loss. This leakage error out of the code space (span$\{|01\rangle, |10\rangle\}$) into the error state $|00\rangle$ can be efficiently converted to an erasure error (i.e., an error in a qubit the location of which is heralded) by QND detection of the parity of the joint photon number [233,240,394] $\hat{P} = e^{i\pi(\hat{n}_1+\hat{n}_2)}$ (which changes from -1 in the code space to $+1$ in the error state). This is important because erasure errors are much more readily corrected (by concatenation of the dual-rail qubits into a next-level code such as the surface code) than errors the location of which is unknown [245,246], thus leading to substantially higher error-correction thresholds if erasure errors are dominant.

4. Oscillator-to-oscillator encoding

Rather than encoding DV logical qubits in CV oscillators, an interesting but extremely challenging alternative would be to attempt to protect the entire bosonic Hilbert space from Gaussian errors by encoding an oscillator into many oscillators. If this were possible, one could carry out, e.g., error-corrected boson sampling. It has been shown that Gaussian errors cannot be corrected using purely Gaussian quantum resources [323,461]. This no-go theorem stands in contrast to the result that in DV error-correction codes, Clifford operations and measurements are sufficient to correct both Clifford and non-Clifford errors. To circumvent this no-go theorem, there has been a proposal to encode an oscillator into many oscillators by using GKP states as a non-Gaussian resource [462], and to use GKP codes as auxiliary qubits. In Ref. [463] this code has been generalized to a GKP two-mode-squeezing code. The problem of optimal oscillator-to-oscillator encoding has been addressed in Refs. [463,464]. However, Ref. [465] shows that oscillator-oscillator code families do not have a threshold for the classical displacement channel, a noise model assumption used in Ref. [462]. This point is discussed further in Sec. IX.

5. Code concatenation

Among the leading candidates for DV codes in practice today are surface codes, which have been shown to possess good error-rate thresholds for FTQC [466]. A high threshold enables the multi-qubit code to achieve error-correction gain with higher physical-level error rates. Recent demonstrations of beyond break-even error correction with GKP and binomial bosonic code qubits [65,68] suggest that it would be fruitful to concatenate bosonic codes with DV codes. Using bosonic codes at the physical level could achieve physical error rates well below the threshold for the next-level DV code, allowing exponential reduction in the logical error rate with increasing size and distance of the DV code, thereby yielding reduced resource overhead for fault-tolerant error correction. Reference [461] has proposed a decoding technique for GKP qubits concatenated with the surface code, while Refs. [456,467,468] have studied how to increase the fault-tolerance threshold for GKP-surface encoding. In Refs. [302,469], the authors lay out perspectives and detailed approaches toward all GKP-surface codes and hybrid GKP-surface codes where the GKP states used at the base layer are stabilized by a biased-noise auxiliary qubit.

This can further be extended to concatenation of cat codes with rectangular surface codes [470]. Another advantage of CV-DV encoding has been studied in Ref. [471] where, using biased noise Kerr-cat qubits in concatenation with the $XZZX$ surface codes, the authors show a significant reduction in overhead for magic state distillation, a key ingredient for FTQC with discrete-variable encoding. Concatenation of the biased-noise cat code into a repetition code has been studied in Refs. [427, 472–474]. Reference [428] has shown that concatenation of GKP code with generic quantum LDPC code can surpass the Calderbank-Shor-Steane (CSS) Hamming bound by exploiting the analog information in the GKP code. Reference [475] has shown the advantages of using cat codes with classical LDPC codes.

Despite the infinite Hilbert space of bosonic modes, the GKP code requires only two stabilizers to reduce this code subspace to dimension to $d = 2$. This is because the stabilizers (which are translations by the lattice constant in phase space) have a continuum of eigenvalues on the unit circle in the complex plane. Although CV codes seek to encode DV information into CV modes, the intrinsic analog character of CV error syndromes could be exploited further in CV-DV code concatenation to improve the overall performance of the hybrid system. Reference [476]

has demonstrated a first example of this where analog information on the GKP qubits is used in concatenation with a CSS code using a maximum-likelihood method. The authors have shown that the resulting CV-DV hybrid code can correct double-bit-flip errors (as opposed to only single-bit-flip errors in usual DV codes) in a three-qubit system. They have also shown superior performance on quantum channel capacity with this concatenation scheme. Reference [461] has connected the analog character of the maximum-likelihood decoding problem for the GKP-toric code to a Feynman path integral for lattice gauge theory.

In Ref. [477] CV-DV code concatenation has been studied and compared with other types of concatenation using bosonic codes concatenated into smaller second-level codes such as the repetition code and the Steane code. The authors have found that to encode quantum information in a way that can correct both logical X and Z errors, the above-mentioned CV-DV–type encoding performs the best. Such code concatenation is yet to be compared with the practical multimode encoding discussed in the previous subsection. An analysis of this nature yields the optimal strategy for scaling up toward fault-tolerant error correction using oscillator codes.

C. Logical ISA for bosonic QEC

In this section, we present compilation schemes for selected bosonic QEC codes using the various hybrid instruction sets. Particular codes are most naturally implemented using particular ISAs. For example, translation symmetric codes, such as the GKP codes shown in Fig. 16(a), are most natural in the phase-space ISA, while the rotation symmetric codes, like four-legged cat codes, shown in Fig. 16(b) are most natural in the Fock-space ISA. In addition to the four-legged cat codes we also discuss the binomial codes and dual-rail codes which fall under the umbrella of rotationally symmetric codes. We note that our classification of dual-rail codes as rotationally symmetric is not conventional; our reasoning for this classification lies in the fact that the logical code words of dual-rail codes are based on Fock states which are rotationally symmetric states.

Importantly, throughout this section, we compile logical instructions using the hybrid CV-DV architecture where the qubits are used as ancillary systems and the oscillator as our encoding for data within the code space. Preliminary information required for the discussion in this section is given in Sec. V C 1. We present examples illustrating all components of bosonic QEC necessary for universal quantum computation with logical qubits: preparation (Sec. V C 2), readout (Sec. V C 3), stabilization (Sec. V C 4), arbitrary single-qubit rotations (Sec. V C 5), and entangling operations (Sec. V C 6). The definitions of logical code words for each bosonic QEC code are given in Sec. V C 2.

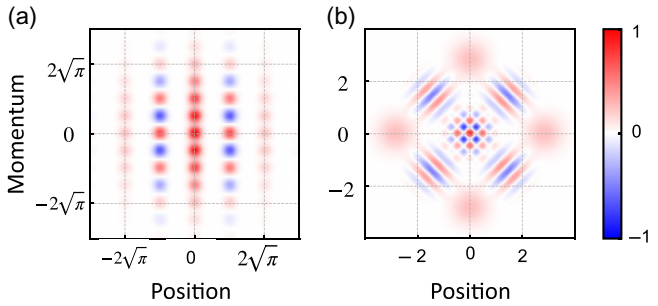


FIG. 16. Wigner plots of two types of single-mode bosonic stabilizer codes, (a) a translationally symmetric GKP code and (b) a rotationally symmetric code cat code. The horizontal axis represents the oscillator position (in standard units $\hat{x} = (a + a^\dagger)/\sqrt{2}$) and the vertical axis represents the oscillator momentum ($\hat{p} = (a - a^\dagger)/\sqrt{2}i$). The color bar represents the magnitude of the quasiprobability distribution in phase space. Details on Wigner function tomography can be found in Appendix E 6. (a) The logical $|+Z_L\rangle$ state of the finite-energy GKP code. In the limit of infinite lattice size (and therefore infinite energy), the states are invariant when shifted by $2\sqrt{\pi}$. This property holds for all (ideal) GKP code words, and is used for error correction against translation (displacement) errors. A method to convert a bosonic damping channel to a displacement channel can be found in Ref. [282]. It has been shown that these codes are optimal in correcting for photon loss under ideal recovery [414]. (b) The logical $|+Z_L\rangle$ state of the four-legged cat code comprises a superposition of four coherent states. These codes have a fourfold rotational symmetry. An n-legged cat state possesses an n-fold rotational symmetry. Such codes can efficiently correct for rotational errors (i.e., dephasing). For example, in Ref. [415], the authors have proposed the number-phase code, a rotationally symmetric code that is efficient in correcting for both oscillator dephasing and photon loss, under ideal recovery.

We primarily review compilation schemes based on the phase-space ISA that utilize the methods of non-Abelian QSP developed in Ref. [478]. In addition, we give examples from Ref. [235] of compilation schemes using the Fock-space ISA. For details on these schemes, we direct readers to these key references.

1. Preliminary information

a. Figure of merit. In practice, the conditional displacements of the phase-space ISA generally involve displacing the oscillator by a large range of distances of order $(-\sqrt{n}_{\max}, \sqrt{n}_{\max})$, where $n_{\max}$ is the photon-number cutoff we impose on the oscillator. These gates thus require a large range of times during which errors on the oscillator-ancilla system occur at an approximately constant rate. Hence, the gate error probability is directly proportional to the duration of the gate. A good metric of efficiency for circuits involving such operations should take into account the total duration of the circuit T_{circuit} rather than just the

gate count. This metric incorporates the fact that conditional displacement gates in the layers of Fig. 30(b) are not instantaneous, but require a time proportional to $|\gamma_i|$. Typical time scales [60] for qubit rotations are 24 ns, while the duration of conditional displacements CD(α) is $\approx |\alpha|$ μs. The ratio of the two types of gate times ($T_R/T_{|\gamma_i|}$) is motivated from experimental parameters used in Ref. [60]. Similar arguments apply to other parametrized gates in the various ISAs.

b. Non-Abelian QSP primitives. Many compilation techniques in this section use the entangling sequence described in the cat state preparation example in Sec. VI D 3 b. Here, in addition to cat codes, we consider other codes, such as GKP codes, the code words of which are also represented by a sum of identical displaced Gaussian states $|\mathcal{G}_i\rangle$ with insignificant overlap [see Eqs. (284)–(286) further below]

$$|\psi\rangle = \sum_i c_i \, |\mathcal{G}_i\rangle_{\text{osc}} \, . \tag{280}$$

We define the sequences $\mathcal{U}_v, \mathcal{E}_v$ that, respectively, disentangle and entangle the ancilla qubit with the oscillator state as

$$\mathcal{U}_v \left[\frac{\theta}{\alpha}, \Delta\right] = e^{i\frac{\theta}{\alpha}\hat{v}\otimes\vec{\sigma}\cdot\hat{n}_\phi} \times e^{i\frac{2\theta}{\alpha}\Delta^2\hat{v}_\perp\otimes\vec{\sigma}\cdot\hat{n}_\gamma}, \tag{281}$$

$$\mathcal{E}_v \left[\frac{\theta}{\alpha}, \Delta\right] = e^{i\frac{2\theta}{\alpha}\Delta^2\hat{v}_\perp\otimes\vec{\sigma}\cdot\hat{n}_\gamma} \times e^{i\frac{\theta}{\alpha}\hat{v}\otimes\vec{\sigma}\cdot\hat{n}_\phi}, \tag{282}$$

where

$$\hat{v} = \hat{x}\cos\xi + \hat{p}\sin\xi, \quad \hat{v}_\perp = \hat{p}\cos\xi - \hat{x}\sin\xi. \tag{283}$$

Here, $\Delta^2 = \langle\hat{v}^2\rangle - \langle\hat{v}\rangle^2$ is the variance of the individual Gaussian states in $|\psi\rangle$. In general, for a squeezed state, this depends on the magnitude and orientation of the squeezing. Note that the choice of ξ is also determined by the orientation of the squeezing. The ancilla Bloch sphere vector $\hat{n}_\phi$ depends on the protocol. The vector $\hat{n}_\gamma$ is decided by the final (initial) state of the disentangled ancilla for $\mathcal{U}_v$ ($\mathcal{E}_v$). Formally, for $\mathcal{U}_v$ ($\mathcal{E}_v$) we need $\hat{n}_\gamma$ to satisfy $\vec{\sigma}\cdot\hat{n}_\gamma\,|\zeta\rangle = i\vec{\sigma}\cdot\hat{n}_\phi\,|\zeta\rangle$, given that the ancilla final (initial) disentangled state is $|\zeta\rangle$.

2. Logical state preparation

High-fidelity fast preparation of states encoded in the logical manifold, popularly known as logical states, is a key requirement for QEC. Any error in this preparation leads to an additional burden on error correction. Here, we discuss the preparation of various bosonic code words using gates from the phase-space and Fock-space instruction sets.

In using the phase-space instruction set, we more broadly employ the tricks used in the specific example of cat state preparation, discussed in Sec. VI D 3. The initial hybrid state for the preparation circuits is $|0\rangle_{\text{vac}} \otimes |0\rangle$. We use the definition of target state fidelity in this case as $\mathcal{F} = |\langle\phi|\psi, g\rangle|^2$, where $|\psi\rangle$ is the target oscillator state and ϕ is the hybrid oscillator-ancilla state at the end of the preparation circuit. Thus, the goal of the preparation circuits is to disentangle the oscillator from the ancilla qubit without measurements, such that the qubit is in a known state at the end of oscillator state preparation. Several of these preparation schemes are on par with the numerical circuits given in Sec. VI E in terms of fidelity and circuit duration. In addition, analytical knowledge of these preparation schemes leads to additional applications in QEC, as we discuss below.

a. GKP code. The physically realistic normalizable version of the GKP code words defined in Eqs. (211)–(212) (see the Wigner function plot for $|0\rangle_{\text{GKP}}$ in Fig. 16) is defined as follows (where we have left the required normalization constants implicit):

$$|0\rangle_{\text{GKP}} = \hat{E}_\Delta \sum_{m\in 2\mathbb{Z}} \left|m\sqrt{\pi}\right\rangle_x, \tag{284}$$

$$|1\rangle_{\text{GKP}} = \hat{E}_\Delta \sum_{m\in 2\mathbb{Z}} \left|(m+1)\sqrt{\pi}\right\rangle_x, \tag{285}$$

$$\hat{E}_\Delta = e^{-\Delta^2\hat{a}^\dagger\hat{a}}. \tag{286}$$

Here, $|.\rangle_x$ is the infinitely squeezed position eigenstates. Similarly, the $\pm$ logical basis is represented by the same superposition of momentum eigenstates (not shown here). The $\hat{E}_\Delta$ operator applies a Gaussian envelope on the infinite support of this state in phase space to restrict the energy of the system to a finite average photon number determined by Δ. This operation takes the infinitely squeezed states $|.\rangle_x$ represented by Dirac delta functions to finitely squeezed states represented by Gaussian functions with standard deviation equal to Δ. Thus, in position basis, the GKP states are a superposition of multiple Gaussian wave functions with the same width displaced from the origin by even and odd integer multiples of $\sqrt{\pi}$.

To prepare these code words with the phase-space instruction set defined in Table V, one can perform multiple rounds of the cat state preparation protocol with displacement $\alpha = \sqrt{\pi}$, which has been discussed briefly in this work in Sec. VI D 3. However, for this preparation, the performance of the circuit is better when the protocol is applied to a squeezed vacuum. Reference [478] has given an efficient state-of-the-art construction for the preparation of squeezed states using phase-space ISA, so here we

take the squeezed states as freely accessible. For logical Z eigenstates, the circuit for k^{th} round is given by

$$\mathcal{C}_k = D_c(\sqrt{\pi}, -\sqrt{\pi})\mathcal{U}_x\left[\frac{\sqrt{\pi}}{4k}, \Delta\right]. \tag{287}$$

Odd (even) k yields a state that has high overlap with $|0\rangle_{\mathrm{GKP}}$ ($|1\rangle_{\mathrm{GKP}}$). N applications of this protocol yield $N+1$ Gaussian wave functions in superposition. The Gaussian width Δ of each state in this superposition is given by the squeezing in the initial vacuum state. The circuit duration T_{circuit} for the preparation scheme is decided by the target Δ and $\mathcal{U}_x$ is optimized for each k in Eq. (287). A squeezed vacuum can be achieved using the sideband ISA as well as the phase-space ISA in combination with various primitives described in Sec. VI such as Trotterization, BCH, etc. The squeezing required for the GKP code word is decided by the target envelope size Δ of the Gaussian functions in the superposition of the desired GKP state. For $\Delta = 0.302$ with $N = 3$ peaks, used in recent experiments [60,65], this method yields an infidelity of $\mathcal{O}(10^{-2})$ using a circuit depth of $T_{\mathrm{circuit}} \approx 3\sqrt{\pi}(1+\Delta^2)$ μs $= 5.84$ μs. However, it should be noted that since the ancilla is reset after every large conditional displacement by $\sqrt{\pi}$, postselection can be performed from measurements of the ancilla after every 1.95 μs. This is the only fully deterministic scheme known to prepare GKP states without numerical optimizations. In addition, the authors have shown that this scheme can be modified to detect errors at various points in the circuit and, upon postselection after each C_k, yield higher fidelity in the presence of faults. In Ref. [479], the authors have prepared GKP states using numerically optimized circuits that resemble the structure of this preparation scheme, i.e., alternating conditional displacements and conditional momentum boosts.

b. Four-legged cat code. The four-legged cat code (see Fig. 16) is a superposition of four coherent states positioned equidistant from the origin. The logical Z code words are defined as cat states along the position axis and the momentum axis. Again with implicit normalization

$$|0\rangle_{4\mathrm{C}} = \left|C_{0,\alpha,+}\right\rangle, \quad |1\rangle_{4\mathrm{C}} = \left|C_{0,i\alpha,+}\right\rangle, \tag{288}$$

where $\left|C_{\beta_1,\beta_2,\pm}\right\rangle$ is an even (+) or odd (−) cat state of size β_2 displaced from the origin by β_1:

$$\left|C_{\beta_1,\beta_2,\pm}\right\rangle = |\beta_1+\beta_2\rangle_{\mathrm{osc}} \pm |\beta_1-\beta_2\rangle_{\mathrm{osc}}. \tag{289}$$

The superposition states $|\pm\rangle_{4\mathrm{C}} = |0\rangle_{4\mathrm{C}} \pm |1\rangle_{4\mathrm{C}}/\sqrt{2}$ are the X-basis four-legged cat code words, as shown in Fig. 16.

Let us discuss the preparation of $|+\rangle_{4\mathrm{C}}$ starting with the hybrid initial state of $|0\rangle_{\mathrm{vac}} \otimes |0\rangle$. Using the phase-space instruction set, we can create the four blobs in phase space as seen in Fig. 16 using a conditional displacement followed by a conditional momentum boost as follows:

$$|\psi_1\rangle = e^{i\alpha(\hat{x}-\hat{p})\otimes\sigma_z} e^{i\alpha(\hat{x}+\hat{p})\otimes\sigma_x} |0\rangle_{\mathrm{vac}} \otimes |0\rangle. \tag{290}$$

Now, to disentangle the ancilla qubit from the oscillator, we can apply the controlled rotation of the ancilla,

$$|\psi_2\rangle = e^{-i\frac{\pi}{4\alpha}(\hat{x}+\hat{p})\sigma_x} |\psi_1\rangle. \tag{291}$$

Following this step, $|\psi_2\rangle$ has a small residual entanglement between the two systems analogous to what is shown in Fig. 28 for the cat state. For two-legged cat states, cancellation of this entanglement has been accomplished using two methods: non-Abelian QSP and a single-variable QSP sequence $\mathrm{BB1}_{90}$ (see Fig. 29 for a performance comparison of the two methods). In the present case, non-Abelian QSP may not apply (see Ref. [478] for details). We therefore use single-variable QSP with the $\mathrm{BB1}_{90}$ scheme for disentangling the four-legged cat state from the ancilla. For $\alpha = 5$, we obtain a fidelity to the target state of $\mathcal{F} = 0.99$. The circuit depth for this sequence is $N = 6$. Reference [478] has shown that if the cat sizes of $|0\rangle_{4\mathrm{C}}$ and $|1\rangle_{4\mathrm{C}}$ in Eq. (288) are unequal—say, α_1, α_2 such that $|\alpha_1|/|\alpha_2|$ is an even integer—a non-Abelian QSP correction can be found.

c. Binomial code. The simplest binomial code [416] can correct at most a single photon loss. It has logical code words defined as

$$|0\rangle_{\mathrm{bin}} = \frac{|0\rangle_{\mathrm{osc}} + |4\rangle_{\mathrm{osc}}}{\sqrt{2}}, \quad |1\rangle_{\mathrm{bin}} = |2\rangle_{\mathrm{osc}}. \tag{292}$$

The preparation of these states can be achieved using the Law-Eberly protocol [480] accessible via the sideband ISA. They are also efficiently realized in experiment [60] using numerically optimized phase-space ISA gate sequences.

d. Dual-rail code. The logical code words of dual-rail encoding [227,241] are represented by a single photon being in one of two distinct modes,

$$|0_{\mathrm{L}}\rangle = |0\rangle_{\mathrm{osc}} \otimes |1\rangle_{\mathrm{osc}}, \tag{293}$$

$$|1_{\mathrm{L}}\rangle = |1\rangle_{\mathrm{osc}} \otimes |0\rangle_{\mathrm{osc}}. \tag{294}$$

These can be distinct spatial or temporal modes of flying qubits, one standing mode in each of two separate resonators [233,234,236], or two standing modes of a single resonator [242]. The same type of encoding for pairs of two-level qubits has also recently been explored [243,244]

but we focus here solely on the bosonic mode case. The preparation of these states can also be achieved using the Law-Eberly protocol accessible via the sideband ISA or using non-Abelian QSP via the phase-space ISA [60,478].

3. Logical readout

A logical readout circuit can be realized by maximally entangling the logical code state in the oscillator with an auxiliary qubit such that the measurement of the auxiliary qubit realizes the measurement of the logical code word. For example, consider the following maximally entangled state:

$$|0\rangle_L|0\rangle + |1\rangle_L|1\rangle, \tag{295}$$

where $|0\rangle_L$ and $|1\rangle_L$ denote a logical bosonic qubit. Measurement of the auxiliary $\{|0\rangle, |1\rangle\}$ qubit in the computational basis is then equivalent to a logical Z measurement of the bosonic qubit. Thus, the task of readout can be accomplished by realizing an entangling sequence $\mathcal{E}$ between the bosonic qubit and the auxiliary qubit, followed by measurement of the qubit. And assuming that the measurement of the auxiliary qubit is perfect, the measurement failure probability is the probability of predicting the wrong logical state in the oscillator given a specific ancilla measurement outcome.

The specific circuits that have been experimentally realized for entangling the auxiliary readout qubit with the logical bosonic qubit state for the dual-rail [234,244], binomial [67,68], four-legged cat [63], and GKP codes [64,65] are described in Sec. V C 5.

Here, as an example, we shall discuss the case of GKP logical readout in detail. Importantly, as we shall see, the GKP codes are the only ones which can be read out in all three logical Pauli bases with equal ease. The inverse of the disentangling gate in Eq. (287) for $\mathcal{C}_1$ can be used as the entangling gate between the GKP qubit and the ancilla qubit.

The GKP readout sequence is as follows. As explained in Sec. V C 2, GKP $|0\rangle_L$ ($|1\rangle_L$) is a superposition of multiple Gaussian functions with peaks located at even (odd) multiples of $\sqrt{\pi}$. The logical GKP readout requires the ancilla qubit to rotate to orthogonal qubit states depending on whether the locations of the peaks in the wave function are at even or odd multiples of $\sqrt{\pi}$. Using the idea of conditional momentum boosts as rotations of the ancillary qubit depending on the position eigenvalue of the oscillator state, we can use $e^{\frac{i}{2}\sqrt{\pi}\hat{x}\sigma_x}$ for logical Z readout.

Just as in the case of cat states (Sec. VI D 3), we have rotation errors due to the uncertainty in the position of each Gaussian state in the superposition. These errors can be corrected using single-variable as well as non-Abelian QSP. In this case, since the peaks have a standard deviation of Δ, as fixed by the envelope parameter [see Eq. (286)], the precorrection to reduce the rotation errors during logical Z readout is achieved via $e^{i\sqrt{\pi}\Delta^2\hat{x}\sigma_y}$. Below, we give the non-Abelian QSP sequences for logical Pauli readout of GKP code words with envelope size Δ,

$$X\text{-basis:}\quad \mathcal{E}_p\left[\frac{\sqrt{\pi}}{2}, \Delta\right] = e^{-i\sqrt{\pi}\hat{x}\Delta^2\sigma_y} e^{i\frac{\sqrt{\pi}}{2}\hat{p}\sigma_x}, \tag{296}$$

$$Z\text{-basis:}\quad \mathcal{E}_x\left[\frac{\sqrt{\pi}}{2}, \Delta\right] = e^{i\sqrt{\pi}\hat{p}\Delta^2\sigma_y} e^{i\frac{\sqrt{\pi}}{2}\hat{x}\sigma_x}, \tag{297}$$

$$Y\text{-basis:}\quad \mathcal{E}_{\frac{x+p}{\sqrt{2}}}\left[\frac{\sqrt{\pi}}{\sqrt{2}}, \Delta^2\right] = e^{-i\sqrt{\pi}(\hat{x}-\hat{p})\Delta^2\sigma_y} e^{i\frac{\sqrt{\pi}}{2}(\hat{x}+\hat{p})\sigma_x}. \tag{298}$$

Note that for a square GKP code defined in Eq. (286) and shown in Fig. 16, the displacement required to implement a logical Y operation is longer than the corresponding displacements required for logical X or Z operations by a factor of $\sqrt{2}$. This fact accounts for the difference in the arguments of Y-basis readout in Eq. (298) and X/Z-basis readouts in Eqs. (296) and (297).

Following the sequence in Eq. (297), the $+1$ eigenstates of the GKP logical Z operator rotate the ancilla by 2π, while the -1 eigenstates rotate the ancilla by π with high probability. Thus, for an arbitrary logical GKP state $|\psi\rangle_{\text{GKP}} = \alpha|0\rangle_{\text{GKP}} + \beta|1\rangle_{\text{GKP}}$, we can write

$$\mathcal{E}_x\left[\frac{\sqrt{\pi}}{2}, \Delta\right]|\psi\rangle_{\text{GKP}}|0\rangle = \alpha|0\rangle_{\text{GKP}}|0\rangle + \beta|1\rangle_{\text{GKP}}|1\rangle + \mathcal{O}(\epsilon^3), \tag{299}$$

where $\epsilon = (\sqrt{\pi}/2)\Delta^2$. At this point, measuring the ancilla realizes the measurement of the logical Z operator. The Y-basis readout is different due to the fact that the phase-space displacement associated with the logical Y operator of the (square-grid) GKP code (shown in Fig. 16) is longer by a factor of $\sqrt{2}$ since it is along the hypotenuse of the triangle whose sides form the displacements associated with logical X and Z. The failure probability for this readout is $\epsilon^3 = 10^{-4}$ for a GKP code with $\Delta = 0.302$. The backaction on the oscillator, however, scales as $\mathcal{O}(\epsilon^2)$. In practice, this error is negligible relative to those associated with faults in the auxiliary qubit during the entangling operation and subsequent measurement. In Ref. [479], the authors have numerically optimized the amplitude of this precorrection using numerical methods and found it to be the same as the analytical quantity given here, by the use of non-Abelian QSP (see Sec. VI D 3). To date, this is the only known measurement circuit designed for finite-energy GKP logical readout, and it works better than the simpler circuit designed for readout of the ideal (infinite-energy GKP) code words. For details, we direct the readers to the aforementioned references.

4. Stabilization and error correction

Errors induced by interactions with the environment may increase or decrease the energy of the system (and/or introduce phase errors), leaving the oscillator in a state outside of the code space. Nonunitary channels involving measurements and qubit resets can be used to extract entropy from the system to bring the oscillator back to the code space. In this section, we discuss how two of these *oscillator-error-correction* protocols can be achieved using hybrid instruction sets. Specifically, we highlight the methods of cat state amplification and dissipative GKP stabilization. These circuits extract entropy from the system via auxiliary-qubit measurements (thereby converting the von Neumann entropy to Shannon entropy) or resets (which transfer the von Neumann entropy of the oscillator into the cold bath of the qubit).

a. Cat state amplification. Energy relaxation extracts energy from an oscillator, shrinking its quantum phase space. Recall that two-legged even (+) and odd (−) cat states, with implicit renormalization, are given by

$$\left|C_{0,\alpha,\pm}\right\rangle = |\alpha\rangle_{\rm osc} \pm |-\alpha\rangle_{\rm osc}\,, \tag{300}$$

using Eq. (289). Energy relaxation reduces the cat size $|\alpha|$ to $|\alpha'| < |\alpha|$, bringing the two blobs closer to the origin. This is a deterministic error and one knows how much the cat size has shrunk at any given time since the coherent state amplitude obeys $\alpha(t) = \exp(-(\kappa/2)t)\alpha(0)$.

The task of cat state amplification is to add energy to the system such that the cat state is boosted back to size α. It is well known in quantum optics that in the process of boosting the coherent state amplitude, even perfect quantum limited amplifiers must either add noise (phase-insensitive amplification entangles the signal mode with another idler mode) or cause squeezing (phase-sensitive amplification preserves phase-space volume so that it squeezes the quadrature not being amplified), thereby introducing errors. This fact can be recalled from the phase-space flow pictures for squeezing in Fig. 4 and the desired phase-space flow picture in Fig. 5. We could use QSP methods to create a polynomial approximation to the Hamiltonian of Eq. (66) to approximate the phase-space flow in Fig. 5.

Here, we describe a new method to achieve cat amplification using prior knowledge of the specific cat state to be amplified and the techniques we have already developed to create the cat state. The cat state preparation scheme (Sec. VI D 3) uses non-Abelian QSP to disentangle the qubit from the oscillator state. First, we use the adjoint of that sequence to entangle the qubit with the individual blobs of a two-legged cat state of size α:

$$\mathcal{E}_x\left[-\frac{\pi}{4\alpha}, 1\right]\left|C_{0,\alpha,\pm}\right\rangle|+\rangle = |\alpha\rangle_{\rm osc}|0\rangle \pm |-\alpha\rangle_{\rm osc}|1\rangle\,. \tag{301}$$

We then apply a conditional displacement to move the blobs farther apart by δ. For real α this step is equivalent to

$$\begin{aligned} &e^{-i\delta\sqrt{2}\hat{p}\sigma_z}|\alpha\rangle_{\rm osc}|0\rangle \pm |-\alpha\rangle_{\rm osc}|1\rangle \\ &\quad = |\alpha+\delta\rangle_{\rm osc}|0\rangle \pm |-\alpha-\delta\rangle_{\rm osc}|1\rangle\,. \end{aligned} \tag{302}$$

After this operation, we can disentangle the oscillator from the ancilla, with a circuit the parameters of which depend on the new size of the cat state $\alpha' = \alpha + \delta$,

$$\mathcal{U}_x\left[\frac{\pi}{4\alpha'}, 1\right]\left|\alpha'\right\rangle_{\rm osc}|0\rangle \pm \left|-\alpha'\right\rangle_{\rm osc}|1\rangle = \left|C_{0,\alpha',\pm}\right\rangle|+\rangle\,. \tag{303}$$

Similar to cat state preparation, the amplification scheme can alternatively be carried out using the BB1 scheme of single-variable QSP in place of the non-Abelian QSP sequences $\mathcal{E}$ and $\mathcal{U}$. In Fig. 17 we show a comparison between the two schemes for the amplification of cat state with size α to size $\alpha + \delta$ (here, $\delta = \alpha$).

Amplification of four-legged cat states requires a scheme that moves all four blobs away from the origin. Such an operation could be done using two ancilla qubits (or a qudit with $d = 4$) where each of the four ancilla states is matched with one blob of the four-legged cat state

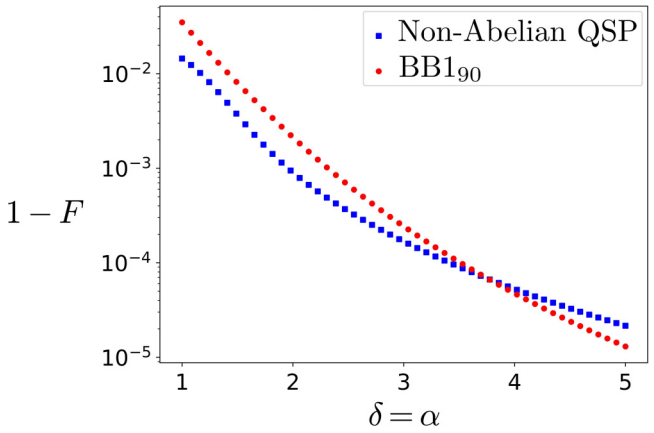


FIG. 17. A direct comparison between the single-variable (red) and non-Abelian QSP (blue) cat state amplification protocols. The plot shows the infidelity of the hybrid output state with the final cat state of size $\alpha + \delta = 2\alpha$ using different cat state amplification protocols. We see that the two curves cross at approximately $\alpha + \delta \approx 7.5$, after which BB1$_{90}$ outperforms non-Abelian QSP. It is useful to contrast the results of this plot with Fig. 29 which compares the two protocols for cat state preparation from vacuum.

in the overall entangled state. The desired flow field for the four-legged cat amplifier can be found in Fig. 6 in Sec. III C 3. Again, QSP could be invoked to achieve a polynomial approximation to the quantum version of the classical Hamiltonian of Eq. (69) to achieve the desired flow field.

Here, we review steps to perform this amplification for the four-legged cat state using a single auxiliary qubit. For simplicity, we consider the case of a four-legged cat state rotated by an angle of $\pi/4$ in phase space. This allows us to see the X-basis code words of the four-legged cat code as a superposition of two two-legged cat states parallel to the position axis or momentum axis of the quantum phase space:

$$|+\rangle_{4C} = \left|C_{\alpha,i\alpha,+}\right\rangle + \left|C_{-\alpha,i\alpha,+}\right\rangle, \tag{304}$$

$$|-\rangle_{4C} = \left|C_{\alpha,i\alpha,+}\right\rangle - \left|C_{-\alpha,i\alpha,+}\right\rangle, \tag{305}$$

where $\left|C_{\alpha,\beta,\pm}\right\rangle$ is defined in Eq. (289). In order to amplify $|-\rangle_{4C}$, we first prepare the entangled state

$$\left|C_{\alpha,i\alpha,+}\right\rangle|+\rangle + \left|C_{-\alpha,i\alpha,+}\right\rangle|-\rangle\,, \tag{306}$$

using the single-variable QSP scheme described in Sec. V C 2 adapted to entangle two displaced cat states (see Ref. [478]). We then amplify the state along the position quadrature to $\alpha' = \alpha + \delta$ using $e^{-i2\delta\hat{p}\otimes\sigma_x}$:

$$\left|C_{\alpha',i\alpha,+}\right\rangle|+\rangle + \left|C_{-\alpha',i\alpha,+}\right\rangle|-\rangle\,. \tag{307}$$

Next, we change the entanglement of the state to entangle cat states parallel to the position axis using non-Abelian QSP [481]:

$$\left|C_{i\alpha,\alpha',+}\right\rangle|+\rangle + \left|C_{-i\alpha,\alpha',+}\right\rangle|-\rangle\,. \tag{308}$$

This allows us to amplify the four-legged cat state along the momentum quadrature using $e^{i2\delta\hat{x}\otimes\sigma_x}$:

$$\left|C_{i\alpha',\alpha',+}\right\rangle|+\rangle + \left|C_{-i\alpha',\alpha',+}\right\rangle|-\rangle\,. \tag{309}$$

Finally, we disentangle the qubit from the oscillator using the method described in Eq. (291) for the preparation of the new four-legged cat state sized α':

$$|+\rangle_{4C} = \left|C_{\alpha',i\alpha',+}\right\rangle + \left|C_{-\alpha',i\alpha',+}\right\rangle. \tag{310}$$

b. Dissipative GKP stabilization. We review here recent theoretical and experimental progress in CV-DV control that has led to successful error-correction gain for quantum memories based on the GKP code [282]. The stabilizers of the GKP states can be viewed as exponentiated modular quadratures—e.g., $e^{i\sqrt{\pi}(\hat{x}_m + i\Delta^2\hat{p}_m)}$, where

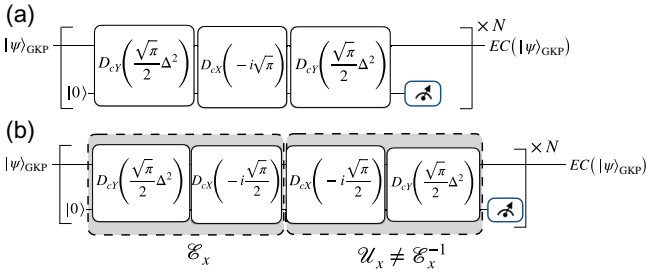


FIG. 18. Dissipative GKP stabilization. In this figure, we have used shorthand notations for symmetric X- and Y-conditional displacements, respectively, as $D_{cX}(\alpha) = e^{(\alpha\hat{a}^\dagger - \alpha^*\hat{a})\otimes\sigma_x}$ and $D_{cY}(\alpha) = e^{(\alpha\hat{a}^\dagger - \alpha^*\hat{a})\otimes\sigma_y}$. The input of the circuits shown here is an arbitrary logical GKP state $|\psi\rangle_{GKP}$ and the output $EC(|\psi\rangle_{GKP})$ is an error-corrected GKP qubit with the same logical information as the input. Here, $EC(*)$ denotes that $*$ has been error corrected. The circuit variable Δ is the finite-energy parameter that represents the envelope size of the GKP states. (a) The circuit representation of the small-big-small (sBs) scheme using the phase-space ISA, repeated for N rounds. Here, higher N yields a closer approximation to the continuous-variable error-correction process. (b) The decomposition of the sBs scheme into an entangling sequence $\mathcal{E}_x$ and a disentangling sequence $\mathcal{U}_x$. Note that $\mathcal{U}_x \neq \mathcal{E}_x^{-1}$ enables the qubit to apply a corrective backaction on the oscillator in the presence of errors. This decomposition will be helpful in Sec. V C 5. Panels (a) and (b) have been adapted from Refs. [326,478], respectively.

$\hat{x}_m = \hat{x} \mod\sqrt{\pi}, \hat{p}_m = \hat{p} \mod\sqrt{\pi}$. Here, Δ is the Gaussian width of the target state, as described in Eq. (286). In Refs. [107,326] the authors have shown that this stabilizer can be extracted by engineering the dissipator $\hat{d} = \hat{x}_m + i\Delta^2\hat{p}_m$. This is possible by using unitary evolution under $e^{i\sqrt{\pi}(\hat{x}\sigma_x + \Delta^2\hat{p}\sigma_y)}$ combined with qubit resets. Using the phase-space ISA, this dissipative quantum channel can be realized by a circuit known as the *small-big-small* (sBs) scheme [326], shown in Fig. 18(a). This circuit uses three conditional displacements such that the first and last conditional displacements have a small amplitude proportional to Δ^2, while the middle gate has a relatively larger amplitude of $\sqrt{\pi}$, and thus the only gate that contributes to the circuit duration significantly is the middle gate with amplitude $\sqrt{\pi} \gg \Delta^2$. This circuit has been used in recent beyond-break-even QEC memory experiments for GKP qubits [65] and qudits [66].

In Ref. [107] the authors have demonstrated the construction of this channel with a trapped-ion system using the sideband ISA. This dissipative channel method is an autonomous stabilization scheme in which no additional correction needs to be applied based on an ancilla measurement outcome. The ancilla is reset at the end of the circuit. With high probability, each reset takes the oscillator closer to the GKP code space, removing (up to) one bit of entropy from the system. In Fig. 18 we show that this

circuit representation is equivalent to a combination of an entangling sequence and a disentangling sequence.

While this circuit corrects errors, it also applies a deterministic logical Pauli operation on the GKP states [482]. The circuit applies an ancilla rotation by $2m\pi$ for each peak (m) in the GKP logical qubit. Hence, the ancilla qubit learns no information about the GKP state of the system. However, if the oscillator is not in a GKP logical eigenstate, the non-Abelian QSP correction for the disentangling operation changes. Thus, the oscillator is not disentangled from the ancilla. The qubit reset in this case applies a corrective backaction on the oscillator in addition to the deterministic Pauli logical operation.

5. Logical single-qubit rotations

This section reviews the construction of single-logical-qubit rotations on the logical Bloch sphere using gates from the various instruction sets in consideration. A majority of the examples discussed below teleport the fast two-level ancilla operations to bosonic logical qubits. Here, we use a special case of teleportation circuits, namely a phase-transfer circuit (see Fig. 19), which can be advantageous in mitigating effects of ancilla errors [235,478,483]. As noted above in Sec. V C 3, the entangling circuits involved in the teleportation are also useful for mapping the logical state onto an ancilla qubit for the purpose of logical readout.

Let $P \in \{X, Y, Z\}$ denote logical Pauli operations of the bosonic code under consideration, and let $|P_\pm\rangle$ be the ± 1 eigenstates of the logical P operation. The phase-transfer circuit can be constructed straightforwardly with access to hybrid entangling operations that flip the ancilla conditioned on the state of the bosonic qubit:

$$C_P X = |P_+\rangle_{\rm osc}\langle P_+|_{\rm osc} \otimes \hat{I} + |P_-\rangle_{\rm osc}\langle P_-|_{\rm osc} \otimes X, \tag{311}$$

as shown in Fig. 19. In this case, $|\psi\rangle = \alpha_1 |0\rangle_L + \alpha_2 |1\rangle_L$ in the top wire is an oscillator state encoded in a bosonic code, while the bottom wire is a two-level ancilla. The first $C_P X$ gate entangles the logical qubit in the oscillator and the ancilla qubit,

$$\alpha_1 |+P\rangle_{\rm osc} |0\rangle + \alpha_2 |-P\rangle_{\rm osc} |1\rangle . \tag{312}$$

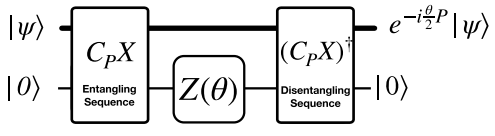


FIG. 19. The phase-transfer circuit. Here, $P \in \{X, Y, Z\}$. The oscillator-controlled Pauli gates $C_P X$ entangle the two wires in the state $|P_+\rangle_{\rm osc}|0\rangle + |P_-\rangle_{\rm osc}|1\rangle$, where the thin line represents the ancilla and $|P_+\rangle_{\rm osc}$ and $|P_-\rangle_{\rm osc}$ represent the ± 1 eigenstates of $\hat{P}$ for the top bosonic qubit. $(C_P X)^\dagger$ is the adjoint of the operator $C_P X$.

The rotation gate $Z(\theta)$ now adds a local phase to this entangled state:

$$\alpha_1 |+P\rangle_{\rm osc} |0\rangle + e^{i\theta}\alpha_2 |-P\rangle_{\rm osc} |1\rangle . \tag{313}$$

Finally, the second $(C_P X)^\dagger$ gate disentangles the oscillator from the auxiliary qubit:

$$(\alpha_1 |+P\rangle_{\rm osc} + e^{i\theta}\alpha_2 |-P\rangle_{\rm osc}) |0\rangle . \tag{314}$$

Thus, the circuit performs a logical $\exp(-i(\theta/2)P)$ operation, realized from the oscillator-controlled operations $C_P X$ and a single-qubit ancilla rotation. One of the key advantages of this circuit is the fact that at the end of the circuit the ancilla qubit is in a deterministic state (here, $|0\rangle$), and thus certain ancilla errors are detectable using ancilla measurement.

Below, we show examples of the construction of the entangling gates $C_P X$ and disentangling gates $(C_P X)^\dagger$ for various bosonic codes using gates available in different selected instruction sets.

a. GKP code. The logical Pauli operations on realistic GKP code words are equivalent to translations in phase space up to errors due to the finite-energy cutoff ($\Delta \neq 0$). Assuming that finite-energy errors do not induce logical errors in the GKP code space, phase-space translations can yield logical Pauli operations with classical controls on the oscillator. However, it is impossible to perform arbitrary Pauli rotation using classical controls since this requires creating a linear combination of identity and a Pauli translation operation. Hence we must teleport the gate into the system using an ancillary qubit and the phase-transfer circuit.

We have discussed the entangling operation between a two-level ancilla and GKP code word in the logical readout section. The sequences used for GKP logical P readout in Eqs. (296) and (297) yield $C_X X$, $C_Y X$, and $C_Z X$. The disentangling operation can be chosen to be the inverse operations $(C_X X)^\dagger$, $(C_Y X)^\dagger$, and $(C_Z X)^\dagger$. An alternative to $(C_Z X)^\dagger$ for disentangling the oscillator and the ancilla qubit is to use the second half of the stabilization scheme, $\mathcal{U}_x$ [see Fig. 18(b)]. The resulting phase-transfer circuit not only teleports gates but also corrects errors in the GKP qubit at the same time. This feature is a consequence of the fact that, in the absence of $Z(\theta)$ rotation, Fig. 19 is equivalent to the GKP stabilization circuit sBs [484]. The stabilization along each quadrature can be used to derive the phase-transfer circuit to teleport universal single-qubit rotations in the GKP Bloch sphere. Importantly, this gate is valid for not only GKP qubits but can also be, *mutatis mutandis*, used for GKP qudits encoding a d-level logical system. The authors have also illustrated a piece-wise circuit construction to reduce the effect of ancilla dephasing.

With this, the protocol can yield high-fidelity gate operation with a biased noise ancilla the same as claimed for the stabilization scheme in Ref. [326].

b. Four-legged cat code. Here, the entangling gate for logical $Z(\theta)$ rotations is given by the controlled rotation gate,

$$C_ZX = e^{i\frac{\pi}{4}a^\dagger a \otimes \sigma_x}. \tag{315}$$

This enables single-qubit logical $Z(\theta)$ rotations. In addition, one can use SNAP gates and displacements (see Sec. IV D 4) from the Fock-space ISA to perform entangling operations which can correct for errors in the ancilla up to first order. In Refs. [422,432] the authors have used this method to perform $Z(\theta)$ and $X(2\phi)$ rotations of the four-legged cat code. Under correctly chosen conditions, the logical rotations can be made fault tolerant by utilizing a three-level (qutrit) ancilla. This technique was first described in [180] for the control of binomial codes with a qubit ancilla.

c. Binomial code. The entangling operation between binomial codes and an ancilla qubit can be carried out using

$$C_ZX = e^{i\frac{\pi}{4}a^\dagger a \otimes \sigma_x}. \tag{316}$$

This operation rotates the ancilla by an angle of 2π (π) to $|0\rangle$ ($|1\rangle$) if the oscillator state is $|0\rangle_{\text{bin}}$ ($|1\rangle_{\text{bin}}$). It should be noted that while the operation C_ZX takes the logical qubit outside of the code space ($|0\rangle + |4\rangle \rightarrow |0\rangle - |4\rangle$), the disentangling operation (also, C_ZX) brings it back to the code space. This circuit in Fig. 19 can be easily modified to use the natively available controlled rotation gate in Box IV.8. The circuit is an exact gate teleportation sequence and is on par with the numerically optimized GRAPE pulses used in Refs. [67,180]. If either an ancilla error or a photon loss error in the binomial qubit occurs, then at the end of this circuit the ancilla would be in state $|1\rangle$, heralding the error. A better protocol from the perspective of fault tolerance uses SNAP gates which are tolerant to single ancilla errors, as described in [67,180]. This work also uses a qutrit ancilla to achieve fault tolerance by enabling ancilla error correction as in the four-legged cat example discussed above. Rotations around other axes on the logical Bloch sphere of the binomial code (required for universal control) are more complex and to date have required numerically optimized sequences of gates from one of the ISAs. Finding efficient analytical gate sequences based on QSP remains an open problem.

d. Dual-rail code. The dual-rail qubit encodes quantum information equivalent to a single qubit in two physical systems. Consider the case in which these two systems are oscillators in Fock states $|n = 0\rangle$ and $|n = 1\rangle$. The preparation of these states can be achieved simply using the sideband ISA. Note that the beam-splitter gate acting on the dual-rail qubit provides universal control, and thus, we already have natively available logical Pauli rotations for this encoding. However, for the sake of completeness, we note that one could also teleport these single-qubit Pauli rotations. A hybrid controlled rotation gate from the Fock-space ISA is defined as (or equivalently, a SNAP gate),

$$C_ZX = e^{i\frac{\pi}{2}a^\dagger a \otimes \sigma_x}, \tag{317}$$

acting on one of the two oscillators. This gate can be used to entangle and disentangle the two-level ancilla from the dual-rail qubit, and teleport single-qubit Z rotation using the circuit in Fig. 19. Again, the circuit in the figure can be easily modified to use the natively available gate in Box IV.8. The same entangling gate can also be used for readout purposes.

6. Entangling bosonic logical qubits

In this section, we focus on entangling operations between two logical qubits (as opposed to the previously discussed entanglement between a logical qubit and an ancilla). These operations are necessary complements to single-qubit logical operations, preparation, and measurement to achieve universal logical control. Primarily, we focus on recent proposals [432,485] (in addition to Refs. [235,478] cited in the introduction to this section) for two-qubit logical operations which use gates from the instruction sets discussed in this work. These operations require the multimode extensions of the ISA, which includes beam splitters between oscillators or controlled-PHASE (CPHASE) gates between ancilla qubits. We have already introduced several ISAs with beam splitters in Table V. The ISA where beam splitters are replaced by CPHASE gates is better suited for Rydberg atoms (see Table XI).

First, we begin with the case in which the ISA only includes CPHASE gates. A simple circuit for the bosonic logical qubit entangling task can be constructed using the entangling-disentangling gates C_PX between two-level ancilla and bosonic code words discussed with respect to arbitrary single-qubit-logical rotations in Sec. V C 5. One can entangle both logical code words $|\psi\rangle_a$ and $|\phi\rangle_b$ that belong to logical qubits encoded in modes a and b, respectively, with one ancilla each,

$$(\alpha_1 |0\rangle_L |0\rangle + \beta_1 |1\rangle_L |1\rangle) \otimes (\alpha_2 |0\rangle_L |0\rangle + \beta_2 |1\rangle_L |1\rangle). \tag{318}$$

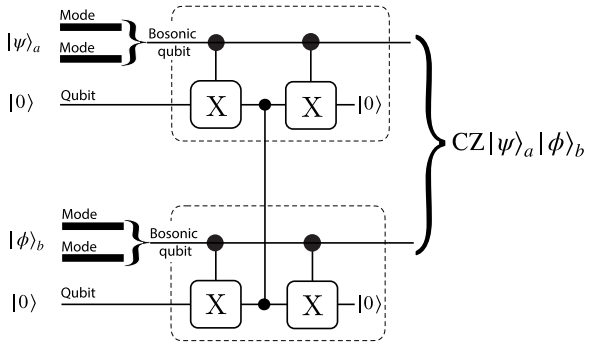


FIG. 20. The entangling operation between two logical bosonic qubits using the oscillator-controlled C_ZX gate on the ancilla and the CPHASE gate to entangle the two ancillas. The top lines in both boxed figures represent bosonic logical qubits, while the bottom lines represent the two-level ancillary qubits. The two thick lines indicate that more than one oscillator can be used to encode a single bosonic logical qubit, e.g., a dual-rail qubit.

We then entangle the two ancillary systems with a CPHASE gate,

$$\alpha_1|0\rangle_L|0\rangle\otimes(\alpha_2|0\rangle_L|0\rangle+\beta_2|1\rangle_L|1\rangle) + \beta_1|1\rangle_L|1\rangle\otimes(\alpha_2|0\rangle_L|0\rangle-\beta_2|1\rangle_L|1\rangle). \quad (319)$$

Finally, we disentangle the ancillary systems using the disentangling gates C_ZX:

$$\alpha_1|0\rangle_L|0\rangle\otimes(\alpha_2|0\rangle_L+\beta_2|1\rangle_L)|0\rangle + \beta_1|1\rangle_L|0\rangle\otimes(\alpha_2|0\rangle_L-\beta_2|1\rangle_L)|0\rangle \quad (320)$$

$$=\alpha_1|0\rangle_L|0\rangle\otimes|\phi\rangle_b|0\rangle+\beta_1|1\rangle_L|0\rangle\otimes(Z_L|\phi\rangle_b)|0\rangle \quad (321)$$

$$\equiv \mathrm{CZ}_L^{a,b}\otimes I\otimes I(|\psi\rangle_a|\phi\rangle_b)\otimes|0\rangle|0\rangle, \quad (322)$$

where Z_L is the logical Z operation on the bosonic code word $|\phi\rangle_b$. This scheme is described in Fig. 20.

This circuit can also detect ancilla errors since post measurement both auxiliary qubits are in a deterministic state $|0\rangle$. Thus, the gate could be made fault tolerant if we assume an upper-layer QEC code above bosonic modes [233,243,302]. While the above can be generalized easily using the discussion of hybrid entangling operations in the previous Sec. V C 5, it requires two extra auxiliary qubits to perform each entangling operation.

Next, we discuss examples for the case in which the ISA includes beam splitters instead of CPHASE gates. For this architecture, logical two-qubit entangling gates for bosonic code words can be realized using a single ancilla qubit and high-fidelity beam splitters [237]. These examples discuss operations which entangle the two bosonic logical qubits with the same qubit ancilla, generalizations of C_PX [see Eq. (311)], which directly entangle the two code words $|\psi\rangle_a$ and $|\phi\rangle_b$ with a single auxiliary qubit in the initial state $|0\rangle$. These operations can then be used directly as shown in Fig. 19, replacing the top line with two thick lines representing the two code words. With the available architecture, this method can yield not only logical entangling gates such as CZ_L or CX_L but also two-qubit logical rotations such as $ZZ(\theta)$ and $XX(\theta)$ gates. All operations discussed below can be realized using conditional beam splitters, compilation of which can be realized natively, or using multimode Fock space or phase-space ISA [see Eq. (337)].

a. GKP code. The logical gates for GKP qubits are obtained via the nonunitary gate $\hat{E}_\Delta A\hat{E}_\Delta^{-1}$, where A is the gate for unrealistic infinite-energy GKP, while $\hat{E}_\Delta$ is the envelope operator defined in Eq. (286). The entangling gate CZ_{L} such that $A=e^{i2\hat{x}\otimes\hat{x}}=e^{i2\hat{x}_1\hat{x}_2}$ takes the following form:

$$\hat{E}_\Delta e^{i2\hat{x}\otimes\hat{x}}\hat{E}_\Delta^{-1} = e^{i(2\cosh\Delta^2\hat{x}+i2\sinh\Delta^2\hat{p})\otimes(2\cosh\Delta^2\hat{x}+i2\sinh\Delta^2\hat{p})} \quad (323)$$

$$\approx e^{i(2\hat{x}_1\hat{x}_2-2\Delta^4\hat{p}_1\hat{p}_2+i2\Delta^2(\hat{x}_1\hat{p}_2+2\hat{p}_1\hat{x}_2))} \times (\text{finite-energy } CZ_L \text{ gate}). \quad (324)$$

These nonunitary gates can be approximated using an auxiliary qubit, where the approximations hold in the small-Δ limit such that $\cosh\Delta^2\approx 1$ and $\sinh\Delta^2\approx\Delta^2$:

$$CZ_L\approx e^{i\Delta^2(\hat{x}_1\hat{p}_2+\hat{p}_1\hat{x}_2)\sigma_y}e^{i(2\hat{x}_1\hat{x}_2-2\Delta^4\hat{p}_1\hat{p}_2)\sigma_x}\times e^{i\Delta^2(\hat{x}_1\hat{p}_2+\hat{p}_1\hat{x}_2)\sigma_y}|\psi\rangle_{\mathrm{GKP}}|0\rangle. \quad (325)$$

Reference [485] has proposed a circuit to perform these nonunitary gates using an auxiliary qubit through the dissipation model described in Ref. [326]. The circuit involves conditional beam splitters. If the middle gate $e^{i(\hat{x}_1\hat{x}_2-\hat{p}_1\hat{p}_2)\otimes\sigma_x}$ is split into two halves just like the GKP stabilization circuit in Fig. 18, then the first half of the circuit can be the entangling operation C_PC_PX, while the second half can be the disentangling operation. Thus, we can introduce a $Z(\theta)$ rotation on the qubit to achieve a $ZZ(\theta)_L$ rotation on the two qubits. The authors have also shown that this circuit can be achieved using non-Abelian QSP, highlighting the generalization of this QSP framework to two modes. In Appendix B, we provide a compilation of the fast $e^{i(\hat{x}_1\hat{x}_2-\hat{p}_1\hat{p}_2)\otimes\sigma_x}$ gates in the weak-dispersive regime, an artifact of phase-space ISA in superconducting circuits [60].

b. Four-legged cat code. We employ the following conditional joint-parity operation to entangle the ancillary qubit with the two modes encoded in the four-legged cat

code.

$$e^{i\frac{\pi}{2}(a^\dagger a+b^\dagger b)\sigma_x}. \tag{326}$$

Again, the circuit can be easily modified to use beam splitters and the SQR gate from the Fock-space ISA discussed in Box IV.9. A realization of this operation in which a qutrit ancilla is used to make the circuit error-detectable against ancilla errors is possible. Alternatively, in Ref. [432] the authors have used SNAP gates and beam splitters from the multimode Fock-space ISA to achieve an $XX(\theta)$ rotation for four-legged cat codes.

c. Binomial code. Again, the trick from Ref. [235] can again be used to construct the conditional joint-parity operation $e^{i\pi/2(a^\dagger a+b^\dagger b)\sigma_x}$. In Ref. [235], the authors show methods to implement two-qubit $ZZ(\theta)$ gates using the same instruction set for all three bosonic codes. We can create arbitrary excitation-preserving two-qubit gates [235,298] which can be achieved using the Fock-space ISA.

d. Dual-rail code. Given two dual-rail qubits with modes $\{a_1, a_2\}$ and $\{b_1, b_2\}$, the required conditional joint-parity operation is given by

$$e^{i\pi({a_1}^\dagger a_1+{b_1}^\dagger b_1)\sigma_x}, \tag{327}$$

which entangles the two dual-rail qubits (comprising of four modes in total) with an ancilla qubit. Note the change in the parameter of this gate from $\pi/4$ to $\pi/2$. This operation can be compiled with the multimode Fock-space ISA using a qutrit ancilla. The authors have also given parameters to construct a controlled-SWAP operation using the same resources for dual-rail qubits and four-legged cat codes as well as binomial codes [235].

VI. COMPILATION METHODS

Having established physical and logical ISAs introduced in Secs. IV E and V C, a vital next question is how to compile programs for the three AMMs established in Sec. II B. By definition, *compilation* means *translation* from one programming language to another. A related concept is *transpilation*, which refers to the compilation of a high-level (compiler-facing) ISA into a low-level (transpiler-facing) ISA, the latter being aware of hardware-specific details such as qubit connectivity and the dominant error rates. For simplicity, we broadly use the word compilation throughout this section for both of these tasks.

The primary challenge is that we cannot directly rely on intuition from Boolean logic to design unitaries acting on qumodes. A second challenge is the significant architectural differences across various hybrid CV-DV platforms in terms of hardware layout (Fig. 2) and native interactions—see Table IX for details. As a result, many compilation and synthesis strategies developed for, e.g., trapped ions cannot be directly adapted to superconducting or neutral-atom platforms.

To jump start efforts toward tackling these challenges, we begin in Sec. VI A with a broad introduction to compilation and lay out a high-level compilation strategy for hybrid CV-DV processors. We then follow in Sec. VI B with two complementary descriptions of an arbitrary single-qubit-oscillator unitary in terms of polynomials of different operators in the hybrid system. In Sec. VI C, we address the exact, analytical synthesis of multi-qubit and multi-oscillator entangling gates and present several novel compilation techniques tailored to the superconducting CV-DV architecture illustrated in Fig. 3(a).

Section VI D is devoted to the use of ideas from quantum simulation and quantum signal processing in DV systems for hybrid CV-DV compilation purposes, particularly toward the systematic control and synthesis of single-qubit–oscillator unitaries. We note that the extension of product formulas (Sec. VI D 1) and LCU techniques (Sec. VI D 2) to bosonic platforms has been previously discussed in Ref. [340] and Ref. [486], respectively, and our treatment on these topics will therefore be brief and pedagogical. In contrast, we will provide a more detailed discussion on the generalization of QSP to hybrid CV-DV platforms (Sec. VI D 3), as this concerns new material. Finally, we discuss numerical compilation strategies in Sec. VI E as a complement to the analytical compilation methods. For the interested reader, we provide further in-depth comparisons between numerical and analytical compilation techniques in Appendix F.

For each compilation technique described in this section, we provide simple illustrative examples and, whenever possible, present them in a "human-readable" fashion to provide physical insights. We hope that the techniques presented in this section (and the examples provided) spur further active research and development on hybrid CV-DV compilation.

A. Introduction and overview of compilation schemes

Given the fact that the development of quantum programming languages is still in its early stage [487], especially for bosonic programming languages, a systematic discussion of bosonic compilation schemes in a rigorous fashion at the level of compilation schemes on classical computers [120] is not yet available. At the core of compiling hybrid oscillator-qubit algorithms lie efficient methods to decompose a multiple-qubit-oscillator unitary operator into one of the hybrid ISAs defined in Sec. IV E. While analytical constructions of arbitrary multi-qubit–oscillator unitaries are challenging [488,489], it is the goal of this section to lay major frameworks to describe qubit-oscillator unitaries as well as ways to

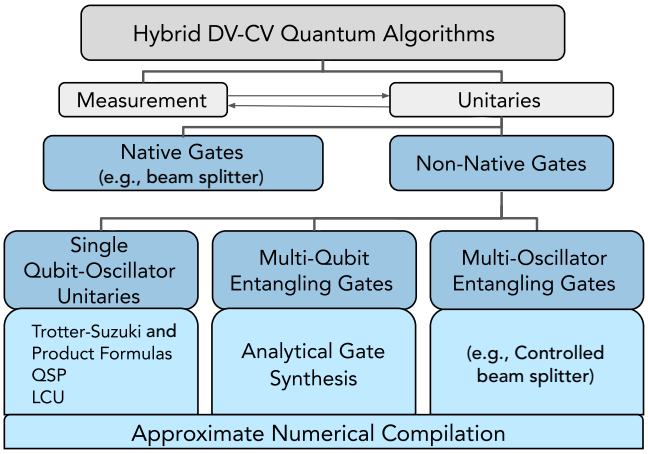


FIG. 21. A compilation flow chart for hybrid CV-DV quantum operations on multiple qubits and oscillators. Starting from the top, a generic quantum operation consists of unitaries interleaved with measurements where feedforward (unitaries conditioned on measurement results) is possible (denoted by the horizontal arrows in between). The multi-qubit–oscillator unitaries are composed of native (e.g., beam-splitter) and non-native gates (e.g., the qubit-conditioned beam splitter). The latter can be further decomposed into native single-qubit oscillator unitaries and entangling gates between either multiple qubits or multiple oscillators.

perform the compilation, with ample examples as demonstrations. For complex algorithms, it becomes increasingly difficult and less efficient to keep track of these compilation rules. Not surprisingly, modern compilers have embraced numerical optimization algorithms to find good (yet usually not globally optimal) circuits, a method which we adopt as well (*mutatis mutandis*, since quantum instructions contain continuous parameters).

The compilation scheme for hybrid CV-DV quantum computers is summarized in Fig. 21. A generic algorithm or operation is composed of multiple unitary operations interleaved with measurements and feedforward to produce new unitaries conditioned on the measurement results. The hybrid architecture has some native multi-qubit–oscillator unitaries such as the beam splitter that can serve as primitives to use in compilation. Non-native multi-qubit–oscillator unitaries may be decomposed into unitary operations on individual qubit-oscillator pairs combined with multi-qubit or multi-oscillator entangling gates, and we present examples to demonstrate compilation methods using these different categories.

We note that in actual systems, there will additionally be practical trade-offs to consider due to the heterogeneous nature of the lifetime and operation speed between qubits and oscillators. See Ref. [167] for a recent work along this line. References [490–492] have discussed related ideas for noise-adaptive compilation schemes for qubit-based NISQ architectures. We also note that while transmon qubits in some superconducting cQED hardware platforms have direct near-neighbor couplings between them, the superconducting-qubit–cavity architecture we have outlined in Fig. 3 intentionally does not for the purpose of avoiding crosstalk issues. Importantly, our novel compilation strategy (Sec. VI C) for multi-qubit entangling gates with cavity-mediated interactions enables high-level algorithms to treat qubits and oscillators on equal footing, using fast, high-fidelity SWAPs to achieve logical all-to-all connectivity despite near-neighbor physical constraints. The compilation strategies in the following sections are novel for superconducting cQED systems.

B. Representations of hybrid CV-DV algorithms

To discuss compilation methods in detail, it is useful to develop a more formal description of generic unitaries on hybrid systems. In this section, we expand the earlier discussion on universality in Sec. IV E and introduce two complementary pictures for understanding a generic multi-qubit–oscillator unitary, namely its Hermitian generator (i.e., the control Hamiltonians in an instruction set), and the unitary itself—accessible via Hamiltonians in the algebra generated by the Lie brackets of the control set.

1. Hermitian generator picture

In the Hermitian generator picture, let us consider the Hermitian operators of a system with n qubits and m oscillators that generate its corresponding unitary evolution. We denote the Pauli operator on n qubits as P_k, where $k = 0, 1, \ldots, 4^n - 1$. Each oscillator term can likewise be described by a finite-degree Hermitian polynomial $h(\hat{x}, \hat{p})$. In particular, we have seen in Eq. (242) that the most general single-qubit and single-oscillator generator is a sum of four different finite-degree polynomials in $\hat{x}, \hat{p}$. Similarly, a generic Hamiltonian on a multi-qubit and multi-oscillator system can be written as

$$H = \sum_{k=0}^{4^n-1} P_k h_k(\hat{x}_1, \hat{p}_1, \hat{x}_2, \hat{p}_2, \ldots, \hat{x}_m, \hat{p}_m), \tag{328}$$

where the subscript k runs over all Hermitian basis P_k of the n-qubit system. In the case of a single qubit $n = 1$, $k = 0, 1, 2, 3$, and P_k reduces to the identity (P_0) and the three Pauli matrices. $h_k(\cdot)$ is a generic polynomial of the $2m$ canonical variables on the m oscillators, which can be of infinite degree (or infinitely smooth such that its arbitrary order derivative concerning $\hat{x}$ and $\hat{p}$ exists). In practice, it is necessary to truncate $h_k(\cdot)$ to a finite degree. More concretely, a degree-d truncated representation can be written as

$$h_k(\hat{x}_1, \hat{p}_1, \hat{x}_2, \hat{p}_2, \ldots, \hat{x}_m, \hat{p}_m) = \sum_{\substack{r_1,\ldots,r_m,s_1,\ldots,s_m=0 \\ \sum_j (r_j+s_j)\le d}}^{d} \left(c_{r_1,\ldots,r_m,s_1,\ldots,s_m} \prod_{j=1}^{m} \hat{x}_j^{r_j} \hat{p}_j^{s_j} + \text{h.c.} \right), \tag{329}$$

for some constant coefficients $c_{r_1,\ldots,r_m,s_1,\ldots,s_m} \in \mathbb{C}$ and where "h.c." means Hermitian conjugate.

This form captures all possible terms of finite degree because $\hat{p}^s\hat{x}^r$ can always be rewritten as a linear combination of $\hat{x}^{r'}\hat{p}^{s'}$ without increasing the power, i.e., $\max\{|r'|,|s'|\} < \max\{|r|,|s|\}$ [see Eq. (D24) in Appendix D 4].

Under this representation, the goal of the compilation algorithm is to find a sequence of gates from the instruction set which results in a Hermitian matrix of the form given in Eq. (328) that generates the target unitary. As alluded to earlier in Sec. III C 3, one advantage of this representation is that the polynomial takes arguments as the canonical positions and momenta, so that the transformation on the oscillator may be viewed as flows in the m-oscillator phase space, in a similar spirit to the phase-space formulation of classical mechanics. It also admits a simple connection to the holomorphic (stellar) representation of CV computation [258], which is useful to quantify the non-Gaussian character of the transformation in terms of zeros of polynomials. We refer to Sec. III C 4 for details.

2. Unitary picture

A complementary representation is to write the unitary matrix generated from Eq. (328) explicitly,

$$U = A_0(\hat{w}_1,\hat{w}_2,\ldots,\hat{w}_m,\hat{v}_1,\hat{v}_2,\ldots,\hat{v}_m) + i\sum_{k=1}^{4^n-1} P_k A_k(\hat{w}_1,\hat{w}_2,\ldots,\hat{w}_m,\hat{v}_1,\hat{v}_2,\ldots,\hat{v}_m), \quad (330)$$

where A_k $(k = 0,\ldots,4^n-1)$ are polynomials acting on the m oscillators which satisfy certain constraints given by the unitarity of U. The arguments $w_j = e^{+i2k_j\hat{x}_j}$ and $\hat{v}_j = e^{-i2\lambda_j\hat{p}_j}$ are the unitaries that, respectively, boost the momentum of the j th oscillator by k_j and displace the position by λ_j (in the dimensionless Wigner units introduced in Sec. III C 2).

The representation in Eq. (330) can be viewed as a generalized notation of QSP for single-qubit dynamics [274,493]. Here, we adopt it to represent unitary dynamics in a coupled multi-qubit–oscillator system. Similar to the Hermitian picture representation, the polynomials A_k are usually truncated to finite degrees. Moreover, an additional truncation is set by the scale of k_j and λ_j. This truncation provides a coarse-grained description of the bosonic algorithm, which is necessary to avoid unphysical requirements during the compilation.

With this representation, the goal of the compilation is to find a sequence of hybrid qubit-bosonic gates from the instruction set such that the resulting unitary admits a given decomposition in Eq. (330). One advantage of this representation is that results from single-qubit QSP may be generalized to bosonic systems. We will see in Sec. VI D 3 that in the case of a single-qubit–oscillator system, Eq. (330) essentially provides a Fourier series decomposition of phase space under plane waves with different wave vectors, similar in spirit to the plane-wave basis used for modeling periodic systems [494].

C. Exact analytical compilation of entangling gates

We begin by discussing useful techniques for exact analytical compilation of entangling gates. Throughout this section, we specialize to the superconducting hybrid oscillator-qubit hardware layout illustrated in Fig. 3(a), though note that many of the techniques discussed are extendable to other hybrid hardware CV-DV platforms. We emphasize that some of these techniques developed in this section, while novel in the context of the superconducting architecture of Fig. 3(a), draw direct inspiration from known methods in the trapped-ion community (the technique for compilation of multi-qubit entangling gates in Sec. VI C 2, reminiscent of Mølmer-Sørenson gates [82,495], being a prime example). We will remark on these connections where appropriate.

1. Useful primitives

a. Beam-splitter SWAP *networks.* In a typical superconducting qubit-based quantum processor, lack of all-to-all connectivity inhibits the ability to natively perform entangling gates between geometrically nonlocal qubits. Instead, such gates are typically realized using SWAP networks, where information is routed to facilitate a geometrically local entangling operation.

A similar challenge persists in the superconducting hybrid oscillator-qubit hardware layout displayed in Fig. 3(a), where entanglement between sites is facilitated by beam splitters between "connected" oscillators. Analogous to the qubit-only case, long-range entangling gates can be synthesized by leveraging a beam-splitter SWAP network, which we now describe. One advantageous property of bosonic SWAP networks is that the operations are agnostic to the contents of the bosonic modes. For example, each bosonic mode can store multiple bits of quantum information in a complex state, or each can encode a single logical qubit in one of a variety of bosonic error-correcting codes. The latter possibility would allow for the correction of errors occurring during the routing of the information across the hardware fabric. Recent experimental progress has dramatically enhanced the speed and fidelity of bosonic SWAP operations [70,234,236,237].

Defining a beam splitter between oscillators at site i and j as

$$\mathrm{BS}^{(i,j)}(\theta,\varphi) = e^{-i\frac{\theta}{2}[e^{i\varphi}a_i^\dagger a_j + e^{-i\varphi}a_i a_j^\dagger]}, \quad (331)$$

the key idea is to first note that SWAP $=$ BS $(\pi,\pi/2)$ can be used to realize the transformation

$$\begin{aligned}\mathrm{BS}^{(i,j)\dagger}\left(\pi,\frac{\pi}{2}\right)a_i\mathrm{BS}^{(i,j)}\left(\pi,\frac{\pi}{2}\right)&=a_j,\\ \mathrm{BS}^{(i,j)\dagger}\left(\pi,\frac{\pi}{2}\right)a_j\mathrm{BS}^{(i,j)}\left(\pi,\frac{\pi}{2}\right)&=-a_i.\end{aligned}\tag{332}$$

We note that the above convention for SWAP differs from the choice of $\varphi=0$ used in the text surrounding Box IV.4. The net effect is a difference in global phase, and either can be used for the purpose of a beam-splitter SWAP network. An alternative option is to define SWAP by incorporating additional single-oscillator rotations to remove the phase entirely, as in Ref. [41]. Here, we opt to use the above definition and track the phase across compilation steps.

To see how this transformation can be leveraged for long-range entangling gates, for specificity let us imagine that we wish to synthesize a phase-space rotation of the kth oscillator controlled on the j th qubit,

$$\mathrm{CR}^{(k,j)}(\theta)=e^{-i\frac{\theta}{2}Z_j a_k^\dagger a_k}.\tag{333}$$

By "stacking" beam-splitter transformations, this is achieved via the sequence

$$\mathrm{CR}^{(k,j)}(\theta)=U^\dagger\mathrm{CR}^{(j,j)}(\theta)U,\tag{334}$$

where $U=\prod_{p=1}^{k-j}\mathrm{BS}^{(j,j+p)}(\pi,\pi/2)$ is a string of SWAPS along a path connecting oscillators j and k. The controlled rotation gate is just one example of the many possible hybrid oscillator-qubit gates that can be achieved using this technique. Furthermore, it is not limited to hybrid oscillator-qubit gates and can also be used to compile nonlocal oscillator-oscillator entangling gates, such as beam splitters, two-mode squeezers, and non-Gaussian gates. As previously noted, multiple SWAP strings can be parallelized and a scheduling protocol invoked to avoid collisions in the $(2+1)$-dimensional space-time fabric.

b. Synthesizing hybrid oscillator-qubit gates. The above SWAP network example illustrates the ability to synthesize hybrid oscillator-qubit gates even when the qubit and oscillator are not natively coupled. Similarly, we can condition natively available oscillator-only gates on a qubit or on multiple (possibly distant) qubits. As an example, let us consider the action of the controlled-parity gate

$$\mathrm{CP}^{(j,k)}a_j\,\mathrm{CP}^{(j,k)\dagger}=iZ_k a_j,\tag{335}$$

where $\mathrm{CP}^{(j,k)}=\mathrm{CR}^{(j,k)}(\pi)$ is the controlled-parity gate between the qubit at site k and oscillator at site j (see Box IV.8). For $j=k$, this gate is natively realizable via the dispersive interaction, and for $j\neq k$, this gate can be synthesized with the aid of beam-splitter SWAPs. Nonetheless, Eq. (335) provides a route for turning unconditional oscillator gates into conditional ones. As described shortly below, this process can be iterated to add higher-weight Pauli strings.

A displacement of mode j controlled by qubit k, realized via the sequence

$$\begin{aligned}\mathrm{CD}^{(j,k)}(\alpha)&=\mathrm{CP}^{(j,k)}D^{(j)}(i\alpha)\mathrm{CP}^{(j,k)\dagger}\\&=e^{Z_k(\alpha a_j^\dagger-\alpha^* a_j)},\end{aligned}\tag{336}$$

is a particularly useful equivalence for converting the Fock-space instruction set (where conditional displacements are not native primitives) into the phase-space instruction set.

Another useful gate is the conditional beam splitter, with applications both toward quantum simulation (see, e.g., Sec. VII C 7), as well as for universal control and measurement of logical dual-rail qubits realized in superconducting cavities [233–236]. It can be compiled as

$$\begin{aligned}\mathrm{CBS}^{(j,k,\ell)}(\theta,\varphi)&=\mathrm{CP}^{(k,\ell)}\mathrm{BS}^{(j,k)}(\theta,\varphi-\pi/2)\mathrm{CP}^{(k,\ell)\dagger}\\&=e^{-iZ_\ell\frac{\theta}{2}[e^{i\varphi}a_j^\dagger a_k+e^{-i\varphi}a_j a_k^\dagger]}.\end{aligned}\tag{337}$$

Alternatively, this gate can be viewed as natively available in the phase-space ISA since it is readily compiled at the hardware level from a short pulse sequence taking advantage of the dispersive qubit-oscillator coupling (see Appendix B). This latter strategy is particularly useful for optimizing gate speed in the weak-dispersive regime in superconducting platforms.

Finally, we note the possibility of repeatedly using parity operations controlled on different qubits to realize oscillator operations conditioned on an arbitrary, high-weight Pauli operator. For example, it is possible to combine single-qubit rotations, controlled parity operations, conditional displacements, and a mediating beam-splitter network to realize a displacement conditioned on a weight-m Pauli operator via the sequence

$$\begin{aligned}\mathrm{CD}^{(k,\{\hat{n}_1,\hat{n}_2,\ldots\hat{n}_m\})}(\alpha)&=U_{\mathrm{seq}}^\dagger D(i^m\alpha)U_{\mathrm{seq}}\\&=e^{(\prod_{j=1}^m\hat{n}_j\cdot\vec{\sigma}_j)(\alpha a_k^\dagger-\alpha^* a_k)},\end{aligned}\tag{338}$$

where $U_{\mathrm{seq}}=\prod_j^m e^{i\pi(\hat{n}_j\cdot\vec{\sigma}_j)a_k^\dagger a_k/2}$. The result is a very valuable resource that can, e.g., be used to compile many-qubit entangling gates (see Sec. VI C 2, where we compile the Toffoli gate as an illustrative example).

c. Composing conditional displacements. The composition of conditional displacement gates offers a powerful resource for the control of hybrid oscillator-qubit systems.

For example, sequences of conditional displacements are integral to the phase-space ISA, and underlie many examples and applications discussed throughout this work, including the readout, stabilization, and logical gates for GKP error-correcting codes in Sec. V C, the compilation of Jaynes-Cummings and conditional squeezing gates in Sec. VI D 1, the non-Abelian QSP framework discussed in Sec. VI D 3, and the realization of random quantum walks in Sec. VII D. However, inextricably tied to their utility, sequences of conditional displacements can combine to create extremely complex hybrid oscillator-qubit operations that are challenging to analyze. Here, our goal is to demystify some of this complexity by providing intuition for how such sequences act in phase space.

To begin, we recall the generalized conditional displacement operation of the form

$$\mathrm{CD}_{\hat{b}\cdot\vec{\sigma}}(\alpha) = e^{\hat{b}\cdot\vec{\sigma}\otimes[\alpha a^\dagger - \alpha^* a]}, \tag{339}$$

where $\hat{b}$ is a unit vector defining an axis on the qubit Bloch sphere. As discussed in Box IV.7, this general form can be synthesized by conjugating the standard controlled displacement operation $\mathrm{CD}(\alpha)$ by an unconditional rotation of the qubit. Next, we note that this conditional displacement can be cast into the form

$$\mathrm{CD}_{\hat{b}\cdot\vec{\sigma}}(\alpha) = \sum_{\eta=\pm1} \left|\eta^{(\hat{b})}\right\rangle\!\left\langle\eta^{(\hat{b})}\right| \otimes D(\eta\alpha), \tag{340}$$

where $\left|\eta^{(\hat{b})}\right\rangle\!\left\langle\eta^{(\hat{b})}\right| = \frac{1}{2}\left[\hat{I} + \eta\hat{b}\cdot\vec{\sigma}\right]$ is a projector onto the eigenstate of $\hat{b}\cdot\vec{\sigma}$ with eigenvalue $\eta = \pm1$. We can compose two such conditional displacements with different displacement amplitudes and qubit axes,

$$U = \mathrm{CD}_{\hat{c}\cdot\vec{\sigma}}(\beta)\mathrm{CD}_{\hat{b}\cdot\vec{\sigma}}(\alpha). \tag{341}$$

Using the form in Eq. (340), this can be decomposed into a Feynman path summation of four distinct terms,

$$\begin{aligned} U &= \sum_{\eta_1,\eta_2=\pm1} W_{\eta_2,\eta_1} \otimes D(\eta_2\beta)D(\eta_1\alpha) \\ &= \sum_{\eta_1,\eta_2=\pm1} e^{2ic\eta_1\eta_2\mathrm{Im}(\beta\alpha^*)} W_{\eta_2,\eta_1} \otimes D(\eta_2\beta + \eta_1\alpha), \end{aligned} \tag{342}$$

where $c = 1$ in standard units and $c = \frac{1}{2}$ in Wigner units (i.e., $[\hat{p},\hat{x}] = -ic$). Here, the qubit transformation operator W_{η_2,η_1} is defined as [496]

$$W_{\eta_2,\eta_1} = \left|\eta_2^{(\hat{c})}\right\rangle \langle\eta_2^{(\hat{c})}|\eta_1^{(\hat{b})}\rangle \left\langle\eta_1^{(\hat{b})}\right|. \tag{343}$$

This decomposition lends itself to an intuitive physical interpretation. In the case in which $\hat{c}$ and $\hat{b}$ are orthogonal, U enacts a superposition of four distinct oscillator displacements, $D(\beta+\alpha)$, $D(\beta-\alpha)$, $D(-\beta+\alpha)$, and $D(-\beta-\alpha)$. Each is paired with a complicated qubit operation that sends eigenstates of $\hat{b}\cdot\vec{\sigma}$ to eigenstates of $\hat{c}\cdot\vec{\sigma}$ with weights given by the overlap between these eigenstates. In contrast, if $\hat{c}$ and $\hat{b}$ correspond to the same direction on the Bloch sphere, U collapses into a single conditional displacement as only terms with $\eta_1 = \eta_2$ contribute.

More generally, a sequence of many conditional displacements yields a larger superposition of the form in Eq. (342), and each conditional displacement tends to double the number of terms. From a phase-space perspective, if the oscillator initially starts in the vacuum state, n conditional displacements split the initial coherent state into 2^n Gaussian phase-space "blobs" (Sec. III C 3).

These blobs can overlap, causing interference fringes in the final state (similar to the four-legged cat state shown in Fig. 16). Generally, the qubit is entangled with the oscillator at the end of this sequence, though it is possible to disentangle the qubit through a carefully chosen series of displacements—see the example of preparing a cat state using a $\mathrm{BB1}_{90}$ composite pulse sequence in Fig. 28.

The simplicity of the above representation for composing conditional displacements allows for the efficient numerical simulation of sequences of many such gates. Rather than representing the controlled displacements numerically as matrices acting on a large Hilbert space, we can simply keep track of how the phase-space "blobs" split and move in phase space. Building upon the intuition presented above, each of the 2^n phase-space blobs is associated with a definite spin state. If the most recent displacement operation was conditioned on, say, the Pauli X operator (i.e., $\hat{b} = \hat{x}$), then every phase space blob corresponds to one of the two X qubit eigenstates, $|\pm\rangle$. However, if two blobs overlap, the phase of their superposition determines the orientation of the spin associated with that region of phase space. Reference [497] discusses an efficient method to analyze linear combinations of Gaussian states which is readily extensible to the present case of entangled qubit-oscillator states produced in the phase-space ISA.

Thus far, our discussion has concerned the composition of conditional displacements in the most general case. To gain further intuition, and as a particular example, we consider a sequence of four conditional displacements,

$$U = \mathrm{CD}_{\hat{y}\cdot\vec{\sigma}}(-\alpha)\mathrm{CD}_{\hat{x}\cdot\vec{\sigma}}(-i\alpha)\mathrm{CD}_{\hat{y}\cdot\vec{\sigma}}(\alpha)\mathrm{CD}_{\hat{x}\cdot\vec{\sigma}}(i\alpha), \tag{344}$$

where α is a real-valued parameter. As will be discussed in Sec. VI D 1, such a sequence approximates a conditional squeezing gate in the limit $\alpha \ll 1$ via a BCH formula [see Eq. (377)]. It is instructive, however, to examine the trajectory of the Wigner function after each conditional displacement for general α. In Fig. 22, we show the phase-space trajectory for three cases: $\alpha = 2$, $\alpha = 1$, and $\alpha = 0.3$. Each case produces qualitatively distinct

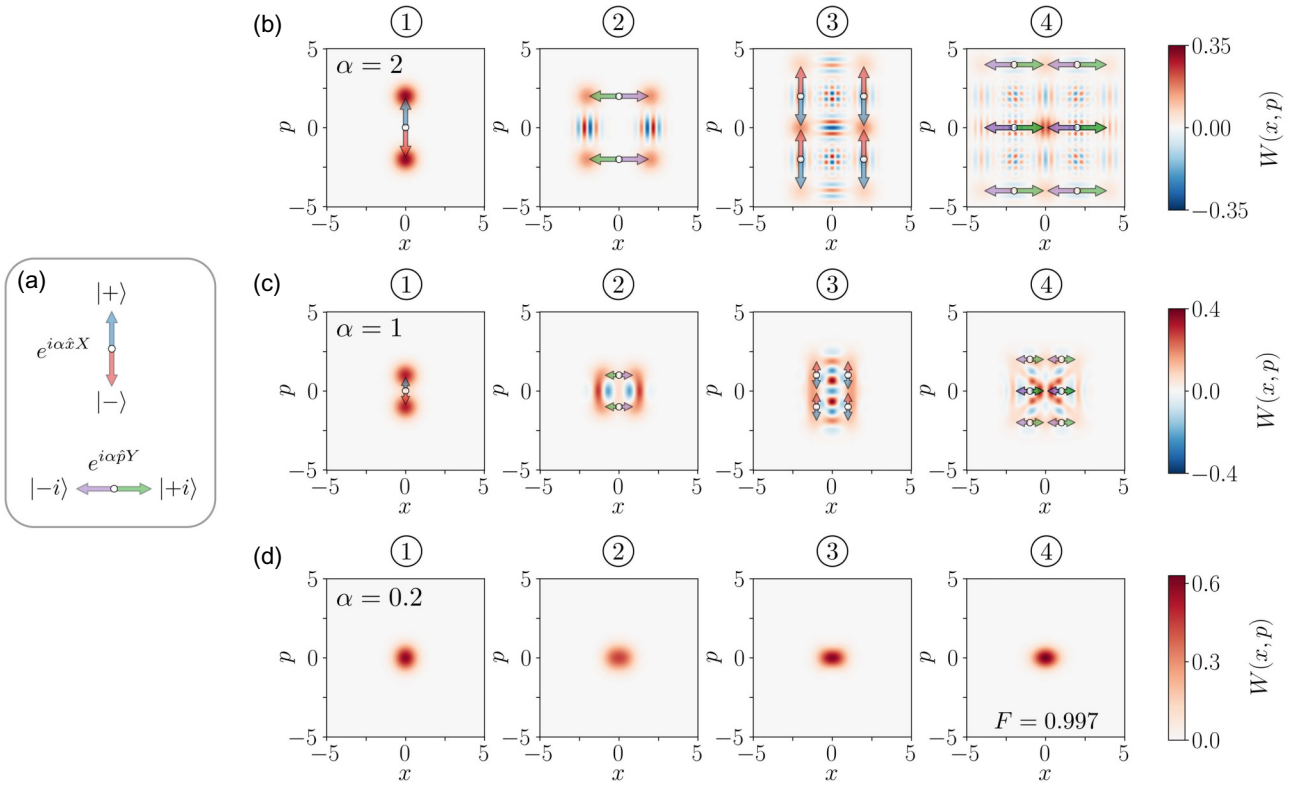


FIG. 22. Intuition for composed conditional displacements. (a) Graphical overlays are used to indicate X-controlled-momentum boosts (top) and Y-controlled coordinate boosts (bottom) in the other panels. The white circle indicates the "starting point" for each blob, while the arrows indicate the displacement direction for each conditioned qubit state. All other panels: the Wigner quasiprobability distributions for the oscillator after each successive displacement (from left to right) of the sequence in Eq. (344) applied to an oscillator and qubit in vacuum and $|0\rangle$, respectively. See also the "sum of paths" decomposition in Eq. (349). Before computing each Wigner distribution, the qubit state is traced out. Interference fringes are therefore visible only between components of definite qubit state. For example, there are no fringes after the initial conditional displacement as the two blobs are entangled with the qubit but, after the second, each displaced two-legged cat is paired with a definite X-basis qubit state (i.e., there are two $|+\rangle$ state blobs that can interfere and two $|-\rangle$ state blobs that can interfere). (b) Large displacement $\alpha = 2$. In this case, distinct blobs are visible after each successive displacement, each corresponding to a unique composite displacement in the sum in Eq. (349). While there are 16 terms in total, only nine contain unique displacements, leading to interference fringes within the blobs along the axes $x = 0$ and $p = 0$. (c) The intermediate displacement $\alpha = 1$. In this case, α is not large enough to resolve distinct blobs past the first displacement, leading to complicated interference patterns in the resulting Wigner function. (d) The small displacement $\alpha = 0.2$. As discussed in Sec. VI D 1, the sequence in Eq. (344) approximates a conditional squeezing gate CS (ζ) gate with $\zeta = -4\alpha^2$ for small α. We compute a fidelity of $F = 0.997$ with exact application of CS (ζ). While the oscillator appears squeezed after the first displacement, we note that the purity $P_i = \mathrm{Tr}(\rho_{\mathrm{osc},i})$ of the oscillator state $\rho_{\mathrm{osc},i}$ after the ith displacement reveals the oscillator and qubit to be entangled until the final step ($P_1 = 0.93, P_2 = 0.75, P_3 = 0.94, P_4 = 0.99$), highlighting the important role of each successive conditional displacement in not only squeezing the oscillator but leaving the hybrid oscillator-qubit system disentangled (in the case in which the qubit is initialized to a Z eigenstate).

behavior, highlighting the expressivity (and complexity) of sequences of conditional displacements.

To provide further intuition, it is helpful to decompose U into a Feynman path summation of unconditional displacements as in Eq. (342),

$$U = \sum_{\vec{\eta}} W_{\vec{\eta}} \otimes D(-\eta_4\alpha)D(-i\eta_3\alpha)D(\eta_2\alpha)D(i\eta_1\alpha), \tag{345}$$

where $\vec{\eta} = (\eta_1, \eta_2, \eta_3, \eta_4)$ is shorthand for a list of indices, each of which is summed over ± 1. Furthermore, as before, we have adopted the shorthand

$$W_{\vec{\eta}} = \left|\eta_4^{(\hat{y})}\right\rangle \langle \eta_4^{(\hat{y})}|\eta_3^{(\hat{x})}\rangle \langle \eta_3^{(\hat{x})}|\eta_2^{(\hat{y})}\rangle \langle \eta_2^{(\hat{y})}|\eta_1^{(\hat{x})}\rangle \left\langle \eta_1^{(\hat{x})}\right| \tag{346}$$

for the weighted qubit transformation operator $W_{\vec{\eta}}$.

To facilitate computation of $W_{\vec{\eta}}$, it is helpful to recognize that the qubit projectors can be conveniently expressed in

the computational basis as

$$\left|\eta_i^{(\hat{x})}\right\rangle\!\left\langle\eta_i^{(\hat{x})}\right| = \frac{1}{2}\begin{pmatrix} 1 & \eta_i \\ \eta_i & 1 \end{pmatrix}, \tag{347}$$

$$\left|\eta_i^{(\hat{y})}\right\rangle\!\left\langle\eta_i^{(\hat{y})}\right| = \frac{1}{2}\begin{pmatrix} 1 & -i\eta_i \\ +i\eta_i & 1 \end{pmatrix}. \tag{348}$$

In addition, similar to Eq. (342), the four unconditional displacements can be combined into a single unconditional displacement: using the convention $[\hat{p},\hat{x}] = -ic$, we may repeatedly use the identities in Box IV.1 to obtain

$$U = \sum_{\vec{\eta}} e^{i\phi_{\vec{\eta}}} W_{\vec{\eta}} \otimes e^{i2\alpha[(\eta_4-\eta_2)\hat{p}+(\eta_1-\eta_3)\hat{x}]}, \tag{349}$$

where $\phi_{\vec{\eta}} = -2c\alpha^2(\eta_1\eta_2 + \eta_2\eta_3 + \eta_3\eta_4 - \eta_4\eta_1)$. Thus, we see that each of the 16 phase-space displacements in the "sum over paths" is paired with an associated phase factor $e^{i\phi_{\vec{\eta}}}$ and qubit transformation operator $W_{\vec{\eta}}$. We also see that some distinct terms contain identical displacements (e.g., $\vec{\eta} = (+1,+1,-1,+1)$ and $\vec{\eta} = (-1,+1,+1,+1)$) and therefore lead to interference effects arising from a superposition of qubit transformation operators associated with the same point in phase space—see, e.g., the center "blob" in the final step of Fig. 22(b).

2. Compilation of entangling gates between physical qubits

A particularly important primitive for any quantum computation is the entangling gate. In the hybrid computational model where both oscillator modes and physical qubits are used for computation, it is imperative to not only have access (either natively or through compilation) to oscillator-oscillator entangling gates such as the two-mode beam-splitter (Box IV.4) and two-mode squeezing (Box IV.5) gates, but also to entangling operations between physical qubits such as the CNOT and Toffoli gates. Ideally, such operations should be achievable independently of the state of the cavity modes to which they are coupled.

Hybrid CV-DV compilation strategies for such entangling operations have a rich history in the trapped-ion platform, where bosonic modes serve as mediators of effective qubit-qubit interactions. One of the earliest and most influential examples is the Mølmer-Sørensen (MS) scheme, which enables high-fidelity two-qubit gates even in the presence of thermally populated motional modes [79–82]. This approach has since been extensively generalized to include the use of local radial modes [87], the synthesis of multi-body spin interactions [83,85], and enhancements through parametric driving [88]. The reverse—mediating oscillator interactions via Berry phases accumulated on spins—has also been realized, leading to programmable, high-fidelity multimode operations [73,75,76]. While these strategies are mature and platform-optimized for trapped ions, they are largely unused or not directly applicable to other architectures, including superconducting, due to the fundamentally distinct native instructions and unique CV-DV connectivity of each physical platform.

To remedy this gap, in the following we present a flexible strategy for compiling arbitrary entangling gates between physical qubits coupled to independent oscillator modes specifically tailored to the superconducting architecture of Fig. 3(a). As each mode is coupled to a single physical qubit, we will heavily rely on the ability to natively realize or compile beam splitters between any pair of modes (using, e.g., the beam-splitter SWAP network described in Sec. VIC 1 a). For convenience, we work in the phase-space ISA, though we note that all assumed native operations can be compiled in any of the ISAs found in Table V to achieve universal hybrid control. We begin by explaining the core principle using a single-qubit gate as an example, and follow this with an extension to two- and multi-qubit gates and comment on optimization strategies. Finally, we conclude with a brief comparison to Mølmer-Sørensen gates [79–81] (in trapped ions) and resonator-induced phase gates [498–500] (in cQED), both of which have provided inspiration for the scheme presented here.

a. Single-qubit gate. To explain the general strategy, it is instructive to first consider the compilation of a single (physical) qubit gate. In particular, we demonstrate how one can synthesize $R_z(\theta) = e^{-i\theta Z/2}$ using only beam splitters and displacements (both conditional and unconditional). We emphasize that this is not the most practical strategy for implementing a single-qubit rotation, an operation that can be much more efficiently realized natively without the need for (relatively time-consuming) oscillator-qubit entangling gates. However, the general idea lays the groundwork for the compilation of multi-qubit gates which cannot be natively realized.

First, recall that a series of displacements forming a closed loop in phase-space composed as

$$D(-\beta)D(-\alpha)D(\beta)D(\alpha) = e^{2iA}, \tag{350}$$

where $A = \mathrm{Im}\{\alpha^*\beta\}$ is the oriented area of the enclosed parallelogram (see Fig. 10 and surrounding text). Consequently, this sequence of unconditional displacements imparts a phase independent of the state of the oscillator. Here, the idea is to condition these displacements on a qubit such that the oriented area, and therefore the imparted phase, depends on its state (but not that of the oscillator). The net effect is a single-qubit rotation.

One such example is the sequence

$$U = D(-i\alpha)D_c(\alpha,-\alpha)D(i\alpha)D_c(-\alpha,\alpha), \tag{351}$$

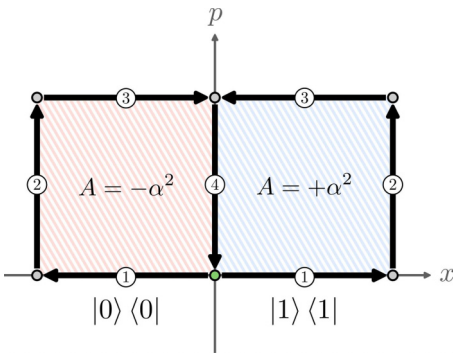


FIG. 23. A conditional oscillator phase-space trajectory that realizes $R_z(\theta) = \exp(-i\theta Z/2)$. The gate sequence is as follows: (1) displacement by α or $-\alpha$, conditioned on the state of the qubit; (2) unconditional displacement by $i\alpha$; (3) displacement by $-\alpha$ or α, conditioned on the state of the qubit; (4) unconditional displacement by $-i\alpha$. In all, the sequence leaves the oscillator unchanged (mapping each point in phase-space back to itself), but imparts a phase that is dependent on the state of the qubit (but independent of the state of the resonator), equivalent to the application of $R_z(\theta)$ for the choice $\alpha = \sqrt{\theta}/2$.

realizing the phase-space trajectory shown in Fig. 23. Associated with each of the two possible "pathways" is an oriented area, the sign of which depends on the Z-basis state of the qubit, i.e., $A = -\alpha^2 Z$. This can equivalently be seen through brute-force simplification:

$$\begin{aligned} U &= |0\rangle\langle 0|\, e^{-2i\alpha^2} + |1\rangle\langle 1|\, e^{2i\alpha^2}, \\ &= e^{-2i\alpha^2 Z}, \end{aligned} \tag{352}$$

where we have used the definition of the conditional displacement (Box IV.7) in combination with Eq. (350). Defining $\theta = 4\alpha^2$, we see that this amounts to a single-qubit rotation about the Z axis, $U = R_z(\theta)$. This gate sequence is routinely used in experiments to calibrate the size of displacements and controlled displacements in phase space (see, e.g., Fig. S4 in the Supplementary Material of Ref. [64]).

b. Two-qubit entangling gates. Building upon the core idea of the previous section, we now demonstrate the general recipe for compiling entangling gates between physical qubits. In particular, we show two examples of two-qubit gates: $R_{ZZ}(\theta) = e^{-i(\theta/2)Z_1\otimes Z_2}$ and $\text{CNOT} = e^{i(\pi/4)[I_1 - Z_1]\otimes[I_2 - X_2]}$. To that end, it is helpful to first define some notation. We envision a system comprising at least two modes (with annihilation operators a_1 and a_2) and two qubits (with Pauli-Z operators Z_1 and Z_2). Working within the phase-space ISA, we assume native access to controlled displacements between the ith mode and ith qubit (denoted $D_c^{(i,i)}(\alpha,\beta)$), a beam splitter between the two cavity modes (denoted $\text{BS}^{(1,2)}(\theta,\varphi)$), and arbitrary single-qubit rotations $R_{\hat{n}\cdot\vec{\sigma}}(\theta)$. Furthermore, we also leverage the native dispersive interaction between the ith qubit and the ith oscillator whenever it is advantageous. As noted earlier, the remote beam splitter can be achieved through a network of SWAPS.

En route to an entangling gate, it is helpful to first compile a conditional displacement between mode 1 and qubit 2 by leveraging mode SWAPs via beam splitters. Recalling the identity in Eq. (332), this is achieved via the sequence

$$\begin{aligned} D_c^{(1,2)}(\alpha,\beta) = \text{BS}^{(1,2)\dagger}\left(\pi, -\frac{\pi}{2}\right) D_c^{(2,2)}(\alpha,\beta) \\ \times \text{BS}^{(1,2)}\left(\pi, -\frac{\pi}{2}\right). \end{aligned} \tag{353}$$

To compile $R_{ZZ}(\theta)$, we can then follow the general strategy of the previous section, but with conditional displacements replacing their unconditional counterparts:

$$\begin{aligned} R_{ZZ}(\theta) = D_c^{(1,2)}(-i\alpha, +i\alpha) D_c^{(1,1)}(\alpha, -\alpha) \\ \times D_c^{(1,2)}(+i\alpha, -i\alpha) D_c^{(1,1)}(-\alpha, \alpha), \end{aligned} \tag{354}$$

with $\theta = 4\alpha^2$. As shown in Fig. 24(a), the four possible eigenstates of $Z_1 \otimes Z_2$ correspond to four distinct phase-space trajectories, enclosing an oriented area proportional to its associated eigenvalue, i.e., $A = -\alpha^2 Z_1 \otimes Z_2$.

Similarly, a CNOT gate can be compiled with the sequence

$$\begin{aligned} \text{CNOT} = D_{cX}^{(1,2)}\left(0, -i\sqrt{\frac{\pi}{2}}\right) D_c^{(1,1)}\left(0, -\sqrt{\frac{\pi}{2}}\right) \\ \times D_{cX}^{(1,2)}\left(0, +i\sqrt{\frac{\pi}{2}}\right) D_c^{(1,1)}\left(0, \sqrt{\frac{\pi}{2}}\right), \end{aligned} \tag{355}$$

where we have employed the shorthand $D_{cX}^{(1,2)}(\alpha,\beta) = e^{-i\pi Y_2/4} D_c^{(1,2)}(\alpha,\beta) e^{+i\pi Y_2/4}$, converting a Z-conditional displacement into an X-conditional displacement. As illustrated in Fig. 24(b), this sequence corresponds to the phase-space trajectory with oriented area $A = \pi[I_1 - Z_1]\otimes[I_2 - X_2]/4$. In total, both the $R_{ZZ}(\theta)$ and CNOT gates require four beam splitters and four conditional displacements.

c. Many-qubit gates. We now demonstrate that this recipe can be extended beyond two-qubit gates and toward more general many-qubit entangling gates. To that end, we use the Toffoli gate as an illustrative example. The core strategy is similar to the previous examples but with additional synthesis of displacements conditioned on high-weight Pauli operators using the technique in Eq. (338).

To compile a many-qubit gate, we follow the same strategy as the two-qubit gates, but replace displacements conditioned on a single qubit by ones conditioned on an arbitrary number of qubits. In particular, in Eq. (338) we have already shown the ability to synthesize a displacement conditioned on an arbitrary, weight-m Pauli

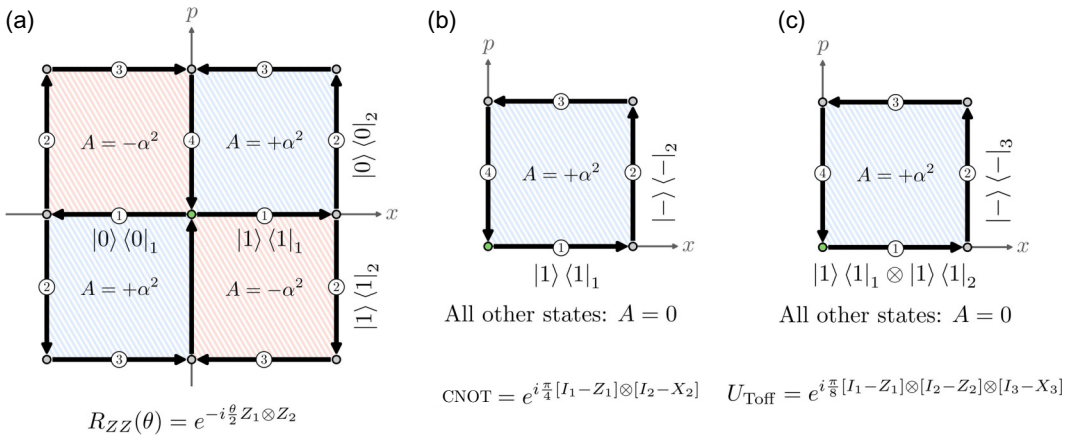


FIG. 24. Phase-space trajectories that realize (a) $R_{ZZ}(\theta)$, (b) CNOT, and (c) Toffoli. In all cases, displacement conditions are indicated by projectors onto particular qubit states, with subscripts labeling the qubit index. We label only the conditions for steps 1 and 2, as steps 3 and 4 consist of applying the Hermitian conjugate operator of the former.

operator. Consequently, this ability facilitates the ability to realize highly complex n-qubit entangling gates.

As a simple example, we now consider the Toffoli gate defined as

$$U_{\text{Toff}} = e^{i\pi[I_1-Z_1]\otimes[I_2-Z_2]\otimes[I_3-X_3]/8}. \tag{356}$$

There is freedom in choosing the particular sequence of conditional displacements. One solution is to rewrite U_{Toff} as a product of displacements conditioned on (commuting) Pauli operators of weight 3 and realize each with Eq. (338). A more efficient solution, though, is to recognize that

$$\begin{aligned} U_{1,2}(\alpha) &= e^{[I_1-Z_1][I_2-Z_2](\alpha a^\dagger-\alpha^* a)/4} \\ &= D_c^{(2,2)}\left(0,\frac{\alpha}{2}\right)\text{CP}^{(2,1)}D_c^{(2,2)}\left(0,-\frac{i\alpha}{2}\right)\text{CP}^{(2,1)\dagger}, \end{aligned} \tag{357}$$

where $\text{CP}^{(j,k)}$ is a controlled-parity operation between mode j and qubit k, natively realizable using the dispersive interaction (see Box IV.8). Here, the general idea is that $U_{1,2}(\alpha)$ realizes a displacement of oscillator 2 conditioned on *both* qubits 1 and 2 being in the state $|1\rangle$, as shown in the first step of Fig. 24(c). This is realized through the combination of the controlled-parity technique described in Eq. (336) and a final conditional displacement such that oscillator 2 is displaced only if both qubits are in the state $|1\rangle$.

Likewise, defining $U_3(\alpha) = D_{cX}^{(2,3)}(0,\alpha)$, the Toffoli gate can then be compiled as

$$U_{\text{Toff}} = U_3^\dagger\left(i\sqrt{\frac{\pi}{2}}\right)U_{1,2}^\dagger\left(\sqrt{\frac{\pi}{2}}\right)U_3\left(i\sqrt{\frac{\pi}{2}}\right)U_{1,2}\left(\sqrt{\frac{\pi}{2}}\right), \tag{358}$$

with the corresponding phase-space trajectory illustrated in Fig. 24(c).

d. Strategies for optimization. As a final remark on the compilation of multi-qubit gates, we note some additional strategies that can be implemented to reduce overall runtime and minimize decoherence. For one, in the above examples we have used a series of conditional displacements to accumulate a geometric phase on a single oscillator mode—a choice largely made for ease of explanation. A more efficient strategy is to "symmetrize" the compilation such that all involved oscillator modes accumulate a geometric phase in parallel. In particular, involving n modes allows us to impart a total phase of A by enclosing a phase-space area of A/n with each mode in parallel, leading to a $1/\sqrt{n}$ reduction in both the size and duration of each displacement relative to the single-mode realization. Thus, the benefit is a reduction in the multi-qubit gate time by a factor of $\sqrt{n}$, while simultaneously reducing the total photon number per cavity throughout the gate. This last feature helps reduce the impact of nonidealities such as cavity self-Kerr, though the total number of photons across all modes remains the same, leaving the rate of photon loss unchanged in comparison. As a simple illustration of this strategy in practice, Fig. 25 compares two circuits realizing a CNOT gate: one using the single-mode displacement sequence in Eq. (354), and the other using two modes in parallel.

In each of the previous examples, we have used a series of displacements of equal magnitude, thereby enclosing a square region of phase space. However, in certain contexts (and depending on the particulars of the experimental platform), it is advantageous to alter this approach and instead enclose a nonsquare region. To see this, let us take as an example the CNOT gate and note that there are multiple strategies for compilation, with Fig. 24(b) providing one such case in which both position and momentum boosts are conditional. Alternatively, we could have performed coordinate boosts conditioned on both $|1\rangle\langle 1|_1$ and $|-\rangle\langle -|_2$

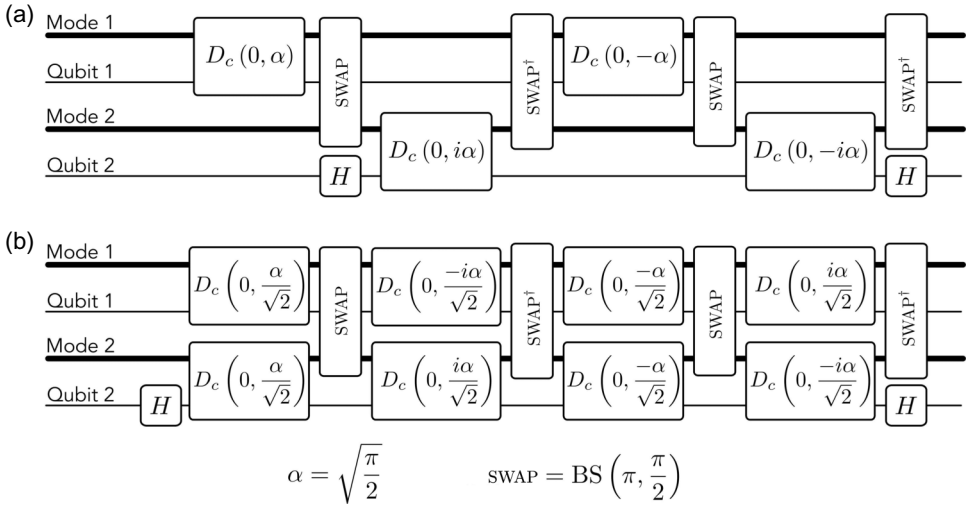


FIG. 25. Two quantum circuits that realize a CNOT gate between qubits 1 and 2. (a) The circuit corresponding to the phase-space trajectory in Fig. 24(b), where the entirety of the requisite geometric phase is achieved through the phase-space area enclosed by a single mode. (b) A more efficient strategy, where two modes are displaced in parallel, enclosing an area such that each contributes half the requisite geometric phase. The benefit of such an approach is that the size (and consequently duration) of each displacement is reduced by a factor of $\sqrt{2}$.

(synthesized similarly to the first step of our U_{Toff} compilation) with momentum boosts *unconditional*. As large unconditional displacements can be carried out much faster than their conditional counterparts [60], the overall runtime can be decreased by enclosing a rectangular region of phase space using short conditional displacements alternated with long unconditional displacements. The trade-off of this strategy is a relative increase in photon number (and, correspondingly, photon loss) during the gate. The optimal balance of these competing effects will in general depend on the specifics of the experimental platform and the gate being compiled, and constitutes an interesting future direction.

e. Advantages over MS and RIP gates. In superconducting platforms, analogs to MS gates known as resonator-induced phase (RIP) gates have been proposed [498–500]. The core idea is to leverage an architecture similar to that of trapped-ion platforms in Fig. 3(b), where multiple qubits are statically coupled to a common microwave bus resonator. Implementation of the gate then follows in analogy to MS gates—the resonator is driven in a qubit-conditioned closed phase-space loop, thereby enacting an entangling operation on the qubits. However, generalizing this gate beyond two-qubit gates to the multi-qubit case is a challenge, involving complex microwave pulse shapes [498].

A major weakness of MS and RIP gates is that both rely on a mediating motional or bus microwave mode that is *statically* coupled to a particular subset of qubits, as in Fig. 3(b), leaving the limited ability to selectively entangle arbitrary combinations of qubits. In contrast, our proposed approach exploits the local oscillator-qubit connectivity of the proposed superconducting architecture in Fig. 3(a) to realize a general multi-qubit entangling scheme with comparatively greater flexibility. While it relies on the same general principle as MS and RIP gates—i.e., an oscillator that is driven through closed, qubit-state-dependent loops in phase space—it is not limited to the use of motional or microwave modes with particular mode-qubit couplings and, instead, one can *selectively* couple particular qubits to a chosen oscillator. In some sense, it may be viewed as a digital analog of MS/RIP gates with a programmable common mode, where one can choose the number and locations of qubits to couple to through judicious use of beam splitters. Furthermore, as noted previously, the 2D connectivity graph of Fig. 3 can facilitate the implementation of many entangling gates in parallel given sufficient care is taken to avoid "collisions" between photons traveling via beam-splitter SWAPS.

Finally, we note that our scheme has potential advantages over conventional entangling gates in superconducting platforms that rely on direct coupling between superconducting qubits and are thus limited by crosstalk (see, e.g., Ref. [501] for a discussion regarding the cross-resonance gate). In contrast, here we rely only on the dispersive coupling to individual oscillators and use beam splitters to controllably route information, avoiding the problem of crosstalk while offering the flexibility of oscillator-mediated gates. An in-depth comparison between conventional entangling gates and the proposed oscillator-mediated scheme is beyond the scope of this work, however, and is left as a future direction.

3. Compilation of entangling gates between oscillators

While the superconducting architecture in Fig. 3(a) includes beam splitters as a native oscillator-oscillator entangling operation, systematic techniques to compile more complex entangling gates are desirable. In this section, we briefly review such a technique developed in Ref. [41] in the context of simulating lattice gauge theories. We refer to this work for details and merely summarize the basic concept here.

The core principle is to leverage an ancillary qubit to mediate interactions between two or more oscillators. Due to the mismatch in Hilbert space between the oscillators and the ancilla, this interaction is synthesized in a sequential, bit-wise fashion. To explain further, let us first assume access to the gate

$$U_i(\theta) = e^{-i\theta Z_{\text{anc}} O_i}, \tag{359}$$

where the subscript "anc" denotes an ancillary qubit in a known initial eigenstate of Z_{anc} and O_i is any Hermitian operator acting on the ith mode. For example, U_i can be a conditional displacement ($O_i \to e^{i\phi} a_i^\dagger + e^{-i\phi} a_i$), a controlled phase-space rotation ($O_i \to a_i^\dagger a_i$), or even a SNAP gate (with $O_i \to \sum_n \varphi_n |n\rangle \langle n|$). We assume $U_i(\theta)$ to be native or easily synthesized using the primitives discussed above.

Next, we can synthesize a variant of this gate that is conditioned on information contained in mode j. As this information is mediated by an ancillary qubit, it is passed one bit at a time. In that vein, let us define the projector P_j onto a set of Fock states in the subspace $\mathcal{P}$,

$$P_j = \sum_{n\in\mathcal{P}} |n\rangle \langle n|_j \,. \tag{360}$$

We can realize a qubit rotation conditioned on this projector with an SQR gate (Box IV.9),

$$\text{SQR}_j(\vec{\theta}_P, \vec{0}) = e^{-i\frac{\pi}{2} X_{\text{anc}} P_j} \equiv V_{P_j}, \tag{361}$$

where $\vec{0} = \{0, 0, 0, \ldots\}$ and $\vec{\theta}_P = \{\theta_0, \theta_1, \theta_2, \ldots\}$ with

$$\theta_n = \begin{cases} \pi, & \text{if } n \in \mathcal{P}, \\ 0, & \text{otherwise.} \end{cases} \tag{362}$$

Leveraging this gate, it can then be shown that

$$\text{CU}_i^{P_j}(\theta) = e^{-i\theta Z_{\text{anc}} O_i P_j} = V_{P_j} U_i(\theta/2) V_{P_j}^\dagger U_i(-\theta/2). \tag{363}$$

Therefore, we can synthesize the "P-conditional" unitary $\text{CU}_i^{P_j}$ using two SQR gates and two applications of U_i.

While $\text{CU}_i^{P_j}$ enacts an operation on the ith mode conditioned on a *single* bit of information concerning the jth mode, it is possible to synthesize more complex entangling operations by iterating this gate for different choices of P_j. As a simple example, let us choose for the unconditioned unitary,

$$U_i(\theta) = \text{SNAP}_i(\vec{\varphi}(f, \theta)) = e^{-i\theta Z_{\text{anc}} f(\hat{n}_i)}, \tag{364}$$

where $f(\hat{n}_i) = \sum_k (\varphi_k/\theta) |k\rangle \langle k|_i$ is an arbitrary polynomial in $\hat{n}_i$, the form of which depends on the chosen vector of angles $\vec{\varphi} = \{\varphi_0, \varphi_1, \varphi_2, \ldots\}$. Furthermore, let the projector for the kth iteration be that of the kth Fock state, $P_j^{(k)} = |k\rangle \langle k|_j$. Then, for an appropriately chosen cutoff N_{max}, one can realize

$$\begin{aligned} \text{CU}_i^{h(\hat{n}_j)} &= e^{-i\theta Z_{\text{anc}} f(\hat{n}_i) h(\hat{n}_j)} \\ &= \prod_{k=0}^{N_{\text{max}}-1} V_{|k\rangle\langle k|_j} U_i(\theta_k/2) V_{|k\rangle\langle k|_j}^\dagger U_i(-\theta_k/2). \end{aligned} \tag{365}$$

Here, $h(\hat{n}_j) = \sum_k (\theta_k/\theta) |n\rangle \langle n|_j$ is an arbitrary function of the number operator $\hat{n}_j$. The depth of this sequence scales as $O(N_{\text{max}})$. However, for certain functions $h(\hat{n}_i)$, a more efficient scheme is possible that requires $O(\log(N_{\text{max}}))$ iterations (see Ref. [41] for details). Moreover, we emphasize that while the product $f(\hat{n}_i)h(\hat{n}_j)$ encompasses the synthesis of extremely complex, nonlinear entangling two-mode gates, it is a specialized form of the more general function $f(\hat{n}_i, \hat{n}_j)$. If the aim is to compile the latter more general form, this can also be done. In particular, it can be realized by iterating Eq. (365) with $f_\ell(\hat{n}_i) = |\ell\rangle \langle \ell|_i$ for $\ell \in [0, N_{\text{max}})$, requiring a total depth $O(N_{\text{max}}^2)$ sequence.

Finally, we note that the examples above merely scratch the surface of the possible gates that can be analytically synthesized using the primitive in Eq. (363). For example, it is possible to iterate this scheme for more than two oscillators, enabling the systematic synthesis of multi-oscillator entangling gates mediated by an ancilla qubit—a powerful resource for the OCMM (see Sec. II B 2). We note that unlike the oscillator-mediated multi-qubit gates of Sec. VI C 2, the qubit-mediated multi-oscillator gates proposed here require the ancillary qubit to be in a known eigenstate of Z_{anc} if one wishes to leave the qubit and oscillators disentangled. However, in line with the HMM (see Sec. II B 3), there may be instances in which this family of gates is useful for hybrid multi-oscillator-qubit operations where the qubit is not ancillary but instead used as a computational degree of freedom. Both of these possibilities are fruitful areas for further research.

D. Systematic compilation of single-qubit-oscillator unitaries

In the previous section, we have presented exact analytical compilation schemes for some key compilation tasks.

In this section, we present systematic compilation techniques for single-qubit oscillator unitaries by leveraging methods in DV quantum simulation for compilation purposes. These techniques are analytical and formally exact, but only in the limit of an infinite number of infinitesimal compilation steps. Thus, while the gate sequences are "human-readable," they are not necessarily efficient. More concretely, Sec. VI D 1 discusses how to compile an algorithm into the bosonic ISA using Trotter-Suzuki discretization and product formulas [336–340]. We then present a generalized version of LCU techniques [502] on hybrid CV-DV systems in Sec. VI D 2 that can be applied on top of any unitaries compiled from product formulas or QSP to generate more sophisticated algorithms. We focus our discussion in Sec. VI D 3 on hybrid CV-DV QSP by generalizing the original QSP theorem developed for compiling single-qubit dynamics [49,274].

1. Trotter-Suzuki and product formulas

The first technique we discuss is the Trotter decomposition method and the product formulas, in general [335–339,341]. This section is a continued discussion of Ref. [340] and is largely pedagogical. The idea is to use Hamiltonian simulation and group-commutator expansions to construct particular unitaries. In the case of hybrid oscillator-qubit operators, special attention is needed to take care of the unbounded operator norm for typical operators of the infinite-dimensional oscillator (e.g., the position, momentum, and photon-number operators).

Specifically, imagine that we wish to add two operators with support on a qubit and an oscillator $A = \hat{b}_1 \cdot \vec{\sigma} h_1(\hat{x},\hat{p})$ and $B = \hat{b}_2 \cdot \vec{\sigma} h_2(\hat{x},\hat{p})$. We can approximate the sum of the two operators using a first-order Trotter formula with r time steps via

$$e^{-i(A+B)} = (e^{-iA/r}e^{-iB/r})^r + O(\max(\|A\|,\|B\|)^2/r), \tag{366}$$

or a second-order Trotter formula

$$e^{-i(A+B)} = (e^{-iA/2r}e^{-iB/r}e^{-iA/2r})^r + O(\max(\|A\|,\|B\|)^3/r^2), \tag{367}$$

where the approximation error is given by the norm of the operator (while a more accurate commutator norm may be used to improve the bound [339,503]).

Note that $\hat{x}$ and $\hat{p}$ have infinite operator norm, so $\|A\|$ and $\|B\|$ are in general unbounded. In practice, however, the unitaries are applied to quantum states with bounded energy, thus producing states with bounded energy; therefore, we can loosely view the various relevant operator norms as effectively bounded. Furthermore, while $\hat{x}$ and $\hat{p}$ are unbounded, we do not apply them directly. Rather, we apply exponentials of $\hat{x}$ and $\hat{p}$ which, despite the unbounded nature of their generators, are themselves well-defined unitaries the eigenvalues of which lie on the unit circle in the complex plane. Products of such unitaries obey a simple composition law, the analytical form of which is derived in Sec. VIC 1 c and this can be useful both for circuit synthesis and for estimation of errors.

In addition to the Trotter-Suzuki method to *add* Hamiltonian terms, BCH formulas are also useful to create commutators of two Hamiltonians [335–341]:

$$e^{itA}e^{itB}e^{-itA}e^{-itB} = e^{-t^2[A,B]} + O(t^3). \tag{368}$$

Using this, the *product* between two operators can be approximated by exploiting the anticommutation properties of Pauli operations. This leads us to the following expansion:

$$e^{-i[h_1(\hat{x},\hat{p})h_2(\hat{x},\hat{p})+h_2(\hat{x},\hat{p})h_1(\hat{x},\hat{p})]\otimes\sigma_z} = \left(e^{-ih_1\otimes\sigma_x/\sqrt{r}}e^{-ih_2\otimes\sigma_y/\sqrt{r}}e^{ih_1\otimes\sigma_x/\sqrt{r}}e^{ih_2\otimes\sigma_y/\sqrt{r}}\right)^r + O(\max(\|h_1\|,\|h_2\|)^2\min(\|h_1\|,\|h_2\|)/\sqrt{r}). \tag{369}$$

Higher-order variants of both the formula for the sum and the formula for the product exist [337–339,503,504] and can be used to achieve subpolynomial scaling of the error with $1/r$ in both cases. These formulas are of particular interest in our case because gates such as conditional displacements and SNAP gates are controlled by an auxiliary qubit that can be manipulated through the use of single-qubit rotations, enabling us to transform from σ_z-controlled operations into σ_x- or σ_y- controlled operations. This ability allows us to leverage a single qubit to build arbitrary polynomial interactions in $\hat{x}$ and $\hat{p}$ (or $a^\dagger$ and a) on the qumode and thus achieve universality, as discussed in Sec. IV E.

More formally, given an arbitrary degree-d Hermitian polynomial $h_d(\hat{x},\hat{p})$ as in Eq. (329), we would like to find a sequence of L_d controlled displacements $\{U_j\}$ ($j = 1,2,\ldots,L_d$) such that

$$U_{L_d}U_{L_d-1}\cdots U_2U_1 \approx e^{i\hat{b}_d\cdot\vec{\sigma}h_d(\hat{x},\hat{p})}, \tag{370}$$

where $U_j = e^{i\hat{c}_j\cdot\vec{\sigma}(\alpha_j a^\dagger - \alpha_j^* a)}$, and $\hat{b}_d$ is an arbitrary unit vector. One may construct the sequence of $\{U_j\}$ in reverse by multiplying a controlled displacement on the left of both sides of Eq. (370), such that the resulting unitary has a simpler generator.

a. Example: Jaynes-Cummings interaction. As a simple example, let us consider the approximate synthesis of the Jaynes-Cummings and anti-Jaynes-Cummings gates (see

Table IV),

$$\mathrm{JC}(\theta) = e^{-i\theta(\sigma_- a^\dagger + \sigma_+ a)}, \tag{371}$$

$$\mathrm{AJC}(\theta) = e^{-i\theta(\sigma_+ a^\dagger + \sigma_- a)}, \tag{372}$$

where for simplicity we have restricted the parameter θ to be real. (The reader is also directed to Sec. VII C 4 which discusses the Rabi model, which is the parent of these models.) The microscopic physical Rabi model Hamiltonian in cQED (see Appendix A) contains the sum of the JC and AJC Hamiltonians, and the corresponding unitaries are individually accessible within separate rotating-wave approximations (RWAs). Hence we list these as natively available within the sideband ISA.

Our goal here is to approximately synthesize these sideband unitaries within the phase-space ISA and the Fock-space ISA. To that end, we first note that they can be cast into the form

$$\mathrm{JC}(\theta) = e^{-i\theta(B-A)}, \tag{373}$$

$$\mathrm{AJC}(\theta) = e^{-i\theta(A+B)}, \tag{374}$$

where

$$A = \hat{p}\sigma_y, \tag{375}$$

$$B = \hat{x}\sigma_x. \tag{376}$$

Importantly, both A and B generate conditional displacements that, up to single-qubit rotations, can be directly realized either natively in the phase-space ISA or through a combination of controlled-parity gates and unconditional displacements in the Fock-space ISA. Thus, both $\mathrm{JC}(\theta)$ and $\mathrm{AJC}(\theta)$ can be approximately implemented within both of these ISAs via Eq. (366) through an alternation of coordinate and momentum displacements, conditioned on orthogonal qubit bases. We give a comparison of first-order and second-order Suzuki-Trotter approximations to the unitary $\mathrm{AJC}(\theta)$ using the phase-space ISA gates in Appendix F.

b. Example: Conditional squeezing operation. Using Eq. (368) with $t = \sqrt{\zeta}$ and A, B as defined in Eqs. (375) and (376) yields (for the case of ζ real and small)

$$\begin{aligned} e^{i\sqrt{\zeta}A} e^{i\sqrt{\zeta}B} e^{-i\sqrt{\zeta}A} e^{-i\sqrt{\zeta}B} &\approx e^{-\zeta[A,B]} \\ &\approx e^{i\zeta Z(\hat{x}\hat{p}+\hat{p}\hat{x})}, \end{aligned} \tag{377}$$

which matches the controlled-squeezing (CS) unitary $\mathrm{CS}(\zeta)$ defined in Table IV. For an illustration of how the Wigner function evolves after each successive conditional displacement, see Fig. 22(d). See also Sec. VIC 1 c, where we provide an analytical decomposition for sequences of conditional phase-space displacements of the form in Eq. (377), and provide intuition on the resulting composite operation when ζ is not infinitesimal. Separately, we note that in Ref. [505], the authors have numerically optimized circuits for unconditional squeezing using a sequence of alternating conditional displacements and momentum boosts. Reference [60] has experimentally demonstrated conditional squeezing using a numerically optimized sequence of conditional displacements and momentum boosts. References [506,507] have presented an analog scheme for producing conditional squeezing by driving a qubit-resonator system at twice the resonator frequency, while Ref. [508] has presented a scheme based on driving an ancilla qubit at the resonator frequency.

c. Example: (Conditional) cross-Kerr interaction. While Trotter-Suzuki and BCH formulas are particularly useful for the compilation of unitaries acting on a single-qubit–oscillator block, we emphasize that these methods are quite general and can also be used to synthesize nontrivial entangling gates. As a simple demonstration, we approximately synthesize a conditional cross-Kerr interaction between a pair of oscillators labeled a_1 and a_2. In particular, we envision the connectivity illustrated in Fig. 3(a), with the ith mode dispersively coupled to the ith qubit, and a tunable beam splitter between the two modes of interest. Utilizing Eq. (368), we make the choice $A = X_1 a_1^\dagger a_1$ and $B = Y_1 a_2^\dagger a_2$, with X_1 and Y_1 Pauli operators for qubit 1, $U_A = \exp(-i(\eta/2)A)$ corresponding (up to single-qubit rotations) to a conditional cavity-rotation gate $\mathrm{CR}(\eta)$, and similarly $U_B = \exp(-i(\eta/2)B)$ is independently compiled by conjugating U_A by appropriate single-qubit rotations and by a beam splitter that SWAPs the two modes. In all, this yields the desired unitary up to an error scaling as η^3,

$$U_A^\dagger U_B^\dagger U_A U_B = e^{-i\frac{\theta}{2} Z_1 a_1^\dagger a_1 a_2^\dagger a_2} + O(\eta^3), \tag{378}$$

where $\theta = \eta^2$.

By setting the auxiliary qubit to be in $|0\rangle$ (the +1 eigenstate of Z), this cross-Kerr interaction can be made useful in quantum simulations of, e.g., the real-time dynamics (Hamiltonian evolution) of repulsively interacting bosons. The fact that the interaction is conditional on the auxiliary qubit state means that we can use the phase kickback onto the qubit to measure the mean value of the interaction Hamiltonian term in a trial variational ground state produced by a variational quantum eigensolver (VQE) [509] or the quantum approximate optimization algorithm (QAOA) [510]. This conditional cross-Kerr interaction can also be useful in compiling the fSim gate given in Eq. (454).

2. Linear combination of unitaries (LCU)

The linear combination of unitaries (LCU) technique has been used in the context of DV systems for efficient Hamiltonian simulation [502]. In the standard LCU approach, a DV register composed of qubits is used to store the amplitudes defining the LCU. A final measurement on the auxiliary-qubit register imposes a linear combination of multiple unitaries on a system DV register. Note that since an LCU is in general not a unitary operation, it can be realized only nondeterministically through the measurement of auxiliary variables. In some cases, it may be possible to use amplitude amplification techniques [274,511,512] to bring the measurement success probability close to unity.

Here, we consider a generalization of LCU to hybrid bosonic quantum processors, where either the qubits or the oscillators can serve as auxiliary registers to store the linear combination coefficients in the LCU. In Fig. 26(a) we illustrate such a scheme for a single-qubit–oscillator system, and the multi-qubit generalization for the auxiliary register immediately follows.

Using oscillators as auxiliary registers enables the implementation of a continuous integral of unitaries on the system register, instead of a discrete sum of unitaries as in the original formulation of LCU. In fact, we have already seen a prototypical example of this concept for the conditional displacement gate in Fig. 13. Reference [513] has given another powerful example of a continuous-variable LCU block-encoding of a Gaussian imaginary time evolution spectral filter and applied this to the task of eigenstate preparation in quantum spin models.

In Fig. 26(b) we show that in the case of a single qubit and oscillator, a state preparation unitary V can be applied, followed by an oscillator-controlled qubit rotation and the inverse preparation unitary $V^\dagger$. A linear combination of a continuous family of qubit rotations can be applied provided that a measurement on the oscillators is performed for postselection. In the case in which the oscillator is projected to a position eigenstate $|x''\rangle$, the resulting qubit state is given by

$$|\psi_{\text{final}}\rangle = \int_{-\infty}^{\infty} \mathrm{d}x\, c_{x''}(x)\hat{R}(x)\,|0\rangle\,, \tag{379}$$

where

$$c_{x''}(x) = \Psi(x)\,\langle x''|\,V^\dagger\,|x\rangle \tag{380}$$

is an amplitude parametrized by a position x when the oscillator is projected to x'', $\hat{R}(x)$ is a single-qubit rotation parametrized by x, and $\Psi(x) = \langle x|\,V\,|0\rangle_{\text{osc}}$ is the wave function of the oscillator resulting from the preparation unitary V. Note that the oscillator-controlled qubit rotation in the middle of Fig. 26(b) applies the qubit rotation $\hat{R}(x)$ conditioned on the position x of the oscillator:

$$|x\rangle_{\text{osc}}\,|0\rangle \rightarrow |x\rangle_{\text{osc}}\,\hat{R}(x)\,|0\rangle\,. \tag{381}$$

As a simple example, let us consider the nondeterministic creation of an oscillator cat state by the application of the controlled unitary given in Eq. (389). The circuit is illustrated in Fig. 27. Measurement of the qubit in the Z basis produces an ideal even- or odd-parity cat depending on the measurement result. Measurement of the qubit in the X basis produces a coherent state $|\pm\alpha\rangle$ instead of a cat state. These hybrid CV-DV circuits have been implemented, and the resulting theoretical Wigner functions shown in Fig. 27 have been obtained using BOSONIC QISKIT [52,53]. For experimental implementation of this protocol in circuit QED, see Ref. [212].

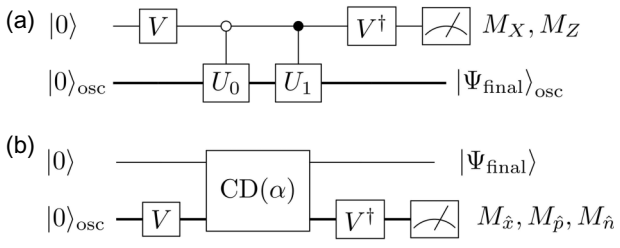


FIG. 26. LCU in hybrid bosonic quantum processors using (a) qubits or (b) oscillators as auxiliary registers to store the linear combination amplitudes. The scheme can be easily generalized to multiple qubits and oscillators. In (a), the measurement on the qubit that postselects the oscillator final state $|\Psi_{\text{final}}\rangle_{\text{osc}}$ can, e.g., be made in either the Z (M_Z) or X (M_X) basis. Similarly in (b), the postselection measurement on the oscillator can, e.g., be made via homodyne measurement in one of the oscillator quadrature bases $M_{\hat{x}}, M_{\hat{p}}$ or the photon-number basis ($M_{\hat{n}}$), each of which yields a different final state $|\Psi_{\text{final}}\rangle$ on the qubit.

3. Quantum signal processing for hybrid systems

Quantum signal processing (QSP) and its cousin, the quantum singular-value transformation (QSVT) [49,274,514], are the second family of techniques that we consider for programming hybrid devices via unitary synthesis. The intellectual origins of the QSP concept are found in the robust spin-rotation control techniques developed by the nuclear magnetic resonance (NMR) community. Today, QSP is the unifying principle behind many existing quantum algorithms for qubit-based platforms [49] with many applications (see, e.g., Ref. [515]).

In this section, we describe how to extend these concepts from control problems where the spin-rotation angles are *classical* control parameters to the case in which they are *quantum* operators—namely the position and momentum operators of harmonic oscillators. Ultimately, this enables the extension of QSP unitary synthesis techniques to hybrid qubit-oscillator systems. We emphasize that this section assumes knowledge of qubit-based QSP; for a brief introduction, see Appendix G.

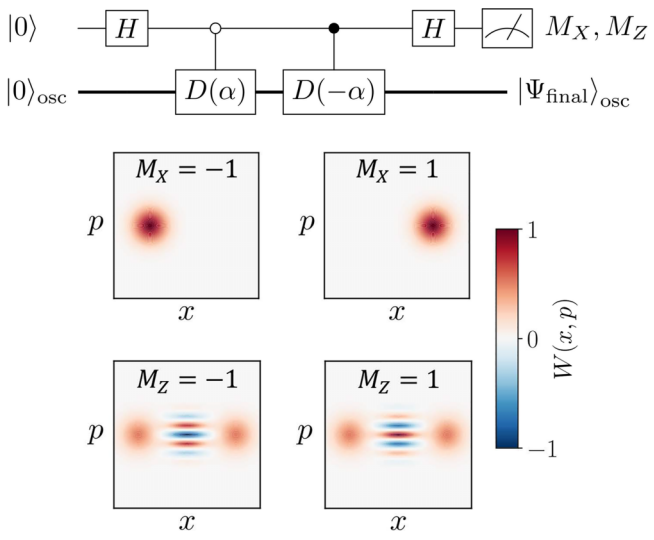


FIG. 27. An example of the LCU circuits in Fig. 26(a) that simply produce a displaced coherent state $|\pm\alpha\rangle$ conditioned on the ancilla-qubit measurement result in the X basis (Wigner functions shown in the middle row), or produces a cat state if the qubit measurement is performed in the Z basis (Wigner functions shown in the lower row). Note (from the color of the central interference fringe) that the parity of the cat is even (lower right panel) if the Z measurement yields $+1$ and odd (lower left panel) if the Z measurement yields $+1$.

Before formally defining hybrid CV-DV QSP, we will, as a first example, use a well-known composite pulse sequence from NMR, namely $\text{BB1}_{90,y}$ [49,516] to illustrate how QSP can be applied to oscillator-qubit control. The $\text{BB1}_{90,y}$ (see also Appendix G) composite pulse can be written as

$$\text{BB1}_{90,y} = Z(\phi_0)\prod_{j=1}^{d} W_x(\theta)Z(\phi_j) \tag{382}$$

$$= e^{-i\frac{\pi}{2}\vec{h}(\theta)\cdot\vec{\sigma}}, \tag{383}$$

where $\vec{h}(\theta) = \big(h_0(\theta), h_x(\theta), h_y(\theta), h_z(\theta)\big)$ is a vector the components of which are four functions of the rotation angle θ and $\vec{\sigma} = (\sigma_0, \sigma_x, \sigma_y, \sigma_z)$. The striking similarity between Eq. (383) and the general qubit-oscillator hybrid Hamiltonian in Eq. (242) for controlling a hybrid qubit system inspires us to think about a method to promote θ from a noisy classical variable to a quantum operator $\hat{\theta}$ such as $\hat{x}$ acting on an oscillator.

Consider for example, the following controlled displacement operation (Box IV.7, Table V):

$$e^{-i\frac{k}{2}\hat{x}\sigma_z} = \begin{bmatrix} e^{-i\frac{k}{2}\hat{x}} & 0 \\ 0 & e^{i\frac{k}{2}\hat{x}} \end{bmatrix} := W_z, \tag{384}$$

which imparts a qubit-dependent momentum boost on the oscillator. Importantly, we can also interpret this as a rotation of the qubit about the z axis by an angle linearly proportional to the position of the oscillator,

$$W_z(\hat{\theta}) = e^{-i\frac{\hat{\theta}}{2}\sigma_z}, \tag{385}$$

$$\hat{\theta} \equiv k\hat{x}. \tag{386}$$

Since the oscillator position is a quantum operator, it has quantum fluctuations analogous to the classical fluctuations of the control parameter θ that we discussed in connection with composite pulses above. Remarkably, as long as we deal only with a single operator $\hat{x}$ or $\hat{p}$ (or a fixed linear combination of them) then all of the rotation angles commute with each other, and the machinery of QSP developed for classical control variables can be applied without modification.

The noise model that we used in the classical BB1_{90} composite pulse example was that we have perfect control over the axis of rotation on the Bloch sphere, but imperfect control over the rotation angle. For the quantum case, we assume that we have perfect control over all qubit rotation axes and angles when they are created using standard classical control pulses directly applied to the qubit. However, when we rotate the qubit by an angle $\hat{\theta}$ associated with the oscillator position, the quantum fluctuations of the oscillator play the role of noise in the rotation angle. Within this model, we can convert the controlled-momentum boost $W_z(\hat{\theta})$ into the form compatible with the $\text{BB1}_{90,y}$ protocol in Eq. (382):

$$W_x(\hat{\theta}) = R_{+\frac{\pi}{2}}\left(\frac{\pi}{2}\right) W_z(\hat{\theta}) R_{+\frac{\pi}{2}}\left(-\frac{\pi}{2}\right)$$

$$= e^{-i\frac{k}{2}\hat{x}\sigma_x} = \begin{bmatrix} 0 & e^{-i\frac{k}{2}\hat{x}} \\ e^{i\frac{k}{2}\hat{x}} & 0 \end{bmatrix}. \tag{387}$$

As a simple illustration of the hybrid CV-DV QSP procedure, we use the $\text{BB1}_{90,y}$ protocol within the phase-space ISA (Sec. IV E 2) to deterministically create a Schrödinger cat state in an oscillator. Two-legged Schrödinger cat states are superpositions of coherent states (defined in Sec. III C 1) $|\pm\alpha\rangle$ equidistant from the origin in phase space,

$$|C\rangle_\alpha = \mathcal{N}(|\alpha\rangle \pm |-\alpha\rangle), \tag{388}$$

where $\mathcal{N} = 1/\sqrt{2 + 2e^{-2|\alpha|^2}}$ and the $\pm$ sign determines whether the photon-number parity is even $(+)$ or odd $(-)$. This is useful for encoding logical information and for quantum sensing of displacements. For simplicity, we have chosen α to be real so that the cat is aligned along the position axis of phase space.

Suppose that we would like to prepare the two-legged cat state from the ground state of the qubit and oscillator

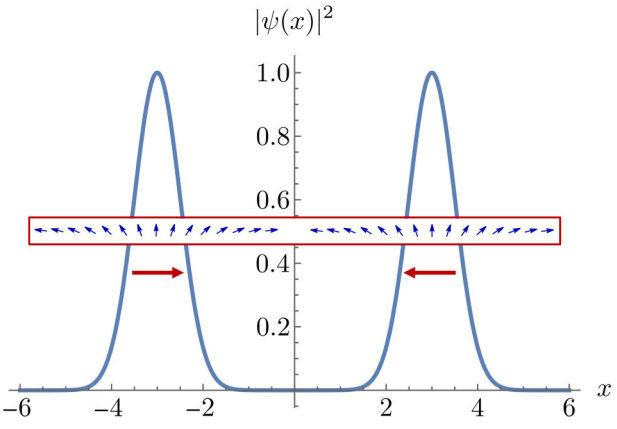


FIG. 28. The probability density $|\psi(x)|^2$ for a cat of size $\alpha = 3$, illustrating the quantum fluctuations of the position about the peak positions $x = \pm\alpha$. The large red arrows indicate the initial spin entanglement with position based on Eq. (389). The blue arrows in the red box indicate the spin configuration after the application of the single-qubit-controlled-momentum boost (or equivalently, an x-dependent qubit rotation) given in Eq. (391). In the position regime where the probability density is high, the spin orientation is approximately constant, indicating that the spin and the oscillator are largely (but not perfectly) disentangled.

without any measurement. Using conditional displacement ($e^{+i2\alpha\hat{p}\otimes\sigma_x}$, in Wigner units) from the phase-space instruction set, the initial hybrid state can be entangled with the qubit as follows:

$$e^{-i2\alpha\hat{p}\otimes\sigma_x}\,|0\rangle_{\rm vac}\otimes|0\rangle = \frac{1}{\sqrt{2}}\left[|\alpha\rangle_{\rm osc}|-\rangle + |-\alpha\rangle_{\rm osc}|+\rangle\right]. \tag{389}$$

Here, $|\pm\rangle = \frac{1}{\sqrt{2}}(|0\rangle \pm |1\rangle)$ are the qubit σ_x eigenstates. For large enough α, we can safely assume that the Gaussian wave functions of $|\pm\alpha\rangle$ do not overlap, as shown in Fig. 28. We can also see directly from this plot that the probability distribution for the position of the oscillator is strongly peaked near $= \pm\alpha$, but as expected, has quantum fluctuations around those nominal positions.

Now, to create an even-parity cat state, the qubit (see the red arrows in Fig. 28) needs to be disentangled from the oscillator state by a rotation about the y axis by an angle $+\pi/2$ if the oscillator position, $\hat{x}$, is positive and $-\pi/2$ if the position is negative to create the target cat state

$$\frac{1}{\sqrt{2}}\left[|\alpha\rangle_{\rm osc} + |-\alpha\rangle_{\rm osc}\right]\otimes|0\rangle. \tag{390}$$

This transformation is (essentially) unitary so long as the coherent states have negligible overlap. As noted above, conditional momentum boosts can be viewed as a conditional rotation of the qubit based on the position of the CV state. Thus, as a first approximation, we could use

$$e^{-i\frac{k}{2}\hat{x}\otimes\sigma_y} = Z\left(+\frac{\pi}{2}\right)W_x(\hat{\theta})Z\left(-\frac{\pi}{2}\right), \tag{391}$$

$$\text{with}\quad k = \frac{\pi}{2\alpha} \tag{392}$$

to rotate the qubits correctly at positions $x = \pm\alpha$ corresponding to the peaks in the oscillator position probability distribution. The blue arrows inside the red box in Fig. 28 show how the initial spin polarization of the qubit (shown as red arrows) changes upon applying this oscillator-position-dependent rotation to a (moderately) large cat state.

We see that the controlled-momentum boost has approximately disentangled the qubit from the oscillator, thus moving the state closer to the target. The disentanglement is not perfect, however, because for $|x| \neq \alpha$, there is some under-rotation or over-rotation due to the (quantum) uncertainty in the value of the position of the oscillator in the coherent states $\pm|\alpha\rangle$. The position uncertainty visible in Fig. 28 is set by the vacuum fluctuations and is independent of the displacement α, while we see from Eq. (392) that the rate of spin rotation with position is inversely proportional to α. Thus we expect that the rotation errors will decrease and the fidelity to the target state should improve with increasing α. Indeed, we find that the fidelity obeys

$$F_{\rm cat} \approx 1 - \frac{\pi^2}{64\alpha^2} \tag{393}$$

for asymptotically large α.

To increase this fidelity, we can directly apply the $\mathrm{BB1}_{90,y}$ protocol in Eq. (382) to (largely) eliminate the over- or under-spin-rotation errors. The only change that needs to be made is to replace the classically noisy rotation $W_x(\theta)$ in Eq. (382) with the quantum noisy rotation $W_x(\hat{\theta})$ in Eq. (387). The classical $\mathrm{BB1}_{90,y}$ protocol still works because the only quantum rotation angle is $\hat{\theta}$ [defined in Eq. (386)] and so there are no noncommuting operators in the quantum version of the $\mathrm{BB1}_{90,y}$ sequence.

In Fig. 29 we show a comparison of the fidelities to the target state for three different state preparation protocols: the simple protocol in Eq. (392), the $\mathrm{BB1}_{90,y}$ composite pulse QSP protocol, and a bivariable or "non-Abelian QSP" protocol [478] which is discussed further below. We see that the $\mathrm{BB1}_{90,y}$ protocol is very efficient at correcting the spin-rotation errors associated with quantum fluctuations in the oscillator position.

a. Formal exposition of QSP for hybrid CV-DV systems. We now turn to a more formal exposition of the extension of QSP to the case in which the rotation angles can be quantum operators acting on an oscillator, focusing on the

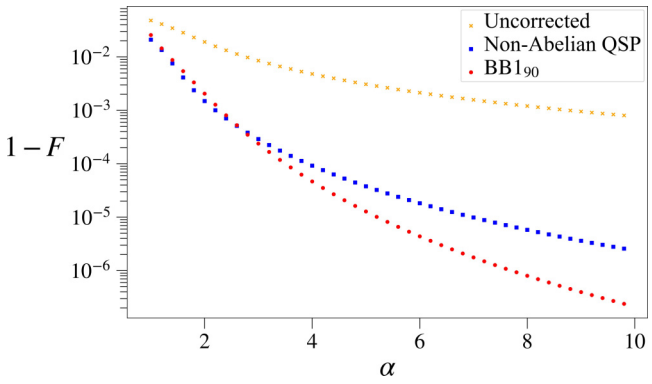


FIG. 29. The infidelity of cat state preparation by unentangling the qubit from the oscillator state in Eq. (389) using the BB1$_{90}$ sequence of commuting variables and non-Abelian QSP correction. For various cat sizes α, we plot the infidelity of the output hybrid state with the desired state $|C\rangle_\alpha \otimes |0\rangle$. Here, "uncorrected" refers to the simple protocol using the operation described in Eq. (391). The "non-Abelian QSP" and "BB1$_{90}$" curves, respectively, denote corrections over the simple protocol using operations described in Eq. (G3) and Sec. VID 3 b. Figure adapted from Ref. [478].

unitary picture (Sec. VI B 2) rather than the Hamiltonian picture of Eq. (383) and Sec. VI B 1. Here, we explore the single commuting variable case (e.g., all rotation angles proportional to $\hat{x}$) and in the next subsection, we consider how to extend QSP to the case of noncommuting rotation angles in which both variables $\hat{x}$ and $\hat{p}$ are involved. One key intuition for extending qubit QSP to hybrid qubit-oscillator systems is to realize that the controlled displacement operation (conditional rotation) in the phase-space ISA in Table V on the qubit-oscillator system can be viewed as a block-encoding of the displacement operator on the 2×2 matrix of the qubit subspace. In the spirit of QSP and QSVT, combining any block-encoding and single-qubit rotations that *process* the block-encoded operator, a polynomial transformation on the block-encoded operator can be realized.

Viewing the controlled displacement W_z in Eq. (385) as a block-encoding of the operator $w = e^{-i(k/2)\hat{x}}$, the following QSP sequence comprising d steps (described in more detail in Appendix G) produces an arbitrary polynomial transformation on w,

$$e^{i\phi_0\sigma_x}\prod_{j=1}^{d} W_z e^{i\phi_j\sigma_x} = \begin{bmatrix} F(w) & iG(w) \\ iG(w^{-1}) & F(w^{-1}) \end{bmatrix}, \tag{394}$$

where here we work in the z-rotation convention in which the roles of x and z rotations are interchanged relative to the BB1$_{90,y}$ example above. We again take the operator ordering to be left to right and

$$\begin{aligned} F(w) &= \sum_{n=-d}^{d} f_n w^n = \sum_{n=-d}^{d} f_n e^{-i\frac{nk}{2}\hat{x}} := f(\hat{x}), \\ G(w) &= \sum_{n=-d}^{d} g_n w^n = \sum_{n=-d}^{d} g_n e^{-i\frac{nk}{2}\hat{x}} := g(\hat{x}), \end{aligned} \tag{395}$$

are Laurent polynomial expansions of arbitrary functions $f(\hat{x})$ and $g(\hat{x})$ in canonical position space. The coefficients $f_n, g_n \in \mathbb{R}$ can be computed by Fourier series of $f(x)$ and $g(x)$ under the discrete plane-wave basis $\{e^{-i(nk/2)x}\}$ for $-d < n < d$ and for a given fixed k. In addition, unitarity requires $F(w)F(w^{-1}) + G(w)G(w^{-1}) = I$, and $F(w), G(w) \in \mathbb{R}[w, w^{-1}]$ with degree d, where $\mathbb{R}[w, w^{-1}]$ represents the set of polynomials of w, w^{-1} with real coefficients.

Note that even though $k\hat{x}$ is a quantum operator rather than a classical rotation angle, this construction is the same as the W_z convention of single-qubit QSP, and therefore the QSP theorem holds. The only difference is that the position variable of the oscillator is hidden inside a complex variable w. Such block-encoding of a unitary operator has also been discussed in Ref. [517] to prepare the ground state and estimate the energy in standard DV quantum computation by transforming the eigenvalues of the encoded unitary matrix.

More concretely, the form of $F(w)$ and $G(w)$ as given in Eq. (395) can be directly verified by explicitly evaluating the left-hand side of Eq. (394). Conversely, given any Laurent polynomial $F(w)$ of degree d written in the form in Eq. (395), there always exists a set of phase angles $\{\phi_0, \phi_1, \ldots, \phi_d\}$ such that the gate sequence as constructed in Eq. (394) block-encodes $F(w)$. See Ref. [518] for a constructive proof on the existence of the QSP phases for an arbitrary Laurent polynomial $F(w)$. By comparing Eq. (394) and (330), we may readily identify the polynomials $A_0(w) = (F(w) + F(w^{-1}))/2, A_1(w) = (G(w) + G(w^{-1}))/2, A_2(w) = i(G(w) - G(w^{-1}))/2, A_3(w) = (F(w) - F(w^{-1}))/2$, corresponding to I, σ_x, σ_y, and σ_z, respectively.

This simple generalization of QSP to hybrid bosonic systems allows us to implement an arbitrary unitary operation on the oscillator if that unitary is the exponential of a purely imaginary (polynomial) function of $\hat{x}$. This can immediately be changed to a function only of momentum $\hat{p}$, by simply using a qubit-dependent position kick $W_z = e^{-i(\lambda/2)\hat{p}\sigma_z}$ in the construction of Eq. (394). More generally, for an operator taking a fixed linear combination of $\hat{x}$ and $\hat{p}$, $(k/2)\hat{x} + (\lambda/2)\hat{p}$, one may use $W_z = e^{-i((k/2)\hat{x}+(\lambda/2)\hat{p})\sigma_z}$ in the construction. The ability of hybrid CV-DV QSP to prepare a wide range of oscillator states is useful for a variety of applications. For example, in Ref. [278], hybrid

QSP has been used to construct an interferometer for quantum sensing applications, where Heisenberg-like scaling is achieved for performing binary decisions on a displacement channel in the single-shot limit. It is also shown that such decision protocols can be chained together for efficient parameter estimation.

b. Non-Abelian QSP. The QSP on hybrid CV-DV systems in the previous section is, of course, a limited transformation, as in general, the position and momentum of the oscillator may be transformed independently by different polynomials.

In this section, we provide a construction of a non-Abelian QSP, discussed in [478], for the same task of cat state preparation and compare it to the single-variable case to highlight the advantages of processing both phase-space quadratures. Despite recent progress on multivariate QSP in the commuting case [275,519], a complete theory on polynomial transformation on the phase space for two noncommuting quadratures is yet to be developed [478,520].

Consider the controlled displacements along the momentum and position directions,

$$e^{-i\frac{k}{2}\hat{x}\cdot\sigma_z} = \begin{bmatrix} e^{-i\frac{k}{2}\hat{x}} & \\ & e^{i\frac{k}{2}\hat{x}} \end{bmatrix} = \begin{bmatrix} w(\hat{x}) & \\ & w^{-1}(\hat{x}) \end{bmatrix} := W_z^{(k)}(w), \tag{396}$$

$$e^{-i\frac{\lambda}{2}\hat{p}\cdot\sigma_z} = \begin{bmatrix} e^{-i\frac{\lambda}{2}\hat{p}} & \\ & e^{i\frac{\lambda}{2}\hat{p}} \end{bmatrix} = \begin{bmatrix} v(\hat{p}) & \\ & v^{-1}(\hat{p}) \end{bmatrix} := W_z^{(\lambda)}(v), \tag{397}$$

which can be viewed, respectively, as a block-encoding of a momentum boost $W_z^{(k)}(w)$ or a position displacement $W_z^{(\lambda)}(v)$.

We may now extend this to obtain the following unitary constructed from the above unitary operator alternating with x rotations on the qubit:

$$\begin{aligned} U_d &= e^{i\phi_0\sigma_x}\prod_{j=1}^{d} W_z^{(k)} e^{i\phi_j^{(k)}\sigma_x} W_z^{(\lambda)} e^{i\phi_j^{(\lambda)}\sigma_x} \\ &= \begin{bmatrix} F_d(w,v) & iG_d(w,v) \\ iG_d(v^{-1},w^{-1}) & F_d(v^{-1},w^{-1}) \end{bmatrix}, \end{aligned} \tag{398}$$

parametrized by the set of phases $\{\phi_j^{(k)},\phi_j^{(\lambda)}\}$ $(j = 1,2,\ldots,d)$. It can be shown by direct multiplication of terms in Eq. (398) that the above construction implements a bivariable Laurent polynomial transformation on noncommuting w and v with the following form:

$$F_d(w,v) := \sum_{r,s=-d}^{d} f_{rs} w^r v^s, \quad G_d(w,v) := \sum_{r,s=-d}^{d} g_{rs} w^r v^s, \tag{399}$$

where f_{rs} and g_{rs} are complex coefficients determined by the phase angles $\{\phi_j^{(k)},\phi_j^{(\lambda)}\}$. Note that because w and v do not commute, their order in the above expression matters. We define (somewhat arbitrarily) the way of writing out the bivariable Laurent polynomial to be *canonical* when all w are written to the left of v. Again, it can be seen that the unitary in Eq. (394) indeed complies with the form given in Eq. (330), but in the non-Abelian case, reduction to the Hamiltonian form in Eq. (242) may not be straightforward (since it constitutes the most hybrid general circuit compilation problem). We can, however, use the rules for the composition of controlled displacements given in Sec. VIC 1 c to obtain the overall unitary as a sum of terms.

To make this construction useful for compilation, it is desirable to show that for an arbitrary function F_d given in Eq. (399), the QSP phases $\{\phi_j^{(k)},\phi_j^{(\lambda)}\}$ $(j = 1,2,\ldots,d)$ always exist. The bivariable nature makes the qubit QSP theorem in Appendix G inapplicable. The noncommutative nature of $\hat{x}$ and $\hat{p}$ distinguishes the current context from the recent development of multi-variable QSP for commuting variables [275], and makes it challenging to prove a theorem governing efficient universality of non-Abelian QSP. However, numerical evidence suggests that the non-Abelian QSP construction can indeed efficiently approximate an arbitrary transformation of the phase space, with resolution set by the momentum kick k and the position kick λ. Reference [478] has shone a light on the analytical understanding of the non-Abelian QSP construction toward control of bosonic modes.

To gain some intuition about the capabilities of bivariable non-Abelian QSP, let us revisit the two-legged cat state preparation example discussed above. There is no classical analog for this framework since the bivariable non-Abelian QSP scheme makes full use of the noncommuting position and momentum variables in addition to the noncommuting Pauli matrices. In a recent work [478], a basic model of composite pulses in phase space has been introduced to correct for errors in Gaussian uncertainty of position variables using momentum variables, and vice versa. The authors begin with the Schrödinger entangled cat state in Eq. (389), but generalize it to allow for the possibility of squeezed states that appear in both squeezed cat [434,439,521,522] and GKP error-correction code states:

$$\psi(x) = \frac{1}{\sqrt{\pi\Delta^2}}\left[e^{-\frac{(x-\alpha)^2}{\Delta^2}}|-\rangle + e^{-\frac{(x+\alpha)^2}{\Delta^2}}|+\rangle\right], \tag{400}$$

where Δ^2 is the position variance of the squeezed coherent state relative to the vacuum value of $\sigma_x^2 = \frac{1}{4}$ found in Eq. (42).

In the Abelian QSP protocol, the next step would be to apply the controlled-momentum boost in Eq. (391) which rotates the qubit by an angle linear in the oscillator position to (nearly) disentangle the spin from the cavity. The quantum fluctuations in the position leave small residual over- or under-rotation errors, illustrated in Fig. 28. Near the peak positions $x = \pm\alpha$ the rotation error around the y axis of the Bloch sphere has the form

$$\delta\hat{\theta}_{\pm} = k(\hat{x} \mp \alpha), \tag{401}$$

with the constant k defined in Eq. (392). In the Abelian QSP protocol, we have substantially reduced these rotation errors by applying the three additional conditional momentum boosts of the $\mathrm{BB1}_{90,y}$ protocol.

In the non-Abelian scheme, Ref. [478] has shown that we have access to new spin-rotation angles that are linear in the oscillator momentum. We can use this new freedom to preemptively correct this error to first order in the quantum fluctuations of the oscillator position. We see from Fig. 28 that the residual rotation errors are the same for both peaks in the wave function. However, a little thought shows that if the correction is to be applied preemptively before the controlled-momentum boost in Eq. (391), then we must apply the small corrective rotations in *opposite* directions for positive and negative oscillator positions,

$$\phi(x) = \frac{1}{\sqrt{\pi\Delta^2}}\left[e^{+i\frac{\delta\hat{\theta}_+}{2}\sigma_y}e^{-\frac{(x-\alpha)^2}{\Delta^2}}|-\rangle + e^{-i\frac{\delta\hat{\theta}_-}{2}\sigma_y}e^{-\frac{(x+\alpha)^2}{\Delta^2}}|+\rangle\right]. \tag{402}$$

To achieve this (to the lowest order in k, we can take advantage of the fact that the derivative of a Gaussian has the right form (again in Wigner units)

$$\hat{p}e^{-\frac{(x\mp\alpha)^2}{\Delta^2}} = i\frac{(x\mp\alpha)}{\Delta^2}e^{-\frac{(x\mp\alpha)^2}{\Delta^2}}, \tag{403}$$

and further, use $\sigma_y = -i\sigma_z\sigma_x$ to obtain the necessary sign difference for the two peaks

$$\sigma_y|\pm\rangle = \mp i\sigma_z|\pm\rangle. \tag{404}$$

Using these results we see that we can obtain the preemptively corrected state (accurate to first order in k) by

$$\phi(x) \approx e^{i\frac{\pi\Delta^2}{4\alpha}\hat{p}\sigma_z}\psi(x). \tag{405}$$

Applying the same controlled-momentum boost in Eq. (391) as before then yields the desired nearly disentangled state with the residual rotation errors eliminated to first order in $\Delta^2/4\alpha$. While the $\mathrm{BB1}_{90,y}$ scheme also corrects second-order errors, this non-Abelian QSP preparation scheme requires fewer operations than $\mathrm{BB1}_{90,y}$ and avoids the associated large momentum boosts, making it more robust to ancilla and cavity decay errors and residual cavity anharmonicity. Essentially the same circuit is already in use experimentally for GKP state generation, stabilization, and measurement [65,326] and has been independently discovered by numerical optimization in [479]. An interesting extension of the method would be to use non-Abelian QSP techniques to modify this "big-small" protocol to also cancel the rotation errors to second order while avoiding the large momentum boosts of $\mathrm{BB1}_{90,y}$.

In Fig. 29 we show a comparison of the fidelities for the Abelian and non-Abelian QSP protocols presented above. A more general discussion on the framework of composite pulses in phase space has been described in Ref. [478], where the authors have also discussed robust protocols for other cases, including cat states represented by overlapping Gaussian functions, GKP code states, and Fock states.

Recall that the cat states realized by the LCU circuits in Sec. VI D 2 were exact but nondeterministic. The Abelian and non-Abelian QSP circuits here are deterministic but approximate.

Finally, we note that Ref. [277] has proposed a different generalization of QSP to purely CV systems (rather than hybrid) based on the $\mathrm{SU}(1,1)$ group, which can be realized via beam-splitter and parametric amplifier interactions (single- or two-mode squeezing) between two bosonic modes and is therefore potentially useful for compiling multimode unitary operations. This reveals the richness of the synergy between QSP and CV quantum systems, and we anticipate more developments in the future.

E. Approximate numerical compilation

As a pedagogical example of numerical compilation for bosonic unitaries, we detail here the particular approach from Ref. [60] and further extend this framework to any universal instruction set developed in this work. In many cases, a target bosonic operation (such as a unitary or state preparation) can be approximately synthesized from a universal gate set using numerical optimization. For example, Ref. [523] presents benchmarking techniques for variational quantum optimizers for bosonic state preparation. Separately, in Ref. [524], model-free reinforcement learning has been applied to approximately synthesize target bosonic operations using SNAP and displacement gates. The model-free reinforcement learning method is particularly well suited for *in situ* optimization, where unknown experimental microwave transfer functions can be included. Other numerical SNAP optimization strategies have been explored in Refs. [177,267]. Similarly, in Ref. [60], gradient-based optimization methods have

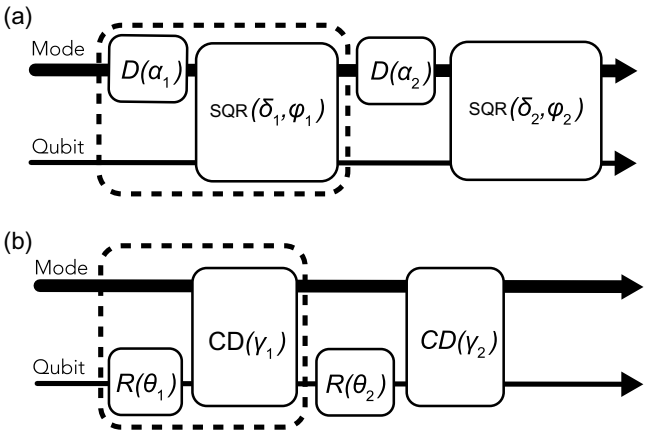


FIG. 30. Examples of circuit constructions used in numerical optimization with the (a) Fock-space and (b) phase-space instruction sets. For each, a layer is shown in the dotted box, and contains enough operations for universality. The circuit depth (the number of layers) and circuit parameters are optimized to realize a target unitary or state transfer. The SNAP construction (not shown) is similar, comprised of alternating SNAP gates and oscillator displacements. Recall from Eq. (239) that the SNAP gates discussed in the text can be compiled from pairs of SQR gates.

been developed that make use of automatic differentiation and gradient back-propagation to find gate parameters that approximately realize a target operation.

To outline the optimization problem, we first define the fundamental unit of each bosonic circuit to be a "layer." Each layer should consist of at least one application of each gate in the generating set of a universal set. Examples of circuits used for the Fock-space and phase-space gate sets are shown as the dotted boxes in Fig. 30. The circuit depth, N_{circuit}, is the number of layers. For circuits acting on multiple oscillators with some well-defined connectivity, the layer should also include one application of entangling gates between each connected oscillator (e.g., a beam splitter between each pair of neighboring oscillators).

To aid in computational efficiency, it is advantageous to optimize multiple randomly initialized circuits in parallel; this set of circuits is called a batch. Let B be the total number of independent circuits in a batch. We can denote b_{jk} as the unitary associated with layer $k \in \{1, 2, \ldots, N_{\text{layer}}\}$ in circuit $j \in \{1, 2, \ldots, B\}$ of the batch. With this, the unitary generated by the jth circuit is written as

$$U(\vec{\theta}_j) = b_{jN} \ldots b_{j1} b_{j0}, \tag{406}$$

where $\vec{\theta}_j$ is the set of parameters specifying the jth circuit. This parametrized unitary can then be used to define a cost function to be minimized. By calculating gradients using back-propagation, the cost function can be chosen to handle different situations. In the case of a target unitary operation, the cost function can be defined using the trace fidelity. For the jth circuit, this is given by

$$\mathcal{F}_j = \left| \frac{1}{\text{Tr}\{P\}} \text{Tr}\left\{ \mathcal{P} U_{\text{target}}^\dagger U_j \right\} \right|^2, \tag{407}$$

where $\mathcal{P}$ is a projector onto a subspace of interest [525, 526]. For example, $\mathcal{P}$ could consist of the first l Fock levels to implement a Fock truncation, $\mathcal{P} = \sum_{n=0}^{l-1} |n\rangle \langle n|$.

If the goal is to instead prepare a target pure state $|\psi_t\rangle$, the fidelity of the jth circuit can be defined using the overlap,

$$\mathcal{F}_j = \langle \psi_t | \rho_j | \psi_t \rangle, \tag{408}$$

where ρ_j is the result of the optimized circuit applied to a specified initial state. For example, consider the setting of a single oscillator coupled to a qubit, with the goal being to prepare a target cavity state $|\psi_t\rangle$ of the oscillator starting from the composite oscillator-qubit state $|0\rangle_{\text{vac}} \otimes |0\rangle$. If we do not care about the final state of the qubit (only that it is disentangled from the oscillator), the qubit can be traced out to obtain $\rho_j = \text{Tr}_q \left(U_j [|0\rangle_{\text{vac}} \langle 0|_{\text{vac}} \otimes |0\rangle \langle 0|] U_j^\dagger \right)$. The reduced density matrix ρ_j can then be used to evaluate the fidelity $\mathcal{F}_j$ of the resulting oscillator state.

From the fidelity $\mathcal{F}_j$ of the jth circuit, a total cost function can be defined using all B circuits in the batch, allowing for parallel optimization. As an example, one option that has proven useful in practice is

$$C = \sum_{j=1}^{B} \log(1 - \mathcal{F}_j). \tag{409}$$

Since the cost function C is a simple sum of independent logarithmic cost functions, gradient descent of C realizes independent gradient descent of each circuit realization in parallel. The logarithm of $1 - \mathcal{F}_j$ is used to aid with the problem of vanishing gradients as the solution infidelity approaches zero.

Once the parametrized cost function is defined, it is optimized by calculating gradients using back-propagation. Gradient-descent methods can then be applied using techniques that are well established. For example, the optimizer called *Adam* [527], commonly used to train neural networks, can be directly applied, making use of graphics processing units (GPUs). The optimization is stopped when any circuit fidelity $\mathcal{F}_j$ reaches the target fidelity, and the parameters from that circuit are selected. As an example of this, in Fig. 31, we show results of optimizing the preparation of Fock [Fig. 31(a)] and GKP [Fig. 31(b)] states using the phase-space or Fock-space ISAs.

It is not always guaranteed that random quantum circuits will converge to realize a target unitary operation. A complete formal complexity-theoretic analysis of hybrid

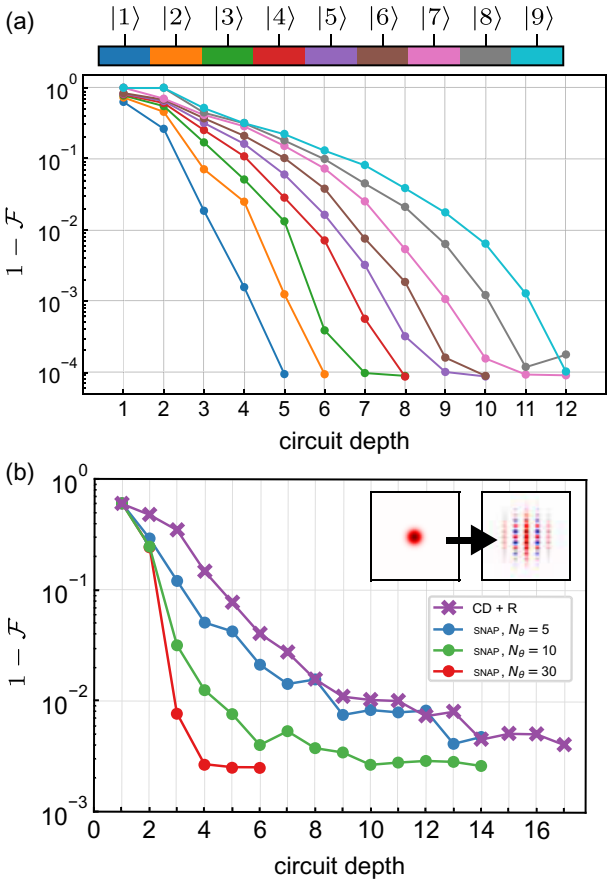


FIG. 31. Examples of numerical circuit optimization for an oscillator coupled to a qubit. (a) The infidelity of Fock states optimized using the phase-space ISA as a function of circuit depth (N) [528]. (b) The optimization of oscillator GKP $|+Z\rangle$ logical states with squeezing $\Delta = 0.32$ using either the phase-space ISA (crosses) or the Fock-space ISA via SNAP (circles). For the SNAP gates, the first N_θ phases are optimized, where N_θ is swept in the figure. As shown, including more phases in each SNAP gate leads to shorter-depth circuits.

CV-DV circuit synthesis for tasks such as state transfer and general unitary synthesis does not yet exist. In many cases, optimization fails when the variance of the cost-function gradient (calculated over the ensemble of randomly initialized circuits in the batch) becomes exponentially small, and the circuit optimization becomes stuck due to so-called "barren plateaus."

Reference [529] has made progress in analyzing the barren-plateau problem for numerical optimization of CV variational quantum circuits, demonstrating that the variance of the cost-function gradient decays as $E^{-\nu M}$, and is thus exponential in the number of modes M but polynomial in the (per-mode) circuit energy E (i.e., the mean boson number). The exponent is $\nu = 1$ for shallow circuits and $\nu = 2$ for deep circuits. The authors have proved these results for state preparation of general Gaussian states and number states. Separately, Ref. [530] has studied the barren-plateau problem in the more general setting of quantum channels that include midcircuit measurements and feedforward. One possible approach to the barren-plateau problem is to replace the global fidelity cost function with local cost functions [531,532]. Therefore, further study of the best local cost functions for hybrid CV-DV circuit synthesis is a promising direction for future study. In this section, we have introduced strategies that we have developed to aid compilation, and in the next section, we introduce applications of these strategies in quantum algorithms.

VII. ALGORITHMS AND APPLICATIONS

Major quantum algorithms including the Deutsch-Josza [533], Grover's search [534], and Shor's quantum factoring algorithm have been generalized from the DV setting to the CV setting [535], and experimental implementations of these are now emerging [536]. In addition, novel CV-based quantum algorithms such as boson sampling protocols [154,537,538] have been devised and experimentally demonstrated in the optical domain [156,158,159]. Apart from certain classically simulable regimes [161–163], these algorithms open the door to scenarios in which predicting experimental outcomes is difficult or intractable for present-day classical computers [539,540]. Small-scale demonstrations of boson sampling [205], Gaussian boson sampling (i.e., including squeezing), and efficient photon-number-resolving detection [147] have also been made in the microwave domain using cQED techniques.

The development of algorithms that uniquely take advantage of CV quantum systems has not yet flourished for several reasons. A common drawback of many CV algorithms is the lack of strict orthogonality in the computational basis (e.g., oscillator position states are not infinitely squeezed). As a result, only quasiorthogonal states are available for use in computations, which complicates the analysis of the algorithm. Another issue is that there is no fault-tolerant threshold in CV systems unless a finite-dimensional logical qubit or qudit subspace is encoded into the infinite-dimensional CV system. This limits our ability to do fault-tolerant operations on oscillators (as full oscillators rather than qudits). Moreover, from an operational point of view, much less is known about how to control and implement quantum algorithms in CV systems. For these reasons, it is unclear how to implement such CV algorithms and evaluate their efficiency relative to their DV counterparts.

The systematic establishment of CV-DV ISAs (Secs. II C and IV E) allows for a more elaborate discussion of CV and hybrid CV-DV algorithms, addressing many of the challenges previously faced by algorithm developers. In this section, we illustrate potential applications of quantum

CV-DV ISAs. We begin by presenting new algorithms for arbitrary CV-to-DV and DV-to-CV state transfer, alongside a review of an existing protocol (Sec. VII A), and then introduce a new implementation of the quantum Fourier transform based on state transfer and oscillator dynamics (Sec. VII B). Section VII C provides pedagogical examples of the application of bosonic ISAs to the simulation of time-evolution dynamics of Hamiltonians containing bosonic degrees of freedom that are otherwise costly to simulate with standard methods using qubit-based fault-tolerant quantum computers. Finally, we briefly review quantum random walks, pointing out the role of the phase-space ISA in the implementation of such walks (Sec. VII D).

A. Arbitrary qubit-oscillator state transfer

We first present how to transfer a multi-qubit DV state to a CV state of a single oscillator, leaving the qubits and the oscillator disentangled at the end. We present two different methods for such DV-to-CV state transfer, using single-variable hybrid QSP (Sec. VII A 1) and non-Abelian QSP (Sec. VII A 2). In these protocols, each qubit basis state is directly mapped to a coherent state located at the position corresponding to the integer representation of the qubit basis state. As these protocols make use of both qubits and oscillators, they serve as useful subroutines for each of the qubit-centric, oscillator-centric, and hybrid CV-DV AMMs. The presentation here is kept brief; for more details on these state-transfer protocols, see their full presentation in Ref. [541].

Formally, in the state-transfer problem, we begin with an n-qubit state $|\psi\rangle_Q = \sum_{\mathbf{x}} c_{\mathbf{x}}|\mathbf{x}\rangle_Q$, where $\mathbf{x} = (x_1, x_2, \ldots, x_n)$ is a Boolean-valued vector and $|\mathbf{x}\rangle_Q = |x_1\rangle|x_2\rangle\cdots|x_n\rangle$ is the corresponding qubit state. Using the binary representation of integers, let us denote the integer corresponding to $\mathbf{x}$ as $x := \sum_{j=1}^{n} x_j \cdot 2^{n-j}$. We wish to transfer $|\psi\rangle_Q$ to a corresponding oscillator state $|\psi\rangle_B = \sum_{\mathbf{x}} c_{\mathbf{x}}|x, \Delta\rangle_B$, where the basis states $|x, \Delta\rangle_B$ now exist in continuous space and depend on x and a spacing parameter Δ. One can think of the state $|x, \Delta\rangle_B$ as being localized around the position $q = x\Delta$, such that adjacent basis states (i.e., $|x \pm 1, \Delta\rangle_B$) are separated by Δ; the exact form of these basis states depends on the specific state-transfer protocol, as discussed below. Ultimately, we desire an n-qubit DV-to-CV *state-transfer unitary* $U_{\text{st}}^{(n)}(\Delta)$, which obeys

$$U_{\text{st}}^{(n)}(\Delta)|\psi\rangle_Q|0, \Delta\rangle_B = |\mathbf{0}\rangle_Q|\psi\rangle_B, \tag{410}$$

where $\mathbf{0}$ is a vector of all zeros and $|0, \Delta\rangle_B$ is the initial state of the oscillator (e.g., the vacuum state or a squeezed vacuum state).

1. State-transfer method 1: Single-variable QSP

Our first DV-to-CV state-transfer protocol employs quantum signal processing (QSP). This approach works by first applying a series of conditional displacements to the initial state, and then disentangling this intermediate state by applying a series of strategically chosen QSP sequences. We depict the circuit that implements this protocol in Fig. 32.

For this protocol, we would ideally want to take our oscillator basis states to be position eigenstates $|q\rangle_B$, with eigenvalues $q = x\Delta$ that are integer multiples of the spacing parameter. However, we cannot perform infinite squeezing to prepare a perfect position space eigenstate. Instead, we take our basis states to be Gaussians of width σ centered around the value $x\Delta$:

$$|x, \Delta\rangle_B = \frac{1}{\sqrt{\sigma}(2\pi)^{1/4}} \int dq\, e^{-(q-x\Delta)^2/4\sigma^2}|q\rangle_B, \tag{411}$$

which reduces to a position eigenstate as $\sigma \to 0$. In addition, we assume that the spacing and squeezing parameters are such that these basis states are nearly orthonormal:

$${}_B\langle y, \Delta|x, \Delta\rangle_B = e^{-\frac{\Delta^2}{\sigma^2}\frac{(x-y)^2}{8}}, \tag{412}$$

which approaches the Kronecker $\delta_{x,y}$ as $\sigma/\Delta \to 0$.

a. Protocol. In more detail, the state-transfer protocol works by first preparing the qubits in the initial state $|\psi\rangle_Q$ and the oscillator in the state $|0, \Delta\rangle_B$. Then, in the first stage, we apply the following series of controlled displacements to the oscillator:

$$\prod_{j=1}^{n} D_{c_j}(0, \Delta 2^{n-j}), \tag{413}$$

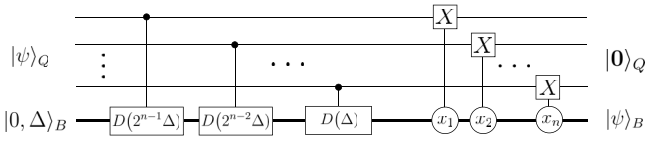


FIG. 32. A circuit that implements DV-to-CV state-transfer with single-variable QSP. The initial n-qubit state is $|\psi\rangle_Q$, and the initial oscillator state is a squeezed vacuum as per Eq. (411). The first stage applies a series of controlled displacements $D(2^{n-j}\Delta)$. The second stage disentangles the state by flipping the j th qubit conditioned on the j th bit of the oscillator position. We illustrate the corresponding operations as X gates conditioned on the bits x_j. These conditional bit flips are implemented using a QSP approximation to a square-wave filter in position space, which is depicted in Fig. 33 and implemented through the procedure described in Eq. (416).

where $D_{c_j}(0, \Delta 2^{n-j})$ is a conditional displacement as in Eq. (224), controlled by the j th qubit. Upon application to the initial state, this produces

$$|\psi\rangle_Q|0,\Delta\rangle_B = \sum_{\mathbf{x}} c_{\mathbf{x}}|\mathbf{x}\rangle_Q|0,\Delta\rangle_B \mapsto \sum_{\mathbf{x}} c_{\mathbf{x}}|\mathbf{x}\rangle_Q|x,\Delta\rangle_B. \tag{414}$$

If we can disentangle the qubits and oscillator as $|\mathbf{x}\rangle_Q|x,\Delta\rangle_B \mapsto |0\rangle_Q|x,\Delta\rangle_B$, then, up to the small errors created by nonorthogonality (see Eq. (412)), we perform the ideal transfer from the qubits to the oscillator.

We perform this disentangling procedure with single-variable QSP. The idea behind this is to devise a series of operations that flips the j th qubit conditioned on the j th bit x_j. This behavior requires that we determine the bit x_j from the binary representation $|x,\Delta\rangle_B$ of the bosonic mode. Accordingly, observe that the bit x_j of the position $q = x\Delta = \Delta\sum_{j=1}^{n} 2^{n-j}x_j$ can be read out by a "square-wave" function:

$$S_j(\hat{x}) := \Theta\left(\cos\left[\frac{\pi}{2^{n-j}}\left(\frac{\hat{x}}{\Delta} - 2^{n-j-1} + \frac{1}{2}\right)\right]\right) = 1 - x_j, \tag{415}$$

where $\Theta(\cdot)$ is the Heaviside step function, and $\hat{x}$ denotes the position operator. For visual intuition, we depict this function in Fig. 33, which illustrates how such a square-wave function can output the bits x_j.

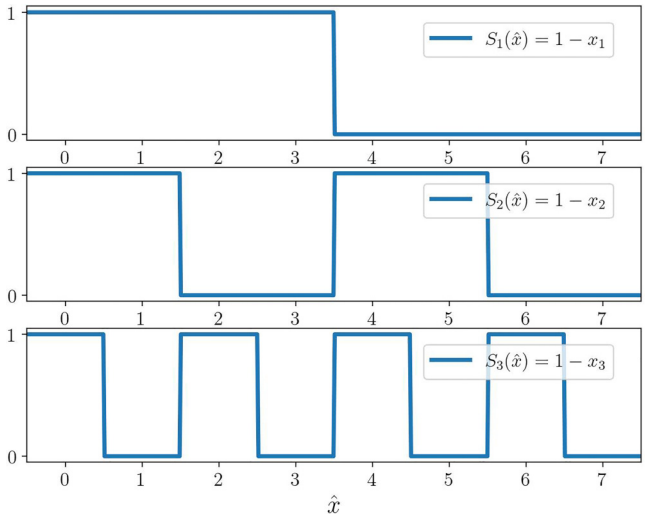


FIG. 33. An illustration of the square-wave functions $S_j(x\Delta)$ from Eq. (415) for $j = 1,2,3$, with $n = 3$ and $\Delta = 1$. Observe how at integer values x, these square waves are equal to $1 - x_j$, which enables one to read out the bits $\{x_j\}$.

We may use this observation to our advantage by constructing the unitary

$$W_j(\hat{x}) = \begin{pmatrix} S_j(\hat{x}) & \sqrt{1 - S_j(\hat{x})^2} \\ -\sqrt{1 - S_j(\hat{x})^2} & S_j(\hat{x}) \end{pmatrix} = \begin{cases} I, & x_j = 0, \\ iY, & x_j = 1, \end{cases} \tag{416}$$

where here, $Y = \sigma_y$. This operation correctly flips the j th qubit conditioned on x_j: $W_j(\hat{x})|x_j\rangle|x,\Delta\rangle_B = |0\rangle|x,\Delta\rangle_B$ [542]. Therefore, the sequence $\prod_{j=1}^{n} W_j(\hat{x})$ correctly disentangles all n qubits.

Our strategy is therefore to approximate each $W_j(\hat{x})$ with a QSP sequence.

To achieve this, we may apply QSP to the rotation $e^{iX\cdot(\pi/2^{n-j})((\hat{x}/\Delta)-2^{n-j-1}+\frac{1}{2})}$, which block-encodes $\hat{R}_j = \cos\left[\pi/2^{n-j}\left((\hat{x}/\Delta) - 2^{n-j-1} + \frac{1}{2}\right)\right]$ in its $|0\rangle\langle 0|$ component. We appropriately select QSP phases to construct a real polynomial $P(\hat{x})$ that approximates the Heaviside step function. Upon conjugating the corresponding QSP unitary (see Eq. (G10)) by a phase gate S, we obtain

$$\begin{pmatrix} P(\hat{R}_j) & \sqrt{1 - P(\hat{R}_j)^2} \\ -\sqrt{1 - P(\hat{R}_j)^2} & P(\hat{R}_j) \end{pmatrix} =: \tilde{W}_j(\hat{x}). \tag{417}$$

This approximates $\tilde{W}_j(\hat{x}) \approx W_j(\hat{x})$ in Eq. (416) because $P(\hat{R}_j) \approx S_j(\hat{x})$, where the accuracy in this approximation is dictated by how well a polynomial can approximate the Heaviside step function. It is well established that a QSP polynomial can approximate the step function to within some error ϵ, except within a region of width w centered around the discontinuity, and that the corresponding polynomial has degree $\mathcal{O}(1/w\log(1/\epsilon))$ [276,512].

Using this construction, the second stage of the protocol applies the series of QSP sequences $\prod_{j=1}^{n}\tilde{W}_j(\hat{x})$. This disentangles the intermediate state of Eq. (414) to (approximately) produce the desired final state $\sum_{\mathbf{x}} c_{\mathbf{x}}|\mathbf{0}\rangle_Q|x,\Delta\rangle_B = |\mathbf{0}\rangle_Q|\psi\rangle_B$.

b. Performance. We illustrate the entire circuit that implements the state-transfer protocol in Fig. 32, showcasing its decomposition into a stage of controlled displacements followed by a series of QSP sequences. Let us now analyze its gate complexity and fidelity.

The first stage requires $\mathcal{O}(n)$ displacement gates of size $\Delta 2^{n-j}$. This translates to a total time complexity that scales as $T_{\mathrm{CD}} = \sum_{j=1}^{n}\Delta 2^{n-j} = \mathcal{O}(\Delta 2^n)$. This implies an overall time complexity that scales linearly with 2^n when we implement the controlled displacements with a fixed coupling between the qubits and oscillator. However, this need

not be exponential in general as $\Delta \in \mathcal{O}(2^{-n})$ is also a possibility if sufficient squeezing is available on our quantum device.

In the second stage (the QSP sequences), we take the polynomial implemented by the jth QSP sequence to be an approximation to the step function that suffers error at most ϵ, outside of a region of width w_j centered about the discontinuity. Each such QSP sequence requires $N_{\mathrm{QSP},j} = \mathcal{O}\left(1/w_j \log(1/\epsilon)\right)$ gates on the qubit register [49,274]. To ensure that the jth QSP sequence can discern the correct bit x_j when acting on a state located at position $q = x\Delta$, we require that the width of the approximate step function needed in the jth bit in our encoding of the binary number scales as $w_j \leq \mathcal{O}(1/2^{n-j})$. Therefore, adding together the gate complexities of all the QSP sequences, we obtain a total gate complexity of $\sum_j N_{\mathrm{QSP},j} = \mathcal{O}\left(\sum_j 1/w_j \log(1/\epsilon)\right) = \mathcal{O}(2^n \log(1/\epsilon))$.

Next, consider the time complexity of this algorithm. Each step of a QSP sequence at fixed j requires application of the block-encoding $\hat{R}_j$ and a qubit rotation, which takes time $\mathcal{O}(1)$. Therefore, as a QSP sequence consists of $N_{\mathrm{QSP},j}$ of these steps, its time complexity is $T_{\mathrm{QSP},j} = \mathcal{O}(N_{\mathrm{QSP},j})$. Accounting for all n QSP sequences, we obtain a total time complexity of the QSP stage: $T_{\mathrm{QSP}} = \sum_j T_{\mathrm{QSP},j} = \mathcal{O}(2^n \log(1/\epsilon))$. The time complexity of the overall algorithm is then

$$T_{\mathrm{Tot}} = T_{\mathrm{CD}} + T_{\mathrm{QSP}} = \mathcal{O}\left(2^n \left(\Delta + \log(1/\epsilon)\right)\right). \quad (418)$$

If Δ is at most a constant, this scales exponentially with the number of bits that we wish to encode in the oscillator. This is perhaps unsurprising as we are not using a multi-modal binary encoding, but instead are storing the information in the position of a single oscillator. In principle, one could use binary encoding to more efficiently encode large numbers in multiple oscillator modes, but for simplicity, we focus on the case in which all the information is encoded in a single mode.

Lastly, the values of Δ and ϵ also influence the fidelity of the state-transfer. The protocol suffers from two possible types of errors: first, the bosonic basis states are not exact position eigenstates but instead Gaussians of finite width σ; and second, the QSP polynomial approximation is only accurate to within ϵ, and fails within the region of width w_j about the discontinuity of the step function. A careful analysis of these error modes, presented in Ref. [541], indicates that the total fidelity scales as

$$1 - \mathcal{O}(n\epsilon) - \mathcal{O}(e^{-\mathcal{O}(\Delta^2/\sigma^2)}). \quad (419)$$

Note that this depends on the ratio σ/Δ, which represents the relative spacing between the basis states, $|x, \Delta\rangle_B$. Therefore, one can improve fidelity by either squeezing the initial state to decrease σ, or selecting a larger separation Δ. Numerical examples of the DV-to-CV state-transfer protocol for various three-qubit states are given in Fig. 34.

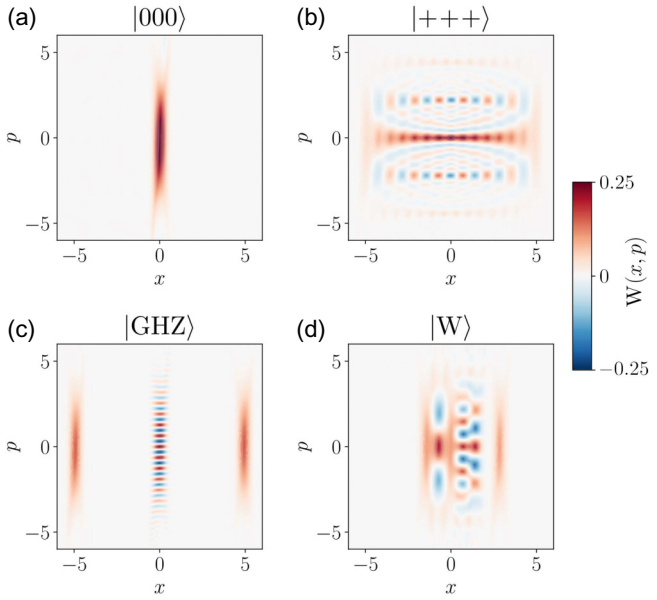


FIG. 34. DV-to-CV state-transfer for various three-qubit states, including $|\mathrm{GHZ}\rangle = (|000\rangle + |111\rangle)/\sqrt{2}$ and $|\mathrm{W}\rangle = (|001\rangle + |010\rangle + |100\rangle)/\sqrt{3}$. In all cases, we have applied the protocol in Sec. VII A 1, using a single-variable QSP sequence of depth 60 with $w = 0.2$, $\Delta = 1$, and $\sigma = \pi/16$. The fidelities have been computed against the ideal state-transfer, yielding (a) $F = 0.970$, (b) $F = 0.980$, (c) $F = 0.972$, and (d) $F = 0.982$. All simulations have been carried out in BOSONIC QISKIT [53].

2. State-transfer method 2: Non-Abelian QSP

Recently, Refs. [479,543] have proposed an approach to quasideterministically transfer an oscillator state to an n-qubit state by repeatedly performing two controlled displacement operations between the oscillator and each qubit. This defines a CV-to-DV state-transfer protocol, which in reverse furnishes a DV-to-CV state-transfer protocol. For completeness, we depict the circuit that implements the DV-to-CV state-transfer protocol in Fig. 35.

Below, we first review the protocol of Ref. [543] and provide bounds on its performance. In addition, we show how this protocol can be viewed as a special case of non-Abelian QSP as discussed in Sec. VI D 3.

a. Protocol. Let us denote the CV-to-DV state-transfer operation presented in Ref. [543] by $U_{\mathrm{st}}^{(n)}(2\lambda)^\dagger$. This uses a spacing parameter 2λ (analogous to Δ above) and transfers an oscillator state to a state on n qubits. Explicitly, $U_{\mathrm{st}}^{(n)}(2\lambda)^\dagger$ is the unitary operation defined by

$$U_{\mathrm{st}}^{(n)}(2\lambda)^\dagger = \prod_{j=n}^{1} W_j V_j = W_n V_n \cdots W_1 V_1, \quad (420)$$

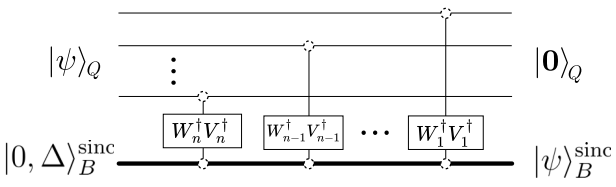


FIG. 35. The circuit that implements DV-to-CV state-transfer, adapted from Ref. [543]; crucially, the order of W_n and V_n here are flipped relative to Fig. 1 of Ref. [543], which we believe to be a typographical error. The initial state of the qubits is $|\psi\rangle_Q$, and the initial oscillator state is a sinc state $|0,\Delta\rangle_B = 1/\sqrt{\Delta}\int dq\ \text{sinc}(\pi q/\Delta)|q\rangle_B$ with $\Delta = 2\lambda$. Then, one applies a series of operations $W_j^\dagger V_j^\dagger$ between the oscillator and the j th qubit [see Eq. (421)], which maps the initial state of the qubits to an equivalent oscillator state $|\psi\rangle_B$ encoded in a basis of displaced sinc states, as per Eq. (425).

where

$$V_j = e^{i\frac{\pi}{2^{j+1}\lambda}\hat{x}\hat{\sigma}_y^{(j)}}, \quad W_j = \begin{cases} e^{i\lambda 2^{j-1}\hat{p}\hat{\sigma}_x^{(j)}}, & j < n, \\ e^{-i\lambda 2^{j-1}\hat{p}\hat{\sigma}_x^{(j)}}, & j = n, \end{cases} \tag{421}$$

are momentum boosts and displacements of the oscillator conditioned on the j th qubit. Note that we have defined this operation in terms of the Hermitian conjugate $U_{\text{st}}^{(n)}(2\lambda)^\dagger$, such that its inverse $U_{\text{st}}^{(n)}(2\lambda)$ provides a DV-to-CV state-transfer protocol, as per our conventions.

The input state to this protocol is $|\mathbf{0}\rangle_Q|\psi\rangle_B$, where $|\psi\rangle_B = \int dq\ \psi(q)dq|q\rangle_B$ is a state on the oscillator to be transferred to the n qubits. Intuitively, the arrangement of momentum boosts and displacements in Eq. (420) is carefully chosen to map the CV wave function $\psi(q)$, evaluated at a discrete set of positions $q_{\mathbf{s}}$, onto the amplitudes of a DV quantum state, thus performing state-transfer. As shown in Ref. [543], application of $U_{\text{st}}^{(n)}(2\lambda)^\dagger$ maps the initial state to

$$\begin{aligned} &U_{\text{st}}^{(n)}(2\lambda)^\dagger\,|\mathbf{0}\rangle_Q\,|\psi\rangle_B \\ &\approx \sum_{\mathbf{s}\in\{-1,+1\}^n} \sqrt{2\lambda}\psi(q_{\mathbf{s}})|\phi_{\mathbf{s}}\rangle_Q \otimes \frac{1}{\sqrt{2\lambda}}\int dq\ \text{sinc}\left(\frac{\pi q}{2\lambda}\right)|q\rangle_B, \end{aligned} \tag{422}$$

where the basis states here are

$$\begin{aligned} |\phi_{\mathbf{s}}\rangle &= (-1)^{\gamma_{\mathbf{s}}}\cdot\bigotimes_{j=1}^{n}\left[Z^{(1-s_j)/2}|+\rangle\right], \\ \gamma_{\mathbf{s}} &= \sum_{j=1}^{n-2}\frac{1}{2}(s_j + s_{j+1}) + \frac{1}{2}(s_{n-1} - s_n). \end{aligned} \tag{423}$$

For instance, if $\mathbf{s} = (1,-1,1,-1)$, then $\gamma_{\mathbf{s}} = 1$, and $|\phi_{\mathbf{s}}\rangle = (-1)\cdot|+\rangle|-\rangle|+\rangle|-\rangle$. Finally, the value $q_{\mathbf{s}}$ in Eq. (422) is

$$q_{\mathbf{s}} = \sum_{j=1}^{n-1} s_j w_j - s_n w_n, \tag{424}$$

where $w_j = \lambda 2^{j-1}$. This quantity takes 2^n discrete values in the range $[-\lambda(2^n-1), \lambda(2^n-1)]$, with each possible value equally spaced by 2λ.

Two approximations are used in Ref. [543] to obtain Eq. (422). It is assumed that the support of $\psi(q)$ is limited to $|q| \le \lambda(2^n-1)$, and that $\psi(q)$ is slowly varying as $|d\psi/dq| \ll 1/\lambda$. This approximately decouples the oscillator from the qubits, thus transferring the initial CV state $\psi(q)$ to a corresponding state of the qubits $\sum_{\mathbf{s}}\sqrt{2\lambda}\psi(q_{\mathbf{s}})|\phi_{\mathbf{s}}\rangle_Q$.

Moreover, this protocol can also be run in reverse to furnish a DV-to-CV state-transfer protocol. In this direction, one first prepares the qubits in the state $\sum_{\mathbf{s}} c_{\mathbf{s}}|\phi_{\mathbf{s}}\rangle$, and the oscillator in the state $|0,2\lambda\rangle_B = 1/\sqrt{2\lambda}\int dq\ \text{sinc}(\pi q/2\lambda)|q\rangle$ [544]. Then, executing $U_{\text{st}}^{(n)}(2\lambda) = \prod_{j=n}^{1} W_j^\dagger V_j^\dagger$ (approximately) outputs the state

$$\begin{aligned} &|\mathbf{0}\rangle_Q \otimes \sum_{\mathbf{s}} c_{\mathbf{s}}\frac{1}{\sqrt{2\lambda}}\int dq\ \text{sinc}\left(\frac{\pi(q-q_{\mathbf{s}})}{2\lambda}\right)|q\rangle_B \\ &\quad = |\mathbf{0}\rangle_Q|\psi\rangle_B. \end{aligned} \tag{425}$$

This has effectively transferred the state of the qubits to an oscillator state, now encoded in the basis of displaced "sinc states" $|q_{\mathbf{s}}/2\lambda, 2\lambda\rangle_B := 1/\sqrt{2\lambda}\int dq\ \text{sinc}\left(\pi(q-q_{\mathbf{s}})/2\lambda\right)|q\rangle_B$. A sinc state represents a state in continuous space that is peaked around $q = q_s$, with peaks of adjacent sinc states separated by 2λ. Moreover, the sinc states are orthonormal: ${}_B\langle q_{\mathbf{s}}/2\lambda, 2\lambda|q_{\mathbf{s}'}/2\lambda, 2\lambda\rangle_B = \delta_{\mathbf{s}\mathbf{s}'}$.

b. Performance. We illustrate the circuit of this DV-to-CV state-transfer protocol in Fig. 35. This showcases the series of $V_j^\dagger W_j^\dagger$ operations acting between the oscillator and the j th qubit. Let us now study the gate complexity and performance of this state-transfer protocol.

Consider first the gate complexity of this protocol. According to Eq. (420), this protocol requires $\mathcal{O}(n)$ gates. However, collectively these gates require a total displacement $\sum_{j=1}^{n}\mathcal{O}(2\lambda 2^j) = \mathcal{O}(\lambda 2^n)$. As in the previous state-transfer algorithm, this implies an overall time complexity of $\mathcal{O}(\lambda 2^n)$ when we implement the displacements with a fixed coupling between the qubits and the oscillator. But, with a sufficiently tunable and strong coupling between the qubits and the oscillator, one can adjust the coupling strength to reduce this exponential overhead.

Next, let us consider the fidelity of this protocol. Reference [543] has presented numerical results on the fidelity

achieved in the CV-to-DV direction. For example, in transferring the Fock state $|3\rangle$ onto n qubits, the protocol achieves an infidelity of ≈ 0.2 for $n = 4$ and $\approx 9 \cdot 10^{-4}$ for $n = 10$, indicating how the performance improves with increasing n. Achieving this performance, however, requires that λ be carefully tuned for each n to maximize the fidelity. While no explicit bounds on the fidelity are provided in Ref. [543] to guide this tuning, in Ref. [541] we derive bounds on the fidelity of both the DV-to-CV and CV-to-DV directions.

First, consider the fidelity of the CV-to-DV direction, in which case the output state is approximately that of Eq. (422). The approximations used in reaching this expression require that the $\psi(q)$ have support limited to $|q| \leq \lambda(2^n - 1)$, and be slowly varying such that $|d\psi/dq| \ll 1/\lambda$. A careful analysis of this protocol, presented in Ref. [541], indicates that the infidelity between these two states, and equivalently, the infidelity of the CV-to-DV direction, is

$$\mathcal{O}\left(\int_{|q|\geq\lambda(2^n-1)} dq|\psi(q)|^2 + \lambda\int_{-\lambda(2^n-1)}^{\lambda(2^n-1)} dq\left|\frac{d}{dq}|\psi(q)|^2\right|\right). \tag{426}$$

Notably, the two contributions to the infidelity stem from the support of $\psi(q)$ outside of $|q| \leq \lambda(2^n - 1)$ and the derivative of $\psi(q)$, arising precisely from the two approximations used in Ref. [543] to obtain Eq. (422).

Moreover, the action of the DV-to-CV state-transfer is defined by Eq. (425). By a similar analysis, also presented in Ref. [541], we find that the infidelity of the DV-to-CV state-transfer protocol is

$$\mathcal{O}\left(\int_{|q|\geq\lambda(2^n-1)} dq|\psi(q)|^2\right), \tag{427}$$

where now $\psi(q) = 1/\sqrt{2\lambda}\sum_{\mathbf{s}} c_{\mathbf{s}}\mathrm{sinc}\,(\pi(q - q_{\mathbf{s}})/2\lambda)$ is the wave function of the state of the qubits transferred to the basis of sinc states. Evidently, in this direction, the fidelity is primarily impeded by the support of $\psi(q)$ outside $|q| < \lambda(2^n - 1)$. The additional term of Eq. (426) that depends on $(d/dq)\psi(q)$ is absent because the approximation that produces this contribution is naturally satisfied in the DV-to-CV direction (see Ref. [541] for details).

c. Recontextualization of DV-to-CV state-transfer as non-Abelian QSP. Let us take a closer look at this DV-to-CV state-transfer protocol, and illustrate how it can be reexpressed in the language of non-Abelian QSP. If we rewrite an individual term $V_j^\dagger W_j^\dagger$ using the conventions of non-Abelian QSP outlined in Sec. VI D 3, the following identity can be verified:

$$\begin{aligned} &e^{-i\frac{\pi}{4}\sigma_y^{(j)}}(V_j^\dagger W_j^\dagger)e^{i\frac{\pi}{4}\sigma_y^{(j)}} \\ &\quad= e^{-i\frac{\pi}{\lambda 2^{j+1}}\hat{x}\hat{\sigma}_y^{(j)}} e^{\mp i\lambda 2^{j-1}\hat{p}\hat{\sigma}_z^{(j)}} \qquad (428) \\ &\quad= e^{i\frac{\pi}{4}\sigma_x^{(j)}} e^{-i\frac{\pi}{\lambda 2^{j+1}}\hat{x}\hat{\sigma}_z^{(j)}} e^{-i\frac{\pi}{4}\sigma_x^{(j)}} e^{\mp i\lambda 2^{j-1}\hat{p}\hat{\sigma}_z^{(j)}}, \qquad (429) \end{aligned}$$

where the "$-$" sign is taken for $j < n$, and the "$+$" sign for $j = n$.

By comparing this expression to the non-Abelian QSP sequence of Eq. (398), it is readily identified that the product $V_j^\dagger W_j^\dagger$, upon conjugation by $e^{-i(\pi/4)\sigma_y^{(j)}}$, corresponds to a degree-1 non-Abelian QSP sequence with momentum and position boosts:

$$W_z^{(k)}(w) = e^{-i\frac{\pi}{\lambda 2^{j+1}}\hat{x}\cdot\sigma_z}, \quad W_z^{(\lambda)}(v) = e^{\mp i\lambda 2^{j-1}\hat{p}\cdot\sigma_z}, \tag{430}$$

for the j th qubit, and with QSP phases

$$\phi_0 = \frac{\pi}{4}, \quad \phi_1^{(k)} = -\frac{\pi}{4}, \quad \phi_1^{(\lambda)} = 0 \tag{431}$$

for all j qubits. In this incarnation, the DV-to-CV state-transfer protocol of this section may be interpreted as a product of n degree-1 non-Abelian QSP sequences, where each sequence acts between the oscillator and the j th qubit. Accordingly, this protocol could admit a generalization by using a higher-degree non-Abelian QSP sequence for each qubit during state-transfer.

B. Quantum Fourier transform

The quantum Fourier transform (QFT) is an important quantum subroutine that is ubiquitous in many quantum algorithms, such as Shor's algorithm [545], phase estimation [1], and quantum gradient estimation [546]. It is defined on an n-qubit state $|\psi\rangle_Q = \sum_{\mathbf{x}} c_{\mathbf{x}}|\mathbf{x}\rangle_Q$ as the unitary transformation

$$U_{\mathrm{QFT}}|\psi\rangle_Q = \sum_{\mathbf{x}}\left[\sum_{\mathbf{y}}\frac{1}{\sqrt{2^n}}c_{\mathbf{y}}e^{2\pi i xy/2^n}\right]|\mathbf{x}\rangle_Q, \tag{432}$$

which effectively implements a discrete Fourier transform on the coefficients of the initial state.

Here, we provide an implementation of the QFT on the qubit register using bosonic operations, incorporating the state-transfer protocols presented in Sec. VII A. In particular, we assume access to a DV-to-CV state-transfer unitary $U_{\mathrm{st}}^{(n)}(\Delta)$ and its inverse, as well as traditional bosonic gates such as free evolution and displacement. Because these state-transfer protocols incorporate qubits and oscillators, this implementation of the QFT is a useful subroutine for each of the three AMMs: qubit-centric, oscillator-centric, and hybrid CV-DV. For more specific details on the implementation of the QFT protocol, see Ref. [541].

To preface our construction, recall that in phase space the free evolution of an oscillator naturally swaps position and momentum, thus enacting a continuous Fourier transform on the underlying wave function. Using this intuition, we show that by mapping an initial DV state to an oscillator, letting it undergo free evolution, and finally transferring the state back to qubits, the QFT of the DV state can be extracted.

1. Continuous-discrete Fourier transform correspondence

A vital component of our construction of the QFT is a correspondence between the continuous Fourier transform and the discrete Fourier transform. Consider a discrete signal c_x for $x \in [0, \ldots, N-1]$, that is made periodic such that $c_x = c_{x \bmod N}$. Let this discrete signal be encoded in a continuous function $f(q)$ as

$$f(q) = \sum_{x\in\mathbb{Z}} c_x g(q - x\Delta), \tag{433}$$

where $g(q)$ is a basis function localized about $q = 0$ and Δ is the spacing between basis states. The continuous Fourier transform of this function evaluates to

$$\tilde{f}(p) = \sum_{x\in\mathbb{Z}} c_x \tilde{g}(p) e^{ipx\Delta} = \tilde{g}(p) \sum_{k\in\mathbb{Z}} e^{ipNk\Delta} \sum_{y=0}^{N-1} c_y e^{ipy\Delta}, \tag{434}$$

where $\tilde{g}(p)$ is the Fourier transform of $g(q)$, and we have split the index x into $x = Nk + y$ for $k \in \mathbb{Z}$ and $y \in [0, 1, \ldots, N-1]$. Noting that the quantity $\sum_{k\in\mathbb{Z}} e^{ipNk\Delta}$ equates to a Dirac comb $\sum_{l\in\mathbb{Z}} \delta((pN\Delta/2\pi) - l)$, this expression simplifies to

$$\sum_{l\in\mathbb{Z}} \frac{2\pi}{\Delta} \tilde{g}\left(\frac{2\pi}{\Delta}\frac{l}{N}\right) \delta\left(p - \frac{2\pi}{\Delta}\frac{l}{N}\right) \cdot \tilde{c}_l, \tag{435}$$

where $\tilde{c}_l = \sum_{y=0}^{N-1} c_y e^{i2\pi yl/N}$ is the discrete Fourier transform of c_y. Therefore, by taking the continuous Fourier transform of a discrete signal that is encoded periodically in some basis function, we obtain a sum over the discrete Fourier transform of the signal, with coefficients proportional to the Fourier transform of the basis function.

This correspondence can be used to perform the quantum Fourier transform by letting c_x be the coefficients of an initial state on qubits. Then, $f(q)$ represents the wave function of an oscillator after transferring the state of the qubits to the oscillator and encoding in the basis $g(q)$. By enacting a continuous Fourier transform on the oscillator (i.e., free evolution), the new wave function given by Eq. (435) plucks out the states $|q = 2\pi l/N\Delta\rangle$ with coefficients proportional to the discrete Fourier transform of c_x. Because the discrete Fourier transform of these coefficients is equivalent to the coefficients of QFT, we find that appropriately transferring this state back to qubits produces the QFT of the initial state. This is of course dependent on choosing a suitable basis function $g(q)$, the Fourier transform of which behaves favorably.

2. QFT method 1: Single-variable QSP

We first present a construction of the QFT using the single-variable QSP state-transfer protocol of Sec. VII A 1 (i.e., the circuit presented in Fig. 32). For illustrative purposes, we present the corresponding QFT circuit in Fig. 36.

a. Protocol. In our construction, we first append the initial state with an appropriate number of ancilla qubits in the $|+\rangle$ state so as to make the coefficients of the initial state periodic over a range $O(2^a)$, so that we can leverage the above continuous-discrete Fourier transform correspondence to generate the QFT. Then, we transfer this state to an equivalent oscillator state with basis spacing Δ, and apply a displacement operation to make the state symmetric about $q = 0$ in position space. Subsequently, the Fourier gate is applied to the oscillator to enact a continuous Fourier transform, and then the state is transferred back to qubits using a reciprocal spacing $\pi/2^n\Delta$. Crucially, this reciprocal spacing is chosen precisely to produce the correct phases of the QFT (i.e., $2\pi xy/2^n$). For more details, we present a pedagogical walk-through of each stage of this protocol in Appendix H.

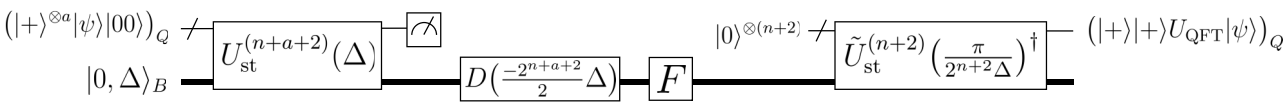


FIG. 36. The circuit used to implement the QFT of an n-qubit state $|\psi\rangle$ using the single-variable QSP state-transfer unitary of Sec. VII A 1, denoted $U_{\rm st}$. Here, $D\left((-2^{n+a+2}/2)\Delta\right)$ is a displacement operation, and F is the Fourier transform operation $F = U(\pi/2)$ from Box IV.2 (i.e., oscillator free evolution). The ancilla qubits $|+\rangle^{\otimes a}$ and $|00\rangle$ are appended to the initial state to increase the fidelity with the exact QFT, as discussed below. The result of this circuit is the state $|+\rangle|+\rangle U_{\rm QFT}|\psi\rangle$ in the qubit register, from which the QFT may be obtained.

b. Performance. The bulk of the gate count in this QFT protocol comes from the use of the single-variable QSP state-transfer of Sec. VII A 1. Copying over the gate count of that protocol, we have that the total gate count of the QFT here is $\mathcal{O}(2^{n+a}\log(1/\epsilon))$. Similarly, the time complexity is $\mathcal{O}(2^{n+a}(\Delta+\log(1/\epsilon)))$.

Next, the steps of this protocol can each suffer errors that decrease the fidelity with the target QFT. By aggregating together the infidelities suffered at each step, which are explicated in Appendix H, we find that the total fidelity of this protocol is

$$F = 1 - \mathcal{O}((n+a)\epsilon) - \mathcal{O}(\sigma/\Delta) - \mathcal{O}(1/2^a), \quad (436)$$

where ϵ is the error in the QSP polynomial approximation to a square-wave function. Understandably, the fidelity is maximized in the limit of small QSP error, large relative spacing between the oscillator basis states, and a large number of ancilla qubits.

3. QFT method 2: Non-Abelian QSP

Analogous to the presentation above, we can also implement the QFT using the non-Abelian QSP state-transfer protocol of Sec. VII A 2. We depict the circuit that implements this algorithm in Fig. 37.

a. Protocol. Our approach is to first append the initial state with ancilla qubits in the $|+\rangle$ state, such that we make the coefficients of the state periodic and can use the continuous-discrete Fourier transform correspondence to generate the QFT. Then, we perform a change of basis from the computational basis to the basis $\{|\phi_{\mathbf{s}}\rangle\}$ used in the non-Abelian QSP state-transfer protocol; this is achieved using an operation denoted $\mathcal{T}$ in Fig. 37, and can be shown to be implemented with simple local operations (see Appendix H). Subsequently, we transfer the state of the qubits to the oscillator, symmetrize the state about $q=0$ in position space, apply the Fourier gate, and finally transfer the oscillator state back to the qubits with an appropriately chosen reciprocal spacing. Upon transforming back to the computational basis, we obtain the QFT of the initial state. We provide a step-by-step analysis of this protocol in Appendix H.

b. Performance. Consider first the gate complexity of this QFT protocol, the bulk of which is due to the non-Abelian state-transfer protocol. According to the performance bounds in Sec. VII A 2, the gate count scales as $\mathcal{O}(n+a)$. However, these gates require a total displacement $\mathcal{O}(\lambda 2^{n+a})$, implying a time complexity of $\mathcal{O}(\lambda 2^{n+a})$ when displacements are implemented with a fixed coupling between the qubits and the oscillator, but this can be reduced with a sufficiently tunable and strong coupling.

Next, consider the fidelity of this protocol. Again, each step of this protocol can incur errors that decease the fidelity with the target QFT, and we analyze these error modes in Appendix H. Upon aggregating these error contributions, we find that the overall fidelity of this QFT protocol is $F = 1 - \mathcal{O}(1/2^a)$.

C. Quantum simulation

In this section, we provide a brief overview of techniques for quantum simulation using hybrid CV-DV hardware. For an in-depth study of the simulation of fermions, bosons, and gauge fields in CV-DV hardware, we refer the reader to Ref. [41].

Quantum simulation aims to solve classically intractable physical problems with controllable quantum devices and has been identified as a likely arena for the first demonstration of *practical* advantage over existing classical methods [547–551]. Exact classical simulation of quantum systems is often intractable due to the exponential growth of Hilbert space with system size. As a result, classical methods often resort to approximations that involve either truncation or finite sampling methods to remove the exponential computational overhead at the price of introducing approximation errors that may or may not be acceptable for a given application. Such approximations include the use of mean-field theory [552,553] and semiclassical approximations of quantum phase space [554] that do not capture particle-particle correlations. Other representative approximations that truncate the Hilbert space, such as many-body perturbation theory, coupled cluster theory [555,556], and tensor network-based methods [557], face difficulties with strongly correlated systems, as they are capable of only limited treatment of many-body correlations. Separately, sampling-based methods, such as

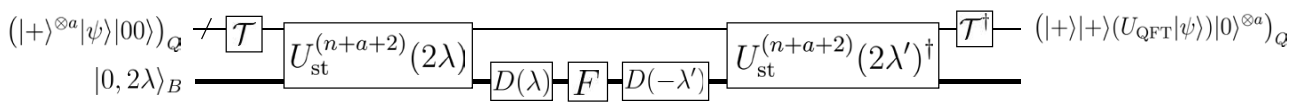


FIG. 37. The circuit used to implement the QFT of an n-qubit state $|\psi\rangle$ using the non-Abelian QSP state-transfer unitary of Sec. VII A 2, denoted $U_{\rm st}$. Here, $\mathcal{T}$ is a basis transformation from the computational basis to the basis $|\phi_s\rangle$, which is easily realized with simple qubit gates, as discussed in the text. The operation $D(\cdot)$ is a displacement operation, and F is the Fourier gate $F = U(\pi/2)$ (oscillator free evolution). Ancilla qubits $|+\rangle^{\otimes a}$ are prepended to the initial state to increase the fidelity with the exact QFT. The circuit outputs the state $|+\rangle|+\rangle U_{\rm QFT}|\psi\rangle|0\rangle^{\otimes a}$ in the qubit register, from which the QFT may be extracted.

quantum Monte Carlo, suffer from sign or phase problems for frustrated spin systems and interacting fermionic matter [558–560].

As originally noted by Richard Feynman [561], quantum hardware, unlike classical hardware, is naturally suited to the simulation of quantum physics [551,562]. To that end, there are two popular approaches: analog quantum simulation and digital quantum simulation. The former corresponds to the *native* realization of a Hamiltonian in a controlled quantum system, while the latter corresponds to the *synthesis* of a Hamiltonian using gate-based instructions on a universal quantum computer (or if not universal, at least a programmable quantum simulator).

While both approaches have proven promising in recent years [563,564], most experiments have thus far focused on "single-species" quantum systems. For example, demonstrations of digital qubit-based quantum simulation have typically concerned either spin [565,566] or fermionic [567] systems (though we note the recent proposal for hybrid fermion-qudit processors in Ref. [568]). Similarly, analog quantum simulators have successfully implemented interacting spin [569,570], all-bosonic [76, 571], and all-fermionic [572] systems. However, only a few experiments have demonstrated the capability to simulate bosons and fermions within a single system [573,574]. However, multi-species quantum Hamiltonians are among the most important and are crucial to the development of fundamental science across a range of energy scales, from the interactions between force-carrying bosons and fermionic elementary particles underlying the standard model to the nuclear vibrational-electron interactions underpinning quantum chemical phenomena such as photosynthesis [575]. Consequently, this motivates the need for flexible approaches to quantum simulation that enable hardware-efficient simulation of different species within a single platform.

We focus on the problem of digital quantum simulation here as it is most pertinent to the scope of this work, but also make a comparison to analog quantum simulation where applicable. Turning first to conventional qubit-based hardware, we note that it is possible to encode fermionic [194], bosonic [200,576], and spin-S [385] degrees of freedom using many qubits, in principle enabling multi-species digital quantum simulation. In practice, however, bosonic degrees of freedom and the corresponding bosonic operations are costly to implement in all-qubit hardware, as discussed in Sec. VIII C. For example, the implementation of a single displacement gate on a qubit-encoded-oscillator containing ten bosons requires on the order of 100 two-qubit gates using the Fock-binary encoding [200] (see Sec. VIII B). This inefficiency is particularly problematic in cases in which it becomes important to represent high mean boson numbers with large density fluctuations, regardless of the choice of encoding. Reference [577] provides an example from cold-atom physics of such a situation where the density fluctuations are large and non-Gaussian.

Hybrid oscillator-qubit platforms can simulate bosonic degrees of freedom natively while retaining the capabilities of qubit-based hardware. Such hybrid platforms show immense promise for the study of both single- and multi-species models containing spin, fermionic, and bosonic degrees of freedom. Furthermore, the universal ISAs presented in this work enable the implementation of arbitrary Hamiltonian terms needed to simulate strongly correlated matter and facilitate measurement of both local and non-local observables, the latter playing a crucial role in the identification of exotic phases of matter [578,579].

In the remainder of this section, we present a high-level survey of the opportunities provided by digital hybrid oscillator-qubit hardware for quantum simulation. For a more in-depth investigation into compilation techniques for the simulation of fermion-boson models and lattice gauge theories, we refer the reader to Ref. [41].

We begin by first describing high-level techniques that enable the study of dynamics (Sec. VII C 1), ground and excited states (Sec. VII C 2), and observables (Sec. VII C 3). Following this, we provide several explicit examples of physically relevant models with bosonic modes, spins, or some combination of the two. These include the Rabi model (Sec. VII C 4), the Bose-Hubbard model (Sec. VII C 5), large-spin models (Sec. VII C 6), a $\mathbb{Z}_2$ lattice gauge theory (Sec. VII C 7), and fermion-boson problems (Sec. VII C 8). In each case, we demonstrate how one can synthesize the Hamiltonian terms of interest for each case using the techniques outlined in Sec. VI. Many of the specific compilation ingredients discussed throughout are summarized in Table VI, which lists a selection of common bosonic Hamiltonian terms and strategies for realizing the unitaries that they generate with native instructions. We then conclude with comments on the potential for simulation of fermion-boson problems.

1. Dynamics

Time dynamics involves the study of a system evolving under the time-evolution operator $U(t) = e^{-iHt}$, where H is some model Hamiltonian of interest. For digital quantum simulation, the goal is therefore to synthesize and enact $U(t)$ through decomposition into a discrete gate set. Many strategies exist to achieve this, including Trotter-Suzuki and product formulas [339,504,580], linear combinations of unitaries [502,511] and multi-product formulas [581], quantum walks [582–584], qubitization [585], and various randomized approaches [341,586,587]. Hybrid oscillator-qubit hardware may facilitate many of these methods. For further discussion of Trotterization in the context of compilation, see Sec. VI D 1.

In the Trotter-Suzuki approach, the unitary time-evolution operator $U(t)$ corresponding to the Hamiltonian

$H = \sum_j H_j$ is written as

$$U(t) = e^{-i\sum_{j=1}^{M} H_j t} = \left(\prod_{j=1}^{M} e^{-iH_j t/r}\right)^r + O\left(\sum_{j,k}^{M} \|[H_j, H_k]\| t^2/r\right), \tag{437}$$

where r is a positive integer introduced to decrease the error in the approximation by increasing the number of time steps, and M is the total number of Hamiltonian terms. This decomposition reduces the problem to the implementation of $e^{-iH_j t/r}$ for each j, via either a single native gate or through compilation using a particular instruction set (see Sec. VI). For example, a chemical potential term $H_j = e^{-i\mu\hat{n}_k(t/r)}$ can be implemented via a phase-space rotation $R(\mu t/r)$ (see Box IV.2), while higher-order interactions such as $H_j = Z_i a_i^\dagger a_i a_k^\dagger a_k$ can be synthesized using BCH methods (see the examples portion of Sec. VI D). Higher-order variants of Eq. (437) exist, such as the second-order (or symmetric) Trotter formula that achieves an improved error scaling at the cost of a higher gate count [339].

2. State preparation

Initial state preparation is a key component for both static and dynamic simulations. For static problems, state preparation is used to construct a prior quantum distribution over eigenstates that can be refined through variational techniques [509] or quantum phase estimation [588]. In general, the question of whether or not quantum computers can efficiently prepare high-fidelity approximations to ground states of physically relevant systems remains an open question; however, if a generic quantum algorithm existed for preparing such states, then a consequence would be that BQP = QMA. This is not believed to be true, as BQP and QMA are the quantum analogs of P and NP for classical computation [1].

For calculating dynamics, quantum states need to be provided that correspond to the initial quantum state $|\psi(0)\rangle$ that we wish to evolve into $|\psi(t)\rangle = U(t)|\psi(0)\rangle$. It is likely, however, that calculating observables of thermal states relevant to material simulations is easier than for ground states [589–591].

Digital methods capable of determining low-energy states using only shallow-depth quantum circuits are of particular interest for simulation of static states in the NISQ era, with the VQE [509] and the QAOA [510] being two examples that are particularly well suited to near-term digital oscillator-qubit hardware [41,529,592]. Annealing, or adiabatic state preparation, approaches can also be used to solve such problems or as intuition for an ansatz for QAOA or VQE [510,593]. However, the cost of these methods scales with the energy gap between the ground state and the first excited state. This means that while adiabatic algorithms can be used to construct arbitrarily accurate approximations to the quantum ground state for gapped systems, the fact that the excitation gap typically shrinks exponentially with the system size causes such approaches to be inefficient in general.

In the fault-tolerant regime, we do not usually assume that fundamental limitations exist on gate depth and number of qubits or qumodes. This allows us to use alternative methods that have exact bounds on success probability and fidelity (or provable speedups) that can be implemented on hybrid oscillator-qubit hardware, with the instruction sets presented in this work enabling the implementation of controlled unitaries implementing $e^{-i\theta H}$ or various energy filter functions. These include quantum phase estimation (QPE) [588] and quantum signal processing (QSP) [49], discussed in Sec. VI D 3 b and applied to Hamiltonian simulation in hybrid qubit-oscillator systems in Ref. [41].

With some modifications, the above preparation methods can be used to prepare either ground or excited states. For example, in the context of variational methods, instead of minimizing $\langle \hat{H} \rangle$ to target the ground state, we can instead minimize $\langle (E - \hat{H})^2 \rangle$ with E the estimated energy of a target excited state. In the context of fault-tolerant algorithms, QPE for example can be used to prepare excited states as long as the initial trial state has non-negligible overlap with the target excited state of interest [594]. In the case of QSP or qubitization, a filter function centering at the excited state energy can be implemented to facilitate excited state preparation [49].

3. Observables

Extracting useful information from a quantum state requires the ability to make measurements and estimate observables. These range from local observables, such as the energy density of a particular state or a time-ordered correlation function of a spin, to nonlocal observables crucial to the identification of exotic phases of matter and topological properties [578,579,595]. For qubit systems, measurements only in the computational basis are inadequate to determine the complex phase factors in superposition states that determine the expectation values of observables that have off-diagonal components. Thus, if measurements in the computational basis are the only ones possible, it is essential to be able to rotate the state before measurement to map off-diagonal observables into the computational basis, thereby effectively creating the ability to measure in other bases. Similarly, in CV systems, one needs to be able to make measurements in various bases [e.g., in the Fock (boson number) basis and in the position and momentum quadrature bases].

Using the gate sets and techniques presented in this work, a wide range of observables can be measured in

hybrid oscillator-qubit hardware. We note that, particularly in the NISQ era, error-mitigation strategies [376,377] such as zero-noise extrapolation (ZNE) and probabilistic error cancellation (PEC) play an important role in the determination of observables from noisy CV [379,596] and DV [566] hardware. Separately, classical shadow methods [597]—a class of techniques to estimate observables from a small number of measurements—have garnered much recent interest in the literature and have recently been extended to CV systems [402,403]. We leave a more complete description of these techniques to the above-cited works, and here simply note their importance for the estimation of observables.

In the following paragraphs, we provide a few illustrative (but nonexhaustive) examples of how one can measure Hamiltonian terms of interest in hybrid oscillator-qubit hardware.

a. Number operator. The occupation of a mode can be measured in multiple ways. First, one can probe whether or not a mode is in a particular Fock state n by, e.g., using an SQR gate (Box IV.9) and encoding the answer in the state of a qubit. However, asking this question for all n (up to some cutoff $n_{\max}$) is costly, requiring a number of measurements that is linear in $n_{\max}$. Alternatively, it has been demonstrated experimentally [147] that one can make a binary search using successive Hadamard tests with the parity operator, where the kth round reveals the kth bit of n (i.e., $\lfloor n/2^k \rfloor \bmod 2$) via phase kickback on the qubit [147]. This latter approach scales logarithmically in $n_{\max}$. For details, see Appendix E. See also Appendix D 1 b, where this approach is recompiled into the phase-space ISA.

Either technique by extension enables the measurement of operators that are functions of n, such as the chemical potential $\langle n \rangle$, on-site interaction $\langle n^2 \rangle$, and other operator strings that include $\hat{n}$, such as the Hamiltonian terms discussed in Sec. VII C 6.

b. Field operators. We can directly synthesize measurements of the resonator field operators $a = \hat{x} + i\hat{p}$ (here in Wigner units) via measurements of $\hat{x}$ and $\hat{p}$ [145]. To achieve this, recall from Eqs. (384)–(386) that a conditional momentum boost CD(ik) (Box IV.7) can be viewed either as a momentum boost $\pm k$ conditioned on σ_z of the qubit or as a qubit rotation about the z axis by an angle $\theta = -4k\hat{x}$ that is linearly dependent on the position of the oscillator. Taking the latter viewpoint, by starting the ancilla in the $|+\rangle$ state and then measuring σ_y after application of the conditional momentum boost, one obtains a one-bit estimate of $\sin(4k\hat{x})$. For sufficiently small k such that the sine can be linearized, this is essentially a homodyne measurement of $\langle \hat{x} \rangle$. This is similar to what has been achieved for flying microwave photons using traditional microwave homodyne techniques (and huge amounts of careful averaging to eliminate amplifier or detector noise) [598]. In Ref. [41], it has been shown that this technique can be leveraged to measure nonlocal observables such as string order correlators.

c. Hopping terms. Many bosonic models include terms of the form $T_{ij} \propto a_i^\dagger a_j + a_i a_j^\dagger$ describing hopping of a single particle from site i to site j. Notably, such a term exactly generates a beam-splitter gate (Box IV.4). To measure this term, we note that conjugation by an appropriately chosen beam-splitter gate maps T_{ij} onto the photon number difference between sites i and j via what is effectively interferometry:

$$\mathrm{BS}^\dagger\left(\frac{\pi}{2}, -\frac{\pi}{2}\right) T_{ij}\, \mathrm{BS}\left(\frac{\pi}{2}, -\frac{\pi}{2}\right) = a_i^\dagger a_i - a_j^\dagger a_j . \quad (438)$$

Consequently, measurement of T_{ij} is achieved by first applying the beam splitter $\mathrm{BS}^\dagger(\pi/2, -\pi/2)$ followed by measurement of the number operator at sites i and j, in close analogy to the combination of a Hadamard gate and Z-basis measurement facilitating an X-basis measurement for qubits.

d. Bosonic interaction terms. Another useful capability is the measurement of interaction terms involving multiple modes. For example, consider an interaction term of the form $V_{ij} \propto a_i^\dagger a_i a_j^\dagger a_j$. To measure V_{ij}, we can proceed in multiple ways. First, we can use the previously described strategy for measuring the number operator to independently determine n_i and n_j and, over repeated shots, correlate the results to determine $\langle n_i n_j \rangle$ (which, importantly, differs from $\langle n_i \rangle \langle n_j \rangle$ when n_i and n_j are correlated). While sufficient as a terminal measurement (i.e., at the very end of a circuit), this strategy extracts more information than needed, collapsing each mode onto a Fock state in the process. In situations in which we wish to measure V_{ij} without disentangling modes i and j (e.g., if we expect V_{ij} to be a conserved quantity and want to check this mid-experiment), the conditional cross-Kerr gate synthesized in Eq. (378) provides an alternative approach strategy, as it rotates an ancilla by an amount proportional to V_{ij}. In principle, one could leverage this to carry out phase estimation to do an efficient binary search for the value of the integer $n_i n_j$. However, in practice, this requires repeated fast and accurate synthesis of the cross-Kerr unitary in Eq. (378) for a variety of different rotation angles.

Next, we provide several explicit examples of physically interesting models that can be implemented efficiently using hybrid oscillator-qubit platforms.

4. One qubit, one mode: The Rabi model

One of the simplest yet most profound models in quantum physics is the Rabi model [359,599]. It describes the interactions between a harmonic oscillator and a single spin-1/2 system. Its applications are wide-ranging, from the fundamental description of quantum optics and molecular physics to real-world technologies such as nuclear magnetic imaging. Along with the Jaynes-Cummings model [356] (to which, as discussed below, the Rabi model reduces in a certain parameter regime), it is widely used to describe light-matter interactions under the dipole approximation and consequently serves as the foundation for the fields of cavity and circuit QED—the latter of which forms the very basis for the superconducting hardware considered in this paper [28,77].

The Rabi model is described by the Hamiltonian

$$H_{\text{Rabi}} = \hbar\omega_m a^\dagger a + \frac{\hbar\omega_q}{2}\sigma_z + g\left(a^\dagger + a\right)\sigma_x, \qquad (439)$$

where ω_m and ω_q are the frequencies of the oscillator and the two-level system, respectively, while g is the coupling strength between the two, here taken (via a gauge choice) to be real. The first and second terms of H_{Rabi} correspond to the free evolution of the oscillator and qubit, while the final term describes interactions between the two. This is, to a good approximation, the natural physical Hamiltonian natively available (and always "on") in cQED settings. However, we may wish to simulate a Rabi Hamiltonian with different (programmable and possibly time-varying) parameters than are natively available in cQED setups, trapped ions, or neutral atoms. Doing so within the instruction sets presented in this work is straightforward: the first two terms can be implemented via oscillator and single-qubit rotations (in software via simple rotating-frame changes), respectively, while the last is exactly the generator of a conditional displacement, $\text{CD}(\alpha)$, up to single-qubit rotations.

As discussed in Appendix A, in the parameter regime in which $g \ll \{\omega_m, \omega_q\}$, one can make the replacement $\sigma_x \to \sigma_+ + \sigma_-$ and discard the rapidly rotating, non-energy-conserving terms, $a^\dagger\sigma_+$ and $a\sigma_-$, under what is known as the rotating-wave approximation (RWA). This results in the comparatively simpler Jaynes-Cummings model,

$$H_{\text{JC}} = \hbar\omega_m a^\dagger a + \hbar\omega_b\sigma_z + g\left(a\sigma_+ + a^\dagger\sigma_-\right), \qquad (440)$$

which reduces light-matter interactions to the number-conserving exchange of mode and qubit excitations, e.g., absorption ($a\sigma^+$) and emission ($a^\dagger\sigma^-$) of a single photon. However, as one increases the coupling such that $g \ll \{\omega_m, \omega_q\}$ no longer holds, the system enters the so-called ultrastrong coupling regime ($g \gtrsim 0.1\{\omega_m, \omega_q\}$) and, for even larger g, the deep strong coupling regime ($g \gtrsim \{\omega_m, \omega_q\}$). It is predicted that these parameter regimes display exotic behavior such as ground state virtual excitations, and as such have been the topic of both theoretical and experimental investigations over the past decade [600,601].

5. Many modes: The Bose-Hubbard model and the fractional quantum Hall effect

Several models in condensed-matter, optical, and nuclear physics are purely bosonic and therefore well suited for simulation using hardware with native bosonic modes. For example, boson sampling, in which photons pass through a linear-optical network before measurement in the number basis, is a candidate for quantum advantage [155,158,159,539,602]. Furthermore, the Bose-Hubbard model has become a test bed for phenomena such as the well-known superfluid-to-Mott-insulator transition, the first observation of which spurred forward the field of quantum simulation with cold atoms in optical lattices [571]. Moreover, it is conjectured that the Bose-Hubbard model hosts topological phases with novel dynamics akin to the ones encountered in the fractional quantum Hall effect (FQHE) [603] for electrons and which could be observed in cQED systems [604–608].

While these models can often be implemented in analog quantum simulators, such approaches are limited in terms of Hamiltonian terms that can be realized and observables that can be measured. For example, the complex phases in the hopping term that incorporate effects of an external magnetic field (required in, e.g., the description of the FQHE) are challenging to implement in cold atoms due to their neutrality. Synthetic magnetic fields have been employed in past work [609], but an FQHE-like state has only so far been observed in a few small-scale experiments in cold-atom [610] and cavity- and circuit-QED platforms [606,611,612]. Furthermore, entanglement-based measures, important for verifying the topological properties of candidate FQHE states [578,579], are challenging to obtain; analog cold-atom experiments have, e.g., typically relied on measurement of occupation numbers, which supplies no information regarding coherences [610] (though we note a recent proposal for measuring arbitrary properties in analog simulators [613]).

Digital simulation with oscillator-qubit hardware presents advantages on all of these fronts. Not only does the hardware natively support bosonic modes and operations, but arbitrary Hamiltonian terms, such as magnetic field terms crucial to the observation of the FQHE, can be synthesized by adjusting the phases of the beam-splitter couplings on each link of the lattice, using the instruction set architecture presented in this work. The programmability of such simulators allows one to easily tune Hamiltonian parameters and synthesize additional terms. Finally,

as previously discussed, digital hardware offers the ability to measure a wide range of observables that otherwise present a challenge in analog simulations.

With the above motivations in mind, we now turn explicitly to the Bose-Hubbard model as an illustrative example. Its Hamiltonian (for zero magnetic field) is given by

$$H_{\mathrm{BH}} = -\mu \sum_i \hat{n}_i + \frac{U}{2} \sum_i \hat{n}_i \left(\hat{n}_i - 1\right) + H_{\mathrm{hop}},$$
$$H_{\mathrm{hop}} = -J \sum_i \left(a_i^\dagger a_{i+1} + \mathrm{h.c.}\right), \tag{441}$$

where μ is the chemical potential, J is the hopping strength, and U is the interaction strength (with $U > 1$ repulsive and $U < 1$ attractive). Time evolution $e^{-iH_i \Delta t}$ under each term H_i for a duration Δt can be directly realized in the Fock-space ISA [41,52]: H_{hop} generates a beam splitter $\mathrm{BS}(-2J\Delta t, 0)$, the chemical potential term is realized using a phase-space rotation gate $R(-\mu\Delta t)$, and the on-site interactions can be implemented using a SNAP gate (Box IV.10) by choosing $\varphi_n = Un(n-1)\Delta t/2$, such that

$$\mathrm{SNAP}(\varphi_0, \varphi_1, \varphi_2, \ldots) = e^{-i\sum_n Un(n-1)\Delta t |n\rangle\langle n|/2}$$
$$= e^{-i\sum_n U\hat{n}(\hat{n}-1)\Delta t/2}. \tag{442}$$

As a simple illustration, Fig. 38 displays the Trotterized evolution of the mode occupation implemented in BOSONIC QISKIT [52].

While here we have simplified to the case in which U, μ, and J are identical across all sites, we note that additional interesting possibilities emerge by varying these parameters across the lattice. For example, it has been shown in cold-atomic platforms that the one-dimensional (1D) Bose-Hubbard model with added disorder in the chemical potential μ_i leads to many-body localization [614]. Separately, alternation of the hopping strength J_i leads to an (interacting) extension of the Su-Schrieffer-Heeger model, known to display topological features such as edge states [358].

Moreover, magnetic fields and strong interactions in the Bose-Hubbard model give rise to fractionally charged quasiparticles, long-ranged entanglement, and anyonic exchange statistics [603,615]. In this direction, the selection of the beam-splitter phase ϕ allows one to implement hopping in a magnetic field, thus facilitating potential investigation of the FQHE. In all, this multitude of possibilities constitutes a rich and promising playground for quantum simulation of many-body bosonic models using hybrid oscillator-qubit hardware.

6. Large-spin models: The transverse-field Ising model

Several large-spin models, i.e., those with spin-S degrees of freedom with $S > 1/2$, are of interest in condensed-matter physics [366,385] and chemistry [616]. While individual qubits are naturally suited to represent spin-1/2 degrees of freedom, encoding a generic spin-S system in qubit-based hardware requires $\log(2S+1)$ qubits. Though not prohibitively expensive in terms of qubit resources, this approach comes at an additional cost in the two-qubit gate count required to implement both single- and multi-spin-S operations, prohibiting straightforward simulation, particularly on near-term devices.

In contrast, in hybrid oscillator-qubit hardware, one can make use of the larger Hilbert space of the bosonic modes to represent d-level systems in a hardware-efficient manner [41,62]. Here, we discuss the Schwinger-boson representation, an encoding that *directly* implements a spin-S and its defining commutation relations in terms of simple combinations of bosonic operators—a possibility first noted by Julian Schwinger [617]. Previously, the Schwinger-boson representation has been proposed for the realization of large-spin $U(1)$ lattice gauge theories using qubit-oscillator hardware [41,618]. Separately, it is the basis for the oscillator-based dual-rail representation of a qubit in, e.g., cQED experiments [233].

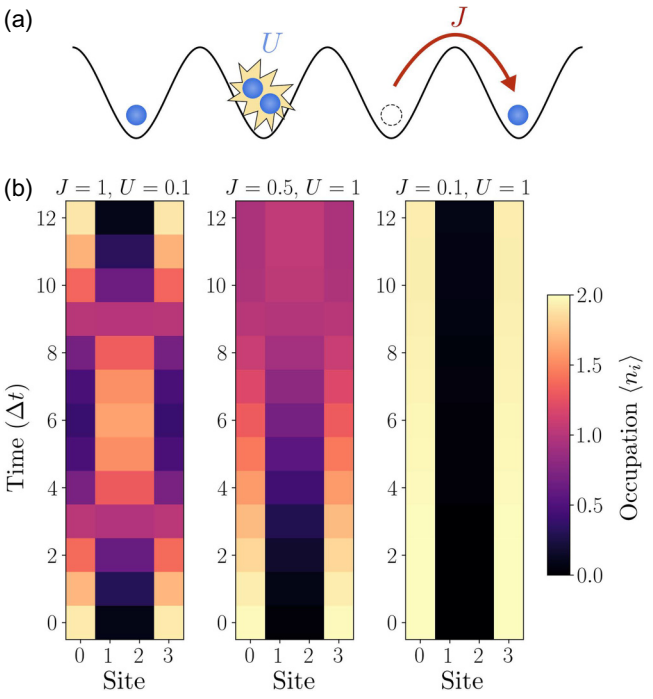


FIG. 38. Simulation of a four-site Bose-Hubbard model using BOSONIC QISKIT. (a) A schematic of the Bose-Hubbard model, illustrating nearest-neighbor hopping with strength J and on-site repulsion of strength $U > 0$. This competition between these Hamiltonian terms leads to a superfluid-to-Mott-insulator phase transition [571]. (b) Occupancy as a function of time for various choices of J and U. In all cases, a time step of $\Delta t = 0.2$ has been used and the system has been initialized to have two bosons at each edge site, with intermediate sites vacant.

The core idea of the Schwinger-boson representation is the following: if one restricts to the subspace of total photon number N (consisting of $N+1$ total states) for two oscillators with annihilation operators a and b, then the combinations $a^\dagger b$ and $ab^\dagger$ act *exactly* as spin-S raising and lowering operators S^+ and S^- within this subspace, with the total spin being given by $S = N/2$. For example, if we take $N = 1$, this subsystem consists of the Fock states $|1,0\rangle$ and $|0,1\rangle$, with $S^+ = a^\dagger b$ and $S^- = ab^\dagger$ acting exactly as raising and lowering operators for an effective spin-1/2 system. Notably, this case is precisely the idea behind the dual-rail qubit in a pair of oscillators (see Sec. V B 3).

More generally, one can represent spin-S systems of arbitrary S by restricting to the subspace of $N = 2S$ photons distributed between a pair of oscillators (representing a spin-S in a single oscillator is also possible [619,620]). Formally, the mapping is as follows:

$$\begin{aligned} S^x &= (a^\dagger b + ab^\dagger)/2, \\ S^y &= (a^\dagger b - ab^\dagger)/2i, \\ S^z &= (a^\dagger a - b^\dagger b)/2, \end{aligned} \tag{443}$$

which exactly obeys the commutation relations

$$[S^x, S^y] = iS^z \quad \text{et cyc.} \tag{444}$$

We also see that the beam splitter (Box IV.4) can be written as

$$\mathrm{BS}(\theta,\phi) = e^{-i\theta[\cos(\phi)S_x - \sin(\phi)S_y]}. \tag{445}$$

Beam splitters alone therefore give us access to two generators of rotations and allow for arbitrary SU(2) rotations of the two-mode qudit. In particular, we note that rotations based on the third SU(2) generator, S^z, can be implemented using a pair of beam-splitter gates:

$$e^{-i\phi S^z} = \mathrm{BS}(\pi,0)\mathrm{BS}(\pi,\phi). \tag{446}$$

This could also be implemented "in software" with phase-space rotations (see Box IV.2).

For the case of a spin-1/2 ($N = 1$), the two modes form a two-level qubit and SU(2) rotations consequently provide universal single-spin operations. However, for $S > 1/2$ ($N > 1$), the most general unitary single-spin operation lives in SU(d) with $d = N + 1 > 2$, and therefore beam splitters alone do not provide universal control. Instead, we need additional hybrid oscillator-qubit operators to synthesize generators that are nonlinear in the spin operators. To that end, one can recognize that any polynomial of S^z is diagonal in the Fock basis. Consequently, it is possible to realize nonlinear unitaries of the form $U = \exp(-i\theta f(S^z))$ via a single SNAP gate on mode a, which has been proposed in Refs. [41,52]. Here, $f(S^z)$ is an arbitrary polynomial in S^z, and we have additionally used the fact that $S^z = (a^\dagger a - b^\dagger b)/2 = a^\dagger a + N/2$. Together with beam-splitter-enabled $SU(2)$ rotations, this nonlinear gate provides universal control for $SU(d)$. This is analogous to the universal control of high-spin atoms in atomic physics utilizing the nonlinear Zeeman effect [460,621,622].

Measurement of the spin-S system in the S^z basis has a natural implementation in the Fock-space ISA—one must simply resolve the number operator $a^\dagger a$ (or $b^\dagger b$), which can then be used to determine S^z (see Appendix E for details on photon-number measurement). By extension, this ability enables the measurement of the spin along any direction, $\hat{n}\cdot\vec{S}$, by first rotating to the $\hat{z}$ basis with an appropriate beam splitter (much like a Hadamard gate prepending a Z-basis measurement enables an X-basis measurement for a qubit). With this primitive, one can more generally carry out quantum state tomography on a spin-S system by measuring $\hat{n}\cdot\vec{S}$ along a minimum of $4S+1$ distinct directions of $\hat{n}$ [623]. To reduce statistical uncertainty in the state construction, however, it is highly beneficial to increase the number of distinct measurement axes beyond the bare minimum [624]. Notably, the tomographic protocols described here require only beam splitters and photon-number measurements, independent of the value of S.

As a simple example of how the Schwinger-boson representation can be leveraged for quantum simulation, let us consider the 1D transverse-field Ising model of spin-S. The model Hamiltonian is

$$H_{\mathrm{TFIM}} = -\sum_i S_i^z S_{i+1}^z - \lambda \sum_i S_i^x, \tag{447}$$

where the first and second terms correspond to spin-spin interactions and coupling to a transverse field, respectively. To implement this model, we envision that the spin at site i is composed of two oscillators with corresponding annihilation operators a_i and b_i. Unitary evolution under the transverse-field terms can be implemented via beam-splitter gates of the form $\mathrm{BS}(\theta,0)$ using Eq. (445). The spin-spin term, on the other hand, requires a non-Gaussian gate due to being quadratic in spin operators. To see how to proceed, we can first recast in terms of bosonic operators,

$$S_i^z S_{i+1}^z = \frac{1}{4}(a_i^\dagger a_i - b_i^\dagger b_i)(a_{i+1}^\dagger a_{i+1} - b_{i+1}^\dagger b_{i+1}), \tag{448}$$

and note that the four above terms commute, such that

$$\begin{aligned} e^{-i\theta S_i^z S_{i+1}^z} = {} & e^{-i\frac{\theta}{4}\hat{n}_{a_i}\hat{n}_{a_{i+1}}} e^{-i\frac{\theta}{4}\hat{n}_{b_i}\hat{n}_{b_{i+1}}} \\ & \times e^{i\frac{\theta}{4}\hat{n}_{a_i}\hat{n}_{b_{i+1}}} e^{-i\frac{\theta}{4}\hat{n}_{b_i}\hat{n}_{a_{i+1}}}. \end{aligned} \tag{449}$$

Taking advantage of the fact that $S = (\hat{n}_{a_j} + \hat{n}_{b_j})/2$ is conserved, we can simplify this to

$$e^{-i\theta S_i^z S_{i+1}^z} = e^{-i\theta(\hat{n}_{a_i} - S)(\hat{n}_{a_{i+1}} - S)}. \tag{450}$$

Therefore, the synthesis of $\exp\left(-i\theta S_i^z S_{i+1}^z\right)$ reduces to that of generators that are products of number operators of distinct oscillators. Although challenging to implement exactly, we can note that each unitary is equivalent to the conditional cross-Kerr gate approximately synthesized using BCH techniques in Eq. (378) if we use an auxiliary qubit in the state $|0\rangle$. Thus, the digital simulation of the transverse-field Ising model is straightforward using the gate sets and compilation techniques presented in this work.

The strategy of using the Schwinger-boson representation to simulate spin systems is of much wider interest than this simple example, with possible applications of classical and quantum spin models spanning the fields of magnetism [625], materials [626], artificial intelligence [627], and even neural processing [628]. Furthermore, because each spin is entirely encoded in a photon-number-preserving subspace, the possibility of detecting errors emerges; similarly to the dual-rail qubit, one can detect photon loss errors by making measurements of the joint parity in the corresponding pair of oscillators [233,235]. Thus, studying spin models using the Schwinger-boson representation in hybrid oscillator-qubit hardware not only provides an avenue for direct study of a wide range of interesting models but does so in a particularly promising NISQ-friendly manner by offering a degree of noise resilience through postselection. For an in-depth example, see Ref. [41] on implementing a $U(1)$ lattice gauge theory in hybrid qubit-oscillator hardware.

7. Many qubits, many modes: $\mathbb{Z}_2$ lattice gauge theory

A variety of interesting Hamiltonians include both qubits and bosonic modes and are thus a natural fit for the hybrid oscillator-qubit hardware. Examples include multispin-boson [629] models, bosonic lattice gauge theories [363], the Jaynes-Cummings-Hubbard model describing many-body states of (dressed) light [630,631], and the exciton-vibrational Hamiltonians governing energy transfer and relaxation pathways in molecules and materials [632,633]. Such Hamiltonians are challenging to implement both on qubit-based hardware, due to the inefficiency of implementing bosonic gates (see Sec. VIII B and Ref. [41]), and analog quantum simulators, due to the multi-species nature of the models and the lack of universality of the operations. Hybrid oscillator-qubit platforms provide a promising hardware-efficient approach for the quantum simulation of these Hamiltonians due to the ease with which multi-oscillator, multi-qubit, and hybrid oscillator-qubit gates can be implemented.

Here, we take as an illustrative example a 1D bosonic $\mathbb{Z}_2$ lattice gauge theory [618]. The implementation of this model in the proposed hybrid oscillator-qubit architecture, along with its extension to 2D, has been studied in Ref. [41]. Here, we provide a high-level review of the general strategy and refer the reader to the aforementioned work for an in-depth discussion.

The $\mathbb{Z}_2$ lattice gauge model describes bosonic matter that hops along a 1D chain of sites alternated by qubits that represent a $\mathbb{Z}_2$ gauge field living on the links of the lattice. The Hamiltonian of this model is given by

$$\begin{aligned} H_{\mathrm{LGT}} &= H_{\mathrm{hop}} + H_{\mathrm{field}}, \\ H_{\mathrm{hop}} &= -J \sum_i a_i^\dagger Z_i a_{i+1} + a_i Z_i a_{i+1}^\dagger, \\ H_{\mathrm{field}} &= -g \sum_i X_i, \end{aligned} \tag{451}$$

where a_i ($a_i^\dagger$) removes (adds) a particle at site i, and X_i and Z_i are Pauli operators for the qubit at link i (which connects site i and site $i+1$). To better understand this model (and the competition between H_{hop} and H_{field} it describes), it is helpful to first imagine that all qubits are initialized in an X eigenstate. H_{field} describes the energy of the field and is minimized by aligning all qubits in the $+X$ direction. In contrast, H_{hop} describes bosons hopping between sites, causing a phase-flip (sending $+X$ to $-X$ and vice versa) at each intermediary link, altering the field energy. Consequently, the model displays drastically different behavior in opposite parameter regimes, with $|J/g| \gg 1$ leading to a "deconfined" dynamics where matter freely hops (or more precisely, the confinement length greatly exceeds the finite sample size), and $|J/g| \ll 1$ leading to "confined" dynamics where the large field inhibits hopping [363]. This is a statement about the dynamics—the ground state phase is always confining for $g/J > 0$ [634].

Much like in electrodynamics, we can define an analog of Gauss's law that relates the number of bosons (or *charges*) at site i to the product of "left" and "right" field values $X_i X_{i+1}$ (akin to a field integral in electrodynamics). In particular, we can identify this gauge symmetry with the generators

$$\hat{G}_j = X_j e^{i\pi a_j^\dagger a_j} X_{j+1}. \tag{452}$$

To see that the generators $\hat{G}_j$ correspond to a conserved quantity, note that when a boson hops from site j to site $j+1$, the factors of $\hat{G}_j$ change as follows: (1) $a_j^\dagger a_j$ is decremented by one, flipping the sign of the photon-number parity $e^{i\pi a_j^\dagger a_j}$; and (2) the presence of Z_j in the hopping amplitude flips the sign of X_j, leaving $\hat{G}_j$ unchanged overall. In other words, this means that $[H_{\mathrm{hop}}, \hat{G}_j] = 0$. Noting

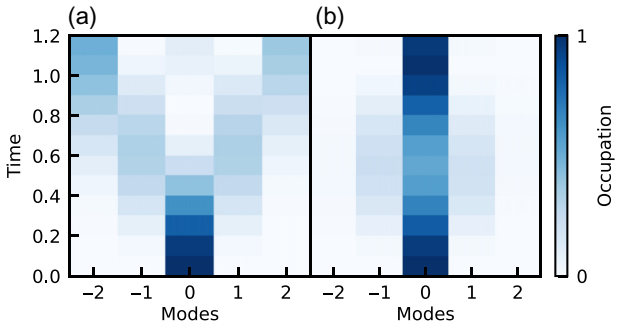


FIG. 39. Implementation of the $\mathbb{Z}_2$ lattice gauge theory in BOSONIC QISKIT [52]. Both panels display the average occupation at each site as a function of time. The initial state is chosen such that there is one boson at the central site, and all others are vacant. For Gauss's law, we fix $G_j = +1$. (a) Results for $g/J = 0.1$ (hopping-dominated): we observe a light-cone spread of the boson. (b) Results for $g/J = 5$ (field-dominated): we observe that the boson is (mostly) trapped at the central site. In both cases, a step size of $0.1J$ has been used. This figure has been adapted from Ref. [41].

that $\hat{G}_j$ trivially commutes with H_{field}, $\hat{G}_j$ is therefore a conserved quantity. Such a feature is advantageous for simulations on noisy hardware, as we can think of the generators $\hat{G}_j$ as stabilizers in an error-detecting code in the sense that we can perform error mitigation by postselection of results that obey the Gauss law constraints.

Both H_{field} and H_{hop} can be implemented directly with gates previously discussed: H_{field} is simply a single-qubit rotation, while H_{hop} corresponds to a controlled beam splitter. As a demonstration, we implement Trotterized time dynamics of H_{LGT} in BOSONIC QISKIT for both hopping-dominated and field-dominated regimes in a five-site system. We initialize a boson at the central site, with all others vacant.

In Fig. 39 we display the resulting time-evolved occupancy at each site. The competing order between J and g is evident, as the hopping-dominated regime gives rise to a light-cone spread of a boson initialized at the central site, whereas the field-dominated regime prohibits hopping due to the prohibitive cost to flip link qubits. For a numerical simulation of dynamics and ground state preparation of this model and others in the presence of noise, see Ref. [41].

8. Outlook: Fermion-boson problems

Classically intractable fermion-boson problems include nonequilibrium dynamics in QCD relevant to heavy-ion collisions [635], calculation of vibrational spectra of large anharmonic molecules coupled to an environment for identification of unknown samples [550], and the study of phonon dynamics in strongly correlated materials such as high-temperature superconductors [636]. Hybrid oscillator-qubit hardware is well suited to such problems, as the native oscillators naturally encode bosonic excitations, while fermionic degrees of freedom can either be mapped onto physical or logical qubits or onto high-Q oscillators [637,638], depending on the specifics of the problem at hand. Furthermore, the instruction set architectures we have presented offer universal control, such that arbitrary fermionic and bosonic gates can be synthesized.

A wide range of encodings can facilitate the mapping of fermions onto qubits. Well-known options include Jordan-Wigner [194], Bravyi-Kitaev [35], Verstraete-Cirac [639], and Derby-Klassen [640] encodings, all of which enable the encoding of fermionic degrees of freedom using either logical qubits derived from CV-to-DV, DV-to-DV, or concatenated codes (see Sec. V). For near-term applications, one could also use the physical qubits themselves. Notably, fermion-to-qubit mappings require nonlocal qubit-qubit gates which, in this case, can be implemented at the physical layer using the oscillator-assisted multi-qubit entangling gate techniques described in Sec. VI C 2. However, in the case of superconducting hardware, it is important to note that transmon-qubit lifetimes are typically on the order of $\sim$100 μs, whereas 3D microwave cavity modes have far longer lifetimes on the order of $\sim$1 ms [24] or more [25,26]. Therefore, even without proper error correction, it may be advantageous to consider encoding the fermions into oscillators using, e.g., the lowest-lying Fock states of cavity modes or dual-rail or Schwinger-boson qubits.

As a simple demonstration, we now illustrate the compilation of the so-called fermionic simulation gate [641,642],

$$\text{fSim}(\theta,\phi) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta) & -i\sin(\theta) & 0 \\ 0 & -i\sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 0 & e^{-i\phi} \end{pmatrix}, \tag{453}$$

a particular application of which is the fermionic-SWAP or fSWAP gate, equivalent to a SWAP gate up to an additional minus sign applied to the $|11\rangle$ state to account for anti-symmetric fermionic exchange statistics (see Table VIII). Specializing to the case in which we encode qubits in the $|0\rangle$ and $|1\rangle$ Fock states of a pair of oscillators indexed by j and k, the fSim gate can be compiled as

$$\text{fSim}(\theta,\phi) = \exp\left(-i\phi\sigma_z a_j^\dagger a_j a_k^\dagger a_k\right) \times \text{BS}\,(2\theta, 0)\,, \tag{454}$$

where the first term is a cross-Kerr gate, synthesized using BCH methods in Eq. (378) and controlled on an ancillary qubit prepared in the state $|0\rangle$. Alternatively, one could replace this approximate cross-Kerr gate with a CPHASE(ϕ) gate realized using the methods of Ref. [235].

The fermionic SWAP operator can be implemented via

$$\text{fSWAP} = \left[R\left(-\frac{\pi}{2}\right) \otimes R\left(-\frac{\pi}{2}\right)\right] \times \text{fSim}\left(\frac{\pi}{2}, 0\right) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}, \tag{455}$$

where $R(\theta)$ corresponds to a phase-space rotation of the oscillator, and we have furthermore displayed only the portion of the unitary acting within the encoded two-qubit subspace.

In summary, hybrid CV-DV hardware offers interesting opportunities for quantum simulation of coupled systems of spins, fermions, and bosons. Here, we have briefly discussed qubits and general spin-S qudits encoded in a pair of cavities using the Schwinger-boson representation and presented a possible compilation of fermionic SWAP gates. For an in-depth manual for using CV-DV hardware for simulating fermions and bosons, including gauge theories in 2D, see Ref. [41].

D. Quantum random walks

Quantum random walks (QRWs) [643] provide a generalization of classical random walks, and serve many purposes in quantum computation, quantum simulation, and algorithm design. They capture the physics of a particle moving in space over a graph or lattice, acting as a quantum analog of diffusion. QRWs also provide a paradigm for quantum algorithms, offering speedups through reduced hit times of random walks, which can be polynomially (or sometimes exponentially) faster than their classical counterparts.

Traditional quantum random walks employ qubits, but using an oscillator in place of qubits can offer unique advantages and access to richer dynamics. Specifically, random walks often involve electromagnetic fields or motional modes which are naturally bosonic. Walk dynamics with oscillators may also capture continuous-valued displacements, rather than discrete changes, making such hybrid quantum walks closer to analog quantum simulators than just qubit-based walks. In addition, a bosonic walker may represent many walker positions (in superposition), compared with a qubit walker.

Here, we briefly review for pedagogical purposes how QRWs with a bosonic random walker can go beyond qubit systems and be realized using a hybrid CV-DV architecture. In this setting, QRWs involve repeated conditional displacements (CD(ϵ); see Box IV.7) of the oscillator, each followed by a multi-dimensional coin toss which decides the basis in which the auxiliary qubit is measured. In other words, the multi-dimensional coin toss decides the auxiliary-qubit rotations about the vectors on the XY plane of the Bloch sphere ($R_\phi(\theta)$; see Table VIII) before measuring in the Z basis. Thus, the phase-space ISA can be exploited to straightforwardly implement QRWs in hybrid CV-DV quantum computing architectures.

Reference [643] has shown that QRWs can be used to study nontrivial effects in quantum mechanics that arise due to quantum interference. For example, let us conditionally displace an oscillator by a very small amount $|\epsilon| \ll 1$. Using the definitions from Eq. (289) (but now including the normalization factors for the cat states), we can write

$$\begin{aligned} &\text{CD}(\epsilon)\,|0\rangle_{\text{osc}}\,|+\rangle_{\text{qubit}} \\ &= \left(|\epsilon\rangle\,|0\rangle + |-\epsilon\rangle\,|1\rangle\right)/\sqrt{2} && (456)\\ &= \frac{\mathcal{N}_{\text{even}}}{2}\left|\mathcal{C}_{0,\epsilon,+}\right\rangle_{\text{osc}}|+\rangle_{\text{qubit}} + \frac{\mathcal{N}_{\text{odd}}}{2}\left|\mathcal{C}_{0,\epsilon,-}\right\rangle_{\text{osc}}|-\rangle && (457)\\ &\approx \frac{\mathcal{N}_{\text{even}}}{2}|0\rangle_{\text{osc}}|+\rangle_{\text{qubit}} + \frac{\mathcal{N}_{\text{odd}}}{2}|1\rangle_{\text{osc}}|-\rangle_{\text{qubit}}\,. && (458)\end{aligned}$$

The final equation uses an approximation in which we have equated a small even (odd) cat state with Fock state $|0\rangle$ ($|1\rangle$). It is important to note that the definitions used here for cat states $\left|\mathcal{C}_{0,\epsilon,\pm}\right\rangle$ include their normalization prefactors, giving rise to the unequal amplitudes for each term in the last two equations. More specifically, the probability of each Fock state depends on the normalization constants of even ($\mathcal{N}_{\text{even}}$) and odd ($\mathcal{N}_{\text{even}}$) cat states of size ϵ,

$$\mathcal{N}_{\text{odd}} = \sqrt{2(1 - e^{-2\epsilon^2})}, \tag{459}$$

$$\mathcal{N}_{\text{even}} = \sqrt{2(1 + e^{-2\epsilon^2})}. \tag{460}$$

The fidelity of even and odd cat states with the Fock states $|0\rangle$ and $|1\rangle$ improves with a decreasing value of ϵ as long as $\epsilon \neq 0$, as shown by the following equations:

$$F_0 = |\langle 0|\mathcal{C}_{0,\epsilon,+}\rangle|^2 = 4\frac{e^{-\epsilon^2}}{\mathcal{N}_{\text{even}}^2}, \quad \alpha \in \mathbb{R}, \tag{461}$$

$$F_1 = |\langle 1|\mathcal{C}_{0,\epsilon,-}\rangle|^2 = 4\epsilon^2\frac{e^{-\epsilon^2}}{\mathcal{N}_{\text{odd}}^2}, \quad \alpha \in \mathbb{R}. \tag{462}$$

Note that, for $\epsilon \to 0$, the small odd cat has a high overlap with the Fock $|1\rangle$ state [see Eq. (462)] in contrast to the small even cat, which has a large overlap with $|0\rangle$ [see Eq. (461)], largely resembling the initial vacuum state. Thus, Eq. (458) shows that there can be a nontrivial unit increase in the photon number of the oscillator upon qubit measurement after an extremely small conditional displacement. We show the marginal distribution in position $|\psi(x)|^2$ and the phase-space quasiprobability distribution $\mathcal{W}(x,p)$ in Fig. 40 for pictorial intuition. A qubit measurement in the X basis collapses the oscillator onto one of the two states shown in this figure.

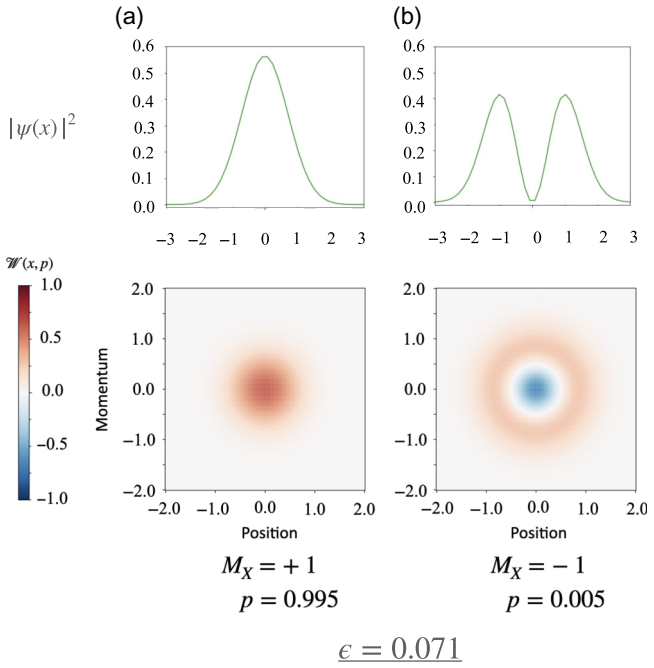


FIG. 40. The marginal probabilities ($|\psi(x)|^2$) and Wigner distribution ($\mathcal{W}(x,p)$) of the oscillator states entangled with orthogonal qubit states $|\pm\rangle$ after small conditional displacements CD($\epsilon = 0.071$), as shown in Eq. (458). (a) The oscillator state entangled with the $|+\rangle$ qubit state. This state closely resembles the vacuum state $|0\rangle$ and occurs for qubit measurement $M_X = +1$ with a high probability $p = 0.995$. (b) The oscillator state entangled with the $|-\rangle$ qubit state. This state is well approximated by Fock state $|1\rangle$, and occurs for qubit measurement $M_X = -1$ with a low probability of $p = 0.005$. The large change in the position of the peaks, shown in the marginal probability distribution, between the two states is achieved merely by a small displacement of $\epsilon - 0.071$. This phenomenon is an artifact of quantum interference and can be observed using quantum random walks, as described in Ref. [643]. This figure has been adapted from Ref. [41].

Let us discuss the consequences of this on a position basis. The even cat [Fig. 40(a)] resembles a vacuum and has its peak centered at $x = 0$, while the odd cat [Fig. 40(b)] has two peaks centered at $x = \pm 1$. If the qubit measurement were to collapse the oscillator onto the odd cat state, the position measurement would yield $|x| = 1$ with maximum probability, hence yielding a net displacement of $|\delta x_-| \sim 1$. Here, the subscript "$-$" denotes displacement upon position measurement of the oscillator state conditioned on the qubit measurement yielding the $|-\rangle$ state.

Thus, it is evident that upon measuring the position of the oscillator there is a nonzero probability that this measurement yields an effective displacement of $\delta x_- \sim \mathcal{O}(1)$ in position, much larger than the original displacement of ϵ. This is a unique and powerful quantum mechanical effect that is only possible due to a local phase of π incurring a minus sign between the two slightly displaced coherent states $|\pm\epsilon\rangle$. The effect of this minus sign is to subtract the wave functions corresponding to the two displaced coherent states. While highly nontrivial, the collapse of the oscillator to this "$-$" superposition is also an extremely rare event because the ratio,

$$\mathcal{N}_{\text{odd}}/\mathcal{N}_{\text{even}} = \sqrt{\frac{1 - e^{-2\epsilon^2}}{1 + e^{-2\epsilon^2}}} = \epsilon + O(\epsilon^5), \tag{463}$$

vanishes as $\epsilon \to 0$.

The probability of observing a large state change (displacement in phase space) induced by the measurement backaction can be increased by changing the qubit measurement axis, a task addressed by a quantum random walk. To see this, let us rewrite the state in Eq. (458) in terms of an arbitrary auxiliary-qubit basis $\vec{\sigma}_\theta = \cos\theta\sigma_z + \sin\theta\sigma_x$, the ± 1 eigenstates of which are denoted $|\pm\theta\rangle$. Alternatively, we can rotate the qubit by $R_Y(\theta/2)$ and then measure it in the Z basis. In this case, the composite oscillator-qubit state just before measurement is equal to (with implicit normalization)

$$\begin{aligned}|\psi\rangle = &(|\epsilon\rangle + \tan(\theta/2)\,|-\epsilon\rangle)\,|+\theta\rangle \\ &+ (\tan(\theta/2)\,|\epsilon\rangle - |-\epsilon\rangle)\,|-\theta\rangle\,.\end{aligned} \tag{464}$$

Reference [643] shows that under the assumption $\epsilon \to 0$, we can use $e^{\pm i\epsilon\hat{p}} \approx I \pm i\epsilon\hat{p}$ and approximate the final state in Eq. (464) effectively as the state after an asymmetric conditional displacement $D_\theta(\delta x_+, \delta x_-)$ of the vacuum:

$$\begin{aligned}|\psi\rangle &\approx D_\theta(\delta x_+, \delta x_-)\,|0\rangle_{\text{osc}}\,|0\rangle_{\text{qub}} &(465)\\ &= \sqrt{\mathcal{P}_+}\,|\delta x_+\rangle\,|+\theta\rangle + \sqrt{\mathcal{P}_-}\,|\delta x_-\rangle\,|-\theta\rangle\,, &(466)\end{aligned}$$

$$\text{where}\quad \delta x_\pm = \frac{1 \mp \tan(\theta/2)}{1 \pm \tan(\theta/2)}\epsilon, \tag{467}$$

$$D_\theta(\delta x_+, \delta x_-) = e^{-i\delta x_+\hat{p}}\,|+\theta\rangle\langle+\theta| + e^{-i\delta x_-\hat{p}}\,|-\theta\rangle\langle-\theta|\,, \tag{468}$$

and $\mathcal{P}_\pm$ denotes the probabilities of collapsing the oscillator onto coherent states $|\delta x_\pm\rangle$. This approximation holds as long as $|\delta x_\pm| \ll \Delta_x$, where Δ_x is the position uncertainty of the initial state (here assumed to be $\Delta_x = \frac{1}{2}$, in correspondence with an initial vacuum state).

Despite this restriction, it is possible to observe large displacements, either $|\delta x_+| \gg \epsilon$ or $|\delta x_-| \gg \epsilon$, with significant probability for an appropriate choice of θ. For example, choosing $\tan(\theta/2) = 1 + \gamma$ such that $\epsilon/\Delta_x \ll |\gamma| \ll 1$ yields $\delta x_- = -2\epsilon/\gamma$, which satisfies the requirement $\epsilon \ll |\delta x_-| \ll \Delta_x$. That is, the effective displacement of the oscillator $|\delta x_-|$, upon collapse to $|-\theta\rangle$ qubit state, is much larger than ϵ, while being much smaller than the initial position uncertainty Δ_x. The probability of this event

is $P_- = \gamma^2/4$. Note that for a premeasurement rotation of $R_Y(\pi/2)$ such that $\gamma = 0$, the probability of collapsing to $|\delta x_-\rangle$ under these assumptions vanishes. In contrast, by choosing a nonzero γ, and thus rotating the qubit by $R_Y(2\tan^{-1}(1+\gamma))$, the probability of observing a nontrivial displacement of the oscillator is enhanced. Reference [643] has shown that for $\epsilon = 0.01$ and $\gamma = -0.1$, this effect is apparent after repeating the protocol for ten rounds with postselection upon -1 measurement outcomes, and the qubit reinitialized to $|+\rangle$ after each measurement. The authors have also shown another example of an interesting quantum mechanical observation using photon numbers and the sideband ISA: the proposed random walk indicates that the photon number of the oscillator can be changed by a large amount even if the interaction only allows for the addition or removal of a single photon.

Another interesting scenario is to carry out the above procedure but skip the "measurement + postselection" step. In this case, the iterative procedure corresponds to a random walk for the composite system with N pairs of conditional displacements and random qubit rotations in the equatorial plane. Such a sequence can be described by a unitary channel that produces a superposition of 2^N different state trajectories. To this end, we note a close similarity between QRWs and the phase-space trajectory, illustrated in Fig. 22(b). Reference [644] has experimentally demonstrated such quantum walks via unitary channels composed entirely out of the phase-space ISA and used them to explore a novel topological effect, which they have argued represents a potential technique for the study of topological quantum materials.

We note that by swapping the roles of qubits and oscillators in QRWs, the coin toss can instead be performed via an oscillator measurement under the Fock or position or momentum basis (or any other basis). This measurement backaction will induce a "walk" of the qubit state according to an infinite-dimensional configuration space parametrized by the oscillator. A coherent version of this can also be constructed by concatenating iteratively a random coin toss on the oscillator and a qubit rotation. This can produce a channel on the qubit which is a superposition of a set of trajectories with a continuous distribution. In the case of a multi-qubit system coupled to an oscillator such as in Refs. [223,513], this QRW construction is more interesting. The flexibility of measuring the oscillator under various bases can enact nontrivial dynamics on the multi-qubit systems.

In summary, quantum random walks are an application of the phase-space ISA (see Sec. IV E 2) that paves the way for observations of novel quantum mechanical phenomena using hybrid CV-DV architectures.

VIII. ADVANTAGES OVER QUBIT-ONLY ARCHITECTURES

In this final section, we illustrate the advantages of hybrid CV-DV hardware compared to DV-only hardware by comparing the gate count of simulating bosonic systems with CV-DV hardware to DV-only hardware. We first discuss in Sec. VIII A different schemes to map oscillators into qubits, before compiling a displacement gate into qubit gates in a qubit-efficient encoding in Sec. VIII B. We close by discussing the gate complexity in the context of other compilation aspects in Sec. VIII C.

A. Mapping between oscillators and qubits

An oscillator has an in-principle infinite-dimensional Hilbert space, whereas a qubit is a two-level system. We therefore need to use a mapping of the Hilbert space of an oscillator and the Hilbert space of a collection of qubits, which requires a truncation of the oscillator Hilbert space.

TABLE VII. An overview of different ways to map an oscillator state to qubits. In the last row, JLP stands for Jordan-Lee-Preskill, where $\varphi_{\max}$ is the maximum field amplitude (boundary of truncated phase space), and δ_φ is the spacing between nearest grid points in the discrete representation of the phase space. $N_{\max}$ is the maximum Fock state number encoded by the encoding.

Encoding	Advantages	Disadvantages	Qubit count	Representation
Unary	Simplifies hardware implementation due to one-hot representation which could allow for repetition code detection of bit-flip errors.	Highly inefficient for high boson occupation due to linear scaling of qubits. Resulting poor scalability for large Fock states.	$O(N_{\max})$	Fock space
Gray [646]	Bit-flip errors represent small boson occupation differences.	Limited benefit for dense matrices or operators with high Hamming distances.	$O(\log(N_{\max}))$	Fock space
Fock binary [200]	Very compact representation. Efficient for diagonal operators.	Susceptible to bit-flip errors since adjacent Fock states are not adjacent in binary representation.	$O(\log(N_{\max}))$	Fock space
JLP [576,645]	Enables representation of bosonic modes in phase space.	High qubit and gate count overhead.	$O(\log(\varphi_{\max}/\delta_\varphi))$	Phase space

We first give two simple mappings as examples in Secs. VIII A 1 and VIII A 2. Other approaches such as the field basis representation used by Jordan, Lee, and Preskill [576,645] and the Gray code [646] can also be useful and may be more efficient for implementation of certain unitary operations. See Ref. [647] for a detailed comparison of simulating d-level system on qubits using different encodings and Ref. [41] for a comparison of the quantum simulation of qubit-boson models on DV versus CV-DV hardware. Table VII summarizes all mappings discussed in this section.

1. Hardware-efficient mapping (Fock-binary encoding)

Note the important (and not accidental) fact that in the multi-qubit states defined in Eqs. (15)–(18), the ordinal number of the basis vector (numbering the list 0–3 from the top) matches the spin state interpreted as a binary number, e.g., $|11\rangle$ corresponds to the third basis vector (starting from 0). More generally for L qubits, we have for the jth standard basis state,

$$|\psi_j\rangle = |b_{L-1}b_{L-2}\ldots b_1 b_0\rangle, \quad j \in [0, 2^L - 1], \tag{469}$$

where $b_k = 0, 1$ labels the standard basis state of the kth qubit and

$$j = \sum_{k=0}^{L-1} 2^k b_k \tag{470}$$

is the integer corresponding to the binary number labeling the qubit state.

As we have seen in Sec. III B, the Fock basis states of an oscillator are also labeled by an integer (beginning at zero). Thus we can very conveniently choose to view the qubit state $|\psi_j\rangle$ as representing the state of an oscillator containing j excitations (bosons), the Fock state $|j\rangle$. A register of L qubits can thus be used to represent the lowest 2^L states of an oscillator (representing boson numbers 0 to $N_{\max} = 2^L - 1$). This provides a mapping from an oscillator to (logarithmically) many qubits, where every Fock state of the oscillator corresponds to a multi-qubit state. This mapping allows us to simulate bosonic modes in classical simulator codes such as QISKIT that are designed for qubits only, a strategy that we have implemented in BOSONIC QISKIT [52–54]. This representation is very convenient because the Fock basis matrix representations of the boson operators a and $a^\dagger$ given in Eqs. (83) and (84) are *unchanged* when we translate them into the multi-qubit basis.

Conversely, a single resonator with boson number cutoff $N_{\max} = 2^L - 1$ in a hybrid architecture can replace L qubits, leading to considerable hardware efficiency [648]. We thus have two ways to realize a Hilbert space of dimension 2^L–using L physical qubits or 2^L states of a single oscillator. Either way, to have universal control, we must be able to perform arbitrary unitary transformations which are $2^L \times 2^L$ matrices. There are "natural" operations in each of the two physical realizations, and those may be quite unnatural in the other realization. In particular, as hinted at in the introduction, it is important to understand that implementing natural bosonic operations on qubit-only quantum hardware can be very difficult. The qubit cost is only logarithmic in $N_{\max}$ and so is not the problem—it is instead the matrix elements of the raising and lowering operators given in Eq. (25) and (26) containing square-root factors. While these are easy to implement [52–54] in a classical simulator designed for qubits only, they are very difficult to realize on quantum qubit hardware because of the quantum arithmetic needed to produce the square-root factors [199,200] (see also Sec. VIII B for details). It is here that we see the main advantage of replacing qubit-only hardware with hybrid oscillator-qubit hardware where the oscillator raising and lowering operators are natively available.

Separately, it is instructive to compare the impact of amplitude damping on a single oscillator to the impact of amplitude damping on the L qubits representing the oscillator up to a photon-number cutoff of $N_{\max} = 2^L - 1$. Recall that the physical oscillator with damping parameter κ, the probability of being in an initially prepared Fock state $|n\rangle$ decays at a rate $n\kappa$, and thus the mean photon number obeys Eq. (271):

$$\frac{d}{dt}\bar{n} = -\kappa\bar{n}. \tag{471}$$

Let us assume that all L qubits used to represent the oscillator have the same decay rate γ out of their excited states. The qubit representation of the oscillator seems advantageous since the error rate in the qubit register representing the oscillator will scale only logarithmically, whereas in an actual oscillator the dominant error (photon loss) scales linearly with the number of photons. For example, the rate of decay out of Fock state $n = 2^k$ is only γ, since only one qubit is in the excited state. In the worst case ($n = N_{\max}$) where all L qubits are in the excited state, the decay rate is still only $L\gamma$, suggesting possible advantages for the qubit encoding.

On the other hand, one might suspect that the loss rate of the mean boson number of the qubit encoding is significantly worse than the one of an oscillator, as 2^j bosons are lost at once if the jth qubit decays. However, a little thought shows that the mean occupation number in the qubit representation still obeys Eq. (471), except with the cavity damping rate κ replaced by the qubit decay rate γ. This is because the mean boson number is a linear function of the qubit excitations [see Eq. (470)]. Typically, the lifetime of state-of-the-art qubits ($\gamma^{-1} \sim 100$–500 μs) is substantially shorter than that of state-of-the-art cavities

($\kappa^{-1} \sim 10^4$–10^6 μs). This demonstrates that while the rate of leaving a given state for the qubits is bounded above by γL (which can in some cases be much smaller than $\kappa\langle\hat{n}\rangle$), the rate of change of the mean photon number represented by the qubits is always worse than $\kappa\langle\hat{n}\rangle$ (assuming $\gamma > \kappa$). For example, in Ref. [26] a transmon controller with a lifetime of $T_1 = 110$ μs has been coupled to a cavity with a lifetime of $T_1 = 35$ ms and used to produce a 10^3-photon cat state. The phase coherence time of this cat state due to photon loss (which produces cat parity jumps) is $\sim T_1/\langle\hat{n}\rangle \approx 35$ μs. A superposition of 0 and 10^3 photons would decohere with precisely the same lifetime. It would take a register of approximately $L = 10$ qubits to simulate this bosonic state and the corresponding lifetime would be only $1/(\gamma L) \sim T_1/10 \sim 11$ μs. We emphasize that beating these constraints on the qubit representation of oscillators would require the use of error-corrected logical qubits that are capable of performing deep and complex circuits with multi-qubit gates.

These considerations show that the way in which decay affects oscillators and qubit representations of oscillators is significantly different. In particular, if the goal were a simulation of an open (i.e., damped) bosonic system, the use of a physical oscillator would be much more natural as the qubit noise processes are a bad representation of bosonic noise. We expect that these considerations concerning noise change substantially for other encodings of oscillators into qubits (e.g., Jordan-Lee-Preskill [576,645]), but the essential conclusion remains the same, namely that qubit noise processes produce a bad representation of the physical oscillator noise processes.

Finally, we stress that to thoroughly address the question of whether a single noisy oscillator holds an advantage over logarithmically many noisy qubits, it is crucial to not only consider the timescale at which errors occur (τ_{error}), but also how this timescale relates to gate duration for the 2^L system (τ_{gates}), as it is comparison of the ratio $\tau_{\text{error}}/\tau_{\text{gates}}$ which determines which implementation holds a practical advantage. This type of analysis has recently been carried out in Ref. [649] to compare noisy d-level qudits to logarithmically many qubits undergoing dephasing errors. There, the authors have found that for certain regimes of τ_{error}, τ_{gates}, and d, single noisy qudits indeed hold an advantage, despite the increased complexity of their dephasing error channel over their many-qubit counterparts. We expect this advantage to be even more pronounced for oscillators due to (i) the simplicity of their dominant error channel and (ii) their substantial advantage in gate complexity (and consequently τ_{gates}) for applications where bosonic gates are desired (see Sec. VIII B). Nonetheless, it would be an interesting direction to rigorously extend the analysis of Ref. [649] for noisy oscillators versus many qubits for the case of amplitude damping. We leave this as a direction for future work.

2. Mapping with redundancy (unary encoding)

A different approach to mapping from oscillators to qubits is to forego the notion of hardware efficiency, and instead introduce redundancy. This may seem inefficient at first glance, but such redundancy is crucial to preserve quantum information in noisy quantum systems. Redundancy can be introduced to both the qubit states and the oscillator states; see Sec. V for a thorough discussion of oscillator redundancy for bosonic quantum error correction. In this section, we instead discuss a mapping of oscillators to qubits that uses redundancy in the qubit states.

One simple mapping with redundancy on the qubits uses the unary representation in a register of $d = N_{\max} + 1$ qubits,

$$|0\rangle_{\text{b}} = |0\ldots0001\rangle_{\text{q}}, \tag{472}$$

$$|1\rangle_{\text{b}} = |0\ldots0010\rangle_{\text{q}}, \tag{473}$$

$$|2\rangle_{\text{b}} = |0\ldots0100\rangle_{\text{q}}, \tag{474}$$

$$|3\rangle_{\text{b}} = |0\ldots1000\rangle_{\text{q}}, \tag{475}$$

where the subscript b labels the bosonic mode Fock states and the subscript q denotes the qubit register. This means that d qubits are needed to represent the lowest d levels of an oscillator. We number the qubits from right to left, starting with zero. In this encoding, only a single qubit is in the 1 state and the location label of that qubit represents the number of photons in the cavity.

Note that this is an exponentially redundant mapping in the sense that the qubit Hilbert-space dimensionality is 2^d but the relevant oscillator levels only have dimension d. However, this mapping has some fault tolerance built into the qubit states. For example, for Fock state $|3\rangle_b$, if the qubit state is changed from $|0\cdots1000\rangle_q$ to $|0\cdots1010\rangle_q$ due to some noise by flipping the second qubit from 0 to 1, the resulting state $|0\cdots1010\rangle_q$ is outside of the computational state space and thus can be detected. Conversely, a spin flip leading to $|0\cdots0000\rangle_q$ cannot be detected. Simple repetition encoding of this representation would allow for bit-flip error correction and, of course, more sophisticated and powerful encodings could be used.

B. Compilation of bosonic operators into and from standard qubit gate sets

In this section, we present a quantitative analysis on the overhead of simulating bosonic hardware on qubit hardware by estimating the CNOT cost for simulating a single displacement operator. For a comparison of different mappings, we refer the reader to Ref. [647]. Specifically, we introduce a novel compilation technique suited for arbitrary bosonic operations which relies on calculating the square-root matrix elements using Newton iterations. We

explicitly calculate the gate count of this method for a displacement and compare to brute-force compilation (which is not scalable beyond a handful of qubits).

The destruction operator a is a natural operator for bosonic systems. Let us consider its matrix elements in the qubit representation. For example, the simple Fock basis expression

$$a|8\rangle = \sqrt{8}|7\rangle \tag{476}$$

tells us that applying the destruction operator to the state with eight bosons yields the state with seven bosons. This is equally simple on the standard multi-qubit basis as well:

$$a|\psi_8\rangle = \sqrt{8}|\psi_7\rangle. \tag{477}$$

This should be all one needs to tell a qubit compiler [by giving it the matrix representation of a in Eq. (83)]. However, it is important to note that a becomes a nontrivial multi-qubit operation in the binary basis that labels the states of the individual qubits. For example,

$$a|1000\rangle = \sqrt{8}|0111\rangle. \tag{478}$$

Thus a is an unnatural weight-4 operator (in this particular example) for qubits because it flips all four qubits. The full matrix representation of the destruction operator in terms of direct products of single-qubit operators is complex because computing which bits have to be flipped requires knowing all the bits. For example, the particular term in a shown in Eq. (478) is given in terms of qubit operators by a complicated (and difficult to physically implement) weight-8 expression,

$$X_3X_2X_1X_0[\hat{I}_3 - Z_3][\hat{I}_2 + Z_2][\hat{I}_1 + Z_1][\hat{I}_0 + Z_1]/16, \tag{479}$$

and there are 15 other similar terms needed to fully represent the operator. From this, we see immediately the huge potential efficiency of having boson degrees of freedom available natively in the hardware, if we are interested in simulating a physical system with bosonic modes.

Conversely, if we would like to use resonators to replace multiple qubits in a quantum simulator or computer, we need to reckon with the fact that simple qubit gates typically correspond to nontrivial unitaries acting on the oscillator. As we discuss further below, oscillators coupled to transmon qubits are "controllable" in the formal language of quantum control theory [266] and recent experiments have demonstrated an impressive level of control of such oscillators [63,64,67,193,266,650].

A natural alternative to the Fock-binary encoding is the Gray code [651], which is a permutation of the binary encoding with the property that only one bit flip at a time as the number the code represents increases or decreases by 1. This would seem to eliminate the multi-qubit flips we saw in Eq. (478). However, which single bit flips is a complicated function of all the bits, so we still have a high-weight operator and there is no simplification to be gained by using the Gray code. Hence, we restrict our attention to the binary code for simplicity.

The implementation of Trotter steps in Hamiltonian simulation on a qubit-based quantum computer requires mapping of the bosonic Hamiltonian to the qubit spin-subspace, mapping of the bosonic state to the qubits, and then mapping the resulting state back to a bosonic state (during classical postprocessing) for interpretation of the results.

Within the context of Trotter-Suzuki–based simulations, one of the most efficient ways of simulating bosonic systems in qubits is via the decomposition of the Hamiltonian into a sum of one-sparse Hamiltonians. One-sparse Hamiltonians are arguably the broadest family of quantum Hamiltonians that can be directly compiled into quantum circuits with nominally zero error. These Hamiltonians correspond to matrices that have at most one nonzero matrix element in every row and column. Most of the common matrices that we use in quantum computing are one-sparse, including the Pauli operators, single-qubit Z rotations (S, P, and T gates), controlled-NOT, controlled-Z, SWAP, and Toffoli—but not the Hadamard gate.

For bosonic systems, however, many of the gates are not one-sparse when compiled into qubit-only hardware. In the following, we calculate the number of qubit gates required to implement a single bosonic displacement operation in a scalable fashion in an all-qubit device.

Assuming that a boson number cutoff of $K = 2^n - 1$ is used, our strategy is to encode such a bosonic state using n qubits while using one additional qubit to ease the calculation of the square-root factor from bosonic dynamics. As we see in the following, the additional qubit allows us to encode coefficients of adjacent Fock levels into it for easy computation of the square-root operation as multi-qubit controlled single-qubit rotations.

The intuition of why the additional qubit helps to compute the square root is as follows. For the displacement operator, the Hamiltonian of interest is

$$a + a^\dagger = \begin{bmatrix} 0 & 1 & 0 & 0 & \\ 1 & 0 & \sqrt{2} & 0 & \\ 0 & \sqrt{2} & 0 & \sqrt{3} & \cdots \\ 0 & 0 & \sqrt{3} & 0 & \\ & & \vdots & & \ddots \end{bmatrix}, \tag{480}$$

which we would like to encode on qubits. Observing the self-repeating (though overlapping) 2×2 block structure,

$$\begin{bmatrix} 0 & \sqrt{k} \\ \sqrt{k} & 0 \end{bmatrix}, \tag{481}$$

and that, e.g.,

$$X \oplus \sqrt{3}X = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & \sqrt{3} \\ 0 & 0 & \sqrt{3} & 0 \end{bmatrix}, \tag{482}$$

we can write $a^\dagger + a$ as follows:

$$\begin{aligned} a + a^\dagger &= [0] \oplus \bigoplus_{j=1}^{\infty} \begin{bmatrix} 0 & \sqrt{2j} \\ \sqrt{2j} & 0 \end{bmatrix} \\ &\quad + \bigoplus_{j=0}^{\infty} \begin{bmatrix} 0 & \sqrt{2j+1} \\ \sqrt{2j+1} & 0 \end{bmatrix} \qquad (483) \\ &= H_{\text{even}} + H_{\text{odd}}. \qquad (484) \end{aligned}$$

Here, the "even" and "odd" subscripts denote the parity of the value inside the square root. We can therefore simulate the quantum dynamics using gates analogous to x rotations $R_0(-2it\sqrt{2j})$ or $R_0(-2it\sqrt{2j+1})$ (where the single-qubit rotation $R_\varphi(\theta)$ is a rotation through an angle θ about an axis ϕ, formally defined in Table VIII of Appendix A).

In the following, we describe in detail the steps necessary to use the above idea to simulate the displacement operator on qubits, the overview of which is given in Algorithm 1. We denote bosonic Fock states by $|k\rangle_{\text{B}}$, where the non-negative integer k is the boson number. The corresponding representation of this Fock state in the n qubit register is denoted $|\text{binary}(k)\rangle$ without the subscript B.

1. Mapping the bosonic state to 2D subspaces and back

Let us write the truncated bosonic state as a sum over even and odd Fock states:

$$|\psi\rangle_{\text{B}} = \sum_{j=0}^{(K-1)/2} \left(a_{2j} |2j\rangle_{\text{B}} + a_{2j+1} |2j+1\rangle_{\text{B}}\right). \tag{485}$$

First, we assume that it is possible to encode this truncated bosonic state onto a register of n qubits using a binary representation of the Fock states:

$$|\psi\rangle = \sum_{j=0}^{(K-1)/2} \left(a_{2j} |\text{binary}(2j)\rangle + a_{2j+1} |\text{binary}(2j+1)\rangle\right). \tag{486}$$

For example, an arbitrary state with a maximum of three bosons,

$$|\psi\rangle_{\text{B}} = \alpha |0\rangle_{\text{B}} + \beta |1\rangle_{\text{B}} + \gamma |2\rangle_{\text{B}} + \delta |3\rangle_{\text{B}}, \tag{487}$$

would require a register of $n = 2$ qubits to represent it:

$$|\psi\rangle = \alpha |00\rangle + \beta |01\rangle + \gamma |10\rangle + \delta |11\rangle. \tag{488}$$

ALGORITHM 1. The Fock-Binary Simulation of Bosonic Displacement Gate with Qubits.

Input : A bosonic state defined in the Fock basis by $|\psi\rangle_B = \sum_{k=0}^{K} a_k |k\rangle_B$ and an n-qubit register augmented by a single additional qubit (hereafter called the main qubit), prepare the initial state on the register of n qubits as $\left(\sum_{k=0}^{K} a_k |\text{binary}(k)\rangle\right) \otimes |0\rangle$, a displacement $\alpha \in \mathbb{R}$ and an error tolerance ϵ

Output: A quantum state $V(\alpha)|\psi_b\rangle$ such that $\|V(\alpha) - D(\alpha)\| \leq \epsilon$

1 Subdivide Hamiltonian into $H = H_{\text{even}} + H_{\text{odd}}$
2 Provide a Trotter decomposition of the form $V(\alpha) = \prod_{j=1}^{N_{\text{exp}}/2} e^{-i\alpha H_{\text{odd}} \tau_{2j-1}} e^{-i\alpha H_{\text{even}} \tau_{2j}}$ such that $\|V(\alpha) - D(\alpha)\| \leq \epsilon/2$
3 **for** $e^{-iH_p t_p} \in$ terms of $V(\alpha)$*:* **do**
4 Simulate $e^{-iH_p t_p}$ by transferring the quantum amplitudes to the main qubit (parameters given for H_{odd}): $\sum_{j=0}^{(K-1)/2} |2j+1\rangle (a_{2j}|0\rangle + a_{2j+1}|1\rangle)$ for odd cutoff K
5 Compute square-roots using reversible operations within error $\epsilon/(2N_{\text{exp}}\alpha)$;
6 Act with a series of register-controlled ancilla rotations to implement Hamiltonian
7 Transfer amplitudes back to the qubit register
8 **end for**

The preparation of this qubit state is in general not a trivial task; however, we do not consider the complexity of this task here.

Second, by coupling one additional ("main") qubit in the last position, we store the coefficients on the main qubit:

$$|\psi'\rangle = \sum_{j=0}^{(K-1)/2} |\text{binary}(2j+1)\rangle \otimes |\psi'_{2j+1}\rangle, \tag{489}$$

$$|\psi'_{2j+1}\rangle \equiv a_{2j} |0\rangle + a_{2j+1} |1\rangle, \tag{490}$$

such that we can then evolve the system under the displacement Hamiltonian using a single unitary acting in parallel on all 2×2 subspaces, each labeled by the projector $|2j+1\rangle\langle 2j+1|$:

$$\sum_{j=0}^{(K-1)/2} |2j+1\rangle\langle 2j+1| \otimes R_0(-i2t\sqrt{2j+1}), \tag{491}$$

where the first part of the tensor product acts on the n qubit register and the R_0 gate (Table VIII) rotates the main qubit around the x axis.

To prepare $|\psi'\rangle$ in Eq. (490) from $|\psi\rangle \otimes |0\rangle$ for $|\psi\rangle$ given in Eq. (486), we use the knowledge that in binary, an odd number ends in 1: we use a CNOT gate the target of which is the main qubit and the control of which is the least significant bit (LSB) in the n-qubit register. This flips the main qubit if the state in the first n-qubit (system register) represents an odd state. We obtain

$$a\,|2j\rangle\,|0\rangle + b\,|2j+1\rangle\,|0\rangle \mapsto a\,|2j\rangle\,|0\rangle + b\,|2j+1\rangle\,|1\rangle\,. \tag{492}$$

Applying an incrementer operation on the system register, controlled on the main qubit being in $|0\rangle$, we obtain

$$a\,|2j\rangle\,|0\rangle + b\,|2j+1\rangle\,|1\rangle \mapsto |2j+1\rangle\,(a\,|0\rangle + b\,|1\rangle). \tag{493}$$

For the example initial state in Eq. (488), we now have

$$|01\rangle\,(\alpha\,|0\rangle + \beta\,|1\rangle) + |11\rangle\,(\gamma\,|0\rangle + \delta\,|1\rangle)\,. \tag{494}$$

In the following, we calculate the two-qubit entangling gates needed to operate Eqs. (492) and (493) for encoding the coefficients to the main qubit. We need one CNOT gate and a single-qubit controlled n-qubit incrementer.

An incrementer, in which we do not reset the additional carry qubits required for the computation of the incrementation, takes $(n-2)$ Toffoli gates and n CNOT gates (with an $(n-1)$-qubit overhead) [652]. There are multiple ways to implement Toffoli gates, but the standard decomposition uses six CNOT gates [1]. Assuming that decomposition, an incrementer takes $7n-12$ CNOT gates.

To implement a controlled incrementer, we can use an extra register of qubits: we carry out a controlled SWAP between the system register and the extra register, then implement the controlled adder on the extra register, and carry out another controlled SWAP to return the incremented value to the system register. A controlled SWAP gate between the two registers requires n Toffoli and $2n$ CNOT gates. This is because a SWAP gate between qubits i and j can be implemented using $\mathrm{CNOT}_{ij}\,\mathrm{CNOT}_{ji}\mathrm{CNOT}_{ij}$ and this circuit can be transformed into a controlled SWAP by only controlling the CNOT_{ji} gate. This implies that two CNOTs and one Toffoli are needed to do a controlled SWAP, which corresponds to eight CNOT gates. Therefore, with the given construction, a controlled incrementer requires $23n-12$ CNOT gates.

Because we need to map the amplitudes to the main qubit to make the computation and back to the system register to interpret the results, this requires two controlled-NOT gates, one controlled incrementer and one controlled decrementer, and the aforementioned controlled rotation for the implementation of the Hamiltonian. The complexity of a decrementer is the same as that of an adder (up to single-qubit gates).

Therefore, for this construction, the total number of CNOT gates required to transfer two neighboring Fock state amplitudes from the bosonic state (represented in binary in an n-qubit register) onto the main qubit, and return that information to the binary qubit register representing the Fock state is

$$\begin{aligned} N_{\mathrm{CNOT/transfer}} &= 2 \times (23n - 12 + 1) \\ &= 46n - 22. \end{aligned} \tag{495}$$

The system register requires n qubits, where n is the number of qubits required to represent the Fock state up to Fock level K in binary. Although all amplitudes can be stored in a single qubit, each incrementer requires $n-1$ extra qubits, and each controlled SWAP operation requires n extra qubits. We also count the main qubit. Therefore, a total of $n + n - 1 + n + 1 = 3n$ qubits are needed.

2. Computation of the square roots

One of the main costs regarding the entangling gate counts for computation of the dynamics relates to the implementation of the $R_0(-i2t\sqrt{2j+1})$ gate, which maps

$$\begin{aligned} &|\mathrm{binary}(2j+1)\rangle \left|\psi'_{2j+1}\right\rangle \\ &\quad\mapsto |\mathrm{binary}(2j+1)\rangle\, R_0(-i2t\sqrt{2j+1}) \left|\psi'_{2j+1}\right\rangle. \end{aligned} \tag{496}$$

Similarly, we can address the states within the even subspaces by computing instead the square root of $2j$ in analogy to the above mapping. Since the square root needs to be computed for each different value of $2j$ and $2j+1$ in the superposition, we cannot construct a fixed circuit for it. It would be possible to calculate the square roots classically in advance and store them in a register; however, given the number of values we need in cases with large cutoffs, the cost may become prohibitive. Several approaches are commonly taken to implement functions such as this. The first approach directly converts classical networks for these operations into quantum gates through the use of Bennett pebbling [653]. This approach is time-consuming, but more importantly, often requires hundreds of qubits to maintain the result.

The cost of performing the square-root transformation can be upper bounded using the cost of a reciprocal square-root operation. The reciprocal square root can be readily computed using Newton iterations, which converge quadratically, meaning that the number of iterations needed scales only doubly logarithmically with the desired

error tolerance. The square root can then be computed via

$$\sqrt{x} = \frac{x}{\sqrt{x}}. \tag{497}$$

Thus we can use this tactic to take advantage of the rapid convergence of the Newton iterations to the reciprocal square root by following the reversible computation of the reciprocal square root with a single multiplication.

We represent the Fock state value (an integer) using n bits. We denote m to be the number of iterations needed to compute the reciprocal square root; e.g., results with 10^{-8} error have been obtained using $m = 3$ Newton iterations [654]. In the following, the iterates are no longer integers; therefore, we denote p as the point precision and choose to set it to n such that there are as many qubits representing the number before the floating point as there are qubits representing the number after. As in fixed-point arithmetic, we now write the formulas with the new number of bits $n' = 2n$.

The cost of computing a square root in terms of the number of Toffoli gates (without uncomputing the iterates) is [654]

$$n'^2\left(\frac{15}{2}m + 3\right) + 15n'pm + n'\left(\frac{23}{2}m + 5\right) - 15p^2m + 15pm - 2m. \tag{498}$$

Rewritten in terms of n, we now have

$$45n^2m + 12n^2 + 38nm + 10n - 2m. \tag{499}$$

The cost of a multiplication of two n-bit numbers with point precision p is [654]

$$N_{\mathrm{CNOT}/\times} = \frac{3}{2}n'^2 + 3n'p + \frac{3}{2}n' - 3p^2 + 3p, \tag{500}$$

which in terms of n is

$$9n^2 + 6n. \tag{501}$$

If we use the standard six CNOTs per Toffoli gate construction [1], then the number of CNOT gates needed to implement the square-root calculation

$$|x\rangle|0\rangle \mapsto |x\rangle\left|\sqrt{x}\right\rangle|\psi_{\mathrm{anc}}\rangle \tag{502}$$

is

$$N_{\mathrm{CNOT/SQRT}} = 6(45n^2m + 12n^2 + 38nm + 10n - 2m) + 6(9n^2 + 6n) \tag{503}$$

$$= (270m + 126)\,n^2 + (228m + 96)\,n - 12m. \tag{504}$$

If m iterations are needed for the computation of the reciprocal square root, m qubit registers will be needed to store the values of the reciprocal square root in addition to any ancillas needed in the implementation of any binary adders that appear in the construction. Each Newton iteration requires three ancilla registers (which are cleaned up after each round) as well as a register to hold the initial guess, totaling $(m + 4)n'$ qubits. Therefore, the total number of ancillas needed to compute the square root to n' bits of precision is

$$n_{\mathrm{qubits}} = n' + (m + 4)n' \tag{505}$$

$$= 2n(m + 5), \tag{506}$$

as an additional register is required to hold the initial value to carry out the multiplication giving the square root.

3. Controlled rotations (CRs)

Having calculated the square roots, controlled rotations must be employed to apply the rotation. That is, if $(\sqrt{x})_j$ represents the value of the jth bit in the fixed-point representation of the square root in increasing order of significance, then

$$e^{-i\sqrt{x}\sigma_x t} = He^{-i\sqrt{x}\sigma_z t}H = H\prod_{j=0}^{n-1} e^{-i(\sqrt{x})_j 2^j \sigma_z t}H. \tag{507}$$

Thus we can implement $e^{-iH_{\mathrm{even}}t}$ or $e^{-iH_{\mathrm{odd}}t}$ by applying this circuit to the ancilla qubit in Eq. (490). A controlled single-qubit rotation can be implemented using two controlled NOT gates and single-qubit rotations. Thus the cost of the simulation of this term would also involve another $2n$ single-qubit rotations and CNOT gates,

$$N_{\mathrm{CNOT/CR}} = 2n. \tag{508}$$

This cost is, however, negligible in most settings compared to the $O(n^2)$ Toffoli gates needed to compute the square root.

4. Cost of simulating a single exponential

Combining Eqs. (495), (504), and (508), the number of two-qubit operations sufficient to implement a single step $e^{-iH_p t_p}$ in Algorithm 1 is

$$\begin{aligned} N_{\mathrm{CNOT}} &= N_{\mathrm{CNOT/H_{even}}} + N_{\mathrm{CNOT/H_{odd}}} \\ &= 2N_{\mathrm{CNOT/transfer}} + 4N_{\mathrm{CNOT/SQRT}} + 2N_{\mathrm{CNOT/CR}} \\ &= (270m + 126)n^2 + (144 + 228m)n - 12m - 22, \end{aligned} \tag{509}$$

where n is the size of the qubit register required to represent the largest initial Fock state in binary, and m is the number of Newton iterations required to calculate the inverse square root.

The number of qubits required for a system with the largest Fock state N, requiring n bits for binary representation, and m Newton iterations would be

$$N_{\text{qubits}} = 2n(m+5) + 3n. \tag{510}$$

If we wish to represent all occupations up to a Fock state cutoff of, e.g., $K+1 = 64$, a seven-qubit register is required for the binary representation: $n = 7$. In the case in which two Newton iterations $m = 2$ are sufficient for our precision goals—the level of approximation error that this corresponds to due to Newton iterations is on the order of 10^{-4} [654]—this evaluates to

$$N_{\text{CNOT}} = 36\,788 \tag{511}$$

and

$$N_{\text{qubits}} = 119. \tag{512}$$

In order to calculate the full cost of a single displacement, we need to find the number of exponentials needed in the Trotter-like decomposition in Algorithm 1. The number of exponentials is determined by the Trotter error, which is difficult to estimate accurately, as the error depends sensitively on the input state and the signs of the commutators that arise in a Trotter-Suzuki decomposition. As a first (most likely grossly overestimated) estimate, the number of exponentials N_{exp} needed in a Trotter-Suzuki formula of order $2k$ to ensure that the error is $\epsilon/2$-small is [655]

$$N_{\text{exp}} \leq \left\lceil 3K\alpha k \left(\frac{25}{3}\right)^{k} \left(\frac{2\alpha K}{\epsilon}\right)^{1/2k} \right\rceil. \tag{513}$$

From Box 4.1 of Nielsen and Chuang [1], we then see that if we denote the idealized case in which the square roots are taken to be exact as $V_0(\alpha)$, and take the error in the evaluation of the square roots to be $\epsilon/(2N_{\text{exp}}\alpha)$, we have that the total error for approximating a displacement $D(\alpha)$ by a Trotter-approximated and square-root-approximated unitary $V(\alpha)$ is bounded above by ϵ from our above error choices and the following argument:

$$\begin{aligned} \|V(\alpha) - D(\alpha)\| &\leq \|V(\alpha) - V_0(\alpha)\| + \|V_0(\alpha) - D(\alpha)\| \\ &\leq \frac{\epsilon}{2} + N_{\text{exp}}\alpha \left(\frac{\epsilon}{2\alpha N_{\text{exp}}}\right) \leq \epsilon, \end{aligned} \tag{514}$$

where the first $\epsilon/2$ is the Trotter error and the second term is the square-root approximation error. For the latter, we have multiplied the per-exponential error $\epsilon/(2N_{\text{exp}}\alpha)$ by the number of Trotter exponentials N_{exp} and the magnitude of the displacement α.

5. Cost of simulating the displacement operator

Putting everything together, we now calculate the CNOT cost of a single displacement. It is given by the product of our estimate of the number of exponentials in Eq. (513) with the number of CNOT gates per exponential in Eq. (509). The latter implicitly depends on the number of exponentials due to the fact that the number of Newton iterations per square root needs to be smaller than $\epsilon/(2N_{\text{exp}}\alpha)$. We therefore calculate the total CNOT cost as follows. First, we calculate the number of exponentials for a given cutoff, displacement amount α, target error ϵ, and Trotter order k from Eq. (513). We then determine the target square-root error by $\epsilon/(2N_{\text{exp}}\alpha)$. Finally, we determine the number of Newton iterations m from the numbers shown in Ref. [654]: for a square-root error of below 10^{-4}, 10^{-6}, 10^{-12}, $m = 2, 3, 4$ is required. Finally, we total number of CNOT gates is the product of our estimate of the number of exponentials in Eq. (513) with the number of CNOT gates per exponential in Eq. (509).

We show the resulting estimate in Fig. 41 as a function of the cutoff K. We find that more than a million CNOT gates are needed in our estimate even for small cutoffs, and an approximately linear increase in the number of gates with cutoff. We also find a significant gap between our gate estimate and the worst-case brute-force decomposition CNOT upper bounds in Ref. [656]. This implies that our estimate is far from optimal. Most importantly, our estimate for the number of Trotter steps needed most likely is a gross overestimation and a proper use of commutator bounds would improve on this [339].

Finally, our gate count estimate does not imply that this represents a fundamental gap between qubit-only and qubit-oscillator gate sets. For example, Ref. [647] has shown more efficient methods which, however, were not scalable and did not result in significantly better results than brute-force decomposition. Instead, our estimate shows that there are some examples where a straightforward application of qubit compilation techniques leads to a cost that is far greater than that needed by the qubit-oscillator system. Regardless, the large gap between the two suggests that applications likely exist where the costs of performing a computation in this hybrid instruction set are lower than that of a traditional qubit-based computer.

C. Efficiency of bosonic instruction set relative to qubit-only simulations

As we have seen in Sec. VIII B, the number of CNOT gates required to carry out a single displacement on a single oscillator with maximum Fock state 63 (requiring six qubits), in the case in which two Newton-Raphson iterations $m = 2$ (which corresponds to errors on the order of 10^{-4} [654]) are sufficient for our precision goals, is $N_{\text{CNOT}} = 36{,}788$, using the described procedure. In contrast, the number of qubit-controlled operations needed in

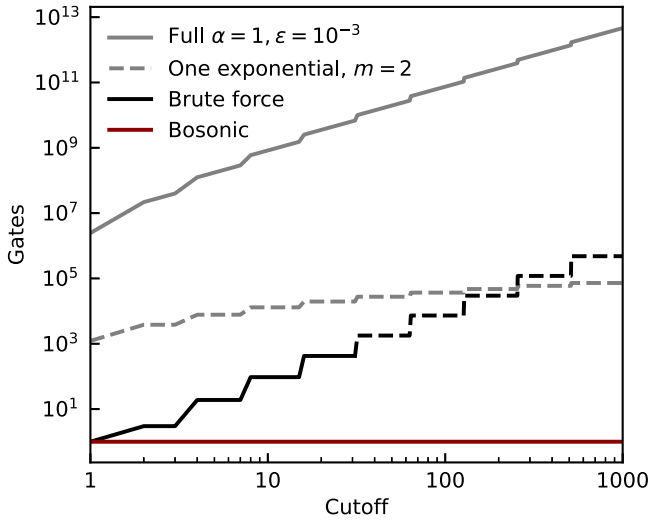


FIG. 41. The CNOT gate count of implementing a displacement operation with qubit-only quantum computers in the Fock-binary encoding. The dashed gray line represents the CNOT count per exponential in Eq. (509) which appears in our estimate for the displacement. It consists of transferring the amplitude from the n-qubit system register to a single main qubit and back, Eq. (495), and the cost of calculating square roots using Newton iterations, in this case for $m = 2$ iterations. The gray line is our full gate estimate for a displacement for target error $\epsilon = 0.1\%$ and $\alpha = 1$. The black line is the currently best-known brute-force decomposition algorithm (see Ref. [656]). It constructs the $2^n \times 2^n$ unitary corresponding to the displacement and decomposes it into two and one gates. From about five qubits, the classical resources to actually perform the compilation are prohibitively expensive, but in this regime the algorithm still provides an upper bound to how many gates are actually required, indicated by a dashed line. The red line is at cost 1, representing the single displacement gate needed in bosonic hardware. This unit circuit-depth metric does not take into account the fact that the time cost of the displacement is linear in the size of the displacement. We also neglect here the subtle effects of noise in the complexity considerations (see discussion in Sec. VIII A 1).

the bosonic instruction set is 1, because this operation corresponds to a native gate in the instruction set. In addition to the circuit complexity issues discussed here, the reader is reminded of the important complications associated with qubit damping errors in the representation of oscillators with qubits previously discussed in Sec. VIII A 1.

The above analysis demonstrates a significant overhead to simulate a hybrid instruction set using qubit-only hardware. However, it is by no means optimal. Alternative mappings exist that may privilege efficiency in terms of gate count rather than, e.g., qubit number (see Table VII). In contrast to the Fock encoding used above, some encodings (e.g., Jordan-Lee-Preskill [576]) privilege representation of oscillator position and momentum and thus might be better for the particular case of simple displacements, but worse for representing interactions best described in the Fock basis. Many encodings have been compared in a comprehensive review [647] with optimized qubit gate counts, and a method for decomposing arbitrary continuous-variable quantum operations via linear combinations of commutation relations is described in Ref. [335]. Reference [657] has discussed block-encoding bosonic operators into qubits using the Jordan-Lee-Preskill representation. Reference [658] has discussed optimizing the efficiency of qubit simulations of bosonic lattice gauge theories.

To better reveal the value of hybrid CV-DV instruction sets, a systematic study of protocols to simulate bosonic instruction sets using qubits in terms of entangling gates, qubit count, and number of samples required, and conversely using bosonic instruction sets to simulate qubits, needs to be performed. Furthermore, theory-experiment collaboration to include realistic gate infidelities in these analyses is another important aspect to facilitate device engineering. Lastly, cross-platform simulation at the system and algorithm level (beyond the gate level) targeting specific applications such as Hamiltonian simulation will deliver a more practical understanding of the advantages and limitations of such hybrid qubit-bosonic hardware.

IX. CONCLUSIONS AND OUTLOOK

The primary focus of this work is pedagogical introduction of explicit and experimentally relevant instruction set architectures (ISAs) and abstract machine models (AMMs) for hybrid CV-DV quantum processors composed of qubits and quantized bosonic modes. We start from the physical behavior of bosonic states and systematically develop bosonic gate models, and hybrid architectures including error-correction strategies, compilation theory and techniques, as well as high-level proxy algorithms and applications. We have presented a few examples illustrating how such hybrid quantum processors can be utilized to solve practical problems, including DV-CV state-transfer, the QFT, bosonic quantum error correction, and quantum simulation of matter with bosonic degrees of freedom. We show for the latter that hybrid digital architectures may provide substantial computational advantages over existing approaches that use only qubits because of the ability to avoid costly arithmetic operations in favor of equivalent native bosonic operations. Below, we close with some perspective on challenges and opportunities and thoughts on directions for the future.

A. Challenges and opportunities for AMMs and ISAs

The challenges of, and opportunities for, defining meaningful AMMs and designing effective ISAs for hybrid quantum CV-DV arise concerning five aspects: (i) the choice of how quantum information is represented using bosonic states, (ii) the universal bosonic gate set employed,

(iii) the variety of hardware and instruction set architectures available for hybrid oscillator-qubit systems, (iv) effective compilation methods, and (v) use cases for hybrid quantum processors. In this work, we have presented and analyzed AMMs and ISAs for hybrid CV-DV quantum hardware by addressing these five aspects. We touch upon each of these below.

Unlike a qubit, which admits a natural representation analogous to a binary bit, bosonic systems offer many different potential representations for encoding state information. Since the bosonic system is a simple harmonic oscillator, the position or momentum of the state can be natural state variables, with continuous values. Alternatively, since the energy of the oscillator is quantized, the Fock basis given by the discrete number of energy quanta may also be a natural choice. Understanding and taking advantage of these state representation options can be complex, particularly when it is desirable to be able to map quantum information between oscillators and qubits, as would be natural within a hybrid CV-DV AMM.

Section III has discussed how quantum logic gate sets acting on bosonic systems are much richer than those for qubits. Available "analog" bosonic gates include not just operations on single oscillators such as displacement and squeezing, but also mixing operations such as beam splitters, between multiple oscillators. And despite their analog nature, the simplest bosonic gates generated by quadratic oscillator Hamiltonians can be understood as being directly analogous to Clifford gates, in that they form a discrete subspace of unitary transforms governed by the Gottesman-Knill theorem and are easy to simulate classically on Gaussian input states. But many bosonic gates, and especially gates that couple bosonic systems to a qubit, are non-Clifford.

Noise models and QEC constructs also depend on the expressibility of a chosen gate set. It is thus challenging to choose appropriate gate sets to make meaningful ISAs for hybrid quantum CV-DV systems. The ISAs that we present are quite small and can each be viewed as a kind of RISC ISA [659] having a small number of instructions defined with a set of continuous parameters accompanying them. Such RISC instruction set architectures greatly ease the problem of translation down the stack from applications at the top to the calibrated microwave pulses that need to be applied to the hardware at the bottom.

Efficient compilation of desired unitary transformations into the instructions for a hybrid CV-DV architecture is also challenging, and two routes are generally available. First, numerical optimization of the parameters for sequences of instructions can be performed to effect the circuit compilation. This is computationally feasible for defining "assembly language" style instructions at small scales for portions of the hardware but is completely infeasible for compiling entire algorithms, and results in circuits that are not "human-readable" and provide little intuition for generalization of hybrid CV-DV hardware utility. A second route is via analytical circuit synthesis. This is possible for specific families of hybrid CV-DV circuits, such as QSP, as we have shown. While the resulting circuits are not necessarily as efficient as numerically synthesized circuits, they are human-readable and provide significant insights into the power of hybrid hardware. They also may prove useful as "seed" circuits and motifs for further numerical optimization.

The fifth challenge for these AMMs and ISAs is their use case: for what end-applications might hybrid quantum CV-DV be particularly well suited? As we have discussed, systems composed of spins and oscillators are natural candidates for quantum simulation with hybrid CV-DV hardware, due to the expected efficiency of mapping possible, compared with using solely qubit-based hardware, or solely oscillator-based hardware. The natural physical dynamics of oscillators also present intriguing possibilities to be harnessed, as the Hamiltonian for a simple harmonic oscillator is known to evolve position eigenstates into momentum eigenstates, embodying a Fourier transform. We have studied both quantum simulation and the quantum Fourier transforms using hybrid hardware in this paper. More applications will arise, and we expect their discovery to be accelerated by the formulation of accessible abstract machine models and instruction set architectures for hybrid quantum CV-DV hardware—precisely our goal here.

The physical-level interactions that enable the proposed ISAs and AMMs can be realized on quantum hardware containing both qubits and bosonic degrees of freedom, and we have discussed here superconducting, trapped-ion, and neutral-atom platforms. Dissipation and quantum noise on such hardware can be corrected or suppressed with bosonic quantum error-correcting codes (potentially further concatenated with qubit codes). These physical interactions allow us to define several instruction sets summarized in Table V, including among others, the Fock-space, phase-space, and sideband ISAs. We have also shown explicitly how instructions from these ISAs can be composed to achieve universal hybrid CV-DV quantum computation and further demonstrate how to characterize the quantum device with a CV analog of gate set tomography.

Systematic compilation techniques compose these universal gates into useful primitives and algorithms that solve practical problems, as we have discussed in Sec. VI. These compilation techniques include strategies for compiling hybrid single-qubit oscillator unitaries, which are obtained by generalizing the well-known Trotter-Suzuki and product formulas, QSP, and LCU techniques from DV systems to hybrid CV-DV systems. We have presented exact analytical gate synthesis methods for multi-qubit entangling gates and multi-oscillator entangling gates in superconducting systems that are inspired by the Mølmer-Sørensen gate for ion traps.

At the system level, we defined three key quantum AMMs: (1) qubit-centric, (2) oscillator-centric, and (3) hybrid oscillator-qubit, which all have universal computation capabilities. These AMMs can represent abstractions of the machine for a high level of the stack that is user-facing, defining the mathematical abstract Hilbert space where the computation happens. Alternatively, they can be lower-level abstractions that are transpiler-visible, presenting to the transpiler what physical degrees of freedom are available as computation components and what their connectivity and allowed gates are. The formal establishment of these AMMs should prove particularly useful for resource estimation on hybrid CV-DV quantum processors.

We have further examined applications of quantum computing, such as quantum simulation, illustrating how quantum electrodynamics and related lattice gauge Hamiltonians can be compiled into hybrid qubit-bosonic operations. We find evidence of advantage for near-term applications of hybrid systems over standard methods for purely qubit-based simulations. In addition, we have presented extensions of quantum signal processing to the non-Abelian case, for which the qubit rotation angles are functions of the noncommuting position and momentum variables of bosonic modes. With this tool in hand, we have examined applications of quantum signal processing within these architectures and discussed explicit approaches to implement the quantum Fourier transform.

In conclusion, the ISAs and AMMs proposed here significantly expand our understanding of how distinct quantum resources can be leveraged in a unified computational framework and demonstrate the potential of DV-CV hybrid quantum processors for the upcoming post-NISQ era of early-stage fault tolerance.

B. Open questions and directions for future work

Despite these new understandings and achievements, there are still many open questions ranging from the fundamental side to more practical aspects of hybrid DV-CV quantum systems that remain to be addressed in the future.

First and foremost, as we have seen in Sec. VII, hybrid quantum hardware opens many possibilities for interesting applications, but it is worth approaching the potential with careful consideration. After all, classical digital computation has advanced just fine without needing assistance from embedded classical harmonic oscillators. But quantum computing hardware is much less mature, and early applications might benefit from hardware capabilities which are closer to application models. Identifying more use cases in which CV-DV hardware can really deliver an advantage over qubit-only or classical computers will be important. Given NISQ needs, examples of such applications likely include quantum simulation of coupled nuclei-electron processes [660,661] or light-matter interactions [662,663] relevant to chemical reactions and materials science applications. As another example, the hybrid CV-DV quantum signal processing framework established in the current work serves as a potentially ideal tool for processing quantum signals. Classical signal processing techniques from electrical engineering for filter design and classical CV-DV conversion [664] can provide much intuition and guidance for further developments in the quantum case.

And while the major quantum algorithms become united in the framework of quantum signal processing and quantum singular-value transforms, the newly emerging hybrid CV-DV quantum signal processing techniques extend possibilities to the non-Abelian case, for which a complete theory remains to be discovered. From a NISQ viewpoint, further development is required for variational CV-DV algorithms beyond Sec. VI E that can be run on such hybrid processors [529]. Such a set of tools would dramatically increase the scope of the signal-processing transformations that can be implemented on input parameters and allow more powerful analytical compilation methods for general multi-variable transformations based on unitary transformations of bosonic modes. While we describe some QSP methods in this work, the presentation is mathematically abstract, and higher-level CV-DV quantum programming languages such as BOSONIC QISKIT [52] are needed to simplify the programming process for a broader community of users.

We have presented various CV-DV ISAs based on discrete sets of control Hamiltonians but continuously parametrized gates. An open question in this direction is to address the complexity of generating arbitrary unitaries in CV systems. While the Trotter-Suzuki–based approaches considered here and in Ref. [110] can provide subpolynomial scaling with the error tolerance for the synthesized gates, existing approaches fall short of the scaling provided by circuit synthesis methods such as the Solovay-Kitaev theorem [665,666] which achieves polylogarithmic scaling in the error tolerance when compiling over a discrete gate set. The development of better CV gate synthesis methods may not only open up new possibilities in compiling unitaries but also open up the door to the development of bosonic error-correcting codes where the physical and logical operations are natural bosonic operations such as discrete displacement or squeezing operations.

The non-Abelian QSP techniques used in this work involve composite pulses in phase space designed to cancel first-order errors. Generalizations of these composite pulses similar to Refs. [274,493,512] are required to develop optimal sequences for arbitrary CV Hamiltonian construction. Formally, this task is equivalent to designing a constructive algorithm that can find optimal non-Abelian QSP sequences for arbitrary polynomials $f(\hat{x},\hat{p})$ [278]. In Ref. [667] the authors find a constructive algorithm to generate arbitrary polynomials for the case of single-variable QSP. The generalization of this scheme to QSP using two

commuting variables has been shown in Ref. [514]. Such a construction for non-Abelian QSP extends to the case of multi-variable QSP with noncommuting variables, and thus, remains an open question. The multiple noncommuting variables for the case of the phase-space ISA will be the numerous quantum phase-space vectors $\hat{v} = \alpha\hat{x} + \beta\hat{p}$. The examples from Ref. [478] discussed in Secs. VI D 3 and V C belong to a subclass of this framework, QSP with just two noncommuting variables. A hierarchy of various QSP frameworks developed with an eye toward the hybrid phase-space instruction set can be found in Ref. [478].

A related aspect of algorithm design is to identify new useful universal instruction sets for such hybrid quantum processors. In particular, we have seen in Sec. III C 4 that the stellar representation provides an intuitive picture of how CV states and operations can be characterized in terms of the roots of complex analytical functions and their associated time dynamics. Designing an ISA that works directly within the stellar representation picture would be very interesting and useful.

From an architectural perspective, we need more complete ISAs that include a memory hierarchy and sophisticated quantum algorithmic and arithmetic logic units (ALUs). In terms of memory, a mixed system of qubits and bosonic modes provides opportunities by balancing the fast gates and high connectivity of qubits with the long lifetimes and novel gate sets of bosonic modes. The needed high connectivity can be obtained through quantum communication using bosonic modes, e.g., by the cavity SWAP gates discussed in Sec. VI C and/or through use of the LAQCC model described in Secs. II B and IV F 1 to efficiently produce qubit GHZ states that can be used for quantum teleportation.

At the simplest level, recent architectures [167] have proposed using bosonic modes as small register memories, embedding them in architectures specifically designed to match the connectivity requirements of algorithms. Recent work with multimode bosonic systems has shown their utility as memories for factoring [668] and unique error-correction architectures [669]. This is particularly important if proposed memory hierarchies are to be realized [141] as in classical systems. Building upon these memory hierarchies, more sophisticated quantum ALUs that take advantage of the hybrid CV-DV gates (Tables III and IV) and the CV nature of oscillator dynamics need to be developed. Section VII has demonstrated key quantum primitives including the Fourier transform and Hamiltonian simulation, but creating more advanced primitives and proxy applications, such as modular exponentiation and matrix arithmetic, remains to be done. Instruction- and processor-level parallelism often used in modern classical computers [43] could also be deployed based on a more sophisticated ISA and memory hierarchy.

Qubit gate fidelities can be efficiently measured using randomized benchmarking and related techniques. To date, however, there are no corresponding systematic methods for CV operations, due to the absence of unitary 2-designs for CV systems. While we provide in Sec. IV F 2 an approach to gate set tomography using Wigner function tomography (see Appendix E 6), this approach cannot be generalized to arbitrarily high dimensions without making strong assumptions about the quantum channels being learned. Such channel tomography estimates will be necessary for us to assess the degree to which these systems will need quantum error correction.

Regarding the topic of noise and reliability, we have discussed examples of bosonic error-correction codes that encode discrete logical qubits into oscillators. Creation and quantum control of such bosonic logical qubits requires non-Gaussian operations that in turn require the presence of a transmon or other auxiliary qubit or anharmonic element which typically has poorer coherence properties than the oscillator they are controlling. This raises fundamental fault-tolerant control and measurement challenges which are only now beginning to be addressed [180,233,239,422,423,670] and many opportunities remain for further progress. An interesting future direction would be to build fault tolerance into the cost function when designing numerically optimized quantum control sequences.

To bring out the role of bosonic error correction in terms of fault tolerance, we have discussed this direction in Sec. V. We discuss how oscillators can be useful in improving the cost of magic state distillation, a resource-intensive component of FTQC. So far, all versions of QEC have deemed magic state distillation to be a necessary component to implement non-Clifford gates, an essential requirement for quantum advantage. We have pointed out in Sec. V B that oscillator-based codes can quadratically reduce the overhead cost associated with this FTQC component by changing the noise profile of the underlying qubit [471]. In addition, we have highlighted that there are efficient ways to realize an error-corrected qudit in an oscillator. Qudits can be helpful in various ways, including reduction of the computation cost and one may expect that low error qudits will be high in demand in the future. The ability to engineer high-fidelity qudits via error correction in the oscillator also takes improvement in magic state distillation several steps further. It has been shown that qudits can directly improve the resource overhead of this preparation by several orders of magnitude by directly enabling more efficient distillation schemes [671,672].

We also briefly discussed the more challenging question of error correction of oscillators (*qua* oscillators, and not as logical qubits). So far there has been only one attempt at such a construction, given by Ref. [462], which corrects errors under a classical displacement channel. Soon after, in Ref. [465] it has been shown that code families designed under this noise model assumption do not have a threshold. Assuming that the true noise model of an oscillator is dominantly photon loss, it remains an open

question if there is an oscillator code family that can protect oscillators against this more realistic noise. Formally, if we introduce a cutoff on the boson number, the problem reduces to qudit error correction for large qudit dimension d, but it seems extremely challenging to efficiently preserve the oscillator states and operator algebra for large d. If such a construction exists, it would be a revolutionary step in reducing the complexity of many quantum simulations and CV algorithms potentially yielding a quantum advantage. However, as analogous ideas for classical computing have failed to deliver advantage thus far [673–675] a complete investigation of all resources needed to achieve a computational advantage in near-term quantum devices will likely be needed to more fully understand the benefits and drawbacks of this approach.

Finally, while this effort seeks to build bridges translating between the languages and worlds of physics and computer science, the core ideas of compilation techniques need to evolve further beyond those presented here. We have demonstrated how techniques such as the Trotter-Suzuki formula and an LCU can be used to achieve effective compilation of single-qubit–oscillator operations and multi-qubit–oscillator entangling gates (Sec. VI). However, more work is required to develop additional compilation strategies. Recently, ZX-calculus—a useful graphical tool for reasoning about DV circuits [676,677]—has been formulated for CV quantum circuits [678]. An extension to hybrid CV-DV circuits constitutes an interesting future direction. Separately, as already noted, a complete theory of non-Abelian hybrid QSP does not yet exist, but its formulation would be very helpful for analytical compilation. More generally, a complete formal complexity-theoretic analysis of hybrid CV-DV circuit synthesis for tasks such as state-transfer and general unitary synthesis does not yet exist. In Appendix D we present cross-compilation techniques among different pairs of ISAs, but much work remains to be done to improve the efficiency and practicality of cross-compilation.

At a fine-grain level, compilation techniques capable of identifying long idle times and inserting SWAPS to memory are crucial. This may be performed in conjunction with the use of flag qubits [679] or single-shot error mitigation [680]. Furthermore, hybrid architectures will require transpilation techniques that map programmatic elements with long idle times to qubits near memories. Compilers should also be adapted to take advantage of the new sets of native, low-cost gates afforded by bosonic hardware. More practically, it would be very useful to create a software library or toolkit built on top of these analytical and numerical compilation techniques that can perform application-tailored resource estimation in terms of appropriate metrics such as the gate count in hybrid ISAs.

Looking further into the future, optimization of quantum compilation using quantum computers [531] is an exciting possibility. Another potentially important direction that would apply to algorithm development for both traditional and hybrid hardware would be the extension of the computation model for quantum computers beyond algorithms based on purely unitary operations and final state measurements to more general quantum channel operations. Such more general quantum channels for algorithms could be created by including engineered dissipation and/or midcircuit measurements and feedforward.

The rate of current experimental and theoretical progress in the field is astounding, yet many problems, both fundamental and engineering, remain to be solved before we fully escape the NISQ era and move into the era of useful, practical, and fault-tolerant quantum computation and simulation. We hope that the hybrid architectures and compilation methods described in this work will help accelerate progress across multiple platforms toward this important goal.

ACKNOWLEDGMENTS

This research was funded by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Co-design Center for Quantum Advantage (C^2QA) under Contract No. DE-SC0012704 (Basic Energy Sciences, PNNL FWP 76274). C^2QA led this research. S.S. and S.G. acknowledge support by the Army Research Office (ARO) under Grant No. W911NF-23-1-0051. J. S.-S. was supported in part by the Army Research Office under the CVQC project W911NF-17-1-0481. J.S.-S. and Y.L. were supported in part by NTT Research. Y.L. acknowledges startup funding support from North Carolina State University. A.S. acknowledges funding from the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, and Quantum Systems Accelerator. J.M.M. acknowledges support from the National Science Foundation Graduate Research Fellowship under Grant No. 2141064. M.B.S. was supported by the Yale Quantum Institute Postdoctoral Fellowship and the Office of the Vice Chancellor for Research and Graduate Education at the University of Wisconsin-Madison with funding from the Wisconsin Alumni Research Foundation. I.L.C. was supported in part by the NSF Frontier Center for Ultracold Atoms (Grant No. PHY-2317134). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Office (ARO), or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein. The authors acknowledge helpful discussions with Jacob Curtis, Michael DeMarco, Shruti Puri, Yongshan Ding, Robert Schoelkopf, Cindy Regal, Baptiste Royer, Timothy Stavenger, and Ang Li, and thank Luke Bell for helpful discussions and verification of some of the key equations presented in the compilation section.

External interest disclosure: S.M.G. receives consulting fees and is an equity holder in Quantum Circuits, Inc.

DATA AVAILABILITY

The data that support the findings of this article are openly available [681].

APPENDIX A: PHYSICAL IMPLEMENTATIONS

Here, we present a physical-level discussion of quantum control of hybrid systems consisting of bosonic modes and qubits, which serves as background information for the hybrid gates presented in Sec. IV D. Table VIII summarizes all qubit-only gates available in superconducting circuits and ion-trap systems and Table IX presents the microscopic Hamiltonian terms freely available in these systems.

1. Superconducting and cavity QED

a. Fundamentals

Let us consider a single oscillator coupled to a single transmon qubit via the Hamiltonian [28,29,77,682]:

$$H = H_0 + H_1 + V_0, \tag{A1}$$

$$H_0 = \bar{\omega}_R a^\dagger a + \bar{\omega}_Q b^\dagger b, \tag{A2}$$

$$H_1 = -\frac{\bar{K}}{2} b^\dagger b^\dagger b b, \tag{A3}$$

$$V_0 = g(a^\dagger + a)(b^\dagger + b). \tag{A4}$$

Here, H_0 is the Hamiltonian of the resonator with "bare" (i.e., uncoupled) frequency $\bar{\omega}_R$ and the transmon with bare frequency $\bar{\omega}_Q$ and H_1 describes the (bare, negative) anharmonicity $-\bar{K}$ of the transmon. The microscopic Hamiltonian of the transmon involves a cosine of the phase difference across the Josephson junction. H_1 is an approximation to the transmon Hamiltonian that results from expanding the cosine to fourth order and making a RWA [682]. Because of the RWA, H_1 is diagonal in the Fock (excitation number) basis and the eigenstates of $H_0 + H_1$ are unchanged by H_1. We still have harmonic oscillator eigenstates, but for the transmon, the energy-level spacing is not uniform due to the presence of the anharmonicity from H_1. The spectrum of the exact Hamiltonian is, of course, bounded from below. The approximate Hamiltonian containing H_1 is not bounded from below, implying that it must eventually break down at large excitation numbers where the truncation of the cosine potential of the Josephson junction ceases to be a valid approximation. We ignore this complication here.

V_0 is the coupling between the charge polarization (dipole moment) of the transmon and the cavity electric field. The coupling constant g is known in the cavity- and circuit-QED literature as the vacuum Rabi coupling [28,29,77,78].

If we assume that the transmon is always within the manifold spanned by its ground state $|g\rangle$ and its first excited state $|e\rangle$, we can make the two-level qubit approximation, mapping the transmon operators onto spin operators. The mapping of the lowest two transmon states

TABLE VIII. Common qubit-only gates. The single-qubit gates are typically natively available. If not natively available, the multi-qubit gates can be readily synthesized through their coupling to the bosonic modes, as discussed in Sec. VI C 2.

		Gate name	Parameters	Definition
Qubit-only	Single-	Qubit x, y rotation[a]	$\theta \in [0, 4\pi), \varphi \in [0, 2\pi)$	$R_\varphi(\theta) = \exp\left(-i\frac{\theta}{2}\sigma_\varphi\right)$ $\sigma_\varphi = \sigma_x \cos\varphi + \sigma_y \sin\varphi$
		Qubit z rotation[a]	$\theta \in [0, 4\pi)$	$R_z(\theta) = \exp\left(-i\frac{\theta}{2}\sigma_z\right)$
	Multi-	ZZ interaction	$\theta \in [0, 2\pi)$	$ZZ(\theta) = \exp\left(-i\frac{\theta}{2}\sigma_z \otimes \sigma_z\right)$
		CNOT (aka CX)		$\frac{I+\sigma_z}{2} \otimes I + \frac{I-\sigma_z}{2} \otimes \sigma_x$
		CPHASE	$\varphi \in [0, 2\pi)$	$\frac{I+\sigma_z}{2} \otimes I + \frac{I-\sigma_z}{2} \otimes e^{i\varphi\frac{I-\sigma_z}{2}}$
		CZ	$\varphi = \pi$	$\frac{I+\sigma_z}{2} \otimes I + \frac{I-\sigma_z}{2} \otimes \sigma_z$
		iSWAP	$\theta \in [0, 2\pi)$	$\exp\left(i\frac{\theta}{2}(\sigma_x\sigma_x + \sigma_y\sigma_y)\right)$
			$\theta = \pi/2$	$\exp\left(i\frac{\pi}{4}(\sigma_x\sigma_x + \sigma_y\sigma_y)\right)$
		fSim	$\theta \in [0, 2\pi), \varphi \in [0, 2\pi)$	iSWAP$(-\theta) \cdot$ CPHASE(φ)
		fSWAP		$i\left[R_z\left(\frac{\pi}{2}\right) \otimes R_z\left(\frac{\pi}{2}\right)\right]$ fSim$\left(\frac{\pi}{2}, 0\right)$

[a]In general, qubit rotations about any axis of the Bloch sphere are given by $R_{\hat{n}}(\theta) = \exp\left(-i(\theta/2)\hat{n} \cdot \vec{\sigma}\right)$, where $\vec{\sigma} = \left(\sigma_x, \sigma_y, \sigma_z\right)$ and $\hat{n}$ is a unit vector, and this gate can be implemented directly on some physical qubits with x-y-z control. However, for many qubits (such as fixed-frequency transmons), only x-y control is native, and the z rotations are implemented *in software* [299]; hence our choice to separate qubit rotations into either x-y or z rotations. fSim (Sec. VII C 8) is the fermion simulation gate [641], while fSWAP (Sec. VII C 8) is the fermionic SWAP gate, which swaps the state of two qubits and adds a minus sign to the state $|11\rangle$ to account for antisymmetric fermionic statistics.

TABLE IX. Natively available interactions on various quantum computing platforms. Not listed are the Gaussian oscillator-only gates available in the superconducting platform: displacement, single-mode squeezing, two-mode squeezing, two-mode SUM gates, and two-mode beam-splitter operations are all natively available through classical microwave drives applied to directly to the oscillators and/or to parametric elements (i.e., three- and four-wave mixing elements) coupled to the oscillators.

Interaction	Hamiltonian	Platform
Dispersive coupling	$a^\dagger a\sigma_z$	Superconducting circuits
Sideband interaction	$a\sigma_+ + a^\dagger\sigma_-$	Trapped ions, superconducting circuits, neutral atoms
Linear drive	$\alpha a^\dagger - \alpha^* a$	Superconducting circuits, trapped ions
Parametric drive	$\zeta^* a^2 - \zeta a^{\dagger 2}$	Superconducting circuits, trapped ions
Controlled squeezing	$\sigma_x(a^{\dagger 2} + a^2)$	Trapped ions, neutral atoms
Free oscillator evolution	$a^\dagger a$	Trapped ions, superconducting circuits
Controlled beam splitter	$\sigma_x(a^\dagger b + b^\dagger a)$	Trapped ions

onto a qubit spin is accomplished by

$$b^\dagger b \to |e\rangle\langle e| = \frac{\hat{I} + \sigma_z}{2}, \tag{A5}$$

$$b \to \sigma^-; \quad b^\dagger \to \sigma^+, \tag{A6}$$

$$H_1 \to 0, \tag{A7}$$

$$H \to \bar{\omega}_R a^\dagger a + \frac{\bar{\omega}_Q}{2}\sigma_z + g\left(a + a^\dagger\right)\sigma_x, \tag{A8}$$

where we have neglected an irrelevant constant in the energy in the last expression. Equation (A8) describes the so-called Rabi Hamiltonian.

The Rabi Hamiltonian can be further simplified if we assume that the cavity-qubit detuning $\Delta = \bar{\omega}_R - \bar{\omega}_Q$ is small enough relative to the cavity frequency $\bar{\omega}_R$ that the RWA can be made to further simplify V_0 to

$$V_0 = g(a^\dagger b + a b^\dagger), \tag{A9}$$

which yields the celebrated Jaynes-Cummings Hamiltonian [77]

$$H = \bar{\omega}_R a^\dagger a + \frac{\bar{\omega}_Q}{2}\sigma_z + g\left(a\sigma^+ + a^\dagger\sigma^-\right). \tag{A10}$$

This "bare" Jaynes-Cummings Hamiltonian has its parameters fixed at the time of hardware fabrication and is not programmable, though in some cQED architectures, the transmon frequency $\bar{\omega}_Q$ can be tuned by application of a magnetic field that threads flux through a superconducting quantum interference device (SQUID) loop. Typical values for the anharmonicity of the transmon are $\bar{K}/(2\pi) \sim$ 50–200 MHz and the vacuum Rabi coupling can be up to $g/(2\pi) \sim 150$ MHz. Typical values for the operating frequencies of transmons and cavities are in the range $\bar{\omega}_Q/(2\pi), \bar{\omega}_R/(2\pi) \sim$ 5–10 GHz.

An alternative approach to the Hamiltonian of a qubit coupled to a cavity that is especially useful when higher levels of the transmon come into play, or when the cavity and transmon are detuned from each other ("dispersive regime"), is described below. In the dispersive regime, we find that it is possible to apply continuous microwave drives to the qubit and the cavity which yields an effective Jaynes-Cummings model with adjustable parameters.

$H_0 + V_0$ is the harmonic part of the total Hamiltonian and is, of course, exactly soluble. We use the strategy of expressing the anharmonic term, H_1 based on exact normal modes of the harmonic part of the Hamiltonian. For simplicity, we make the RWA for V_0 so that the harmonic Hamiltonian can be written

$$H_0 + V_0 = \begin{pmatrix} a^\dagger & b^\dagger \end{pmatrix} \begin{pmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix}, \tag{A11}$$

where

$$M = \begin{pmatrix} \bar{\omega}_R & g \\ g & \bar{\omega}_Q \end{pmatrix} = \frac{\bar{\omega}_R + \bar{\omega}_Q}{2}\hat{I} + \begin{pmatrix} \frac{\Delta}{2} & g \\ g & -\frac{\Delta}{2} \end{pmatrix}. \tag{A12}$$

From this, we see that the two normal modes of the coupled oscillators have frequencies

$$\omega_\pm = \frac{\bar{\omega}_R + \bar{\omega}_Q}{2} \pm \sqrt{(\Delta/2)^2 + g^2} \tag{A13}$$

and the normal mode operators are

$$\begin{pmatrix} A \\ B \end{pmatrix} = \begin{pmatrix} \cos\frac{\theta}{2} & +\sin\frac{\theta}{2} \\ -\sin\frac{\theta}{2} & \cos\frac{\theta}{2} \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix}, \tag{A14}$$

where $\tan\theta = 2g/\Delta$, and

$$\sin\frac{\theta}{2} = \sqrt{\frac{1 - (\Delta/2)/\sqrt{g^2 + (\Delta/2)^2}}{2}}, \tag{A15}$$

$$\cos\frac{\theta}{2} = \sqrt{\frac{1 + (\Delta/2)/\sqrt{g^2 + (\Delta/2)^2}}{2}}. \tag{A16}$$

To express the anharmonic term H_1 in the normal-mode basis, it is convenient to have the inverse transform

$$\begin{pmatrix} a \\ b \end{pmatrix} = \begin{pmatrix} \cos\frac{\theta}{2} & -\sin\frac{\theta}{2} \\ +\sin\frac{\theta}{2} & \cos\frac{\theta}{2} \end{pmatrix} \begin{pmatrix} A \\ B \end{pmatrix}. \tag{A17}$$

Making the RWA, we have

$$H_1 = -\frac{\bar{K}}{2} b^\dagger b^\dagger b b, \tag{A18}$$

$$\approx -\left[\chi_{BB} B^\dagger B^\dagger B B + \chi_{AB} B^\dagger B A^\dagger A + \chi_{AA} A^\dagger A^\dagger A A\right], \tag{A19}$$

where the renormalized transmon self-Kerr (anharmonicity), the transmon-cavity cross-Kerr (dispersive interaction), and the self-Kerr induced in the cavity are, respectively,

$$\chi_{BB} = \frac{K}{2}\cos^4(\theta/2), \tag{A20}$$

$$\chi_{AB} = 2K\cos^2(\theta/2)\sin^2(\theta/2), \tag{A21}$$

$$\chi_{AA} = \frac{K}{2}\sin^4(\theta/2). \tag{A22}$$

In the so-called dispersive regime, $g \ll |\Delta|$, $\theta \ll 1$ and hence $\chi_{BB} \gg \chi_{AB} \gg \chi_{AA}$. In the so-called "strong-dispersive regime," the cross-Kerr coupling is much larger than κ, the energy dissipation rate of the cavity and γ, the energy relaxation rate of the qubit $\chi_{AB} \gg \kappa, \gamma$ [238,683].

If we now make the two-level qubit approximation for the transmon using the analog of Eqs. (A5) and (A6), we obtain the Hamiltonian (up to an irrelevant constant)

$$H = H_{\rm D} + V, \tag{A23}$$

$$H_{\rm D} = \omega_{\rm R} A^\dagger A + \frac{\omega_{\rm Q}}{2}\sigma_z + \frac{\chi}{2}\sigma_z A^\dagger A, \tag{A24}$$

$$V = -\frac{K_A}{2} A^\dagger A^\dagger A A, \tag{A25}$$

where $\chi \equiv \chi_{\rm AB}$ is the dispersive shift and $K_A \equiv 2\chi_{\rm AA}$ represents the weak negative anharmonicity inherited by the cavity normal mode due to the coupling to the transmons. Here, the renormalized cavity frequency is given by Eq. (A13), with $\omega_{\rm R} = \omega_+$ if the detuning $\Delta > 0$, and $\omega_{\rm R} = \omega_-$ if $\Delta < 0$. The renormalized qubit frequency is the remaining eigenvalue, $\omega_{\rm Q} = \omega_-$ if the detuning $\Delta > 0$, and $\omega_{\rm Q} = \omega_+$ if $\Delta < 0$.

Because $K_A \ll \chi$, we make the simplifying assumption that we can neglect the induced cavity self-Kerr term V, and work only with the dispersive Hamiltonian $H_{\rm D}$. The dispersive Hamiltonian V_0 can be interpreted two different ways. First, it tells us that the resonant frequency of the cavity takes on different values, $\omega_{\rm R} \pm \chi/2$, depending on the state of the qubit. The frequency of the cavity is thus a proxy that can be used to make a high-fidelity QND dispersive measurement of the state of the qubit [77,78].

In the so-called "strong-dispersive regime," [238,683] where χ greatly exceeds the damping rate of the cavity and the transmon, the two resonance peaks of the cavity are spectrally well resolved. We see shortly that this allows us to make operations on the cavity that are conditioned on the qubit state, thus creating entanglement between the cavity and qubit.

The second interpretation of the dispersive Hamiltonian is that the qubit frequency shifts by χ for each photon that is added to the cavity. We will see shortly that this allows us to make rotations of the qubit conditioned on the precise number of photons in the cavity, giving us a second way to entangle the qubit and cavity. In particular, it gives us the ability to apply selective number-dependent arbitrary phase (SNAP; see Sec. IV D 4) gates [266] to each Fock state of the cavity.

Together, these two resources, qubit-conditioned cavity displacements and cavity-conditioned qubit rotations, grant us full quantum control of the combined system [59,266,329]. While some gates require only simple and intuitive microwave control pulses, the execution of more general unitaries at the fastest possible rates requires numerical determination of the control pulses using, e.g., the GRAPE algorithm from optimal control theory (OCT), which is difficult to interpret. For example, Heeres *et al.* [329] have carried out a highly non-Gaussian operation that executes a state-transfer from the $n = 0$ Fock state of the cavity to the $n = 6$ Fock state, using OCT pulses applied simultaneously to the transmon and to the cavity. Efficiency gains from bypassing qubit and gate abstractions to directly compile quantum algorithm steps at the control pulse level have been discussed recently in Ref. [684].

b. Conditional cavity gates

i. Conditional displacements. Qubit-controlled cavity displacement operations have been previously discussed in Refs. [60,64,327,342]. In this subsection, we aim to sketch at a high level the implementation of the conditional displacement gate CD(α), an integral gate to the phase-space ISA defined in Box IV.7.

In Ref. [60,64,342], a clever echo sequence is used to make rapid conditional displacements with a strong drive and simultaneously eliminate unwanted Stark shift rotations of the ancilla. Adding a drive to the cavity centered on frequency $\omega_{\rm R}$ and having an envelope $\epsilon(t)$, the Hamiltonian in Eq. (A24) in the frame doubly rotating at the qubit and cavity frequencies becomes

$$H = \frac{\chi}{2}\sigma_z a^\dagger a + i\epsilon(t)a^\dagger - i\epsilon^*(t)a, \tag{A26}$$

where for notational simplicity and to match the usage in the main text, we shall henceforth substitute $A \to a$. For the case in which the complex drive phase is time independent, we can remove the drive using a time-dependent unitary transformation

$$U(t) = e^{-\beta(t)a^\dagger + \beta^*(t)a} \tag{A27}$$

that shifts the oscillator to a displaced frame

$$a \to a + \beta(t), \tag{A28}$$

where

$$\beta(t) = \int_0^t d\tau\, \epsilon(\tau). \tag{A29}$$

The Hamiltonian in the new frame becomes

$$H = \frac{\chi}{2}\sigma_z[a^\dagger + \beta^*(t)][a + \beta(t)] \tag{A30}$$

$$= \frac{\chi}{2}\sigma_z\left(\left[a^\dagger a\right] + \left[\beta^*(t)a + \beta(t)a^\dagger\right] + \left[|\beta(t)|^2\right]\right). \tag{A31}$$

The middle term is the generator of the desired conditional displacement. The first and third terms are spurious but can be eliminated by a "double echo" sequence in which we first drive β to a large positive value, then flip the ancilla state and drive β to a large negative value. The details are described in the Supplementary Material of Ref. [60], where the authors call the sequence an *echoed-conditional-displacement (ECD) gate*. For the appropriate choice of $\beta(t)$, this gate is equivalent (up to a bit flip on the ancilla) to the conditional displacement CD(α) defined in Box IV.7 and used throughout this work.

ii. Conditional oscillator-oscillator entangling gates in the weak-dispersive regime. In Sec. VI C 1, we have demonstrated a useful trick for compiling conditional oscillator-oscillator entangling gates using controlled-parity gates (see the controlled beam-splitter gate in Eq. (337) as an example). Importantly, controlled parity gates are realized by allowing the dispersive interaction $e^{-i(\chi/2)ta^\dagger a\sigma_z}$ to act for time $t = \pi/\chi$. However, in the weak-dispersive regime (where dispersive coupling χ is extremely small), this strategy is undesirable because it results in a long gate duration. Here, we show an alternative method to compile a conditional oscillator-oscillator entangling gate, namely a conditional two-mode squeezing, in the weak-dispersive regime. To achieve this, we leverage uncontrolled two-mode squeezing operations, TMS(r,ϕ), as defined in Table III.

Using the definition of TMS(r,ϕ), we first note that

$$\mathrm{TMS}(\alpha,\pi)a\mathrm{TMS}^\dagger(\alpha,\pi) = a\cosh\alpha + b^\dagger\sinh\alpha. \tag{A32}$$

This in turn implies

$$\mathrm{TMS}(\alpha,\pi)e^{-i\chi ta^\dagger a\sigma_z}\mathrm{TMS}^\dagger(\alpha,\pi) = \exp\left[-i\chi t\left(\cosh^2(\alpha)a^\dagger a + \sinh^2(\alpha)bb^\dagger + \frac{1}{2}\sinh(2\alpha)(a^\dagger b^\dagger + ab)\right)\sigma_z\right]. \tag{A33}$$

The next steps of the method are inspired by the construction of echoed-conditional displacements (ECD) described above. Note that the first two terms in Eq. (A33) do not change signs with α, whereas the last two will. Hence, running the pulse shown below yields an echoed two-mode squeezing, in close analogy with the echoed displacement gate previously discussed,

$$\mathrm{TMS}(\alpha,\pi)e^{-i\chi ta^\dagger a\sigma_Z}\mathrm{TMS}^\dagger(\alpha,\pi)\times\sigma_x \times \mathrm{TMS}(-\alpha,\pi)e^{-i\chi ta^\dagger a\sigma_Z}\mathrm{TMS}^\dagger(-\alpha,\pi) = e^{-i\chi t\sinh 2\alpha(a^\dagger b^\dagger + ab)\sigma_z}. \tag{A34}$$

The speed of the conditional two-mode squeezing gate, in this case, is decided by $\chi t\sinh 2\alpha$ instead of χt as in the case in which this gate is synthesized with controlled parity gates. As $\sinh 2\alpha$ is an unbounded function, we can in principle increase it to extremely large values by varying α. Thus, in the weak-dispersive regime, we can achieve fast conditional oscillator-oscillator entangling gates by leveraging unconditional two-mode squeezing with large α as a resource. Hamiltonian amplification techniques [685] have been used in ion traps as well (see Ref. [88,686]).

c. Cavity-controlled qubit gates

i. SQR *gates.* A microwave electric field applied to the qubit couples to the electric dipole moment of the qubit and causes transitions between the ground and excited state. If the drive was applied at the qubit frequency it would coherently rotate the qubit around an axis that is determined by the phase of the microwave drive. Note from the dispersive Hamiltonian H_D for the coupled qubit-cavity system in Eq. (A24) that the qubit transition frequency suffers a quantized light shift depending on the photon number in the cavity:

$$\omega_Q + \chi\hat{n}. \tag{A35}$$

Hence we are able to execute a qubit gate conditioned on the photon number in the cavity. If we apply the qubit drive at frequency $\omega_Q(m) = \omega_Q + m\chi$, it is on resonance and therefore effective at rotating the qubit on the Bloch sphere if and only if there are exactly m photons in the cavity (assuming χ is sufficiently large relative to the drive amplitude). We now derive this in detail.

Starting from the driven-dispersive Hamiltonian in the corotating frame of the oscillator and qubit,

$$H = \frac{\chi}{2}a^\dagger a\sigma_z + \Omega^*(t)\sigma_- + \Omega(t)\sigma_+, \tag{A36}$$

we perform a unitary frame transformation to eliminate the dynamics of the dispersive term by using the unitary $U_\chi =$

$\exp\left(-i(\chi/2)ta^\dagger a\sigma_z\right)$. The transformed Hamiltonian $\tilde{H} = U_\chi^\dagger H U_\chi + i\left(\partial_t U_\chi^\dagger\right)U_\chi$ is

$$\tilde{H} = \sum_{n=0}^{\infty} |n\rangle\langle n| \left(\Omega^*(t)\sigma_- e^{-i\chi nt} + \Omega(t)\sigma_+ e^{+i\chi nt}\right). \quad \text{(A37)}$$

For selective qubit rotation gates up to Fock state N, we drive the qubit resonant with each number-shifted transition. This drive takes the form $\Omega(t) = \sum_{k=0}^{N}\Omega_k(t)e^{-ik\chi t}$, where the $\Omega_k(t)$ are slowly varying (complex-valued) envelopes. With this drive, the Hamiltonian becomes

$$\tilde{H} = \sum_{n=0}^{\infty} |n\rangle\langle n| \times \sum_{k=0}^{N}\left(\Omega_k^*(t)\sigma_- e^{-i\chi(n-k)t} + \Omega_k(t)\sigma_+ e^{+i\chi(n-k)t}\right). \quad \text{(A38)}$$

Under the condition of slow pulses (total time $T_{\text{gate}} > 2\pi/\chi$), $\Omega_k(t)$ will have a small bandwidth compared to χ, and the first-order RWA can be invoked. The implications of this approximation, along with mitigation techniques for faster pulses, are discussed in Ref. [524]. The RWA results in

$$\tilde{H}_{\text{RWA}} = \sum_{n=0}^{N} |n\rangle\langle n| \left(\Omega_n^*(t)\sigma_- + \Omega_n(t)\sigma_+\right). \quad \text{(A39)}$$

To realize an SQR gate, we pick $\Omega_n(t) = s(t)(\theta_n/2)e^{-i\varphi_n}$, where $s(t)$ is a real-valued envelope function with $\int_0^{T_{\text{gate}}} s(\tau)d\tau = 1$ so that

$$\text{SQR}\left(\vec{\theta},\vec{\varphi}\right) = \mathcal{T}\exp\left\{\left(-i\int_0^{T_{\text{gate}}}\tilde{H}_{\text{RWA}}dt\right)\right\}. \quad \text{(A40)}$$

By transforming the operator back to the initial frame, the total unitary after time T_{gate} is

$$U = U_\chi\ \text{SQR}\left(\vec{\theta},\vec{\varphi}\right) \quad \text{(A41)}$$

$$= \text{CR}\left(\chi T_{\text{gate}}\right)\ \text{SQR}\left(\vec{\theta},\vec{\varphi}\right). \quad \text{(A42)}$$

As indicated, the full gate in the laboratory frame includes an additional conditional cavity rotation, which accounts for the dispersive interaction during the gate. This CR can be canceled by applying a counter CR gate after the evolution, by constructing a suitable qubit echo sequence similar to the ECD construction [60] or, in cases where the initial and final qubit states are known eigenstates of σ_z (such as when implementing the SNAP gate), can be handled simply by tracking the cavity frame change in software.

ii. SNAP *gates.* As described in Box IV.10, two SQR gates can be combined to produce the photon-number-selective phase (SNAP) gate. One interesting choice of phases is

$$\varphi_n = n\varphi, \quad \text{(A43)}$$

which simply rotates the cavity state in phase space by angle $-\varphi$. Since the Berry phase is equal and opposite if the qubit starts and ends in the excited state, one can achieve the qubit-dependent cavity-rotation gate defined in Box IV.8,

$$\text{CR}(2\varphi) = e^{-i\varphi\hat{n}\sigma_z}, \quad \text{(A44)}$$

which can be used either to enhance or to compensate for the naturally occurring time evolution under the dispersive Hamiltonian in Eq. (A24).

For quantum simulation and control purposes, another interesting choice of phases is

$$\theta_n = \lambda n^2, \quad \text{(A45)}$$

which yields

$$U_{\text{Kerr}}(\lambda) = e^{i\lambda\hat{n}^2\sigma_z}, \quad \text{(A46)}$$

which can be used to dynamically cancel the undesired residual Kerr interaction in the cavity found in Eq. (A25). This gate can also be used to enhance the cavity self-Kerr to provide a programmable interaction for simulations of the boson Hubbard model, as discussed in Sec. VII C 5.

2. Trapped-ion quantum processors

In ion traps [16], one can produce the phase-space instructions set using the interaction between ions and lasers or microwaves. As discussed in Fig. 3, the mechanical oscillation modes of the ions form the bosonic modes. Unlike in the superconducting case in which the microwave resonators typically have gigahertz frequencies, the mechanical modes in ion traps typically have frequencies in the megahertz range. Beyond the various gates we discuss below, multi-chromatic laser drives that simultaneously couple many qubits to many mechanical modes have been developed [687].

a. Qubit rotations

The qubits in ion-trap systems can be encoded in optical transitions of the ions, but can also be encoded in hyperfine levels resulting from the coupling between electron and nuclear spins in the ground state manifold of the ions. The hyperfine energy-level spacing is usually on the order of megahertz and the transitions can be driven using microwaves to achieve single-qubit rotations much as in

NMR experiments [688]. Given their long wavelength, microwaves can be difficult to focus to address individual ions, but by engineering near-field microwave gradients, individual addressability can be achieved [689]. An alternative way to drive the qubit hyperfine transition is to use two detuned laser beams which have a frequency difference that matches the hyperfine splitting to drive the qubit transitions via a stimulated Raman process (STIRAP). Because the stimulated Raman transition is second order (i.e., requires one-photon absorption and one emission), it generally requires relatively strong laser intensities.

b. Sideband gates

A sideband gate $\mathrm{JC}(\theta,\varphi)$ (see Table IV) that flips one of the qubits and emits or absorbs one mechanical quantum (phonon) is (for the case of a red sideband)

$$\mathrm{JC}(\theta,\varphi)=\exp\left[-i\theta\left(e^{i\varphi}\sigma_{-}a^{\dagger}+e^{-i\varphi}\sigma_{+}a\right)\right], \qquad \text{(A47)}$$

which is universal when combined with an arbitrary single-qubit rotation $R_{\varphi}(\theta)$ in the X-Y plane (see Table VIII) [61,62,72]. This gives the ISA listed in Table X.

Similarly, a blue sideband gate $\mathrm{AJC}(\theta,\varphi)$ (Table IV) could be used, and when combined with a red sideband gate produces the controlled displacement gate $\mathrm{CD}(\alpha)$ of the phase-space ISA, as discussed below. We note that by use of a pulse shaping technique, a uniform sideband gate has been realized wherein the Rabi frequency between nearest-neighbor Fock levels does not scale as the Fock number but is a constant [343,344].

In the ion-trap setting, the finite momentum kick from the absorption or emission of an optical photon allows for the possibility of higher-order sideband transitions in which a spin flip is coupled to the creation or destruction of two or more mechanical oscillation quanta (phonons). While such higher sideband Hamiltonians are weaker than the lowest sideband coupling by powers of the Lamb-Dicke parameter, they present interesting possibilities for additional control schemes and instruction sets. However, we leave these for future work and do not discuss them further here.

We are not aware of a SNAP gate realization on trapped-ion platforms because of the lack of dispersive coupling between the qubits and the mechanical modes of the system. Such a coupling can, however, be synthesized from sideband transitions [72]. Beam splitters can also be synthesized outside this sideband gate set without using any lasers by parametrically modulating the electric fields used to create the trapping potential seen by the ions [72,74]. Alternatively, the second sidebands include beam-splitter interactions, which are, however, weaker by an additional power of the Lamb-Dicke factor. Moreover, arbitrary all-to-all connected beam-splitter Hamiltonians can be implemented using multi-tone driving [76].

c. Conditional displacements

Recall that the red sideband produces a coupling Hamiltonian in the form of $\sigma^{+}a+\sigma^{-}a^{\dagger}$, whereas a blue sideband produces coupling in the form of $\sigma^{-}a+\sigma^{+}a^{\dagger}$. If a bichromatic laser drive [328] is applied so that both sideband drives are switched on simultaneously and with equal strengths, we can realize the following Hamiltonian as a sum of the two:

$$\begin{aligned}&(\sigma^{+}a+\sigma^{-}a^{\dagger})+(\sigma^{-}a+\sigma^{+}a^{\dagger})\\&\quad=\sigma^{+}(a+a^{\dagger})+\sigma^{-}(a+a^{\dagger})\propto\sigma_{x}\hat{x}.\end{aligned} \qquad \text{(A48)}$$

This is nothing but a conditional displacement Hamiltonian, conditioned on the eigenstates of the qubit Pauli-X operator. We may further recast this into the familiar controlled displacement, $\mathrm{CD}(\alpha)$ (Table IV), by conjugating the qubit using a Hadamard gate since $\sigma_z=H\sigma_x H$.

Unconditional displacements can be achieved by using one of the qubits as an ancilla that is not involved in the circuit. Initializing this spare qubit in $|+\rangle$ and driving a conditional displacement realizes an unconditional displacement.

TABLE X. A minimal instruction set for a trapped-ion platform. This is the sideband ISA defined in Table V, with the important difference that any qubit can have JC coupling with any motional mode in which it participates. Application of the sideband gates to different qubits allows them to be entangled through the exchange of quanta of one of the bosonic modes. A beam-splitter gate (see Box IV.4) to couple different mechanical modes, a and b, can be synthesized from this gate set or can be created using rf control of the trap parameters as noted in the caption of Fig. 3. Blue sideband gates (AJC, Table IV) are equally easy to perform and are a convenient natively available gate that can be added to this minimal set.

Trapped-ion ISA	
Red sideband gate (Table IV)	$\mathrm{JC}(\theta,\varphi)$
Qubit rotation (Table VIII)	$R_{\varphi}(\theta)$

3. Neutral-atom quantum processors

The motion of single atoms in optical tweezers [7,113,114] is also approximately that of a quantum harmonic oscillator. In contrast to ions, each atom is coupled to a maximum of three modes, where in practice one of those modes is coupled to the atomic qubit, as illustrated in Fig. 3. This means that there is no mode-crowding effect restricting the scaling to larger systems, but this comes at the cost of reduced connectivity among the modes. Moreover, due to the much weaker dipole force confining the atoms compared to the electrostatic force confining

TABLE XI. A minimal instruction set for neutral-atom platforms. Sideband gates couple a single qubit to a single mode. Beam-splitter gates (see Box IV.4) to a couple of different mechanical modes, a and b, can be synthesized via the qubit-qubit Rydberg CZ gate.

Neutral-atom ISA	
Red sideband gate (Table IV)	$\mathrm{JC}(\theta,\varphi)$
Qubit rotation (Table VIII)	$R_\varphi(\theta)$
Qubit-qubit gate (Table VIII)	CZ

ions, the motional mode frequency ω_r is smaller, usually in the tens of kilohertz. Qubits may be encoded in hyperfine levels (in alkali atoms such as rubidium) or optical clock qubits (in strontium). By qubit-state-selective excitation of the atoms into Rydberg states, the Rydberg blockade effect can be used to create two-qubit entangling gates [690]. Recently, combined qubit-boson operations have been demonstrated in a tweezer array [115], showing coherence times of the modes over 100 ms and nonclassical motion of the atoms due to Fock-state preparation has been shown [691]. There have been recent theory proposals for realization of bosonic QEC codes in mechanical vibration states of neutral atoms [118,214].

Contrary to ions, the tweezer potential has a sizeable nonlinearity. The maximum boson number occupation is given by the trap depth V_0, which in current implementations is on the order of 1–10 MHz such that the maximum accessible Fock state is 100–1000. Another effect of the finite trap depth is a finite Kerr nonlinearity of the potential, which is 0.1%–1% of the mode frequency (coincidentally similar to that for superconducting transmon qubits). The nonlinearity is in principle tunable by the power of the trapping laser. However, despite this sizeable nonlinearity, qubit-boson gates are not substantially altered since the qubit-boson coupling is so large.

a. Single-qubit gates

Single-qubit gates are implemented just as in qubit-only applications of neutral atoms by either Raman transitions [9] or microwaves combined with ac Stark shifts [692].

b. Multi-qubit gates

Multi-qubit gates are implemented by using a Rydberg level of the atoms. High-fidelity two-qubit CZ and even n-qubit multi-controlled Z gates have been demonstrated [690] by using the blockade mechanism, which forbids two atoms (within the blockade radius of each other) to be transferred to the Rydberg state simultaneously [6]. When two atoms are simultaneously in the Rydberg state, they produce forces that modify the bosonic state, as discussed in the next subsection. However, during the Rydberg blockade entangling gates, the atoms do not simultaneously occupy Rydberg states, thus preventing any impact on the motional states.

c. Single-mode gates

Phase-space rotations are naturally implemented by the free evolution under the tweezer potential. Moreover, the nonlinearity enables universal single-mode control by modulating the tweezer potential [119].

d. Multimode gates

i. Using two-qubit gates. Intermode gates such as beam splitters can be compiled using sideband gates and CZ gates. We employ the BCH formula

$$e^{iA\theta}e^{iB\theta}e^{-iA\theta}e^{-iB\theta}=e^{-[A,B]\theta^2+O(\theta^3)}. \tag{A49}$$

First, we create a doubly controlled displacement by choosing $A=Z_1Z_2$ and $B=X_1(a+a^\dagger)$. In this way we get $-[A,B]\theta^2=2iY_1Z_2(a+a^\dagger)\theta^2\equiv iC\tilde{\theta}$ with $C=Y_1Z_2(a+a^\dagger)$ and $\tilde{\theta}=\sqrt{2}\theta$. Then, we can repeat the BCH trick, conjugating C with $D=X_2(a+a^\dagger)$. From this concatenation, we obtain

$$e^{-[C,D]\tilde{\theta}^2}=e^{iY_1Y_2(a+a^\dagger)(b+b^\dagger)4\tilde{\theta}^4}. \tag{A50}$$

Hence, we have realized a doubly qubit-controlled beam splitter–two-mode-squeezing gate. The two-mode-squeezing or beam-splitter terms can be singled out by applying the same sequence with different phases and employing the Trotter formula.

ii. Using dipolar forces in Rydberg state. Two neutral atoms excited to Rydberg states exert strong dipolar forces on each other. By "dressing" the ground state with the Rydberg state and modulating the tweezer potentials, beam splitters as well as two-mode squeezing can be realized [693].

iii. Using tunneling. Another approach is to use tunneling between tweezers, generalizing the approach discussed in the context of qubits encoded in the lowest two levels of each tweezer [694].

e. Sideband gates

Qubit-boson coupling is implemented using sideband gates as in ions, which in the Lamb-Dicke regime $\eta\sqrt{n}\ll 1$ is given by

$$H=\frac{1}{2}\Omega\eta\left(a^\dagger\sigma^-+a\sigma^+\right), \tag{A51}$$

where Ω is the Rabi frequency, and the Lamb-Dicke factor η is given in terms of the momentum transfer δk of the

Raman beams and the atomic mass m by $\eta = \delta k/\sqrt{2m\omega_r}$. Due to the essentially unlimited Rabi frequency in this context, the main restriction to the coupling $\Omega\eta$ of these gates is given by the boson mode frequency—to resolve the sideband transitions it needs to fulfill $\Omega\eta < \omega_r$.

f. Alternative architectures

i. Motional states in optical lattices. Alternatively to the tweezer platform, the motional modes in optical lattices can be similarly used. Lattice shaking can then be used to excite oscillators, which may also be done in a qubit-dependent way. Superlattices can be used to engineer tunneling interactions for beam splitters.

ii. Bosonic atoms. A completely different approach is to use bosonic atoms in optical lattices or tweezers, similar to the recent proposal for fermion-qubit digital computation [37]. Recently, boson sampling has been realized in this way [117]. This approach is most naturally suited for problems conserving the total number of bosons. However, displacements may be engineered by coupling the tweezers to a Bose-Einstein condensate.

APPENDIX B: BLOCH-MESSIAH DECOMPOSITION

In the last appendix, we learned that a unitary Gaussian transformation on the infinite Hilbert space of N bosonic modes corresponds to a transformation of the $2N$ mode positions and coordinates (or equivalently, the raising and lowering mode operators), which must be symplectic to preserve the commutation relations. The Bloch-Messiah decomposition [222,300,301] is essentially a singular-value decomposition for these $2N \times 2N$ symplectic matrices. Any $2N \times 2N$ real symplectic matrix $\mathbf{M}$ can be decomposed into the form

$$\mathbf{M} = \mathbf{O_1 Z O_2}, \tag{B1}$$

where $\mathbf{O_1}$ and $\mathbf{O_2}$ are $2N \times 2N$ orthogonal symplectic matrices, and $\mathbf{Z}$ is a $2N \times 2N$ diagonal squeezing matrix containing the singular values in the form

$$\mathbf{Z} = \mathrm{Diag}\left(e^{r_1}, e^{-r_1}, e^{r_2}, e^{-r_2}, \ldots, e^{r_N}, e^{-r_N}\right). \tag{B2}$$

Physically, this decomposition can be realized by having a system of beam splitters applied before ($\mathbf{O_2}$) and after ($\mathbf{O_1}$) a middle step $\mathbf{Z}$ involving only single-mode squeezing (by varying amounts on each mode). We make use of this decomposition in the following derivations.

There are two reasons that this decomposition is important for experiments. The first is that it allows synthesis of two-mode squeezing using single-mode squeezing, which is often substantially easier to realize. The second is that very high-fidelity beam-splitter gates between neighboring qubits now existing in superconducting qubit systems and can be used to synthesize the $\mathbf{O}_1$ and $\mathbf{O}_2$ operations which in general involve beam splitters between distant resonators. This in turn requires SWAP networks (again based on beam splitters; see Sec. VIC 1 a) to be able to achieve spatially nonlocal beam-splitter gates.

We now turn to some specific examples.

1. SUM gate

The two-mode SUM gate given in Eq. (192) is a unitary operation that can be represented by the corresponding 4×4 symplectic matrix (in the mode operator basis),

$$\mathbf{SUM}(\lambda) = \begin{pmatrix} 1 & 0 & -\lambda/2 & +\lambda/2 \\ 0 & 1 & +\lambda/2 & -\lambda/2 \\ +\lambda/2 & +\lambda/2 & 1 & 0 \\ +\lambda/2 & +\lambda/2 & 0 & 1 \end{pmatrix}, \tag{B3}$$

acting on the vector of mode operators

$$\begin{pmatrix} a \\ a^\dagger \\ b \\ b^\dagger \end{pmatrix}. \tag{B4}$$

The transformation obeys Eq. (101),

$$\begin{aligned} &\mathbf{SUM}^{\mathrm{T}}(\lambda) \begin{pmatrix} 0 & +1 & 0 & 0 \\ -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & +1 \\ 0 & 0 & -1 & 0 \end{pmatrix} \mathbf{SUM}(\lambda) \\ &= \begin{pmatrix} 0 & +1 & 0 & 0 \\ -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & +1 \\ 0 & 0 & -1 & 0 \end{pmatrix}, \end{aligned} \tag{B5}$$

and is therefore symplectic, proving that it preserves the commutation relations among the mode operators as required.

The Bloch-Messiah decomposition given in Eq. (196) can be written in this form using the symplectic matrix representation of the two-mode beam splitter, the unitary of which is given in Eq. (169) as

$$\begin{aligned} &\mathbf{BS}(\theta, \varphi) \\ &= \begin{pmatrix} \cos\frac{\theta}{2} & 0 & -ie^{i\varphi}\sin\frac{\theta}{2} & 0 \\ 0 & \cos\frac{\theta}{2} & 0 & ie^{-i\varphi}\sin\frac{\theta}{2} \\ -ie^{-i\varphi}\sin\frac{\theta}{2} & 0 & \cos\frac{\theta}{2} & 0 \\ 0 & ie^{+i\varphi}\sin\frac{\theta}{2} & 0 & \cos\frac{\theta}{2} \end{pmatrix}. \end{aligned} \tag{B6}$$

Similarly, the symplectic matrix representation of the pair of single-mode squeezers [with the upper arm of the interferometer being antisqueezed by $s(-r)$ and the lower arm squeezed by $s(r)$] is

$$\mathbf{s}(r)\otimes\mathbf{s}(-r)=\begin{pmatrix}\cosh r & +\sinh r & 0 & 0\\ +\sinh r & \cosh r & 0 & 0\\ 0 & 0 & \cosh r & -\sinh r\\ 0 & 0 & -\sinh r & \cosh r\end{pmatrix}. \tag{B7}$$

Plugging these into Eq. (196) yields the following expression:

$$\mathbf{SUM}(\lambda)=\begin{pmatrix}-\cosh(r)\sin(2\theta) & 0 & -\cosh(r)\cos(2\theta) & \sinh(r)\\ 0 & -\cosh(r)\sin(2\theta) & \sinh(r) & -\cosh(r)\cos(2\theta)\\ \cosh(r)\cos(2\theta) & \sinh(r) & -\cosh(r)\sin(2\theta) & 0\\ \sinh(r) & \cosh(r)\cos(2\theta) & 0 & -\cosh(r)\sin(2\theta)\end{pmatrix}. \tag{B8}$$

Choosing the parameter values for θ and r as in Eqs. (197)–(199) yields the desired transformation given in Eq. (B3).

2. Two-mode squeezing

The two-mode squeezing gate in Box IV.5 has the 4×4 symplectic representation given in Eq. (182):

$$\mathbf{TMS}(r,\varphi)=\begin{pmatrix}\cosh(r) & 0 & 0 & e^{i\varphi}\sinh(r)\\ 0 & \cosh(r) & e^{-i\varphi}\sinh(r) & 0\\ 0 & e^{i\varphi}\sinh(r) & \cosh(r) & 0\\ e^{-i\varphi}\sinh(r) & 0 & 0 & \cosh(r)\end{pmatrix}. \tag{B9}$$

It can be synthesized in a similar manner using the same, rather than opposite, single-mode squeezing operations in each arm of the interferometer, as illustrated in Fig. 11 and Eq. (186). These squeezing operations are represented by the symplectic matrix

$$\mathbf{s}(r)\otimes\mathbf{s}(r)=\begin{pmatrix}\cosh r & -\sinh r & 0 & 0\\ -\sinh r & \cosh r & 0 & 0\\ 0 & 0 & \cosh r & -\sinh r\\ 0 & 0 & -\sinh r & \cosh r\end{pmatrix}. \tag{B10}$$

Plugging this expression and Eq. (B6) into Eq. (186) correctly yields

$$\mathbf{TMS}(r,\pi/2)=\begin{pmatrix}\cosh(r) & 0 & 0 & i\sinh(r)\\ 0 & \cosh(r) & -i\sinh(r) & 0\\ 0 & i\sinh(r) & \cosh(r) & 0\\ -i\sinh(r) & 0 & 0 & \cosh(r)\end{pmatrix}, \tag{B11}$$

in agreement with Eq. (B9).

The corresponding symplectic transformation $\mathrm{TMS}(r,0)$ of the quadrature coordinates is

$$(\hat{x}'_a+\hat{x}'_b)=e^{+r}(\hat{x}_a+\hat{x}_b), \tag{B12}$$
$$(\hat{p}'_a+\hat{p}'_b)=e^{-r}(\hat{p}_a+\hat{p}_b), \tag{B13}$$
$$(\hat{x}'_a-\hat{x}'_b)=e^{-r}(\hat{x}_a-\hat{x}_b), \tag{B14}$$
$$(\hat{p}'_a-\hat{p}'_b)=e^{+r}(\hat{p}_a-\hat{p}_b), \tag{B15}$$

showing that (for large r and the phase choice $\varphi=0$), the position quadratures of output modes become strongly correlated while the momentum quadratures become strongly anticorrelated.

APPENDIX C: OSCILLATOR NOISE MODELS

In this appendix, we discuss a few common noise models for bosonic systems in different experimental platforms.

1. Kraus operator representation of an error channel

A quantum channel is the most general operation that can be applied to a density matrix that creates another physically allowed density matrix. Because physically allowed density matrices must be positive semidefinite and have unit trace, this mapping on the space of density matrices must be completely positive and trace preserving (CPTP). The Kraus representation of a quantum channel has the form

$$\rho(t+\tau)=\sum_{\ell=0}^{D_K-1}\hat{K}_\ell\,\rho(t)K_\ell^\dagger, \tag{C1}$$

where the different Kraus operators $\hat{K}_\ell$ can be unitary or nonunitary and describe the action of the time evolution of the system under its own Hamiltonian plus the couplings to the dissipative bath. The number of such operators D_K

needed to describe the quantum channel is called the Kraus dimension (or rank) of the map. For a single qubit $D_K \leq 4$, while for a bosonic mode, the number of operators is formally infinite.

The form of the Kraus map guarantees that it is completely positive. To guarantee that it is trace-preserving, the Kraus operators must obey the constraint

$$\sum_{\ell=0}^{D_K-1} \hat{K}_\ell^\dagger \hat{K}_\ell = \hat{I}, \tag{C2}$$

where $\hat{I}$ is the identity operator.

2. Amplitude damping

In Sec. V A, we have briefly introduced the amplitude damping channel for bosonic modes by using the master equation formalism. Here, we briefly discuss its Kraus representation. Following Ref. [416] we can choose to label the Kraus operators by the total number of quanta (photons or phonons) ℓ lost during time τ:

$$\hat{K}_\ell = \sqrt{\frac{(1-e^{-\kappa\tau})^\ell}{\ell!}} e^{-\frac{\kappa}{2}\tau\hat{n}} a^\ell, \quad \ell = 0, 1, 2, 3, \ldots, \tag{C3}$$

where $\hat{n} = a^\dagger a$ is the photon-number operator and κ is the energy relaxation (photon loss) rate from Eq. (271). We see that the a^ℓ term in $\hat{K}_\ell$ destroys ℓ photons, but that even if by chance no photons are lost, the $\ell = 0$ Kraus operator $\hat{K}_0$ enacts a "no-jump backaction," causing the probability amplitude for Fock states with large photon numbers to be reduced. This is essentially a Bayesian probability update—given the additional information that no photons decayed into the environment, we must revise downward our prior estimate of the initial photon number.

For small values of $\kappa\tau$, the probability of losing ℓ photons in a short time interval scales as $(\kappa\tau)^\ell$ and becomes small for large ℓ. Hence we can safely truncate the Kraus dimension to some value $\ell_{\max}$ if needed. However, the truncated series of Kraus operators will not be trace-preserving. We can fix this by modifying the "no-jump" operator $\hat{K}_0$ to be

$$\hat{K}_0' = \sqrt{\hat{I} - \sum_{\ell=1}^{\ell_{\max}} \hat{K}_\ell^\dagger \hat{K}_\ell}. \tag{C4}$$

Because the operator inside the square root is diagonal (and each diagonal element is nonnegative) in the Fock basis, the square root of the matrix can be found by taking the element-wise square root.

3. Dephasing

The action of the bosonic dephasing channel can be represented as a sum over continuously many Kraus operators, each representing a phase-space rotation by some random angle θ via

$$\mathcal{E}(\rho) = \int_{-\pi}^{\pi} d\theta\, p(\theta) e^{i\theta a^\dagger a} \rho e^{-i\theta a^\dagger a}. \tag{C5}$$

The channel randomizes the phase according to the probability distribution $p(\theta)$. When ρ is expressed in the photon-number basis,

$$\mathcal{E}(\rho) = \sum_{m,n} \langle m|\, \rho\, |n\rangle \int_{-\pi}^{\pi} d\theta\, p(\theta) e^{i\theta(n-m)}\, |n\rangle\, \langle m|\,, \tag{C6}$$

we see that the channel preserves the diagonal elements but reduces the magnitude of the off-diagonal elements, indicating decoherence.

The exact solution for the quantum capacity of the bosonic dephasing channel has recently been given in Ref. [695]. For microwave oscillators realized as 3D superconducting cavities, intrinsic dephasing rates (i.e., rates not associated with state changes in the auxiliary qubit coupled to the cavity) are known to be substantially smaller than the energy relaxation rate κ (unlike superconducting qubits, where the rates are often similar) and have been measured to be below $\kappa_\phi \lesssim 1/(10\text{ ms})$ [24,419], and more recently below $\kappa_\phi \lesssim 1/(0.5\text{ s})$ [26].

In ion traps, the mechanical mode frequencies are sensitive to electric field noise, which can lead to significant dephasing. Similarly, the tweezer beams that trap neutral atoms may have amplitude or phase fluctuations, which lead to mode dephasing.

4. Heating

Microwave resonators are generally in good thermodynamic equilibrium with the cryogenic refrigeration system, which means that they are cold enough that the resonator is rarely thermally excited out of the vacuum state. This is not the case for ion traps, as we discuss next. Further below, we also discuss heating in neutral-atom tweezer arrays.

Instead of operating at a typical frequency of gigahertz as in cQED hardware, the ac electric field that traps ions in ion traps produces mechanical resonance frequencies in the megahertz regime. As a result, the dominant quantum noise in the trapped-ion platform is the heating of the motional vibrational modes of the ions [16,696,697], which can be described by the time derivative of the frequency-dependent average Fock number, i.e., the population of a given mode $\dot{\bar{n}}$. The microscopic origin of the heating mechanism remains an open question and has been historically under debate, such that it is often dubbed "anomalous heating."

It has been demonstrated that the heating rate depends on the trap frequency and geometry as well as on the temperature. Because the ions mostly interact with the ac

trapping electric field, it is clear that fluctuations in the electric field can significantly affect the heating rate. From the fluctuation-dissipation theorem and the fact that the ions couple linearly to the fluctuating electric field, the heating rate $\dot{\bar{n}}$ can be quantified by [697]

$$\dot{\bar{n}}(\omega, d, T) = \frac{q^2}{2m_I \hbar \omega} S_E(\omega, d, T), \tag{C7}$$

where S_E and ω ($\sim$MHz) are the noise spectral density and frequency of the trapping electric field, and m_I is the mass of the ion. Additionally, S_E depends on the distance d between the ion and the trap surface ($\sim$10–100 μm) and the temperature of the operation T, which is experimentally found to obey the following empirical expression [696]:

$$S_E(\omega, d, T) \propto \omega^{-\alpha} d^{-\beta} f(T), \tag{C8}$$

where most experiments suggest $\alpha \sim 1$ and $\beta \sim 4$, and $f(T)$ is a function of temperature that typically exhibits thermal activation behavior. The distance dependence d^4 excludes the possibility of Johnson noise, which should have a dependence d^2. The reason why S_E has such a dependence on ω, d, and T is still under debate, but there have been many theoretical studies, ranging from phenomenological models [698–702] to *ab initio* electronic structure and molecular dynamics simulation [703–706], to reveal more microscopic details of the heating mechanism. One possibility is that the molecular absorbates on the trap surface serve as a thin layer of lossy dielectric medium that induces dissipation, where the temperature dependence of the motion of the absorbed molecules can contribute to the temperature scaling. Typical values for the heating rate $\dot{\bar{n}}$ in Eq. (C7) range from $\sim 10^1$ phonons/s at 4 K to 10^3 phonons/s at room temperature, for a trapping frequency of 1 MHz and a trapping distance of 50 μm [707]. Also note that because the electric field fluctuations have a long wavelength, they mostly couple to the center-of-mass mode. Modes other than the center-of-mass mode generally have much lower heating rates [696].

In neutral-atom arrays, heating is mainly caused by atom movement [8], which can be reduced by making the atom movement slower. Moreover, in the future, optimal control movement schemes may be developed to optimize the transport of oscillator states [708]. Finally, atom movement may be avoided altogether by using a fixed architecture and using SWAP gates to entangle distant tweezers.

APPENDIX D: UNIVERSALITY OF ISAs AND CROSS-COMPILATION

Here, we demonstrate explicit cross-compilation techniques between various ISAs, including the phase-space and Fock-space ISAs (Appendix D 1), the sideband and phase-space ISAs (Appendix D 2), as well as the sideband and Fock-space ISAs (Appendix D 3). These cross-compilation schemes also serve as a universality proof for all three ISAs on single-qubit–oscillator systems, since previous work has shown that at least one of them is universal [59,62].

Building upon the single-qubit–oscillator universality, we provide a recursive proof of universality for multi-qubit–oscillator systems (Appendix D 4) by adding the analytical compilation of multi-qubit entangling gates, as described in Sec. VI C. The cross compilation schemes and proofs presented in this appendix are far from optimal, but we hope the material presented here spurs more efficient ways for cross-compilation in future works.

1. Compilation between the Fock-space ISA and the phase-space ISA

In this section, we explore cross-compilation between the two most important ISAs in the current paper, the phase-space and the Fock-space ISAs. Section D 1 a provides physical intuition on compiling the SNAP gate from CD gates. Appendix D 1 b leverages Hamiltonian simulation techniques to compile the SNAP gate in the Fock-space ISA using the phase-space ISA. Finally, Appendix D 1 c presents the reverse direction of compiling a CD gate using the Fock-space ISA.

a. Fock-space ISA from phase-space ISA: Intuition

SNAP gates (described in Box IV.10) are cavity-controlled qubit rotations which apply a separate Berry phase to each Fock state in the cavity. These gates plus unconditional cavity displacements constitute a universal Fock-space gate set which is dual to the phase-space ISA in the sense that in the phase-space ISA, it is the cavity operations that are controlled by the qubit and the qubit rotations are unconditional. As described in Sec. IV D 4), cavity-controlled qubit operations are most easily engineered in physical architectures with large dispersive coupling between the ancilla qubit and the cavity. Conversely, the advantage of the phase-space ISA is that its gates can be rapidly executed (using large displacements) even when the dispersive coupling is relatively small. An interesting computational task is to attempt to compile each of these ISAs in terms of the other and analyze the complexity of this task.

Physically, we know how to implement SNAP gates using the strong-dispersive Hamiltonian (see Appendix A 1 c):

$$H = \chi a^\dagger a \sigma_z. \tag{D1}$$

The phase-space ISA is based on this very Hamiltonian but operates experimentally with a small value of the dispersive coupling χ. However, we could use the phase-space ISA with Trotter-Suzuki to synthesize time evolution

under this Hamiltonian with a larger effective value of χ. We can then use controlled displacements with this to literally do a quantum simulation of the SNAP gate dynamics. Although slow, this strategy is, in principle, sufficient to compile a SNAP gate as we outline below.

b. SNAP *gate compilation by quantum Hamiltonian simulation*

For this compilation, we need to use the phase-space ISA to simulate the dispersive Hamiltonian given in Eq. (D1) with the physical dispersive coupling χ replaced by a larger effective value, χ_{eff}.

The time-evolution unitary corresponds to a conditional cavity-rotation gate,

$$\text{CR}(\theta) = e^{-i\frac{\theta}{2}\hat{n}\sigma_z}, \tag{D2}$$

which we can synthesize from CD gates via Trotter-Suzuki. Importantly, this conditional cavity-rotation gate can also be interpreted as a rotation of the qubit around the z axis by an angle that is proportional to the number of photons in the cavity, $\theta_n = \theta\hat{n}$.

The SNAP gate (Box IV.10) is realized within the Fock state ISA via a pair of SQR gates (Box IV.9). In the phase-space ISA, we make use of a Rabi drive to rotate the qubit gate:

$$R_\varphi(\delta) = e^{-i\frac{\delta}{2}[\cos\varphi\sigma_x + \sin\varphi\sigma_y]}. \tag{D3}$$

To make the qubit rotation be conditioned on the photon number being m, let us choose $\theta, \delta \ll 1$ and repeatedly execute the gate sequence

$$\begin{aligned} U_m(\theta,\delta) &= e^{+i\frac{\theta}{2}m\sigma_z}\text{CR}(\theta)R_\varphi(\delta) \\ &= e^{-i\frac{\theta}{2}[\hat{n}-m]\sigma_z}R_\varphi(\delta). \end{aligned} \tag{D4}$$

For example, if we wish to rotate the qubit through angle π if and only if the photon number is m, then we repeat this π/δ times. If the photon number is not equal to m, then the rapid rotation of the qubit about the z axis simulates the strong dispersive coupling that makes the Rabi drive off-resonant and hence ineffective at rotating the qubit. Thus we have produced the SQR gate that selectively rotates the qubit conditioned on the photon number. A second SQR gate with a different Rabi phase φ then imparts the desired phase on cavity Fock state m yielding the SNAP gate.

If we simply rotate the qubit through 2π conditioned on the photon number being m, we have a SNAP gate that produces a Grover oracle that adds a minus sign to Fock state $|m\rangle$:

$$G = \hat{I} - 2|m\rangle\langle m|. \tag{D5}$$

If we choose special values of θ and δ that, rather than being very small, are rational multiples of π, we can achieve other novel gates that, e.g., execute a "superparity" operation. Executing the unitary N times with $\delta = 2\pi/N$ for selected values of θ applies a minus sign to every Fock state, satisfying the condition $\hat{n} = 0 \mod \ell$ shown in Table XII:

TABLE XII. QSP for "superparity" operations. Applying the unitary $U_0(\theta,\delta)$ found in Eq. (D4) N times, with $m = 0$ and $\delta = 2\pi/N$, applies a minus sign to every Fock state satisfying the condition shown in the first column. The phases applied to the other states are 0 to varying degrees of approximation. The worst-case error is shown in the last column and scales to zero like $1/N$. Note that the errors in the QSP realization of the dispersive Hamiltonian evolution have not been included here.

Condition $n = 0 \mod \ell$	θ	N	Maximum phase error (radians)
mod 2	$\pi/2$	4	0
mod 4	$\pi/4$	32	0.152
mod 4	$\pi/4$	64	0.077
mod 4	$\pi/4$	128	0.0385
mod 8	$\pi/8$	64	0.184
mod 8	$\pi/8$	128	0.093
mod 8	$\pi/8$	256	0.047
mod 16	$\pi/16$	128	0.192
mod 16	$\pi/16$	256	0.097

$$G_\ell = \hat{I} - 2\sum_m |m\ell\rangle\langle m\ell|. \tag{D6}$$

All the gates shown in Table XII disentangle the ancilla nearly perfectly from the cavity and leave the ancilla in its starting state. However, the phases for those Fock states that are supposed to be 0 can be erroneous. The worst-case phase error is shown in the last column. For this simple algorithm, the error scales toward zero only as $1/N$. Remarkably, Ref. [332] has combined QSP with the natively available dispersive interaction to produce an efficient and nearly exact synthesis of the superparity operation and other non-Gaussian hybrid CV-DV operations.

c. Phase-space ISA from Fock-space ISA

We now turn to the problem of compiling the phase-space ISA in terms of the Fock-space ISA. In particular, this requires the compilation of single-qubit rotations $R_\varphi(\theta)$ and conditional displacements $\text{CD}(\alpha)$ in terms of selective qubit rotations $\text{SQR}(\vec{\theta},\vec{\varphi})$ and unconditional displacements $D(\alpha)$. Realizing the first of these is particularly straightforward: one just needs to choose $\theta_n = \theta$ and $\varphi_n = \varphi$, such that the same qubit rotation is applied for every Fock state, reducing the SQR gate to a single-qubit rotation $R_\varphi(\theta)$.

To synthesize a controlled displacement, we next recognize that an SQR gate with $\theta_n = n\pi$ conjugated by single-qubit rotations $R_y(\pi/2)$ realizes a controlled parity operation CP (Box IV.8). We note that one can more

efficiently realize this gate using the native dispersive interaction, though here we fix ourselves to gates within the phase-space ISA to prove cross-compilation capability. With the controlled parity gate as a primitive, we can now use the hybrid oscillator-qubit gate synthesis technique outlined in Sec. VI C 1 [see the text surrounding Eq. (336)] to realize a conditional displacement,

$$\mathrm{CD}(\alpha) = \mathrm{CP}\, D(i\alpha)\, \mathrm{CP}^\dagger, \tag{D7}$$

thus completing the compilation of the phase-space ISA in terms of the Fock-space ISA.

2. Compilation between sideband ISA and phase-space ISA

In this section, we show that cross-compilation between the sideband gates and the controlled-displacement gates is possible via Trotterization.

a. Sideband gates from controlled displacement

We first show that sideband gates can be compiled from controlled displacements. Redefining $\alpha = -i(\theta/2)e^{i\phi}$, the red sideband gate in Eq. (A47) can be rewritten as

$$\mathrm{JC}(\theta,\varphi) = e^{\alpha\sigma^+ a - \alpha^*\sigma^- a^\dagger} \tag{D8}$$

$$= \left(e^{\delta\sigma^+ a - \delta^*\sigma^- a^\dagger}\right)^L, \tag{D9}$$

where we have Trotterized into a product of L smaller sideband gates with $\delta = \alpha/L$. Each small sideband gate admits the following decomposition:

$$e^{\delta\sigma^+ a - \delta^*\sigma^- a^\dagger} = \exp\left[\frac{\sigma_x}{2}(\delta a - \delta^* a^\dagger) + \frac{i\sigma_y}{2}(\delta a + \delta^* a^\dagger)\right]$$
$$= e^{\frac{\sigma_x}{2}(\delta a - \delta^* a^\dagger)} e^{\frac{i\sigma_y}{2}(\delta a + \delta^* a^\dagger)} + O(|\delta|^2), \tag{D10}$$

where the final line corresponds to two controlled-displacement gates. Therefore, we can compile a single sideband gate with parameter α into $2L$ controlled-displacement gates with an error of $L \cdot O(|\delta|^2) = O(|\alpha|^2/L)$. Note that higher-order product formulas can be used in Eq. (D10) to improve the compilation accuracy.

b. Controlled-displacement gates from sideband gates

Next, we show the reverse direction of compiling controlled-displacement gates into sideband gates by again relying on Trotterization and product formulas. Without loss of generality, any controlled-displacement gate can always be written as $e^{i\sigma_x(\alpha^* a^\dagger + \alpha a)}$ upon conjugation by single-qubit rotations. In addition, we can express this conditional displacement as a product of L small conditional displacements via Trotterization,

$$e^{i\sigma_x(\alpha^* a^\dagger + \alpha a)} = \left(e^{i\sigma_x(\delta^* a^\dagger + \delta a)}\right)^L, \tag{D11}$$

where $\delta = \alpha/L$. Using the identity $\sigma_x = \sigma^+ + \sigma^-$ in combination with the first-order product formula, each small conditional displacement by δ can be rewritten as

$$e^{i\sigma_x(\delta a^\dagger + \delta a)} = e^{i\delta\sigma^+ a + i\delta^*\sigma^- a^\dagger} e^{i\delta^*\sigma^+ a^\dagger + i\delta\sigma^- a} + O(|\delta|^2). \tag{D12}$$

We can then note that the first term in Eq. (D12) is a red sideband gate as given in Eq. (A47), while the second term is a blue sideband gate that (if not natively available) can be synthesized via a red sideband gate and single-qubit rotations,

$$e^{i\delta^*\sigma^+ a^\dagger + i\delta\sigma^- a} = e^{i\frac{\pi}{2}\sigma_x}\left[e^{i\delta^*\sigma^- a^\dagger + i\delta\sigma^+ a}\right] e^{-i\frac{\pi}{2}\sigma_x}. \tag{D13}$$

Therefore, we can express "small" controlled displacements by δ as two sideband gates interleaved by π rotations on the qubit,

$$e^{i\sigma_x(\delta a^\dagger + \delta a)} = V_s(i\delta) e^{i\frac{\pi}{2}\sigma_x} V_s(i\delta) e^{-i\frac{\pi}{2}\sigma_x} + O(|\delta|^2). \tag{D14}$$

Consequently, a "large" controlled displacement by α can be compiled as a product of $2L$ red and blue sideband gates interleaved with single-qubit rotations, with an overall error of $O(|\alpha|^2/L)$. We note that the red sideband gates can be alternatively replaced by blue sideband gates, again at the expense of additional single-qubit rotations. We also note the analogy here with controlled-displacement gates realized in ion traps by simultaneous red and blue sideband drives. Here, we have demonstrated cross-compilation using first-order product formulas, but generalization to higher-order formulas is possible for more favorable error scaling.

3. Compilation between the sideband ISA and the Fock-space ISA

In this section, we demonstrate the ability to cross-compile between the sideband ISA and the Fock-space ISA. We note that having already shown how to cross-compile between (i) the Fock-space and phase-space ISAs and (ii) the sideband and phase-space ISAs, it is possible to cross-compile between sideband and Fock-space ISAs by employing the phase-space ISA as an intermediary. While this is a valid approach that proves cross-compilation is possible, in this section we discuss strategies for *direct* cross-compilation between the sideband and Fock-space ISAs to provide further insight.

a. Sideband gates from Fock-space ISA

To compile sideband operations in terms of the Fock-space ISA, it is necessary to synthesize both single-qubit rotations and sideband operations. The first of these is straightforward: by setting all rotation angles $\theta_n = \theta$ and $\phi_n = \phi$ to be identical for all n, the SQR gate $\mathrm{SQR}(\vec{\theta}, \vec{\phi})$ (Box IV.9) naturally reduces to a single-qubit rotation,

$$e^{-i\sum_{n=0}^{n_{\max}} \frac{\theta}{2}\vec{b}\cdot\vec{\sigma}\,|n\rangle\langle n|} = e^{-i\frac{\theta}{2}\vec{b}\cdot\vec{\sigma}\sum_n |n\rangle\langle n|} = e^{-i\frac{\theta}{2}\vec{b}\cdot\vec{\sigma}\otimes I} = R(\theta,\phi), \tag{D15}$$

where $\vec{b}$ is determined by the phase ϕ.

The remaining task is to compile the sideband gate $\mathrm{JC}(\theta,\varphi)$ in Eq. (A47) in terms of the SQR gate and unconditional displacement. We can modify the technique developed in Ref. [59] to achieve this. We begin by defining a particular SQR gate,

$$\hat{U}_n(\varphi) = e^{-i\frac{\theta}{2}\sigma_\varphi \hat{Q}_n}, \tag{D16}$$

$$\hat{Q}_n = \sum_{n'=0}^{n} |n'\rangle\langle n'|. \tag{D17}$$

Now note the following useful commutators:

$$\hat{J}_n = [(a^\dagger - a), \hat{Q}_n] = \sqrt{n+1}\,(|n+1\rangle\langle n| + |n\rangle\langle n+1|)\,, \tag{D18}$$

$$\hat{L}_n = [i(a^\dagger + a), \hat{Q}_n] = i\sqrt{n+1}\,(|n+1\rangle\langle n| - |n\rangle\langle n+1|)\,. \tag{D19}$$

Using unconditional displacements and these commutation relations we obtain from BCH, for small $\alpha \in \mathbb{R}$,

$$\Omega_n = D(\alpha)\hat{U}_n(0)D(-\alpha)\hat{U}_n^\dagger(0) \approx e^{-i\frac{\alpha\theta}{2}\sigma_x \hat{J}_n}, \tag{D20}$$

$$\Xi_n = D(i\alpha)\hat{U}_n\left(\frac{\pi}{2}\right)D(-i\alpha)\hat{U}_n^\dagger\left(\frac{\pi}{2}\right) \approx e^{-i\frac{\alpha\theta}{2}\sigma_y \hat{L}_n}. \tag{D21}$$

Using Trotter to add these two exponents for all n yields the blue sideband Hamiltonian evolution operator,

$$\prod_{n=0}^{n_{\max}} \Omega_n \Xi_n \approx e^{-i\alpha\theta\sum_{n=0}^{n_{\max}}\sqrt{n+1}\left(\sigma^+|n+1\rangle\langle n| + \sigma^-|n\rangle\langle n+1|\right)}. \tag{D22}$$

To realize the red sideband we would use

$$\prod_{n=0}^{n_{\max}} \Omega_n \Xi_n^\dagger \approx e^{-i\alpha\theta\sum_{n=0}^{n_{\max}}\sqrt{n+1}\left(\sigma^-|n+1\rangle\langle n| + \sigma^+|n\rangle\langle n+1|\right)}. \tag{D23}$$

Each gate in Eqs. (D22) and (D23) costs $4n_{\max}$ unconditional displacement gates and $4n_{\max}$ SQR gates to synthesize. Having synthesized unconditional qubit rotations and sideband unitaries we have formally compiled the sideband gate set using only SQR gates and uncontrolled cavity displacements from the Fock space gate set, without going through an intermediate compilation into the phase-space gate set. While this compilation is formally correct and convergent (for a finite cutoff, $n_{\max}$ on the photon number), the accuracy could be improved by using higher-order product formulas. Finding a highly efficient cross-compilation with shallow circuit depth remains an open problem.

b. Fock-space ISA from sideband gates

To compile the SQR gate sets from the sideband gate sets, we may use the intermediate constructions in Ref. [62]. For the SQR gate, note that Ref. [62] has constructed arbitrary clean SU(2) carrier rotations in the subspace $\{|0,n\rangle, |1,n\rangle\}$ for any Fock level n expanded by the upper and lower states of the qubit. Therefore, multiplying many such gates together directly leads to an SQR gate.

For the unconditional displacement gate, note that it is also possible to realize arbitrary SU(2) rotations between any pair of oscillator Fock levels [62]. Therefore, the use of many such gates can immediately realize an unconditional displacement gate on the oscillators.

4. Universality of multi-qubit–oscillator systems

In this section, we show that the phase-space ISA is universal on n-qubit–m-oscillator systems for $\forall\, n, m \in \mathbb{N}^+$ by explicitly constructing all possible generators of it as given in Eqs. (328) and (329). The idea is to start from the one-qubit–one-oscillator case and then repeatedly couple one qubit at a time to first realize an n-qubit–one-oscillator universality. Next, by repeatedly coupling one oscillator at a time to the n-qubit–one-oscillator system, universality can be shown successively on an n-qubit–j-oscillator system for $j = 2, 3, \ldots, m$ until finally the n-qubit–m-oscillator universality is achieved. We outline the details in the following. We note that a recent work has demonstrated universal multimode control by coupling a single qubit to the modes using a conditional NOT displacement gate [342], which is in a similar spirit as the CD gate defined in Box IV.7.

a. From one-qubit–one-oscillator to n-qubit–one-oscillator universality

Consider an arbitrary generator on an n-qubit–one-oscillator system of the form $G_Q h(\hat{x},\hat{p})$, where $h(\hat{x},\hat{p}) = \sum_{r,s} c_{r,s}\hat{x}^r\hat{p}^s + \mathrm{h.c.}$ acts on the oscillator and G_Q acts on the qubits. One way to construct the qubit generators is via the tensor product of individual Pauli matrices $\bigotimes_{j=1}^n \sigma_{j,\alpha_j}$ for $\alpha_j = 0, x, y, z$, where $\sigma_{j,0}$ is the identity matrix for the jth qubit. However, a better set of generators are the generalized Gell-Mann matrices [62,709], which possess cleaner commutation relationships. In the following proof,

the notation G_Q is reserved for these matrices. Since $h(\hat{x},\hat{p})$ can be written as a sum of different terms, we only need to focus on generating one such term—say, $(\bigotimes_{j=1}^{n}\sigma_{j,\alpha_j})\hat{x}^r\hat{p}^s$ for every α_j, r, s.

From the one-qubit–one-oscillator universality, we can first compile the generator $\sigma_{1,\alpha_1}\hat{x}^r\hat{p}^s$ using qubit rotations on qubit 1 and conditional displacements. The qubit-qubit entangling gate in Sec. VI C can be used to first entangle a second qubit with qubit 1 to obtain the generator $(\sigma_{1,\alpha_1}\otimes\sigma_{2,\alpha_2})\hat{x}^r\hat{p}^s$ (proper single-qubit rotation on qubit 2 is needed to precisely realize the desired Pauli σ_{2,α_2}). This process can be continued to entangle more qubits one by one until we achieve the desired $(\bigotimes_{j=1}^{n}\sigma_{j,\alpha})\hat{x}^r\hat{p}^s$. A linear combination of these then gives us the desired generator $G_Q\hat{x}^r\hat{p}^s$.

b. From n-qubit–one-oscillator to n-qubit–m-oscillator universality

We now consider entangling more oscillators with the above n-qubit–one-oscillator system to create the desired generators. Let us first note the following identity when commuting products of high-powers of $\hat{x}$ and $\hat{p}$:

$$\hat{p}^s\hat{x}^r = \hat{x}^r\hat{p}^s - is\sum_{u=0}^{r-1}\hat{x}^{r-1-u}\hat{p}^{s-1}\hat{x}^u. \tag{D24}$$

Recursively applying Eq. (D24) to terms on the right-hand side to push any powers of $\hat{p}$ to the rightmost position, it can be seen that $\hat{p}^s\hat{x}^r$ can be written as a linear combination of $\hat{x}^{r'}\hat{p}^{s'}$, where the powers r' and s' is no larger than r and s.

Now consider adding one more oscillator to achieve an arbitrary generator on the n-qubit–two-oscillator system $G_Q h(\hat{x}_1,\hat{p}_1,\hat{x}_2,\hat{p}_2)$, where $h(\hat{x}_1,\hat{p}_1,\hat{x}_2,\hat{p}_2)$ can be written in the form of Eq. (329). Again, we only need to focus on one term in the series sum of (329)—say, $G_Q(\hat{x}_1^{r_1}\hat{p}_1^{s_1}\hat{x}_2^{r_2}\hat{p}_2^{s_2} + \text{h.c.})$. Our goal is to generate this from the n-qubit–one-oscillator generator with additional gates in the phase-space ISA.

Due to the closure of the n-qubit generator set, it is always possible to find another two elements G'_Q and G''_Q such that $[G'_Q, G''_Q] = iG_Q$. Now consider the following four Hermitian generators for small t:

$$h_1 = tG'_Q(\hat{x}_1^{r_1}\hat{p}_1^{s_1} + \text{h.c.}), \tag{D25}$$

$$h_2 = tG''_Q(\hat{x}_2^{r_2}\hat{p}_2^{s_2} + \text{h.c.}), \tag{D26}$$

$$h_3 = itG'_Q(\hat{x}_1^{r_1}\hat{p}_1^{s_1} - \text{h.c.}), \tag{D27}$$

$$h_4 = itG''_Q(\hat{x}_2^{r_2}\hat{p}_2^{s_2} - \text{h.c.}), \tag{D28}$$

each of which can be realized following from the n-qubit–one-oscillator universality. More specifically, h_1 and h_3 can be realized from universality on n-qubits with oscillator 1, whereas h_2 and h_4 follow from universality on n qubits with oscillator 2. Using product formulas we have

$$\begin{aligned}&(e^{ih_1}e^{ih_2}e^{-ih_1}e^{-ih_2})(e^{ih_4}e^{ih_3}e^{-ih_4}e^{-ih_3})\\ &\quad\approx \left(e^{-it^2G_Q(\hat{x}_1^{r_1}\hat{p}_1^{s_1}\hat{x}_2^{r_2}\hat{p}_2^{s_2}+\hat{x}_1^{r_1}\hat{p}_1^{s_1}\hat{p}_2^{s_2}\hat{x}_2^{r_2}+\text{h.c.})}\right)\\ &\qquad\times\left(e^{-it^2G_Q(\hat{x}_1^{r_1}\hat{p}_1^{s_1}\hat{x}_2^{r_2}\hat{p}_2^{s_2}-\hat{x}_1^{r_1}\hat{p}_1^{s_1}\hat{p}_2^{s_2}\hat{x}_2^{r_2}+\text{h.c.})}\right)\\ &\quad\approx e^{-i2t^2G_Q(\hat{x}_1^{r_1}\hat{p}_1^{s_1}\hat{x}_2^{r_2}\hat{p}_2^{s_2}+\text{h.c.})},\end{aligned} \tag{D29}$$

where the generator on the right-hand side is the desired generator on the n-qubit–two-oscillator system.

More generally, to construct a desired generator on an n-qubit–m'-oscillator system ($m' > 2$) given by the form $G_Q(\prod_{j=1}^{m'}\hat{x}_j^{r_j}\hat{p}_j^{s_j} + \text{h.c.})$, we use the universality of an n-qubit–$(m'-1)$-oscillator system to obtain

$$h_1 = tG'_Q\left(\prod_{j=1}^{m'-1}\hat{x}_j^{r_j}\hat{p}_j^{s_j} + \text{h.c.}\right), \tag{D30}$$

$$h_3 = itG'_Q\left(\prod_{j=1}^{m'-1}\hat{x}_j^{r_j}\hat{p}_j^{s_j} - \text{h.c.}\right) \tag{D31}$$

and use the n-qubit–one-oscillator universality for oscillator m' to obtain

$$h_2 = tG''_Q(\hat{x}_{m'}^{r_{m'}}\hat{p}_{m'}^{s_{m'}} + \text{h.c.}), \tag{D32}$$

$$h_4 = tG''_Q(\hat{x}_{m'}^{r_{m'}}\hat{p}_{m'}^{s_{m'}} - \text{h.c.}). \tag{D33}$$

Combining Eqs. (D30)–(D33), and using the same product formula as in Eq. (D29), we obtain

$$\begin{aligned}&(e^{ih_1}e^{ih_2}e^{-ih_1}e^{-ih_2})(e^{ih_4}e^{ih_3}e^{-ih_4}e^{-ih_3})\\ &\quad\approx e^{-2it^2G_Q(\prod_{j=1}^{m'}\hat{x}_{m'}^{r_{m'}}\hat{p}_{m'}^{s_{m'}}+\text{h.c.})}.\end{aligned} \tag{D34}$$

This accomplishes the induction step of adding one oscillator. When $m' = m$, we have therefore succeeded in constructing an arbitrary n-qubit–m-oscillator generator $G_Q(\prod_{j=1}^{m'}\hat{x}_{m'}^{r_{m'}}\hat{p}_{m'}^{s_{m'}} + \text{h.c.})$ as given in Eq. (328) for any term in Eq. (329).

5. Oscillator-only universality: Controllability and the cubic instruction set

This set comprises the Gaussian instruction set augmented by one operator polynomial that is cubic in $\hat{x}$ and $\hat{p}$ (or $a^\dagger$ and a). Let us see how the presence of a single control term higher-order than quadratic is sufficient to allow universal oscillator control [218,291,334].

Formal discussions of universal control for systems containing oscillators require dealing with the subtleties associated with the (countably) infinite dimension of the

Hilbert space and the fact that essential operators such as position, momentum, and photon number are unbounded. For simplicity, we will not go into these questions here, but rather simply define quantum (operator) controllability [710,711] to mean that for a hybrid system consisting of an oscillator and a qubit, we have access to (i.e., can effectively generate) arbitrary unitary time evolution of the form

$$U(t) = e^{-ih_0(\hat{x},\hat{p})t}, \tag{D35}$$

where the (effective) Hamiltonian h_0 is an arbitrary (Hermitian) polynomial function of position and momentum with a large but bounded degree.

Such a general Hamiltonian is not always natively available, so our task is to understand how to generate the corresponding unitary evolution from a given finite (but continuously parametrized) set of primitive instructions. Let us assume that we can write the physically available Hamiltonian h_c in the analog (time-dependent) control form

$$h_c(\hat{x},\hat{p}) = \sum_{k=1}^{M} \mu_k(t) V_k, \tag{D36}$$

where $G = \{V_k; k = 1,\ldots,M\}$ is a discrete set of (low-order polynomial) oscillator control operators to which we have access and the $\mu_k(t)$ are temporal modulation parameters available to us.

For the qubit case, a set of n qubits has a finite Hilbert-space dimension $D = 2^n$, and the fundamental test of full controllability is whether the Lie algebra generated by the control set G is full rank; i.e., is the Lie Algebra $\mathfrak{su}(D)$ [710,711]. If so, then there exist modulation functions $\{\mu_k(t); k = 1,\ldots,M\}$ such that time evolution under h_c generates the full Lie group SU(D). The connection between the Lie algebra based on operator commutators and the Lie group of unitary transformations can be seen from the presence of a commutator in the Heisenberg equation of motion for an operator. More specifically, it is useful to consider the BCH identity

$$\begin{aligned} e^{i\theta_j V_j} e^{i\theta_k V_k} e^{-i\theta_j V_j} &= \exp\{i\theta_k\{V_k + (i\theta_j)[V_j, V_k] \\ &\quad + \frac{1}{2!}(i\theta_j)^2[V_j,[V_j,V_k]] + \ldots\}\}, \end{aligned} \tag{D37}$$

from which it follows (for small angles θ_j and θ_k) that

$$e^{i\theta_j V_j} e^{i\theta_k V_k} e^{-i\theta_j V_j} e^{-i\theta_k V_k} \approx \exp\{(\theta_k\theta_j)[V_k, V_j]\}. \tag{D38}$$

Such so-called "product formulas" [335–341] allow us to multiply group generators (in the form of commutators) to generate new effective Hamiltonian terms. Equivalently, if the individual unitaries on the left-hand side above are taken from a discrete (but parametrized) instruction set, this equation shows us how to chain together elements of the instruction set to create new and more complex unitaries. Similarly, various Trotter-Suzuki formulas [712] such as

$$e^{i\frac{\theta_j}{2}V_j} e^{i\theta_k V_k} e^{+i\frac{\theta_j}{2}V_j} \approx e^{i(\theta_j V_j + \theta_k V_k)} \tag{D39}$$

allow us to sum Hamiltonian terms corresponding to different elements of the instruction set.

With these methods, expanded in Sec. VI, one can create (approximations to) any unitary in the group SU(D), where $D = 2^n$ is the Hilbert space of n qubits, provided that the set G of control terms (or equivalently, the instruction set) is sufficiently rich to give controllability. That is, being able to create $D^2 - 1$ linearly independent traceless Hamiltonian terms (traceless since adding a constant to the Hamiltonian does not grant additional controllability).

The case of infinite Hilbert-space dimension is more subtle. Nominally, one only needs the algebra generated from the Lie brackets of the control Hamiltonians to be infinite, but this does not preclude the possibility of missing important terms. Even if this algebra is infinite, there could be one or more symmetry operators that commute with all the generators. For example, if the generators happen to commute with photon-number parity, you will have control over only half of the infinite number of states. Since we are defining controllability here as being able to generate effective Hamiltonian terms that are arbitrary polynomials of some bounded degree, this subtlety can be explicitly checked and avoided. For an in-depth discussion of these matters, we refer to Ref. [713], which explicitly considers situations where the operator algebra closes.

As mentioned above, any nonquadratic Hamiltonian term is enough to achieve controllability of an oscillator [218,291,334]. Perhaps the simplest such term is $\hat{x}^3$ which, together with Gaussian operations, defines the Cubic instruction set. To understand why a single term in the Hamiltonian higher than quadratic is sufficient for universality, we can examine the commutation relations for different degrees of polynomial (using standard dimensionless units for which $[\hat{x},\hat{p}] = +i$):

$$[\hat{x}, \hat{x}^m\hat{p}^n] = in\hat{x}^m\hat{p}^{n-1}, \tag{D40}$$

$$\begin{aligned} [\hat{x}^2, \hat{x}^m\hat{p}^n] &= \hat{x}[\hat{x}, \hat{x}^m\hat{p}^n] + [\hat{x}, \hat{x}^m\hat{p}^n]\hat{x} \\ &= in(\hat{x}^{m+1}\hat{p}^{n-1} + \hat{x}^m\hat{p}^{n-1}\hat{x}), \end{aligned} \tag{D41}$$

$$[\hat{x}^3, \hat{x}^m\hat{p}^n] = \hat{x}[\hat{x}^2, \hat{x}^m\hat{p}^n] + [\hat{x^2}, \hat{x}^m\hat{p}^n]\hat{x}, \tag{D42}$$

$$\begin{aligned} &= in(\hat{x}^{m+2}\hat{p}^{n-1} + \hat{x}^{m+1}\hat{p}^{n-1}\hat{x} \\ &\quad + \hat{x}^{m+1}\hat{p}^{n-1}\hat{x} + \hat{x}^m\hat{p}^{n-1}\hat{x}^2). \end{aligned} \tag{D43}$$

As is evident, the degree of a polynomial in $\hat{x}$ and $\hat{p}$ is preserved by its commutator with a quadratic term $\hat{x}^2$.

In contrast, the commutator of a degree-$(m+n)$ polynomial with a cubic term $\hat{x}^3$ yields a polynomial of degree $m+n+1$. Similarly, we can prove the same properties for $\hat{p}, \hat{p}^2, \hat{p}^3$ and $\hat{x}^m\hat{p}^n$. Thus, the algebra generated by the Lie brackets of these terms is infinite in the presence of a cubic or higher-order polynomial. As a specific example, consider

$$[\hat{x}^3, \hat{p}^2] = i3(\hat{p}\hat{x}^2 + \hat{x}^2\hat{p}), \tag{D44}$$

$$[\hat{x}^3, [\hat{x}^3, \hat{p}^2]] = -18\hat{x}^4. \tag{D45}$$

Using this inductive proof, we have shown that the algebra generated by the control Hamiltonian set which includes quadratic Hamiltonians with only one cubic Hamiltonian can generate arbitrary order polynomials enabling universal control.

The above discussion shows us how to generate the oscillator-only control unitary in Eq. (D35). From an experimental physics perspective, we require only the harmonic Hamiltonian $\hat{p}^2 + \hat{x}^2$, linear drives on the oscillator that couple to $\hat{x}$ (and/or $\hat{p}$), and one additional higher-order term such as $\hat{x}^3$ to generate an infinite algebra which can synthesize arbitrary unitary operations. The presence of the $\hat{x}^3$ (or higher) term makes the harmonic oscillator anharmonic so that each transition between adjacent levels has a different frequency, making the transitions individually addressable through resonant drives of selected frequencies, amplitudes, and phases. This physical picture of the origin of controllability gives nice intuition to accompany the mathematics of the infinite algebra generated by the Lie bracket of the control set. Experimental oscillator-only control using cubic interactions has been demonstrated for microwave circuits in Refs. [291,334].

Arguments similar to those above show that multimode universal control only requires the addition of Gaussian beam-splitter gates between physically adjacent oscillators (since the contents of any two oscillators can be made physically adjacent through SWAP gates built from the beam-splitter interaction). The beam splitter allows one to perform coherent rotations between mode operators (see Box IV.4), and hence create arbitrary polynomials of multiple oscillator variables. We have thus (nonrigorously) established the universality of the Cubic instruction set for oscillator-only control.

APPENDIX E: MEASUREMENTS AND TOMOGRAPHY

We have established instruction sets from which arbitrary universal operations on the hybrid bosonic quantum processor can be achieved, and which enable useful computation and information processing tasks to be accomplished. To extract information from the bosonic quantum processor after the computation, measurements are needed. In this appendix, we shall delve more into aspects of measurement and discuss how measurements may be performed on these hybrid systems.

Measurements can be carried out on either qubits or oscillators, although direct measurements on oscillators are difficult because photon-number detectors might be required. Surprisingly, by leveraging the native couplings between the qubit and oscillator, indirect measurement of the oscillator can be achieved by simple qubit measurement. We start by presenting a few basic measurement schemes based on qubit measurement in Appendixes E 1–E 4 to extract oscillator photon numbers and parities. We then discuss a qubit-free homodyne measurement scheme in Appendix E 5 that utilizes a simple oscillator-oscillator gate. In cases in which more information in the phase space of oscillators needs to be read out, Appendixes E 6 and E 7 describe three different ways of performing tomography of the phase space.

Finally, just as in DV system where random Clifford circuits combined with measurement are used to benchmark the performance of the circuit, Gaussian operations combined with the measurement techniques in this appendix can be used to benchmark CV systems. We discuss random benchmarking of CV systems in Appendix E 8. We note that these measurement protocols can be reformulated in the framework of classical shadow tomography for continuous-variable systems [399,403] and optimizations in the case of Gaussian states can also be considered [714]. Benchmarking of displacement operations using a random sequence of oscillator displacements, followed by an inverting displacement that nominally returns the system to its starting state, has been considered in Appendix E 8 and Ref. [715].

1. Hybrid qubit-cavity measurements

The dispersive Hamiltonian for the coupled qubit-cavity system in Eq. (A24) commutes with both σ_z and $\hat{n} = A^\dagger A$ and can therefore be used for QND readout of both the state of the qubit and the state of the cavity. Indeed, one of the major advantages of storing quantum information in the large Hilbert space of an oscillator is that state tomography is much easier and more accurate than it would be (at least with present technology) in a qubit system with the corresponding Hilbert-space dimension.

2. Dispersive readout of the qubit

The dispersive Hamiltonian in Eq. (A24) tells us that the cavity resonance frequency is shifted to

$$\omega_R \pm \frac{\chi}{2}, \tag{E1}$$

depending on the state of the qubit. By measuring the phase shift of a microwave pulse reflected from the cavity one can infer the resonance frequency of the cavity and hence

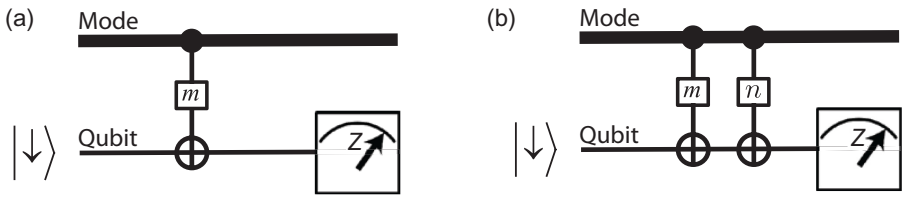


FIG. 42. (a) A quantum circuit using the cavity-conditioned qubit rotation gate $U_m(\pi/2, 0)$ and a Z measurement of the ancilla qubit to measure the operator $\hat{P}_m$ that answers the question "Does the cavity contain m photons?" (b) A circuit using $U_n(\pi/2, 0)U_m(\pi/2, 0)$ and a Z measurement of the ancilla qubit to measure the operator $\hat{P}_m + \hat{P}_n$ that answers the question "Does the cavity contain *either* m or n photons?"

determine the state of the qubit (i.e., measure σ_z) with high ($\sim$99% $-$ $-$99.9%) fidelity [29,77,716–718].

3. Cavity photon-number measurement

We can ask the question "Is the cavity photon number m?" by measuring the value of the projector onto Fock state $|m\rangle$:

$$\hat{P}_m = |m\rangle\langle m|. \tag{E2}$$

This can be achieved in a two-step process [147], illustrated in Fig. 42(a), in which we first apply the SQR gate (Box IV.9),

$$U_m(\pi, 0) = e^{i\frac{\pi}{2}\sigma_x \hat{P}_m} = (\hat{I} - \hat{P}_m) + i\sigma_x \hat{P}_m, \tag{E3}$$

which flips the qubit only if there are m photons in the cavity. We then read out the qubit state using the resulting dispersive shift in the (possibly separate readout) cavity frequency.

Suppose that the initial cavity state is

$$|\psi\rangle = \sum_{j=0}^{N_{\max}} \psi_j |j\rangle. \tag{E4}$$

If the measurement result on the ancilla qubit is $Z = +1$, the state collapses to $|\uparrow\rangle|m\rangle$. If the measurement result on the ancilla qubit is $Z = -1$, the state collapses to

$$|\psi\rangle = |\downarrow\rangle \frac{1}{\sqrt{1 - |\psi_m|^2}} \sum_{j \neq m} \psi_j |j\rangle. \tag{E5}$$

These gates have been used in Sec. V C to present schemes that read out or entangle dual-rail qubits using an ancilla qubit. Interestingly, we can apply multiple versions of the cavity-conditioned bit-flip gate to ask other questions about the cavity state. As illustrated in Fig. 42(b), the compound operator

$$U_n(\pi, 0)U_m(\pi, 0) = e^{i\frac{\pi}{2}\sigma_x[\hat{P}_m + \hat{P}_n]}, \tag{E6}$$

flips the qubit if the photon number is *either* n or m (since it cannot be both). These two unitaries can be executed in parallel by applying two simultaneous drive tones on the ancilla qubit. Generalizing to a product of many unitaries, the most general projector we can measure is

$$\hat{P}_{\vec{c}} = \sum_{m=0}^{N_{\max}} c_m \hat{P}_m, \tag{E7}$$

where $\vec{c} = (c_0, c_1, \ldots, c_{N_{\max}})$ is a binary-valued vector, with $c_j \in \{0, 1\}; j = 0, N_{\max}$.

This is a powerful gate. One could use it to, e.g., flip the ancilla bit if the photon number is less than (or greater than) a selected value. One could then use the qubit-conditioned cavity displacement gate (Box IV.7) to displace the cavity conditioned on the photon number being less than (or greater than) the selected value. As we describe further below, it can also be used in a highly efficient binary search for the photon number.

a. Cavity photon-number parity measurement

A particularly interesting and useful example of the general projectors available in the class $\hat{P}_{\vec{c}}$ is the photon-number parity operator

$$\hat{\Pi} = e^{i\pi\hat{n}} = \sum_m (-1)^m \hat{P}_m, \tag{E8}$$

which can be mapped onto the qubit for measurement using the unitary

$$U_{\vec{c}} = e^{i\frac{\pi}{2}\sigma_x \hat{P}_{\vec{c}}} \tag{E9}$$

with $\vec{c} = (0, 1, 0, 1, 0, 1, 0 \ldots)$ so that the qubit is flipped only if the photon-number parity is odd. The photon parity operator plays an essential role in cavity state tomography (to be described below) and in QEC using bosonic codes [63,416]. The actual pulse sequence used in practice to measure the parity [240] is different from that outlined above but produces the same result (see Sec. V C). The circuit used takes advantage of the conditional cavity rotation afforded directly under time evolution with the dispersive Hamiltonian in Eq. (A24). Time evolution under the

dispersive coupling

$$\frac{\chi}{2}\sigma_z A^\dagger A \tag{E10}$$

for a time $t = -\theta/\chi$ (χ is typically negative) is described by the "conditional rotation" unitary (Box IV.8)

$$\mathrm{CR}(\theta) = e^{-i\frac{\theta}{2}\sigma_z\hat{n}}. \tag{E11}$$

Applying the phase-estimation protocol to this unitary will give us a powerful way to measure the full photon-number distribution along with photon-number parity and "superparity" measurements [147] described further below in Appendix E 4.

To see the connection to the parity operator, consider the conditional rotation gate for $\theta = \pi$ (Box IV.8),

$$\mathrm{CR}\,(\pi) = e^{-i\frac{\pi}{2}\hat{n}}\,\mathrm{CP}, \tag{E12}$$

where

$$\mathrm{CP} = e^{i\pi|e\rangle\langle e|\hat{n}} = |g\rangle\langle g|\hat{I} + |e\rangle\langle e|\hat{\Pi} \tag{E13}$$

is the qubit-controlled cavity parity operation, which is active only if the qubit is in its excited state $|e\rangle$. The additional rotation of the cavity state in Eq. (E12) corresponds for our purposes to an irrelevant frame choice. The circuit shown in Fig. 43, familiar from the phase-estimation problem, takes advantage of the phase kickback on the ancilla from the controlled parity unitary to entangle the qubit state with the cavity parity, thus enabling the cavity parity to be measured by reading out the qubit state. Effectively, the controlled parity gate rotates the qubit state by angle $\hat{n}\pi$ around the z axis, taking it from $|+\rangle$ to $|-\rangle$ if the photon number is odd and returning it to $|+\rangle$ if the photon number is even. The measurement result $z_0 = \pm 1$ after the final Hadamard gate determines the parity bit of the photon number via

$$b_0 = \frac{1 - z_0}{2}. \tag{E14}$$

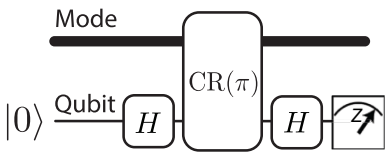


FIG. 43. A quantum circuit using phase kickback from the controlled oscillator rotation, $\mathrm{CR}(\pi)$, to entangle the ancilla-qubit state with the photon-number parity of the mode. Measurement of the qubit state then yields a QND measurement of the parity of the photon number.

4. Cavity photon-number "superparity" measurement

We can use phase estimation to generalize the above result to obtain the "superparities" b_j that yield the full binary number representation of the photon number,

$$n = \sum_{j=0}^{j\max} 2^j b_j. \tag{E15}$$

If the measurement result for b_0 is known, then we can determine the next binary digit by measuring $b_1 = [(n - b_0) \bmod 4]/2$. This utilizes the same circuit as in Fig. 43, with the $\mathrm{CR}(\pi)$ gate replaced by the following unitary:

$$\mathrm{CR}\left(\frac{\pi}{2}\right) R_z\left(-b_0\frac{\pi}{2}\right), \tag{E16}$$

where the R_z is the single-qubit rotation in Table VIII (again executable "in software" via a frame change),

$$R_z(\theta) = e^{-i\frac{\theta}{2}\sigma_z}. \tag{E17}$$

It can be shown that

$$b_1 = \frac{1 - z_1}{2}, \tag{E18}$$

where z_1 is the measurement outcome of the qubit.

Similarly, once the measurement results for b_0 and b_1 are known, b_2 is obtained using the unitary

$$\mathrm{CR}\left(\frac{\pi}{4}\right) R_z\left(-[b_0 + 2b_1]\frac{\pi}{4}\right) \tag{E19}$$

and b_3 is obtained using the unitary

$$\mathrm{CR}\left(\frac{\pi}{8}\right) R_z\left(-[b_0 + 2b_1 + 4b_2]\frac{\pi}{8}\right). \tag{E20}$$

In general, b_k can be obtained using the unitary

$$\mathrm{CR}\left(\frac{\pi}{2^k}\right) R_z\left(-\left[\sum_{j=0}^{k-1} 2^j b_j\right]\frac{\pi}{2^k}\right). \tag{E21}$$

Importantly, since all these unitaries commute with each other we can make a QND measurement of all the binary digits in a single shot (single state preparation) using the quantum circuit with classical feedforward shown in Fig. 44. This is very powerful because it effectively performs a binary search to obtain the value of n in only $\sim \log N_{\max}$ steps. It is exponentially more efficient than using the circuit in Fig. 42(a) to sequentially ask "Is n equal to zero?, Is n equal to 1? Is n equal to two?...?" and so on.

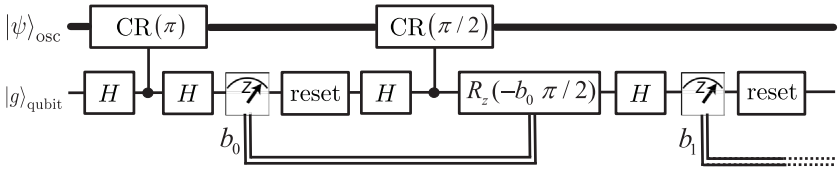


FIG. 44. A quantum circuit with classical feedforward of the ancilla measurement results $b_0, b_1, \ldots$ to carry out the phase-estimation algorithm to measure the binary representation $(\ldots, b_1, b_0)$ of the photon number.

As an alternative to the superparity measurements which determine the least-significant bit (LSB) first, we can perform a more traditional binary search for the photon number value by starting from the most significant bit (MSB). If we consider a range of photon numbers $[0, N_{\max}]$, with $N_{\max}+1$ being an integer power of 2, then we could use the projector in Eq. (E7) to ask "Is the photon number in the upper half of the range?" (i.e., in the interval $[(N_{\max}+1)/2, N_{\max}]$). Given the answer, we then recursively ask"Is the photon number in the upper half of the interval containing the photon number?" This is just like classical binary search.

Proceeding in the spirit of Eq. (E9), let us define a generalized CNOT gate controlled by the state of the cavity using a family of binary vectors

$$\vec{c}_2 = (00001111), \tag{E22}$$

$$\vec{c}_1 = (00110011), \tag{E23}$$

$$\vec{c}_0 = (01010101). \tag{E24}$$

In this simple example, we take the maximum photon number to be $N_{\max} = 7$ and the vector $\vec{c}_j$ to be the jth basis vector of the Walsh-Hadamard transform. Now consider the use of this gate in the circuit depicted in Fig. 45. Note that in analyzing this sequence of measurements, we will assume that the auxiliary qubit is initialized in state $|0\rangle$ at the beginning of the protocol but is *not* reset after each measurement.

From the vectors $\{\vec{c}_2, \vec{c}_1, \vec{c}_0\}$ we see that the measurement result a_2 tells us whether the photon number is in the upper half $[4, 7]$ ($a_2 = 1, |e\rangle_{\text{qubit}}$) or the lower half $[0, 3]$ ($a_2 = 0, |g\rangle_{\text{qubit}}$) of the considered range of $[0, 7]$.

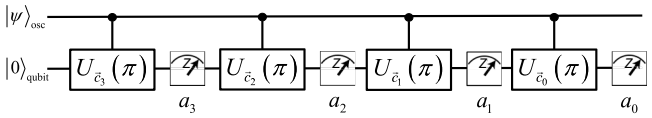


FIG. 45. A quantum circuit for binary search to measure photon number $\hat{n}$. Unlike the phase-estimation protocol, this does not require classical feedforward of the ancilla measurement results $a_0, a_1, a_2, \ldots$. The binary representation of the photon number can be determined directly from the sequence of measurement results. Note that the auxiliary qubit is assumed to be initialized in $|g\rangle = |0\rangle$, but is not reset after each measurement. Reference [147] uses essentially this protocol but *does* reset the qubit after each measurement.

If the measurement result in a_1 is different than a_2, the ancilla spin has flipped and therefore the photon number is in the upper half of the relevant range. Similarly, if the measurement result a_0 is different than a_1, the spin has again flipped and the photon-number parity must be odd. Thus we obtain the binary digits representing the photon number:

$$b_2 = a_2, \tag{E25}$$

$$b_1 = a_2 \oplus a_1, \tag{E26}$$

$$b_0 = a_1 \oplus a_0, \tag{E27}$$

where $\oplus$ corresponds to addition modulo 2.

The fact that active feedforward is not required is also a significant advantage. This is (essentially) the technique used in the boson sampling experiment of Ref. [147], except that those authors have used numerically generated optimal control pulses on the qubit to map the binary bits in the photon number onto an ancilla. The other difference in Ref. [147], is that the authors have reset the ancilla to its ground state after each measurement to minimize errors from the energy relaxation of the ancilla.

5. Homodyne detection

Homodyne detection can be seen as a projective measurement of the phase-space quadratures. The setup for homodyne detection includes a beam splitter with a transmission coefficient (t) and a reflection coefficient (r), and two photodetectors (see Fig. 46). The quantity of interest here is

$$\langle n_3\rangle - \langle n_4\rangle, \tag{E28}$$

the difference of the mean photon numbers detected by the two detectors for a given input state. The state to be measured ($|\psi\rangle$) is sent through one port while the coherent state ($|\alpha\rangle$) is sent through other input port. In the Heisenberg picture, the photon-number operators at the output ports are

$$n_3 = \frac{1}{2}\left[|t|^2 n_1 + |r|^2 n_2 - i(t^* r a_1^\dagger a_2 - r^* t a_2^\dagger a_1)\right], \tag{E29}$$

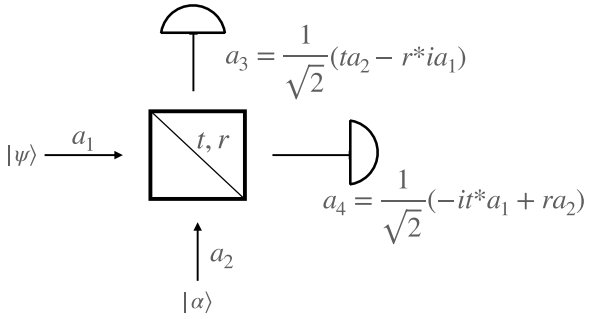


FIG. 46. Homodyne detection involves a beam splitter (BS) with transmission coefficient t and reflection coefficient r, and two detectors at the output ports of the BS. For balanced homodyne detection, $|t|^2 = |r|^2 = \frac{1}{2}$.

$$n_4 = \frac{1}{2}\left[|t|^2 n_2 + |r|^2 n_1 - i(t^* r a_2^\dagger a_1 - r^* t a_1^\dagger a_2)\right], \quad \text{(E30)}$$

and the mean photon number is given by

$$\langle n \rangle = \langle \psi, \alpha | \, n \, | \psi, \alpha \rangle \,, \quad \text{(E31)}$$

$$\langle n_3 \rangle = \frac{1}{2}\Big[|t|^2 \langle \psi | \, n_1 \, | \psi \rangle + |r|^2 |\alpha|^2 - i \langle \psi | \, (tr^* \alpha) a_1^\dagger - (tr^* \alpha)^* a_1 \, | \psi \rangle \Big]. \quad \text{(E32)}$$

Assuming a balanced homodyne detection with $t = r$ and $\alpha = |\alpha| e^{i\phi}$,

$$\langle n_3 \rangle - \langle n_4 \rangle = |\alpha| \, \langle \psi | \, (a^\dagger e^{i\phi} - a e^{-i\phi}) \, | \psi \rangle \,. \quad \text{(E33)}$$

Thus, the mean photon number measures the mean value of a phase-space quadrature along an axis dependent on the phase of the coherent state. This is a multi-purpose method. Homodyne detection on identically prepared states reconstructs the quadrature probability distribution. Repeating these measurements for all different phases is equivalent to computing the marginal integrals of the Wigner function, an important visualization tool for quantum states, which then can be used to obtain the Wigner function (and the density operator) via a Radon transform [719–722].

Homodyne detection of traveling photons at optical frequencies can be relatively efficient, but at microwave frequencies, even nearly quantum limited linear amplifiers add enough noise to complicate state reconstruction, but with sufficient care, this can be done relatively well [598]. For microwave photons trapped in stationary modes in high-Q resonators, it is possible to avoid these difficulties by synthesizing homodyne and heterodyne measurements by mapping the cavity state onto an auxiliary qubit and measuring the qubit [145].

6. Wigner and characteristic function tomography

A key step in validating quantum processes is to estimate the Wigner function for a quantum state. This allows us to assess whether there are errors in our quantum circuits. In this section and the following, we review processes for estimating the characteristic function, the Wigner function, and the Husimi Q-function.

The phase-space density functions W (Wigner), Q, and P are different Fourier transforms of characteristic functions. The characteristic function for a bosonic system is the expectation value of the displacement operator at each point in phase space for a given density matrix ρ,

$$C(\alpha) = \mathrm{Tr}\left[\rho D(\alpha)\right], \quad \text{(E34)}$$

where $\mathcal{D}(\alpha)$ is the (unconditional) displacement operator defined in Box IV.1. It can be seen that for a pure state, this quantity is an autocorrelation function of the quantum state in phase space.

How does one measure the overlap of a state with a displaced version of itself? This can be done by replacing the unconditional displacement operator with a conditional displacement operator controlled by the Z value of an ancilla qubit. Measurement of the phase kickback on the ancilla yields a one-bit estimator expectation value of the displacement operator. Initializing the ancilla in $|+\rangle$ and measuring $\langle X \rangle$ ($\langle Y \rangle$) gives a one-bit estimate of the real (imaginary) part of the characteristic function. Repeated shots of these two types of measurements can be averaged to obtain the complex expectation value $C(\alpha)$ [60,108].

$C(\alpha)$ is often referred to as the symmetric characteristic function to distinguish it from its normal (C_n) and antinormal (C_{an}) order counterparts,

$$C_n(\alpha) = \mathrm{Tr}[\rho e^{\alpha a^\dagger} e^{-\alpha^* a}], \quad \text{(E35)}$$

$$C_{an}(\alpha) = \mathrm{Tr}[\rho e^{-\alpha^* a} e^{\alpha a^\dagger}]. \quad \text{(E36)}$$

The Fourier transform of the symmetric characteristic function $C(\alpha)$ is called the Wigner function and is defined in Eqs. (249) and (250) in the main text and is restated here for convenience:

$$W(\beta) = \frac{1}{\pi^2} \int d^2\alpha \, C(\alpha) e^{\beta\alpha^* - \alpha\beta^*}. \quad \text{(E37)}$$

The density matrix ρ of a system contains the full information needed to predict the (statistical) outcomes of any measurement on that system. The Wigner function is a quasiprobability distribution in phase (position and momentum) space that contains the equivalent information to the density matrix but is easier to measure. Moreover, practical methods for Wigner function estimation have been developed in the literature [723].

In the same dimensionless complex coordinates that we use to describe the amplitude of a coherent state, the relationship of the Wigner function to the density

matrix is given by the following remarkable identity due to Lutterbach and Davidovich [207]:

$$W(\beta) = \frac{2}{\pi}\mathrm{Tr}\left[\rho\mathcal{D}^\dagger(-\beta)\hat{\Pi}\mathcal{D}(-\beta)\right], \tag{E38}$$

where $\hat{\Pi}$ is the photon-number parity operator as defined in Eq. (E8).

This formula tells us that the value of the Wigner function at the point β in the complex plane is determined by displacing that point to the origin and then measuring the expectation value of the photon number parity. This remarkable result tells us that state tomography on the oscillator is relatively easy to obtain, even though the effective dimension of the Hilbert space $N_{\max}+1$ can be quite large (though see the following paragraph for an important caveat regarding sampling complexity). In addition, the basic procedure does not change as $N_{\max}$ increases. It is worth noting that Eq. (E38) shows that determining only the photon-number parity, instead of the full photon number distribution, suffices for Wigner function tomography. However, Refs. [724,725] have demonstrated that the redundant information available in the full photon-number distribution can significantly improve the robustness of state tomography via the Wigner and characteristic functions.

Despite the simplicity of determining $W(\beta)$ for a single point in phase space, the cost of performing full Wigner function characterization using the above relation faces challenges if we wish to learn the entire Wigner function of even a single oscillator. The reason for this is that an infinite number of values of β are needed to ensure that the Wigner function can be accurately estimated at any β. This means that we will ultimately be forced to estimate the value of the Wigner function at a given point by using a scheme such as a Lagrange interpolation, wherein the value of the Wigner function at an arbitrary point is estimated from values learned at points on a grid. Assuming that the Wigner function is differentiable $n+1$ times, from interpolation theory (in essence, Taylor's remainder theorem), the error in estimating $W(\beta)$ at a particular point β from surrounding grid points with a polynomial function $P_n(\beta)$ with constant grid spacing h is

$$|W(\beta) - P_n(\beta)| \leq \frac{\max_{y,\alpha}|D_\alpha^{n+1}W(y)|h^n}{(n+1)!}, \tag{E39}$$

where D_α is the directional derivative operator in the direction α. This suggests that if we wish to estimate the Wigner function at an arbitrary point within error ϵ for an unknown Wigner function that is differentiable at least $n+1$ times, then it suffices to choose

$$h \in O\left(\frac{n\epsilon^{1/n}}{\max_{y,\alpha}|D_\alpha^{n+1}W(y)|^{1/n}}\right). \tag{E40}$$

If we assume that the Wigner function is (within sufficiently small error) supported on a compact region $[-\beta_{\max},\beta_{\max}]^{\otimes 2}$, then since the Wigner function for a single oscillator is a 2D function, the total number of grid points needed within this space scales as

$$N_1 = \left(\frac{\beta_{\max}}{h}\right)^2 \subseteq O\left(\left(\frac{\beta_{\max}\max_{y,\alpha}|D_\alpha^{n+1}W(y)|^{1/n}}{n\epsilon^{1/n}}\right)^2\right). \tag{E41}$$

Repeating the same analysis for a system with $p \geq 1$ oscillators (but taking y, β, and α now to be vectors in a $2p-$dimensional space), we conclude that the number of points needed to extrapolate the Wigner distribution is

$$\begin{aligned} N_p &= \left(\frac{\beta_{\max}}{h}\right)^{2p} \\ &\subseteq O\left(\left(\frac{p\beta_{\max}\max_{y,\alpha}|D_\alpha^{n+1}W(y)|^{1/n}}{n\epsilon^{1/n}}\right)^{2p}\right). \end{aligned} \tag{E42}$$

Choosing for simplicity $n = 2p$, we see that the scaling is

$$N_p = O\left(\frac{\max_{y,\alpha}|D_\alpha^{2p+1}W(y)|}{\epsilon}\left(\frac{\beta_{\max}}{2}\right)^{2p}\right). \tag{E43}$$

We see from this that unless the derivatives of the Wigner function become very small as p increases, the number of points that need to be accurately estimated to approximately learn the Wigner function on p oscillators scales exponentially with p. This can be problematic even for physical states. For example, if we consider a coherent state in one dimension of the form

$$|\psi\rangle^{\otimes j} = \left(\int dq\, e^{-q^2/2\sigma^2}|q\rangle/\sqrt{2\pi}\sigma\right)^{\otimes j}, \tag{E44}$$

then for small σ we have that the norm of the derivatives of the state grows exponentially with the number of copies j. This illustrates the often-observed fact with conventional tomography that the cost grows exponentially with the number of systems.

This exponential increase of the number of measurement settings with the number of qumodes in the system exactly matches what we would expect from qubit tomography of a density matrix (with no promises about the rank of the density operator) that 4^p distinct sets of measurements of Pauli operators would be needed to reconstruct a p-qubit state [726]. This suggests that the situation is analogous, if more complex, to the qubit case, as we have to approximately learn the density matrix because we do not have a tomographically complete set of measurements for our continuous scheme. Thus while individual components of the Wigner function can be easily estimated, learning the

overall Wigner function for a generic p-mode oscillator system will likely be forever out of reach owing to the previous argument about the scaling with the derivatives of the tensor products between coherent states. However, for cases in which prior knowledge is provided about the form of the Wigner function, more efficient reconstruction schemes can be found and the optimization of such reconstruction schemes remains an important open question.

7. The Husimi Q-function

The Fourier transform of the antinormal order characteristic function $C_{an}(\alpha)$ in Eq. (E36) is called the Husimi Q-function, and is defined as

$$Q(\beta) = \frac{1}{\pi^2}\int d^2\alpha\, C_{an}(\alpha) e^{\beta\alpha^* - \alpha\beta^*}. \tag{E45}$$

The Husimi Q-function can be more convenient to use in certain cases since, unlike W, this quantity is always positive (but therefore is less suited to study the nonclassicality of quantum states). In terms of the density matrix, it is given by

$$Q(\beta) = \frac{1}{\pi}\mathrm{Tr}\left[|0\rangle\langle 0|\, \mathcal{D}^\dagger(-\beta)\rho\mathcal{D}(-\beta)\right]. \tag{E46}$$

Thus, it is the average of the projector onto the vacuum state of the state displaced by β.

Similar to the definition of the Q-function, the P-function is the Fourier transform of $C_n(\alpha)$ in Eq. (E35). It is unnormalizable and difficult to use, and hence will not be discussed in the present work.

8. Calibrating controlled displacements

As discussed in Sec. IV F 2, benchmarking hybrid quantum processors is not as easy as qubit systems. However, some simple calibrations on hybrid CV-DV gates can still be performed. One example is that the gate sequence used in Eq. (351) in Sec. VI can be used to calibrate the controlled-displacement gate.

It works as follows. Consider conjugating Eq. (351) using Hadamard gates to construct the following $U_{\text{calibration}}$:

$$U_{\text{calibration}} = H\left[D(-i\alpha)D_c(\alpha,-\alpha)D(i\alpha)D_c(-\alpha,\alpha)\right]H. \tag{E47}$$

As we have seen from Eq. (352) and Fig. 23, the $|0\rangle$ state of the qubit causes the displacements to drive the oscillator around the counterclockwise closed loop in phase space, which means that the $|0\rangle$ state acquires phase $\varphi_{|0\rangle} = -2|\alpha|^2$. Similarly, the $|1\rangle$ state of the qubit causes the displacements to drive the oscillator around the clockwise closed loop in phase space, which means that the $|1\rangle$ state acquires phase $\varphi_{|1\rangle} = +2|\alpha|^2$. The net effect is that the qubit undergoes a rotation around the z axis through an angle $\varphi = \varphi_{|0\rangle} - \varphi_{|1\rangle} = -4|\alpha|^2$. The conjugation by the Hadamard gates converts this to a rotation around the x axis leading to the "Berry fringe" oscillations

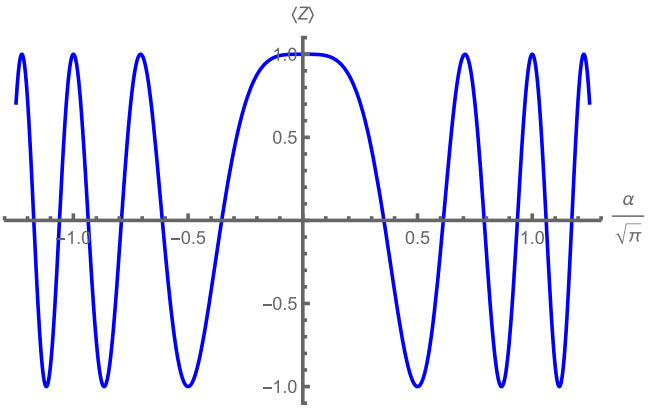


FIG. 47. Average measurement results for the qubit measurement outcome using the calibration given in Eq. (E47), demonstrating "Berry fringe" oscillations due to the phase kickback of the noncommuting displacements in phase space. The known form of these oscillations from Eq. (E48) can be used to accurately calibrate the size of displacements in phase space.

$$\langle Z\rangle = \cos\left(4\alpha^2\right), \tag{E48}$$

as is shown in Fig. 47. Therefore, by accumulating enough shots on the qubit measurement, we can get an estimation of $\langle Z\rangle$ to infer the actual displacement α performed by the displacement gates. Note that if we make the unconditional displacements in Eq. (E47) be constant $\pm i\beta$ independent of α, then we obtain a simple sinusoidal oscillation

$$\langle Z\rangle = \cos\left(4\alpha\beta\right). \tag{E49}$$

Reference [715] has taken a different approach, considering a set of random displacements applied to an initial vacuum state. This has been followed by an inverting displacement that nominally returns the system to the vacuum state. Measurement of the probability of being in the vacuum state (or more generally, measurement of the photon-number distribution) provides information on noise and errors in the displacement operations.

APPENDIX F: COMPARISON BETWEEN ANALYTICAL AND NUMERICAL COMPILATION

In this appendix, we compare the cost of using numerically optimized circuits described in Sec. VI E against analytical compilations using approximate techniques discussed in Sec. VI D. For this comparison, we employ two figures of merit:

(a) The first is the circuit depth or number of layers N_{circuit} used in the numerical optimization scheme in Sec. VI E.
(b) The second is the circuit duration, or the total time taken by the circuit T_{circuit} used in Sec. V C. Note that the duration of the conditional displacements is lower bounded by $T_{|\gamma_i|<0.024} = 48$ ns. This duration includes the necessary components for an echoed conditional displacement, an unconditional displacement $|\alpha_0|$ (24 ns), and a midcircuit qubit rotation (24 ns). For details see Appendix A and Ref. [60].

We use the preparation of a Fock state $|1\rangle$ and a squeezed vacuum as examples to study the efficiency of various approximate methods against numerical methods using phase-space ISA. Such comparisons can be used as a figure of merit to benchmark the analytical schemes.

1. Trotterization

We can prepare Fock states with access to the JC Hamiltonian. For the phase-space ISA, the JC Hamiltonian can be engineered as discussed in Sec. VI D 1 via Trotterization. We show the operator fidelity to achieve this using first-order and second-order Trotter-Suzuki formulas (TS) in Fig. 48. Once available, the JC Hamiltonian can be used to prepare arbitrary states from the vacuum state $|0\rangle_{\text{vac}}$ using the Law-Eberly protocol [480]. The simplest of these tasks is to prepare a Fock $|1\rangle$ state. In Fig. 49(a), we show a comparison between the Law-Eberly protocol using the JC Hamiltonian realized in this manner and numerical optimization. From the plots in in Fig. 49(a), we can see that the numerical scheme outperforms the Trotterization-based scheme in terms of circuit depth. On the other hand, the right-hand plot in Fig. 49(a) shows that the numerical scheme requires a longer circuit duration for the same infidelity. We emphasize that the numerical circuits have been optimized on the circuit depth and not the duration, which is why they do not converge to the Trotterized results. However, it would be an interesting future direction to develop a numerical scheme optimized on the circuit duration, as this more accurately reflects the true cost of the circuit.

2. QSP

Reference [478] has described an analytical scheme to achieve the squeezed vacuum state using non-Abelian QSP within the phase-space ISA. This scheme has been optimized for the circuit duration (T_{circuit}) and performs better than the state-of-the-art partially numerical scheme in Ref. [505]. While we direct the readers to Ref. [478] for the details of this preparation, here we compare it against the squeezed vacuum achieved numerically in Ref. [60] using the method employed to construct Fig. 31.

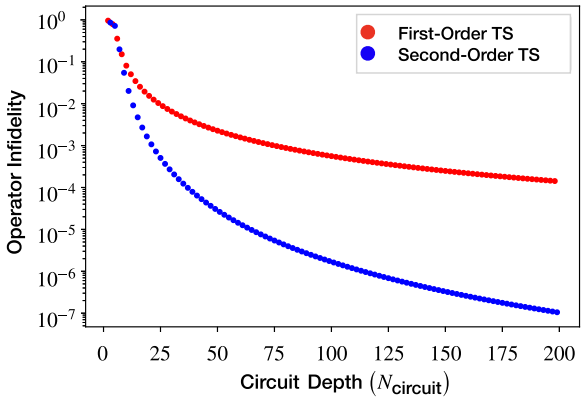


FIG. 48. The fidelity of an approximate compilation of the anti-Jaynes Cummings Interaction using the phase-space ISA. The y axis gives the operator infidelity, defined using Hilbert-Schmidt product: $1 - |(1/2d)\text{Tr}(PU^{\dagger}V)|^2$, between the exact operator $U = \text{AJC}(\theta)$ and the approximate phase-space ISA sequence V. This operator fidelity is computed on the oscillator-qubit subspace with projector $P = \sum_{\ell=0}^{d-1} |\ell\rangle\langle\ell| \otimes \sum_{q=0}^{1} |q\rangle\langle q|$. This is the projector on the joint subspace of a truncated oscillator with $d = 15$ and a qubit. The x axis corresponds to the circuit depth N_{circuit} of the different Trotter-Suzuki (TS) sequences used. See Appendix F for the definition of N_{circuit}. The red curve shows the first-order TS approximation $V = (e^{-i(\theta/r)A}e^{-i(\theta/r)B})^r$. The circuit depth in this case is $N_{\text{circuit}} = 2r$. The blue curve shows the second-order TS [see Eq. (367)] approximation $V = (e^{-i(\theta/2r)A}e^{-i(\theta/r)B}e^{-i(\theta/2r)A})^r$. The circuit depth in this case is $N_{\text{circuit}} = 2r + 1$. The AJC and JC Hamiltonians can be used to prepare arbitrary superpositions of Fock states and hence can be employed for universal oscillator state preparation via the Law-Eberly protocol [480]. The value of $\theta = \pi/\sqrt{2}$ chosen for this comparison is suitable for the preparation of Fock state $|1\rangle$. The Hilbert space of the oscillator used to compute the operators U and V is $N_{\text{dim}} = 50 \gg d$ dimensional and we have checked that the results are unaffected upon a further increase in N_{dim}.

In Fig. 49(b), we plot the circuit depth (left) and the duration (right) required to achieve increased squeezing of the vacuum with fidelity $\mathcal{F} > 0.99$. We can see that the squeezing achieved with this analytical scheme is on par with the numerical scheme in terms of circuit duration. The circuit-depth variation is optimal for the numerical scheme because this is not the metric on which the QSP scheme was optimized. Notably, the protocol in Ref. [505] has used circuit depth as a metric and produced states with worse fidelity compared to the QSP scheme. For example, the QSP-based Ref. [478] has achieved a squeezing of 8.5 dB and an antisqueezing of -8.40 dB with $T_{\text{circuit}} \approx 7.6$ with a fidelity of 0.9978. With comparable T_{circuit}, partial numerical optimization in Ref. [505] based on the circuit depth yields a squeezing of 8.5 dB and an antisqueezing of -9.9 dB with a fidelity of 0.99.

These comparisons indicate the need to optimize the numerical circuits not only using gate counts but also the lengths of conditional displacements. A simple way to

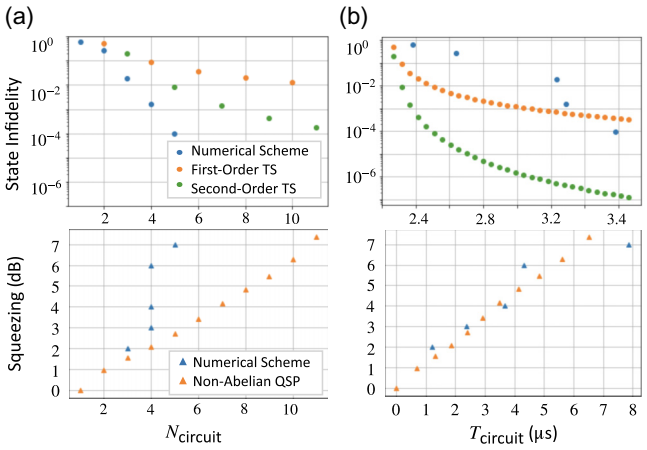


FIG. 49. A comparison between the numerically optimized and analytically derived circuits using techniques such as Trotter-Suzuki (TS) and non-Abelian QSP. (a) The infidelity of Fock state $|1\rangle$ preparation from the Law-Eberly protocol [480] using the JC Hamiltonian synthesized via TS methods, as discussed in Sec. VI D 1. On the x axis, we vary two metrics to improve these quantities. The left-hand plot shows variation in the circuit depth or N_{circuit}, and the right-hand plot shows varying circuit duration or T_{circuit}. (b) The squeezing of vacuum achieved with fidelity $\mathcal{F} > 0.99$ using non-Abelian QSP, as discussed in Ref. [478]. As before, the x axis shows varying N_{circuit} in the left-hand plot and varying T_{circuit} in the right-hand plot.

resolve this issue would be to optimize the circuits using T_{circuit} as a metric, thus taking into account the time cost of performing large conditional displacements $T_{|\gamma_i|}$ as well that of a large number of qubit rotations T_R. Such a metric would improve upon the numerical methods and would facilitate direct comparison with analytical schemes such as QSP-based squeezing. Changing the cost function of the optimizer from the circuit depth to the circuit duration may be more expensive in computational time. However, it would be interesting to use the various analytical schemes presented in this paper as starting templates for numerical optimization of circuit duration. This exercise would help us understand if these analytical methods are close to or at a local minimum.

APPENDIX G: BRIEF REVIEW OF QSP AND QSVT FORMALISMS FOR QUBITS

Quantum signal processing (QSP) combines noncommutative qubit operations to realize some desired polynomial transformation of an unknown signal encoded by a 2×2 unitary matrix [49,274]. The QSVT [514] generalizes QSP by enabling a polynomial transformation of the singular values of larger matrices. For QSVT, this matrix need not be 2×2 and the matrix need not even be square. This is a powerful capability, since such a transformation ordinarily requires the ability to perform a singular-value decomposition, which has a $O(2^{3n})$ cost for an n-qubit operator. In contrast, QSVT realizes transformations of singular values *without* needing to perform the singular-value decomposition itself. The challenge for QSVT is that the matrix to be transformed must be "block-encoded" into a unitary operator. For example, to realize quantum simulation of Hamiltonian H necessitates block-encoding of a positive semidefinite Hamiltonian inside a unitary matrix. The theory of QSP and QSVT comes from composite pulses in NMR, which we review below in Appendix G 1. A formal overview of qubit QSP is presented in Appendix G 2.

1. Intuition from composite pulses

A composite pulse is a sequence of qubit rotations which accomplishes the same thing as an ideal single qubit rotation, but behaves differently away from the ideal control point. For example, composite pulses can be highly insensitive to pulse amplitude errors (e.g., to improve robustness to hardware imperfections).

Thus, composite pulses are frequently used in NMR [516], trapped-ion [727,728], and neutral-atom [729] systems for optical or magnetic spin control to correct for classical errors in these single-qubit rotations associated with fluctuations in the frequency, phase, and amplitude of the control pulses (e.g., due to laser noise in optical spin control).

The mathematical framework for composite pulses is as follows. Let us consider a single-qubit rotation of the form (see Table VIII)

$$R_\varphi(\theta) = e^{-i\frac{\theta}{2}[\cos(\varphi)X + \sin(\varphi)Y]}, \tag{G1}$$

defined by the angle θ (determined by the drive pulse amplitude and duration) and the angle φ (determined by the drive pulse phase relative to a master clock). Suppose that we wish to implement the rotation $\theta = \pi/2$ and $\varphi = 0$, but the control system is imperfect due to control pulse amplitude fluctuations that cause $\theta = (\pi/2)(1+\epsilon)$, where ϵ quantifies the over- or under-rotation error. We may assume for simplicity that the phases φ of the pulses are perfectly controlled, for this pedagogical example.

A composite pulse scheme known as $\text{BB1}_{90,x}$ [516] corrects for the small errors ϵ in the classical parameter θ by appending to the nominal rotation $R_0(\theta)$ (rotation by angle θ about the x axis) three subsequent extra rotations,

$$\text{BB1}_{90,x} = R_{\varphi_1}(2\theta)R_{\varphi_2}(4\theta)R_{\varphi_1}(2\theta)R_0(\theta), \tag{G2}$$

where $\varphi_1 = \cos^{-1}(-1/8)$ and $\varphi_2 = 3\varphi_1$. For the cat state preparation example that we consider below, it is useful to have the analogous composite pulse for performing

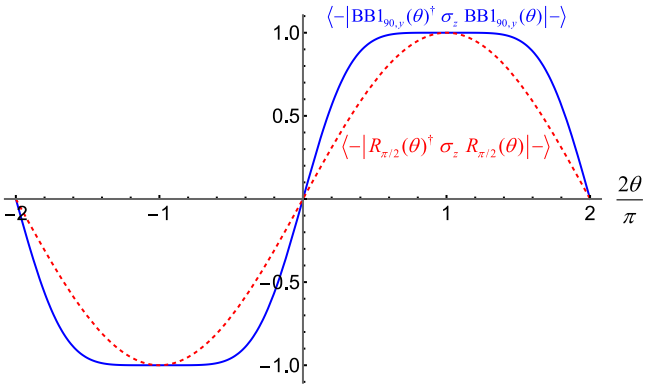


FIG. 50. The $\mathrm{BB1}_{90,y}(\theta)$ composite pulse sequence is robust against large over- and under-rotation errors in the angular parameter θ. An initial state $|-\rangle$ is accurately rotated to the state $|0\rangle$, as evidenced by this plot of the expectation value of σ_z versus the rotation angle θ. The solid (blue) line is for the $\mathrm{BB1}_{90,y}(\theta)$ composite pulse sequence and the dashed (red) line is for the ordinary single pulse rotating the spin by angle θ about the y axis. Even relatively far from the nominal value of $\theta = \pi/2$, the qubit polarization in the z direction on the Bloch sphere is close to unity for the $\mathrm{BB1}_{90,y}(\theta)$ pulse sequence.

rotations about the y axis,

$$\mathrm{BB1}_{90,y} = Z\left(+\frac{\pi}{2}\right)\mathrm{BB1}_{90,x}Z\left(-\frac{\pi}{2}\right), \tag{G3}$$

where $Z(\theta) = e^{-i(\theta/2)Z}$ is a rotation about the z axis by angle ϕ.

It is important to note that the error ϵ is presumed to be sufficiently slowly varying in time that it can be treated as a constant during any given short sequence of pulses. Hence, each of the four pulses in this sequence suffers the same angle scaling factor $1+\epsilon$. It is clear from inspection that if $\epsilon = 0$, the four successively applied rotation angles are $\theta = \pi/2, 2\theta = \pi, 4\theta = 2\pi$, and $2\theta = \pi$ and thus the three additional rotations combine to form the identity operation and the composite sequence of four pulses exactly yields the target unitary $R_0(\pi/2)$. Remarkably, by correctly choosing the azimuthal angles φ_1, φ_2, the errors in the three additional pulses combine to nearly perfectly cancel (for small ϵ) the error in the first pulse, giving an excellent approximation to the target unitary despite the presence of the error.

This strong robustness against fluctuations in the rotation angle θ away from the nominal value $\theta = \pi/2$ is clearly illustrated in Fig. 50. For small ϵ, the polarization (and hence the gate fidelity decay away from unity as $\langle\sigma_z\rangle \approx 1 - 3.7\epsilon^6$, indicating that the rotation angle error is $\mathcal{O}(\epsilon^3)$.

For later more formal purposes, it is useful to note here that inspection of Eq. (G3) shows that it can be rewritten in a "canonical form" as a sequence of $d = 9$ identical rotations $W_x \equiv R_0(\theta)$ by angle θ around the x axis, alternating with varying rotations $Z(\phi_j)$ around the z axis,

$$\mathrm{BB1}_{90,y} = Z(\phi_0)\prod_{j=1}^{d} W_x(\theta)Z(\phi_j), \tag{G4}$$

where we order the product of operators so that $j = 1$ is the leftmost and $j = d$ is the rightmost, and the set of $d+1$ angles for the z rotations is given by

$$\begin{aligned}\vec{v}_\phi &\equiv (\phi_0, \phi_1, \ldots, \phi_d)\\ &= \left(\varphi_1 + \frac{\pi}{2}, 0, -\varphi_1 + \varphi_2, 0, 0, 0, -\varphi_2 + \varphi_1, 0, -\varphi_1, -\frac{\pi}{2}\right),\end{aligned} \tag{G5}$$

where φ_1, φ_2 have been defined below Eq. (G2). Furthermore, since the transformation is a product of qubit rotations it can be written as a single rotation having the generic form

$$\mathrm{BB1}_{90,y} = e^{-i\frac{\pi}{2}\vec{h}(\theta)\cdot\vec{\sigma}}, \tag{G6}$$

where

$$\vec{h}(\theta) = \big(h_0(\theta), h_x(\theta), h_y(\theta), h_z(\theta)\big) \tag{G7}$$

is a vector the components of which are four functions of the rotation angle θ and $\vec{\sigma} = (\sigma_0, \sigma_x, \sigma_y, \sigma_z)$.

As discussed in the introduction, in NMR composite pulse design, a common goal is to make these functions as constant as possible near the target value of θ so that the rotation matches the target and is robust against noise in the parameter θ. One could also imagine the opposite situation in a quantum sensing application where one is interested in detecting small changes in θ by making $\vec{h}(\theta)$ have a strong derivative. This and other types of "filter" response functions are discussed in Refs. [3,49].

2. Qubit QSP

QSP provides a way to generalize the intuition behind the $\mathrm{BB1}_{90,y}$ to transform arbitrary input rotations in SU(2). The simplest case of QSP is the case in which the matrix is an SU(2) matrix such as an ordinary single-qubit rotation. Specifically, assume that $x \in [-1, 1]$ and let $U \in \mathbb{C}^{2\times 2}$ be

$$U := \begin{bmatrix} x & i\sqrt{1-x^2} \\ i\sqrt{1-x^2} & x \end{bmatrix} = e^{iX\arccos(x)}. \tag{G8}$$

The aim of QSP is to find a transformation such that

$$\Lambda : U \mapsto \begin{bmatrix} P(x) & i\sqrt{1-x^2}Q(x) \\ i\sqrt{1-x^2}Q^*(x) & P^*(x) \end{bmatrix} \tag{G9}$$

for polynomial functions $P(x)$ and $Q(x)$.

Although not all polynomial transformations can be directly constructed using this approach, the family of attainable transformations is broad. Specifically, assume that we have a sequence of rotation angles $\{\phi_0, \ldots, \phi_k\}$ such that

$$e^{i\phi_0 Z}\prod_{j=1}^{k} U(x)e^{i\phi_j Z} = \begin{bmatrix} P(x) & i\sqrt{1-x^2}Q(x) \\ i\sqrt{1-x^2}Q^*(x) & P^*(x) \end{bmatrix}. \tag{G10}$$

Such a set of angles can be found if

(1) $\deg(P(x)) \leq k$ and $\deg(Q(x)) \leq k-1$
(2) $\text{Parity}(P(x)) = k \mod 2$ and $\text{Parity}(Q(x)) = k-1 \mod 2$
(3) $|P(x)|^2 + (1-x^2)|Q(x)|^2 = 1$

There exist efficient classical algorithms to find these QSP angles [518,667,730–732]. A generalization of QSP from a single-qubit to multi-qubit settings where the dynamics takes place in direct sums of two-level subspaces can be achieved by the qubitization technique (see Ref. [49] for a tutorial on more details). In the present work, we extend the QSP formulation to bosonic systems and exploit it for compilation on hybrid oscillator-qubit hardware.

In the main text, we discuss an alternative convention for QSP wherein, for any y,

$$U \to U_z(y) := \begin{bmatrix} e^{iy} & 0 \\ 0 & e^{-iy} \end{bmatrix}. \tag{G11}$$

We can use the previous results to see that the Hadamard transform of the original matrix can be written as

$$\begin{aligned} U_{\boldsymbol{\phi}} &= He^{i\phi_0 Z}\prod_{j=1}^{k} U(x)e^{i\phi_j Z}H \\ &= e^{i\phi_0 X}\prod_{j=1}^{k} U_z(\cos^{-1}(x))e^{i\phi_j X}. \end{aligned} \tag{G12}$$

We then have, using the Hadamard transform, for $z = \cos^{-1}(x)$, that

$$U_{\boldsymbol{\phi}} = \begin{bmatrix} i\text{Re}(Q(z))\sqrt{1-z^2} + \text{Re}(P(z)) & \text{Im}(Q(z))\sqrt{1-z^2} + i\text{Im}(P(z)) \\ -\text{Im}(Q(z))\sqrt{1-z^2} + i\text{Im}(P(z)) & -i\text{Re}(Q(z))\sqrt{1-z^2} + \text{Re}(P(z)) \end{bmatrix}. \tag{G13}$$

We can see from the constraints on Q and P that the norm of each column and row is

$$|P(z)|^2 + (1-z^2)|Q(z)|^2 = 1 \tag{G14}$$

and thus the resulting matrix is manifestly unitary. In this form for the polynomials $Q(z)$ and $P(z)$, it is clear that in this formalism we can block-encode a polynomial of degree at most k subject to the above constraints on P and Q.

APPENDIX H: DETAILS ON QFT IMPLEMENTATION

In this appendix, we provide further details on the realization of the QFT on a CV-DV device, as presented in Sec. VII B.

1. Details on QFT method 1: Single-variable QSP

For intuitive purposes, we reproduce the circuit that realizes the QFT here in Fig. 51. Overall, the QFT is achieved in our construction by first transferring the qubit state to an equivalent oscillator state with basis spacing Δ, and then applying a displacement operation to make the state symmetric about $q = 0$ in position space. Subsequently, the Fourier gate is applied to the oscillator to enact a continuous Fourier transform, and then the state is transferred back to qubits by using a *modified* state-transfer unitary $\tilde{U}_{\text{st}}^{(n)}(\pi/2^n\Delta)$ with reciprocal spacing $\pi/2^n\Delta$. Importantly, this operation is modified such that the controlled displacements that comprise it (i.e., the first stage in Fig. 32) are each amplified by a factor of 2, which implements the correct phases corresponding to the QFT. Intuitively, it ensures that the product of the displacements used in each

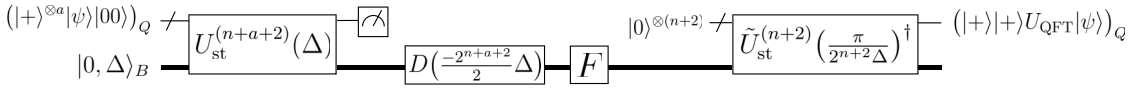


FIG. 51. The circuit used to implement the QFT of an n-qubit state $|\psi\rangle$ using the single-variable QSP state-transfer unitary of Sec. VII A 1, denoted U_{st}. Here, $D\left((-2^{n+a+2}/2)\Delta\right)$ is a displacement operation, and F is the Fourier transform operation $F = U(\pi/2)$ (i.e., oscillator free evolution).

state-transfer step equates to $(x\Delta)\cdot(y\cdot 2\cdot(\pi/2^n\Delta)) = 2\pi xy/2^n$, which is the correct phase of the QFT. Throughout this construction, ancilla qubits are appended to the initial state to make the underlying state periodic and increase fidelity with the exact QFT.

We verify the performance of this QFT construction by considering its action on the input qubit state $|\psi\rangle_Q = \sum_{\mathbf{x}} c_{\mathbf{x}}|\mathbf{x}\rangle_Q$, where $|\mathbf{x}\rangle = |x_1\rangle|x_2\rangle\cdots|x_n\rangle$. Let us go through this verification step by step:

Step 1. Our first step is to prepend $|\psi\rangle_Q$ by a ancilla qubits in the $|+\rangle_Q$ state, generating the state:

$$\left(|+\rangle^{\otimes a}|\psi\rangle\right)_Q = \frac{1}{\sqrt{2^a}}\sum_{k=0}^{2^a-1}\sum_{x=0}^{2^n-1} c_{\mathbf{x}}|2^n\cdot k + x\rangle_Q = \sum_{r=0}^{2^{n+a}-1} C_r|r\rangle_Q. \tag{H1}$$

The coefficients C_r repeat over 2^a periods of size 2^n: $C_r = C_{r+2^n \pmod{2^{n+a}}}$. This makes the initial state coefficients periodic from $r=0$ to $r = 2^{n+a}-1$ and allows us to use the above continuous-discrete Fourier transform correspondence to implement the QFT.

Additionally, we note that in the circuit of Fig. 36, the initial state is also appended by a state $|00\rangle$. As we will see later, this is incorporated to improve the fidelity with the exact QFT. To simplify the current presentation, we exclude these qubits for now and reintroduce them at a later step when they are necessary.

Step 2. The next step is to use the single-variable QSP state-transfer protocol of Sec. VII A 1 to transfer this state to the oscillator, encoded in a Gaussian basis of width σ. Applying the state-transfer unitary $U_{\text{st}}^{(n+a)}(\Delta)$ produces the state $|\mathbf{0}\rangle_Q\sum_r C_r|r,\Delta\rangle_B$ and suffers infidelity $\mathcal{O}((n+a)\epsilon) + e^{-\mathcal{O}(\Delta^2/\sigma^2)}$, where ϵ is the error in the QSP polynomial approximation to the square-wave function.

Step 3. At this stage, the initial state has been mapped to an oscillator state, encoded in a basis of states peaked around $q = r\Delta$, from $q=0$ to $q = (2^{n+a}-1)\cdot\Delta$. Next, we symmetrize these basis states about $q=0$ by applying a displacement $D((-2^{n+a}/2)\Delta)$, producing the state

$$\sum_{r=0}^{2^{n+a}-1} C_r|r - \frac{2^{a+n}}{2},\Delta\rangle_B = \sum_{k=-2^a/2}^{2^a/2-1}\sum_{x=0}^{2^n-1} c_{\mathbf{x}}|x+2^nk,\Delta\rangle_B, \tag{H2}$$

where we have reintroduced the indices x and k, and now symmetrized the sum over k.

Step 4. The next step is to apply the Fourier gate (i.e., free evolution, provided through the bosonic ISA in Box IV.2) to this state, which enacts the continuous Fourier transform on the wave function. As the basis states in Eq. (H2) are Gaussians, this produces the state

$$\sqrt{\sigma\sqrt{\frac{2}{\pi}}}\int dq\sum_{\mathbf{x}} c_{\mathbf{x}}e^{iq\Delta x} \times\left[\frac{1}{\sqrt{2^a}}\sum_{k=-2^{a-1}}^{2^{a-1}-1} e^{iq2^nk\Delta}\right]e^{-\sigma^2q^2}|q\rangle_B. \tag{H3}$$

The factor $\sum_{k=-2^a/2}^{2^a/2-1} e^{iq\Delta 2^nk}$ in brackets is known as the Dirichlet kernel. In the limit of large a, it behaves as the sum of delta functions $\sum_{l\in\mathbb{Z}}\delta((q\Delta 2^n/2\pi)-l)$, i.e., a Dirac comb. At finite a, its behavior under an integral deviates from that of the Dirac comb by an error $\mathcal{O}(1/2^a)$. Taking the large-a limit in Eq. (H3), this factor selects out values $q_l := l(2\pi/\Delta 2^n)$ from the integral, such that we can reexpress the state as

$$\sqrt{\frac{2\pi\sigma}{\Delta 2^n}}\sqrt{\frac{2}{\pi}}\sum_l\sum_{\mathbf{x}} c_{\mathbf{x}}e^{i2\pi lx/2^n}e^{-\sigma^2q_l^2}|\tilde{q}=q_l\rangle_B, \tag{H4}$$

where $|\tilde{q}=q_l\rangle_B$ is a normalized state concentrated around $q=q_l$ (see Ref. [541] for the exact expression), and the infidelity in reaching Eq. (H4) is $\mathcal{O}(1/2^a)$. Also observe that taking this limit has produced a phase $e^{i2\pi lx/2^n}$ resembling that of the QFT, which arises from the continuous-discrete Fourier transform correspondence.

Step 5. Lastly, we introduce n qubits $|\mathbf{0}\rangle_Q^{\otimes n}$ and transfer the oscillator state back to the qubits. For this, we use the inverse state-transfer operation with an appropriately chosen reciprocal spacing $\tilde{U}_{\text{st}}^{(n)}((\pi/2^n)\Delta)^\dagger$. Note that as we mentioned earlier, this unitary is modified such that the controlled displacements that comprise it are each amplified by a factor of 2.

The application of this operation is rather involved, and its analysis is fully presented in Ref. [541]. The final result is that the output is approximately a product state $|\alpha\rangle_Q|\beta\rangle_B$, with an infidelity $\mathcal{O}(\sigma/\Delta)$. Here, the state of the qubits is

$$|\alpha\rangle_Q = \sum_{\mathbf{x}} c_{\mathbf{x}}e^{-i\pi x/2}U_{\text{QFT}}|\mathbf{x}\rangle_Q, \tag{H5}$$

which is nearly the QFT, and the oscillator state is approximately a Gaussian state but will not concern us here.

The state of the qubits is the QFT of the initial state $|\psi\rangle_Q$ up to a phase factor $e^{-i\pi x/2}$. This phase factor arises because the action of free evolution on a state symmetrized about position $q=0$ (as we have here) implements a symmetric QFT, with coefficients given by a symmetrized sum: $\sum_{y=-2^n/2}^{2^n/2-1} c_y e^{2\pi ixy/2^n}$. This can be mapped to the usual QFT by shifting $x\mapsto x+2^n/2$, at the expense of incurring a phase factor $e^{-i\pi x/2}$.

Fortunately, this phase factor is inconsequential and can be removed by appending $|00\rangle$ to the initial state as

$(|\psi\rangle|00\rangle)_Q$. This guarantees that $x \bmod 4 = 0$ for any nonzero $c_{\mathbf{x}}$, such that each phase factor is always $e^{-i\pi x/2} = 1$. Then, performing the inverse state-transfer step onto $n+2$ qubits yields the state

$$U_{\mathrm{QFT}}(|\psi\rangle|00\rangle) = |+\rangle|+\rangle U_{\mathrm{QFT}}(|\psi\rangle), \tag{H6}$$

from which $U_{\mathrm{QFT}}(|\psi\rangle)$ may be extracted. This addresses the additional ancilla qubits $|00\rangle$ depicted in Fig. 36.

a. Fidelity analysis

By aggregating together the infidelities suffered at each step above, we find that the total fidelity of this entire protocol is

$$F = 1 - \mathcal{O}((n+a)\epsilon) - \mathcal{O}(\sigma/\Delta) - \mathcal{O}(1/2^a), \tag{H7}$$

where ϵ is the error in the QSP polynomial approximation to a square-wave function.

2. Details on QFT method 2: Non-Abelian QSP

We again reproduce the figure that realizes the QFT in Fig. 52. Our approach is to first perform a change of basis from the computational basis to the basis $\{|\phi_{\mathbf{s}}\rangle\}$ used in the non-Abelian QSP state-transfer protocol; this is achieved using an operation denoted $\mathcal{T}$ in Fig. 37, which we show can be implemented with simple local operations. Subsequently, we transfer the state of the qubits to the oscillator, symmetrize the state about $q = 0$ in position space, apply the Fourier gate, and finally transfer the oscillator state back to the qubits with an appropriately chosen reciprocal spacing. Upon transforming back to the computational basis, we obtain the QFT of the initial state. We again incorporate ancilla qubits throughout this protocol to improve the fidelity with the exact QFT.

Let us again study this algorithm on an initial n-qubit state $|\psi\rangle_Q = \sum_{\mathbf{x}} c_{\mathbf{x}}|\mathbf{x}\rangle_Q$:

Step 1. As in the previous QFT construction, we first prepend $|\psi\rangle$ by a ancilla qubits in the $|+\rangle$ state to make the coefficients periodic: $\left(|+\rangle^{\otimes a}|\psi\rangle\right)_Q = \sum_{r=0}^{2^{n+a}-1} C_r|r\rangle_Q$, whereas before, the coefficients repeat over 2^a periods of size 2^n: $C_r = C_{r+2^n \,(\mathrm{mod}\, 2^{n+a})}$.

Step 2. The next step is to map the previous state from the computational basis to the basis $\{|\phi_{\mathbf{s}}\rangle\}$ used in the non-Abelian state-transfer protocol.

To specify an appropriate correspondence between basis states, recall that in this context the string $\mathbf{s} \in \{+1, -1\}^{n+a}$ is associated with a position $q_{\mathbf{s}}$ that takes discrete values in the range $[-\lambda(2^{n+a}-1), \lambda(2^{n+a}-1)]$ with spacing 2λ [see Eq. (424)]. It is therefore suitable to index each such possible value by an integer s, defined as

$$s := \frac{1}{2\lambda}\left(q_{\mathbf{s}} + \lambda(2^{n+a}-1)\right) \in \{0, 1, \ldots, 2^{n+a}-1\}. \tag{H8}$$

Thus, we assume access to a unitary change of basis operation $\mathcal{T}$ that acts as $\mathcal{T}|s\rangle = |\phi_{\mathbf{s}}\rangle$, such that its application to the state from step 1 generates the state $\sum_{s=0}^{2^{n+a}-1} C_s|\phi_s\rangle$, where we are now indexing by the associated integer s to simplify analysis. As shown in Ref. [541], $\mathcal{T}$ can be easily implemented as a product of local operations:

$$\mathcal{T} = (HX)^{\otimes(n+a-1)} \cdot H_{n+a} \cdot Z_1 Z_{n+a} \cdot \mathcal{P}_{\leftrightarrow}, \tag{H9}$$

where $\mathcal{P}_{\leftrightarrow}$ is a permutation operation that reverses the order of qubits (e.g., $\mathcal{P}_{\leftrightarrow}|0011\rangle = |1100\rangle$)

Step 3. The next step uses the non-Abelian state-transfer protocol to transfer the state to an oscillator as per Eq. (425). To achieve this, we prepare the oscillator in the sinc state $|0, 2\lambda\rangle_B = 1/\sqrt{2\lambda}\int dq\ \mathrm{sinc}\,(\pi q/2\lambda)\,|q\rangle_B$, and apply $U_{\mathrm{st}}^{(n+a)}(2\lambda)$ to the joint system, producing

$$|\mathbf{0}\rangle_Q \otimes \sum_{s=0}^{2^{n+a}-1} C_s \frac{1}{\sqrt{2\lambda}} \int dq\ \mathrm{sinc}\left(\frac{\pi(q-q_{\mathbf{s}})}{2\lambda}\right)|q\rangle_B. \tag{H10}$$

The infidelity suffered here is quantified by Eq. (426).

Step 4. Next, we symmetrize this state about $q = 0$ by applying a displacement operation $D(\lambda)$. Using the identities $s = 2^n \cdot k + x$ and $q_s = 2\lambda(s - \frac{1}{2}2^{n+a} - \frac{1}{2}) = 2\lambda(2^n \cdot [k - \frac{1}{2}2^a] + x - \frac{1}{2})$, this yields the wave function

$$\sum_{x=0}^{2^n-1} \sum_{k=-2^a/2}^{2^a/2-1} \frac{1}{\sqrt{2^a}} c_{\mathbf{x}} \frac{1}{\sqrt{2\lambda}} \mathrm{sinc}\left(\frac{\pi(q - 2\lambda(2^n \cdot k + x))}{2\lambda}\right). \tag{H11}$$

Step 5. Subsequently, we apply the Fourier gate to the oscillator, which enacts a continuous Fourier transform on the wave function. Noting the Fourier transform

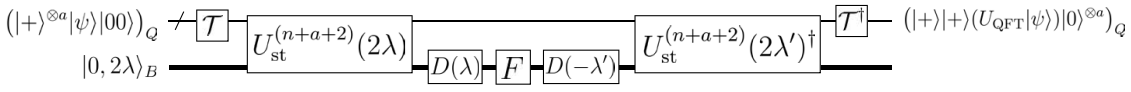


FIG. 52. The circuit used to implement the QFT of an n-qubit state $|\psi\rangle$ using the non-Abelian QSP state-transfer unitary of Sec. VII A 2, denoted U_{st}. Here, $\mathcal{T}$ is a basis transformation from the computational basis to the basis $|\phi_s\rangle$, which is easily realized with simple qubit gates, as discussed in the text. The operation $D(\cdot)$ is a displacement operation, and F is the Fourier gate (i.e., oscillator free evolution).

relation $1/\sqrt{2\pi}\int dq' e^{iqq'}(1/\sqrt{2\lambda})\mathrm{sinc}\left(\pi(q'-b)/2\lambda\right) = \sqrt{\lambda/\pi}\,1_{|q|\le\pi/2\lambda}e^{iqb}$, where $1_{|q|\le\pi/2\lambda}$ is an indicator function (i.e., a box function in this context), this transforms the above wave function to

$$\sqrt{\frac{\lambda}{\pi}}1_{|q|\le\pi/2\lambda}\sum_x c_x e^{iq2\lambda x}\left[\frac{1}{\sqrt{2^a}}\sum_k e^{iq2\lambda 2^n k}\right] =: \tilde{\psi}(q). \tag{H12}$$

Step 6. After free evolution, we perform the inverse of steps 4 and 5, but now with reciprocal spacing $2\lambda' = (2\pi/2^{n+a})1/2\lambda$ chosen to produce the QFT. We first enact the inverse of step 5, which consists of applying the displacement $D(-\lambda')$ to map the wave function to $\tilde{\psi}(q+\lambda')$.

Step 7. Next, we transfer the oscillator state back to $n+a$ qubits by using $U_{\mathrm{st}}^{(n+a)}(2\lambda')^\dagger$. As per Eq. (422), this maps the above state to

$$\begin{aligned}&\sum_{s=0}^{2^{n+a}-1}\int dq\,\tilde{\psi}(q+q'_s+\lambda)\mathrm{sinc}\left(\frac{\pi q}{2\lambda}\right)|\phi_{\mathbf{s}}\rangle_Q|q\rangle_B\\ &\quad=\sum_{s=0}^{2^{n+a}-1}\sum_{x=0}^{2^n-1}\sqrt{\frac{\lambda}{\pi}}\int_{-\pi/2\lambda}^{\pi/2\lambda}dq\,c_{\mathbf{x}}e^{i(q+q'_s+\lambda)2\lambda x}\\ &\qquad\times\left[\frac{1}{\sqrt{2^a}}\sum_k e^{i(q+q'_s+\lambda)2\lambda 2^n k}\right]\mathrm{sinc}\left(\frac{\pi q}{2\lambda'}\right)|\phi_{\mathbf{s}}\rangle_Q|q\rangle_B,\end{aligned} \tag{H13}$$

where $q'_s = 2\lambda' s - \lambda'(2^{n+a}-1)$. Note that we have not used any approximations to simplify this expression, as was done in Sec. VII A 2.

As in the previous QFT protocol, the quantity in brackets is the Dirichlet kernel, which behaves as $\sum_{l\in\mathbb{Z}}\delta(((q+q'_s+\lambda')\Delta 2^n/2\pi)-l)$, with an overall error $\mathcal{O}(1/2^a)$ at finite a. One can take this limit to drastically simplify this state; the full technical analysis of this step is presented in Ref. [541]. The final result is that the state of the system can be simplified to the product state

$$\left[\sum_{l=0}^{2^n-1}\left(\frac{1}{\sqrt{2^n}}\sum_{x=0}^{2^n-1}c_j e^{i2\pi lx/2^n}e^{-i\pi x/2}\right)|\phi_{s=2^a l}\rangle_Q\right]|\tilde{q}=0\rangle_B, \tag{H14}$$

where $|\tilde{q}=0\rangle_B$ is a state concentrated around $q=0$, and the infidelity suffered in reaching this expression is $\mathcal{O}(1/2^a)$. The oscillator and qubits are now decoupled, and the state of the qubits is nearly the QFT evaluated in the basis $|\phi_{s=2^a l}\rangle$. Again, the presence of the phase $e^{i2\pi lx/2^n}$ arises from the continuous-discrete correspondence of the Fourier transform discussed earlier.

Step 8. To extract the quantum Fourier transform from this state, the last step is to transform from the $\{|\phi_s\rangle\}$ basis back to the computational basis using $\mathcal{T}^\dagger$. This acts as $\mathcal{T}^\dagger|\phi_{s=2^a l}\rangle = |2^a l\rangle = |l\rangle|0\rangle^{\otimes a}$, such that its application to the previous state produces

$$\left[\sum_{x=0}^{2^n-1}e^{-i\pi x/2}c_{\mathbf{x}}U_{\mathrm{QFT}}\left(|x\rangle\right)|0\rangle^{\otimes a}\right]_Q|\tilde{q}=0\rangle_B. \tag{H15}$$

The remaining state on the first n qubits is nearly our target QFT, up to a phase factor $e^{-i\pi x/2}$.

As before, this phase factor arises because this protocol really implements a symmetric QFT, which maps to the usual QFT at the expense of a factor $e^{-i\pi x/2}$. Again, this phase factor is inconsequential and can be removed by appending $|00\rangle$ to the initial state, such that $e^{-i\pi x/2}=1$ for all x with nonzero coefficients $c_{\mathbf{x}}$. This then outputs the target QFT as $U_{\mathrm{QFT}}(|\psi\rangle|00\rangle) = |+\rangle|+\rangle U_{\mathrm{QFT}}(|\psi\rangle)$.

a. Fidelity analysis

Consider the fidelity of this protocol. The first source of infidelity occurs in step 3 due to the usage of the non-Abelian QSP state-transfer protocol. The infidelity of this operation is quantified in Eq. (426); a careful analysis of this scenario, presented in Ref. [541], indicates that the resulting infidelity scales as $\mathcal{O}(1/2^a)$. Moreover, the other source of infidelity arises in step 6 when taking the large-a limit, which also contributes an infidelity $\mathcal{O}(1/2^a)$. Lastly, we note that no infidelity is suffered in the inverse state-transfer step, as we do not use any approximation to simplify this step. Therefore, summing these contributions to the infidelity, we find that the overall fidelity of this QFT protocol is $F = 1-\mathcal{O}(1/2^a)$.

[1] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2010).

[2] Y. Alexeev, *et al.*, Quantum computer systems for scientific discovery, PRX Quantum **2**, 017001 (2021).

[3] L. M. K. Vandersypen and I. L. Chuang, NMR techniques for quantum control and computation, Rev. Mod. Phys. **76**, 1037 (2005).

[4] J. A. Jones, Controlling NMR spin systems for quantum computation, Prog. Nucl. Magn. Reson. Spectrosc. **140-141**, 49 (2024).

[5] J. Dijkema, X. Xue, P. Harvey-Collard, M. Rimbach-Russ, S. L. de Snoo, G. Zheng, A. Sammak, G. Scappucci, and L. M. K. Vandersypen, Cavity-mediated iSWAP oscillations between distant spins, Nat. Phys. **21**, 168 (2025).

[6] D. Jaksch, J. I. Cirac, P. Zoller, S. L. Rolston, R. Côté, and M. D. Lukin, Fast quantum gates for neutral atoms, Phys. Rev. Lett. **85**, 2208 (2000).

[7] N. Schlosser, G. Reymond, I. Protsenko, and P. Grangier, Sub-Poissonian loading of single atoms in a microscopic dipole trap, Nature **411**, 1024 (2001).

[8] D. Bluvstein, H. Levine, G. Semeghini, T. T. Wang, S. Ebadi, M. Kalinowski, A. Keesling, N. Maskara, H. Pichler, M. Greiner, V. Vuletić, and M. D. Lukin, A quantum processor based on coherent transport of entangled atom arrays, Nature **604**, 451 (2022).

[9] T. M. Graham, *et al.*, Multi-qubit entanglement and algorithms on a neutral-atom quantum computer, Nature **604**, 457 (2022).

[10] J. I. Cirac and P. Zoller, Quantum computations with cold trapped ions, Phys. Rev. Lett. **74**, 4091 (1995).

[11] C. Monroe, D. M. Meekhof, B. E. King, W. M. Itano, and D. J. Wineland, Demonstration of a fundamental quantum logic gate, Phys. Rev. Lett. **75**, 4714 (1995).

[12] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland, Quantum dynamics of single trapped ions, Rev. Mod. Phys. **75**, 281 (2003).

[13] H. Häffner, C. Roos, and R. Blatt, Quantum computing with trapped ions, Phys. Rep. **469**, 155 (2008).

[14] J. P. Home, D. Hanneke, J. D. Jost, J. M. Amini, D. Leibfried, and D. J. Wineland, Complete methods set for scalable ion trap quantum information processing, Science **325**, 1227 (2009).

[15] C. Monroe and J. Kim, Scaling the ion trap quantum processor, Science **339**, 1164 (2013).

[16] C. D. Bruzewicz, J. Chiaverini, R. McConnell, and J. M. Sage, Trapped-ion quantum computing: Progress and challenges, Appl. Phys. Rev. **6**, 021314 (2019).

[17] S. Moses, *et al.*, A race-track trapped-ion quantum processor, Phys. Rev. X **13**, 041052 (2023).

[18] K. Hémery, K. Ghanem, E. Crane, S. L. Campbell, J. M. Dreiling, C. Figgatt, C. Foltz, J. P. Gaebler, J. Johansen, M. Mills, S. A. Moses, J. M. Pino, A. Ransford, M. Rowe, P. Siegfried, R. P. Stutz, H. Dreyer, A. Schuckert, and R. Nigmatullin, Measuring the Loschmidt amplitude for finite-energy properties of the Fermi-Hubbard model on an ion-trap quantum computer, PRX Quantum **5**, 030323 (2024).

[19] O. Katz, M. Cetina, and C. Monroe, Programmable n-body interactions with trapped ions, PRX Quantum **4**, 030311 (2023).

[20] M. Iqbal, N. Tantivasadakarn, R. Verresen, S. L. Campbell, J. M. Dreiling, C. Figgatt, J. P. Gaebler, J. Johansen, M. Mills, S. A. Moses, J. M. Pino, A. Ransford, M. Rowe, P. Siegfried, R. P. Stutz, M. Foss-Feig, A. Vishwanath, and H. Dreyer, Non-Abelian topological order and anyons on a trapped-ion processor, Nature **626**, 505 (2024).

[21] S. Pezzagna and J. Meijer, Quantum computer based on color centers in diamond, Appl. Phys. Rev. **8**, 011308 (2021).

[22] G. Burkard, T. D. Ladd, A. Pan, J. M. Nichol, and J. R. Petta, Semiconductor spin qubits, Rev. Mod. Phys. **95**, 025003 (2023).

[23] M. Reagor, H. Paik, G. Catelani, L. Sun, C. Axline, E. Holland, I. M. Pop, N. A. Masluk, T. Brecht, L. Frunzio, M. H. Devoret, L. Glazman, and R. J. Schoelkopf, Reaching 10 ms single photon lifetimes for superconducting aluminum cavities, Appl. Phys. Lett. **102**, 192604 (2013).

[24] M. Reagor, W. Pfaff, C. Axline, R. W. Heeres, N. Ofek, K. Sliwa, E. Holland, C. Wang, J. Blumoff, K. Chou, M. J. Hatridge, L. Frunzio, M. H. Devoret, L. Jiang, and R. J. Schoelkopf, Quantum memory with millisecond coherence in circuit QED, Phys. Rev. B **94**, 014506 (2016).

[25] A. Romanenko, R. Pilipenko, S. Zorzetti, D. Frolov, M. Awida, S. Belomestnykh, S. Posen, and A. Grassellino, Three-dimensional superconducting resonators at $t = 20$ mk with photon lifetimes up to $\tau = 2$ s, Phys. Rev. Appl. **13**, 034032 (2020).

[26] O. Milul, B. Guttel, U. Goldblatt, S. Hazanov, L. M. Joshi, D. Chausovsky, N. Kahn, E. Çiftyürek, F. Lafont, and S. Rosenblum, Superconducting cavity qubit with tens of milliseconds single-photon coherence time, PRX Quantum **4**, 030336 (2023).

[27] A. Copetudo, C. Y. Fontaine, F. Valadares, and Y. Y. Gao, Shaping photons: Quantum information processing with bosonic cQED, Appl. Phys. Lett. **124**, 080502 (2024).

[28] A. Blais, A. L. Grimsmo, S. M. Girvin, and A. Wallraff, Circuit quantum electrodynamics, Rev. Mod. Phys. **93**, 025005 (2021).

[29] A. Blais, S. M. Girvin, and W. D. Oliver, Quantum information processing and quantum optics with circuit quantum electrodynamics, Nat. Phys. **16**, 247 (2020).

[30] Z.-L. Xiang, S. Ashhab, J. Q. You, and F. Nori, Hybrid quantum circuits: Superconducting circuits interacting with other quantum systems, Rev. Mod. Phys. **85**, 623 (2013).

[31] U. L. Andersen, J. S. Neergaard-Nielsen, P. Van Loock, and A. Furusawa, Hybrid discrete- and continuous-variable quantum information, Nat. Phys. **11**, 713 (2015).

[32] G. Kurizki, P. Bertet, Y. Kubo, K. Mølmer, D. Petrosyan, P. Rabl, and J. Schmiedmayer, Quantum technologies with hybrid systems, Proc. Natl. Acad. Sci. **112**, 3866 (2015).

[33] S. Lloyd and S. L. Braunstein, Quantum computation over continuous variables, Phys. Rev. Lett. **82**, 1784 (1999).

[34] R. Dutta, B. Allen, N. P. Vu, C. Xu, K. Liu, F. Miao, B. Wang, A. Surana, C. Wang, Y. Ding, and V. S. Batista, Solving constrained optimization problems using hybrid qubit-qumode quantum devices, ArXiv:2501.11735.

[35] S. B. Bravyi and A. Y. Kitaev, Fermionic quantum computation, Ann. Phys. **298**, 210 (2002).

[36] T. E. O'Brien, P. Rożek, and A. R. Akhmerov, Majorana-based fermionic quantum computation, Phys. Rev. Lett. **120**, 220504 (2018).

[37] D. González-Cuadra, D. Bluvstein, M. Kalinowski, R. Kaubruegger, N. Maskara, P. Naldesi, T. V. Zache, A. M. Kaufman, M. D. Lukin, H. Pichler, B. Vermersch, J. Ye, and P. Zoller, Fermionic quantum processing with programmable neutral atom arrays, PNAS **120**, e2304294120 (2023).

[38] A. Rad, A. Schuckert, E. Crane, G. Nambiar, F. Fei, J. Wyrick, R. M. Silver, M. Hafezi, Z. Davoudi, and M. J. Gullans, Analog quantum simulator of a quantum field theory with fermion-spin systems in silicon, ArXiv:2407.03419.

[39] A. Schuckert, E. Crane, A. V. Gorshkov, M. Hafezi, and M. J. Gullans, Fermion-qubit fault-tolerant quantum computing, ArXiv:2411.08955.

[40] S. Girvin, in *Proc. Int. Sch. Phys. "Enrico Fermi," Vol. 209: Quantum Fluids of Light and Matter* (IOS Press Ebook, Amsterdam, 2025), p. 23.

[41] E. Crane, K. C. Smith, A. Schuckert, T. Tomesh, A. Eickbusch, J. M. Martyn, S. Kühn, L. Funcke, N. Wiebe, M. A. DeMarco, I. L. Chuang, and S. Girvin, Hybrid oscillator-qubit quantum processors: Simulating fermions, bosons, and gauge fields, ArXiv:2409.03747.

[42] E. Chiari, W. Makhlouf, L. Pepe, E. Koridon, J. Klein, B. Senjean, B. Lasorne, and S. Yalouz, *Ab initio* polaritonic chemistry on diverse quantum computing platforms: Qubit, qudit, and hybrid qubit-qumode architectures, ArXiv:2506.12504.

[43] J. Hennessy, D. Patterson, and K. Asanović, *Computer Architecture: A Quantitative Approach* (Elsevier Science, Amsterdam, 2012).

[44] M. Davis, R. Sigal, and E. J. Weyuker, *Computability, Complexity, and Languages: Fundamentals of Theoretical Computer Science* (Elsevier, Amsterdam, 1994).

[45] M. E. Beverland, P. Murali, M. Troyer, K. M. Svore, T. Hoeffler, V. Kliuchnikov, G. H. Low, M. Soeken, A. Sundaram, and A. Vaschillo, Assessing requirements to scale to practical quantum advantage, ArXiv:2211.07629.

[46] S. Bravyi, O. Dial, J. M. Gambetta, D. Gil, and Z. Nazario, The future of quantum computing with superconducting qubits, J. Appl. Phys. **132**, 160902 (2022).

[47] D. Hall, *Microprocessors and Interfacing: Programming and Hardware*, McGraw-Hill Computer Science Series (McGraw-Hill, Inc., Glencoe, 1992).

[48] W. Stallings, *Computer Organization and Architecture: Designing for Performance*, William Stallings Books on Computer and Data Communications Technology (Pearson Education, 2003.)

[49] J. M. Martyn, Z. M. Rossi, A. K. Tan, and I. L. Chuang, Grand unification of quantum algorithms, PRX Quantum **2**, 040203 (2021).

[50] N. Killoran, J. Izaac, N. Quesada, V. Bergholm, M. Amy, and C. Weedbrook, STRAWBERRY FIELDS: A software platform for photonic quantum computing, Quantum **3**, 129 (2019).

[51] J. Zhou, Y. Liu, Y. Shi, A. Javadi-Abhari, and G. Li, Bosehedral: Compiler Optimization for Bosonic Quantum Computing, in *2024 ACM/IEEE 51st Annual International Symposium on Computer Architecture (ISCA)* (Buenos Aires, Argentina, 2024), p. 261.

[52] T. J. Stavenger, E. Crane, K. C. Smith, C. T. Kang, S. M. Girvin, and N. Wiebe, C2QA - Bosonic Qiskit, in *2022 IEEE High Performance Extreme Computing Conference (HPEC)* (Waltham, MA, USA, 2022), p. 1.

[53] BOSONIC QISKIT Github Repository, 2022, https://github.com/C2QA/bosonic-qiskit.

[54] K. C. Smith, E. Crane, T. Stavenger, and S. M. Girvin, Introducing BOSONIC QISKIT: A package for simulating bosonic and hybrid qubit-bosonic circuits (2023), https://medium.com/qiskit/introducing-bosonic-qiskit-a-package-for-simulating-bosonic-and-hybrid-qubit-bosonic-circuits-1e1e528287bb.

[55] Z. Liu and K. C. Smith, Simulating a bosonic error correcting code using BOSONIC QISKIT (2024), https://www.ibm.com/quantum/blog/bosonic-error-correction.

[56] Z. Chen, J. Li, M. Guo, H. Chen, Z. Li, J. Bierman, Y. Huang, H. Zhou, Y. Liu, and E. Z. Zhang, GENESIS: A compiler framework for Hamiltonian simulation on hybrid CV-DV quantum computers, ArXiv:2505.13683.

[57] E. Decker, E. Gustafson, E. McKinney, A. K. Jones, L. Goetz, A. Li, A. Schuckert, S. Stein, G. Li, and E. Crane, SYMBOLIC HAMILTONIAN COMPILER for hybrid qubit-boson processors, ArXiv:2506.00215.

[58] X. Pan, P. Song, and Y. Y. Gao, Continuous-variable quantum computation in circuit QED, Chin. Phys. Lett. **40**, 110303 (2023).

[59] S. Krastanov, V. V. Albert, C. Shen, C.-L. Zou, R. W. Heeres, B. Vlastakis, R. J. Schoelkopf, and L. Jiang, Universal control of an oscillator with dispersive coupling to a qubit, Phys. Rev. A **92**, 040303(R) (2015).

[60] A. Eickbusch, V. Sivak, A. Z. Ding, S. S. Elder, S. R. Jha, J. Venkatraman, B. Royer, S. M. Girvin, R. J. Schoelkopf, and M. H. Devoret, Fast universal control of an oscillator with weak dispersive coupling to a qubit, Nat. Phys. **18**, 1464 (2022).

[61] B. Mischuck and K. Mølmer, Qudit quantum computation in the Jaynes-Cummings model, Phys. Rev. A **87**, 022341 (2013).

[62] Y. Liu, J. Sinanan-Singh, M. T. Kearney, G. Mintzer, and I. L. Chuang, Constructing qudits from infinite-dimensional oscillators by coupling to qubits, Phys. Rev. A **104**, 032605 (2021).

[63] N. Ofek, A. Petrenko, R. Heeres, P. Reinhold, Z. Leghtas, B. Vlastakis, Y. Liu, L. Frunzio, S. M. Girvin, L. Jiang, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf, Extending the lifetime of a quantum bit with error correction in superconducting circuits, Nature **536**, 441 (2016).

[64] P. Campagne-Ibarcq, A. Eickbusch, S. Touzard, E. Zalys-Geller, N. E. Frattini, V. V. Sivak, P. Reinhold, S. Puri, S. Shankar, R. J. Schoelkopf, L. Frunzio, M. Mirrahimi, and M. H. Devoret, Quantum error correction of a qubit encoded in grid states of an oscillator, Nature **584**, 368 (2020).

[65] V. V. Sivak, A. Eickbusch, B. Royer, S. Singh, I. Tsioutsios, S. Ganjam, A. Miano, B. L. Brock, A. Z. Ding, L. Frunzio, S. M. Girvin, R. J. Schoelkopf, and M. H. Devoret, Real-time quantum error correction beyond break-even, Nature **616**, 50 (2023).

[66] B. L. Brock, S. Singh, A. Eickbusch, V. V. Sivak, A. Z. Ding, L. Frunzio, S. M. Girvin, and M. H. Devoret, Quantum error correction of qudits beyond break-even, Nature **641**, 612 (2025).

[67] Y. Ma, Y. Xu, X. Mu, W. Cai, L. Hu, W. Wang, X. Pan, H. Wang, Y. P. Song, C.-L. Zou, and L. Sun, Error-transparent operations on a logical qubit protected by quantum error correction, Nat. Phys. **16**, 827 (2020).

[68] Z. Ni, S. Li, X. Deng, Y. Cai, L. Zhang, W. Wang, Z.-B. Yang, H. Yu, F. Yan, S. Liu, C.-L. Zou, L. Sun, S.-B. Zheng, Y. Xu, and D. Yu, Beating the break-even point with a discrete-variable-encoded logical qubit, Nature **616**, 56 (2023).

[69] J. M. Gertler, B. Baker, J. Li, S. Shirol, J. Koch, and C. Wang, Protecting a bosonic qubit with autonomous quantum error correction, Nature **590**, 243 (2021).

[70] Y. Lu, A. Maiti, J. W. O. Garmon, S. Ganjam, Y. Zhang, J. Claes, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, High-fidelity parametric beamsplitting with

a parity-protected converter, Nat. Commun. **14**, 5767 (2023).
[71] J. Leppäkangas, P. Stadler, D. Golubev, R. Reiner, J.-M. Reiner, S. Zanker, N. Wurz, M. Renger, J. Verjauw, D. Gusenkova, S. Pogorzalek, F. Vigneau, P. Yang, W. Kindel, H.-S. Ku, F. Deppe, and M. Marthaler, Quantum algorithms for simulating systems coupled to bosonic modes using a hybrid resonator-qubit quantum computer, ArXiv:2503.11507.
[72] R. T. Sutherland and R. Srinivas, Universal hybrid quantum computing in trapped ions, Phys. Rev. A **104**, 032609 (2021).
[73] H. C. J. Gan, G. Maslennikov, K.-W. Tseng, C. Nguyen, and D. Matsukevich, Hybrid quantum computing with conditional beam splitter gate in trapped ion system, Phys. Rev. Lett. **124**, 170502 (2020).
[74] D. J. Gorman, P. Schindler, S. Selvarajan, N. Daniilidis, and H. Häffner, Two-mode coupling in a single-ion oscillator via parametric resonance, Phys. Rev. A **89**, 062332 (2014).
[75] P.-Y. Hou, J. J. Wu, S. D. Erickson, D. C. Cole, G. Zarantonello, A. D. Brandt, S. Geller, A. Kwiatkowski, S. Glancy, E. Knill, A. C. Wilson, D. H. Slichter, and D. Leibfried, Coherent coupling and non-destructive measurement of trapped-ion mechanical oscillators, Nat. Phys. **20**, 1636 (2024).
[76] O. Katz and C. Monroe, Programmable quantum simulations of bosonic systems with trapped ions, Phys. Rev. Lett. **131**, 033604 (2023).
[77] A. Blais, R.-S. Huang, A. Wallraff, S. M. Girvin, and R. J. Schoelkopf, Cavity quantum electrodynamics for superconducting electrical circuits: An architecture for quantum computation, Phys. Rev. A **69**, 062320 (2004).
[78] A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, R. S. Huang, J. Majer, S. Kumar, S. M. Girvin, and R. J. Schoelkopf, Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics, Nature **431**, 162 (2004).
[79] A. Sørensen and K. Mølmer, Quantum computation with ions in thermal motion, Phys. Rev. Lett. **82**, 1971 (1999).
[80] K. Mølmer and A. Sørensen, Multiparticle entanglement of hot trapped ions, Phys. Rev. Lett. **82**, 1835 (1999).
[81] G. Milburn, S. Schneider, and D. James, Ion trap quantum computing with warm ions, Fortschr. Phys. **48**, 801 (2000).
[82] A. Sørensen and K. Mølmer, Entanglement and quantum computation with ions in thermal motion, Phys. Rev. A **62**, 022311 (2000).
[83] O. Katz, M. Cetina, and C. Monroe, *n*-body interactions between trapped ion qubits via spin-dependent squeezing, Phys. Rev. Lett. **129**, 063603 (2022).
[84] O. Băzăvan, S. Saner, D. Webb, E. Ainley, P. Drmota, D. Nadlinger, G. Araneda, D. Lucas, C. Ballance, and R. Srinivas, Squeezing, trisqueezing, and quadsqueezing in a spin-oscillator system, ArXiv:2403.05471.
[85] O. Katz, L. Feng, A. Risinger, C. Monroe, and M. Cetina, Demonstration of three- and four-body interactions between trapped-ion spins, Nat. Phys. **19**, 1452 (2023).
[86] J. Whitlow, Z. Jia, Y. Wang, C. Fang, J. Kim, and K. R. Brown, Quantum simulation of conical intersections using trapped ions, Nat. Chem. **15**, 1509 (2023).
[87] A. Serafini, A. Retzker, and M. B. Plenio, Manipulating the quantum information of the radial modes of trapped ions: Linear phononics, entanglement generation, quantum state transmission and non-locality tests, New J. Phys. **11**, 023007 (2009).
[88] S. C. Burd, R. Srinivas, H. M. Knaack, W. Ge, A. C. Wilson, D. J. Wineland, D. Leibfried, J. J. Bollinger, D. T. C. Allcock, and D. H. Slichter, Quantum amplification of boson-mediated interactions, Nat. Phys. **17**, 898 (2021).
[89] V. So, M. D. Suganthi, A. Menon, M. Zhu, R. Zhuravel, H. Pu, P. G. Wolynes, J. N. Onuchic, and G. Pagano, Trapped-ion quantum simulation of electron transfer models with tunable dissipation, Sci. Adv. **10**, eads8011 (2024).
[90] M. Mielenz, H. Kalis, M. Wittemer, F. Hakelberg, U. Warring, R. Schmied, M. Blain, P. Maunz, D. L. Moehring, D. Leibfried, *et al.*, Arrays of individually controlled ions suitable for two-dimensional quantum simulations, Nat. Commun. **7**, 11839 (2016).
[91] S. Jain, J. Alonso, M. Grau, and J. P. Home, Scalable arrays of micro-Penning traps for quantum computing and simulation, Phys. Rev. X **10**, 031027 (2020).
[92] S. Jain, T. Sägesser, P. Hrmo, C. Torkzaban, M. Stadler, R. Oswald, C. Axline, A. Bautista-Salvador, C. Ospelkaus, D. Kienzler, *et al.*, Penning micro-trap for quantum computing, Nature **627**, 510 (2024).
[93] F. Hakelberg, P. Kiefer, M. Wittemer, U. Warring, and T. Schaetz, Interference in a prototype of a two-dimensional ion trap array quantum simulator, Phys. Rev. Lett. **123**, 100504 (2019).
[94] R. C. Sterling, H. Rattanasonti, S. Weidt, K. Lake, P. Srinivasan, S. Webster, M. Kraft, and W. K. Hensinger, Fabrication and operation of a two-dimensional ion-trap lattice on a high-voltage microchip, Nat. Commun. **5**, 3637 (2014).
[95] D. Palani, F. Hasse, P. Kiefer, F. Boeckling, J.-P. Schroeder, U. Warring, and T. Schaetz, High-fidelity transport of trapped-ion qubits in a multilayer array, Phys. Rev. A **107**, L050601 (2023).
[96] T. Olsacher, L. Postler, P. Schindler, T. Monz, P. Zoller, and L. M. Sieberer, Scalable and parallel tweezer gates for quantum computing with long ion strings, PRX Quantum **1**, 020316 (2020).
[97] D. Schwerdt, L. Peleg, Y. Shapira, N. Priel, Y. Florshaim, A. Gross, A. Zalic, G. Afek, N. Akerman, A. Stern, A. B. Kish, and R. Ozeri, Scalable architecture for trapped-ion quantum computing using rf traps and dynamic optical potentials, Phys. Rev. X **14**, 041017 (2024).
[98] A. Schuckert, O. Katz, L. Feng, E. Crane, A. De, M. Hafezi, A. V. Gorshkov, and C. Monroe, Observation of a finite-energy phase transition in a one-dimensional quantum simulator, Nat. Phys. **21**, 374 (2025).
[99] M. K. Joshi, F. Kranzl, A. Schuckert, I. Lovas, C. Maier, R. Blatt, M. Knap, and C. F. Roos, Observing emergent hydrodynamics in a long-range quantum magnet, Science **376**, 720 (2022).

[100] W. Chen, J. Gan, J.-N. Zhang, D. Matuskevich, and K. Kim, Quantum computation and simulation with vibrational modes of trapped ions, Chin. Phys. B **30**, 060311 (2021).
[101] W. Chen, Y. Lu, S. Zhang, K. Zhang, G. Huang, M. Qiao, X. Su, J. Zhang, J.-N. Zhang, L. Banchi, M. S. Kim, and K. Kim, Scalable and programmable phononic network with trapped ions, Nat. Phys. **19**, 877 (2023).
[102] S. Saner, O. Băzăvan, D. J. Webb, G. Araneda, D. M. Lucas, C. J. Ballance, and R. Srinivas, Generating arbitrary superpositions of nonclassical quantum harmonic oscillator states, ArXiv:2409.03482.
[103] K. Toyoda, R. Hiji, A. Noguchi, and S. Urabe, Hong-Ou-Mandel interference of two phonons in trapped ions, Nature **527**, 74 (2015).
[104] C.-H. Nguyen, K.-W. Tseng, G. Maslennikov, H. Gan, and D. Matsukevich, Experimental SWAP test of infinite dimensional quantum states, ArXiv:2103.10219.
[105] J. Y. Araz, M. Grau, J. Montgomery, and F. Ringer, Hybrid quantum simulations with qubits and qumodes on trapped-ion platforms, Phys. Rev. A **112**, 012620 (2025).
[106] C. Flühmann, T. L. Nguyen, M. Marinelli, V. Negnevitsky, K. Mehta, and J. P. Home, Encoding a qubit in a trapped-ion mechanical oscillator, Nature **566**, 513 (2019).
[107] B. de Neeve, T.-L. Nguyen, T. Behrle, and J. P. Home, Error correction of a logical grid state qubit by dissipative pumping, Nat. Phys. **18**, 296 (2022).
[108] C. Flühmann and J. P. Home, Direct characteristic-function tomography of quantum states of the trapped-ion motional oscillator, Phys. Rev. Lett. **125**, 043602 (2020).
[109] V. G. Matsos, C. H. Valahu, T. Navickas, A. D. Rao, M. J. Millican, X. C. Kolesnikow, M. J. Biercuk, and T. R. Tan, Robust and deterministic preparation of bosonic logical states in a trapped ion, Phys. Rev. Lett. **133**, 050602 (2024).
[110] M. Kang, H. Nuomin, S. N. Chowdhury, J. L. Yuly, K. Sun, J. Whitlow, J. Valdiviezo, Z. Zhang, P. Zhang, D. N. Beratan, and K. R. Brown, Seeking a quantum advantage with trapped-ion quantum simulations of condensed-phase chemical dynamics, Nat. Rev. Chem. **8**, 340 (2024).
[111] C. H. Valahu, V. C. Olaya-Agudelo, R. J. MacDonell, T. Navickas, A. D. Rao, M. J. Millican, J. B. Pérez-Sánchez, J. Yuen-Zhou, M. J. Biercuk, C. Hempel, T. R. Tan, and I. Kassal, Direct observation of geometric-phase interference in dynamics around a conical intersection, Nat. Chem. **15**, 1503 (2023).
[112] T. Nguyen, D. Lyakh, R. C. Pooser, T. S. Humble, T. Proctor, and M. Sarovar, Quantum circuit transformations with a multi-level intermediate representation compiler, ArXiv:2112.10677.
[113] Y. R. P. Sortais, H. Marion, C. Tuchendler, A. M. Lance, M. Lamare, P. Fournet, C. Armellin, R. Mercier, G. Messin, A. Browaeys, and P. Grangier, Diffraction-limited optics for single-atom manipulation, Phys. Rev. A **75**, 013406 (2007).
[114] A. M. Kaufman, B. J. Lester, and C. A. Regal, Cooling a single atom in an optical tweezer to its quantum ground state, Phys. Rev. X **2**, 041014 (2012).
[115] A. L. Shaw, P. Scholl, R. Finkelstein, R. B.-S. Tsai, J. Choi, and M. Endres, Erasure cooling, control, and hyper-entanglement of motion in optical tweezers, Science **388**, 845 (2025).
[116] A. M. Kaufman, B. J. Lester, C. M. Reynolds, M. L. Wall, M. Foss-Feig, K. R. A. Hazzard, A. M. Rey, and C. A. Regal, Two-particle quantum interference in tunnel-coupled optical tweezers, Science **345**, 306 (2014).
[117] A. W. Young, S. Geller, W. J. Eckner, N. Schine, S. Glancy, E. Knill, and A. M. Kaufman, An atomic boson sampler, Nature **629**, 311 (2024).
[118] L. H. Bohnmann, D. F. Locher, J. Zeiher, and M. Müller, Bosonic quantum error correction with neutral atoms in optical dipole traps, Phys. Rev. A **111**, 022432 (2025).
[119] P. T. Grochowski, H. Pichler, C. A. Regal, and O. Romero-Isart, Quantum control of continuous systems via nonharmonic potential modulation, Quantum **9**, 1824 (2025).
[120] A. V. Aho, M. S. Lam, R. Sethi, and J. D. Ullman, *Compilers: Principles, Techniques, & Tools* (Pearson Education India, Boston, 2007), 2nd ed.
[121] S. Khatri and M. M. Wilde, Principles of quantum communication theory: A modern approach, ArXiv:2011.04672.
[122] P. Murali, D. M. Debroy, K. R. Brown, and M. Martonosi, Architecting noisy intermediate-scale trapped ion quantum computers, in *Proceedings of the ACM/IEEE 47th Annual International Symposium on Computer Architecture, ISCA '20* (IEEE Press, 2020), p. 529.
[123] M. Oskin, F. T. Chong, I. L. Chuang, and J. Kubiatowicz, Building quantum wires: the long and the short of it, SIGARCH Comput. Archit. News **31**, 374 (2003).
[124] S. Xu, C. T. Hann, B. Foxman, S. M. Girvin, and Y. Ding, in *Proceedings of the 56th Annual IEEE/ACM International Symposium on Microarchitecture, MICRO '23* (ACM, New York, 2023), p. 526.
[125] V. Giovannetti, S. Lloyd, and L. Maccone, Quantum random access memory, Phys. Rev. Lett. **100**, 160501 (2008).
[126] V. Giovannetti, S. Lloyd, and L. Maccone, Architectures for a quantum random access memory, Phys. Rev. A **78**, 052310 (2008).
[127] S. Jaques and A. G. Rattew, QRAM: A survey and critique, Quantum **9**, 1922 (2025).
[128] C. T. Hann, G. Lee, S. M. Girvin, and L. Jiang, Resilience of quantum random access memory to generic noise, PRX Quantum **2**, 020311 (2021).
[129] F.-Y. Hong, Y. Xiang, Z.-Y. Zhu, L.-Z. Jiang, and L.-N. Wu, Robust quantum random access memory, Phys. Rev. A **86**, 010306(R) (2012).
[130] J. Liu, C. T. Hann, and L. Jiang, Data centers with quantum random access memory and quantum networks, Phys. Rev. A **108**, 032610 (2023).
[131] O. D. Matteo, V. Gheorghiu, and M. Mosca, Fault-tolerant resource estimation of quantum random-access memories, IEEE Trans. Quantum Eng. **1**, 1 (2020).
[132] D. K. Weiss, S. Puri, and S. M. Girvin, Qram architectures using superconducting cavities, PRX Quantum **5**, 020312 (2024).
[133] K. Phalak, A. Chatterjee, and S. Ghosh, Quantum random access memory for dummies, Sensors **23**, 17 (2023).

[134] H. J. Briegel, D. E. Browne, W. Dür, R. Raussendorf, and M. Van den Nest, Measurement-based quantum computation, Nat. Phys. **5**, 19 (2009).
[135] J. Preskill, Plug-in quantum software, Nature **402**, 357 (1999).
[136] S. Bravyi and A. Kitaev, Universal quantum computation with ideal Clifford gates and noisy ancillas, Phys. Rev. A **71**, 022316 (2005).
[137] QISKIT website, https://www.ibm.com/quantum/qiskit, 2025.
[138] H. Buhrman, M. Folkertsma, B. Loff, and N. M. P. Neumann, State preparation by shallow circuits using feed forward, Quantum **8**, 1552 (2024).
[139] N. M. P. Neumann, Ph.D. thesis, University of Amsterdam, https://hdl.handle.net/11245.1/8b878351-22bb-4fe9-bb5d-591a41d8a027, 2025.
[140] P. Pham and K. M. Svore, A 2D nearest-neighbor quantum architecture for factoring in polylogarithmic depth, Quantum Inf. Comput. **13**, 937 (2013).
[141] D. D. Thaker, T. S. Metodi, A. W. Cross, I. L. Chuang, and F. T. Chong, in *Proceedings of the 33rd Annual International Symposium on Computer Architecture, ISCA '06* (IEEE Computer Society, Washington, DC, 2006), p. 378.
[142] J. Majer, J. M. Chow, J. M. Gambetta, J. Koch, B. R. Johnson, J. A. Schreier, L. Frunzio, D. I. Schuster, A. A. Houck, A. Wallraff, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Coupling superconducting qubits via a cavity bus, Nature **449**, 443 (2007).
[143] C. R. Monroe, R. J. Schoelkopf, and M. D. Lukin, Quantum connections, Sci. Am. **314**, 50 (2016).
[144] P. O. Schmidt, T. Rosenband, C. Langer, W. M. Itano, J. C. Bergquist, and D. J. Wineland, Spectroscopy using quantum logic, Science **309**, 749 (2005).
[145] I. Strandberg, A. M. Eriksson, B. Royer, M. Kervinen, and S. Gasparinetti, Digital homodyne and heterodyne detection for stationary bosonic modes, Phys. Rev. Lett. **133**, 063601 (2024).
[146] B. R. Johnson, M. D. Reed, A. A. Houck, D. I. Schuster, L. S. Bishop, E. Ginossar, J. M. Gambetta, L. DiCarlo, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, Quantum non-demolition detection of single microwave photons in a circuit, Nat. Phys. **6**, 663 (2010).
[147] C. S. Wang, J. C. Curtis, B. J. Lester, Y. Zhang, Y. Y. Gao, J. Freeze, V. S. Batista, P. H. Vaccaro, I. L. Chuang, L. Frunzio, L. Jiang, S. M. Girvin, and R. J. Schoelkopf, Efficient multiphoton sampling of molecular vibronic spectra on a superconducting bosonic processor, Phys. Rev. X **10**, 021060 (2020).
[148] C. T. Hann, S. S. Elder, C. S. Wang, K. Chou, R. J. Schoelkopf, and L. Jiang, Robust readout of bosonic qubits in the dispersive coupling regime, Phys. Rev. A **98**, 022305 (2018).
[149] A. V. Dixit, S. Chakram, K. He, A. Agrawal, R. K. Naik, D. I. Schuster, and A. Chou, Searching for dark matter with a superconducting qubit, Phys. Rev. Lett. **126**, 141302 (2021).
[150] B. Vlastakis, G. Kirchmair, Z. Leghtas, S. E. Nigg, L. Frunzio, S. M. Girvin, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf, Deterministically encoding quantum information using 100-photon Schrödinger cat states, Science **342**, 607 (2013).
[151] B. Yurke, S. L. McCall, and J. R. Klauder, SU(2) and SU(1,1) interferometers, Phys. Rev. A **33**, 4033 (1986).
[152] K. M. Backes, *et al.*, A quantum enhanced search for dark matter axions, Nature **590**, 238 (2021).
[153] C. Oh, S. Chen, Y. Wong, S. Zhou, H.-Y. Huang, J. A. H. Nielsen, Z.-H. Liu, J. S. Neergaard-Nielsen, U. L. Andersen, L. Jiang, and J. Preskill, Entanglement-enabled advantage for learning a bosonic random displacement channel, Phys. Rev. Lett. **133**, 230604 (2024).
[154] S. Aaronson and A. Arkhipov, in *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery, New York, NY, USA, 2011), p. 333.
[155] Y.-H. Deng, *et al.*, Gaussian boson sampling with pseudo-photon-number-resolving detectors and quantum computational advantage, Phys. Rev. Lett. **131**, 150601 (2023).
[156] M. Tillmann, B. Dakić, R. Heilmann, S. Nolte, A. Szameit, and P. Walther, Experimental boson sampling, Nat. Photonics **7**, 540 (2013).
[157] J. B. Spring, B. J. Metcalf, P. C. Humphreys, W. S. Kolthammer, X.-M. Jin, M. Barbieri, A. Datta, N. Thomas-Peter, N. K. Langford, D. Kundys, *et al.*, Boson sampling on a photonic chip, Science **339**, 798 (2013).
[158] G. S. Thekkadath, S. Sempere-Llagostera, B. A. Bell, R. B. Patel, M. S. Kim, and I. A. Walmsley, Experimental demonstration of Gaussian boson sampling with displacement, PRX Quantum **3**, 020336 (2022).
[159] H.-S. Zhong, L.-C. Peng, Y. Li, Y. Hu, W. Li, J. Qin, D. Wu, W. Zhang, H. Li, L. Zhang, *et al.*, Experimental Gaussian boson sampling, Sci. Bull. **64**, 511 (2019).
[160] J. Huh, G. G. Guerreschi, B. Peropadre, J. R. McClean, and A. Aspuru-Guzik, Boson sampling for molecular vibronic spectra, Nat. Photonics **9**, 615 (2015).
[161] C. Oh, M. Liu, Y. Alexeev, B. Fefferman, and L. Jiang, Classical algorithm for simulating experimental Gaussian boson sampling, Nat. Phys. **20**, 1461 (2024).
[162] H. Qi, D. J. Brod, N. Quesada, and R. García-Patrón, Regimes of classical simulability for noisy Gaussian boson sampling, Phys. Rev. Lett. **124**, 100502 (2020).
[163] R. García-Patrón, J. J. Renema, and V. Shchesnovich, Simulating boson sampling in lossy architectures, Quantum **3**, 169 (2019).
[164] R. Babbush, D. W. Berry, R. Kothari, R. D. Somma, and N. Wiebe, Exponential quantum speedup in simulating coupled classical oscillators, Phys. Rev. X **13**, 041041 (2023).
[165] N. Heurtel, A. Fyrillas, G. d. Gliniasty, R. Le Bihan, S. Malherbe, M. Pailhas, E. Bertasi, B. Bourdoncle, P.-E. Emeriau, R. Mezher, L. Music, N. Belabas, B. Valiron, P. Senellart, S. Mansfield, and J. Senellart, PERCEVAL: A software platform for discrete variable photonic quantum computing, Quantum **7**, 931 (2023).
[166] C. Liu, M. Wang, S. A. Stein, Y. Ding, and A. Li, Quantum memory: A missing piece in quantum computing units, ArXiv:2309.14432.

[167] S. Stein, S. Sussman, T. Tomesh, C. Guinn, E. Tureci, S. F. Lin, W. Tang, J. Ang, S. Chakram, A. Li, M. Martonosi, F. Chong, A. A. Houck, I. L. Chuang, and M. Demarco, in *Proceedings of the 56th Annual IEEE/ACM International Symposium on Microarchitecture, MICRO '23* (Association for Computing Machinery, New York, 2023), p. 539.

[168] S. Arunachalam, V. Gheorghiu, T. Jochym-O'Connor, M. Mosca, and P. V. Srinivasan, On the robustness of bucket brigade quantum ram, New J. Phys. **17**, 123010 (2015).

[169] N. Guo, Y. Huang, T. Mai, S. Patil, C. Cao, M. Seok, S. Sethumadhavan, and Y. Tsividis, Energy-efficient hybrid analog/digital approximate computation in continuous time, IEEE J. Solid-State Circuits **51**, 1514 (2016).

[170] Y. Huang, N. Guo, M. Seok, Y. Tsividis, K. Mandli, and S. Sethumadhavan, in *Proceedings of the 50th Annual IEEE/ACM International Symposium on Microarchitecture, MICRO '17* (Association for Computing Machinery, New York, 2017), p. 665.

[171] S. M. Girvin, Introduction to quantum error correction and fault tolerance, SciPost Phys. Lect. Notes, 70 (2023).

[172] J. M. Arrazola, T. R. Bromley, J. Izaac, C. R. Myers, K. Brádler, and N. Killoran, Machine learning method for state preparation and gate synthesis on photonic quantum computers, Quantum Sci. Technol. **4**, 024004 (2019).

[173] M. Krenn, J. S. Kottmann, N. Tischler, and A. Aspuru-Guzik, Conceptual understanding through efficient automated design of quantum optical experiments, Phys. Rev. X **11**, 031044 (2021).

[174] M. Krenn, M. Erhard, and A. Zeilinger, Computer-inspired quantum experiments, Nat. Rev. Phys. **2**, 649 (2020).

[175] A. Cervera-Lierta, M. Krenn, and A. Aspuru-Guzik, Design of quantum optical experiments with logic artificial intelligence, Quantum **6**, 836 (2022).

[176] R. Nichols, L. Mineh, J. Rubio, J. C. F. Matthews, and P. A. Knott, Designing quantum experiments with a genetic algorithm, Quantum Sci. Technol. **4**, 045012 (2019).

[177] T. Fösel, S. Krastanov, F. Marquardt, and L. Jiang, Efficient cavity control with SNAP gates, ArXiv:2004.14256.

[178] U. Chabaud and M. Walschaers, Resources for bosonic quantum computational advantage, Phys. Rev. Lett. **130**, 090602 (2023).

[179] U. Chabaud, M. Joseph, S. Mehraban, and A. Motamedi, Bosonic quantum computational complexity, ArXiv:2410.04274.

[180] P. Reinhold, S. Rosenblum, W.-L. Ma, L. Frunzio, L. Jiang, and R. J. Schoelkopf, Error-corrected gates on an encoded qubit, Nat. Phys. **16**, 822 (2020).

[181] F. T. Chong, D. Franklin, and M. Martonosi, Programming languages and compiler design for realistic quantum hardware, Nature **549**, 180 (2017).

[182] B. Heim, M. Soeken, S. Marshall, C. Granade, M. Roetteler, A. Geller, M. Troyer, and K. Svore, Quantum programming languages, Nat. Rev. Phys. **2**, 709 (2020).

[183] A. JavadiAbhari, S. Patil, D. Kudrow, J. Heckey, A. Lvov, F. T. Chong, and M. Martonosi, ScaffCC, Parallel Comput. **45**, 2 (2015).

[184] Cirq Developers, Cirq, 2024, https://doi.org/10.5281/zenodo.11398048.

[185] K. Svore, A. Geller, M. Troyer, J. Azariah, C. Granade, B. Heim, V. Kliuchnikov, M. Mykhailova, A. Paz, and M. Roetteler, in *Proceedings of the Real World Domain Specific Languages Workshop 2018, RWDSL '18* (Association for Computing Machinery, New York, 2018).

[186] T. Häner, D. S. Steiger, K. Svore, and M. Troyer, A software methodology for compiling quantum programs, Quantum Sci. Technol. **3**, 020501 (2018).

[187] A. Cross, A. Javadi-Abhari, T. Alexander, N. De Beaudrap, L. S. Bishop, S. Heidel, C. A. Ryan, P. Sivarajah, J. Smolin, J. M. Gambetta, *et al.*, OpenQASM 3: A broader and deeper quantum assembly language, ACM Trans. Quantum Comput. **3**, 1 (2022).

[188] Quantum Intermediate Representation, 2020, https://devblogs.microsoft.com/qsharp/introducing-quantum-intermediate-representation-qir/.

[189] Quantum Intermediate Representation Alliance, 2022, https://github.com/qir-alliance.

[190] R. S. Smith, M. J. Curtis, and W. J. Zeng, A practical quantum instruction set architecture, ArXiv:1608.03355.

[191] Superconducting microwave cavity modes have weak amplitude damping (photon loss) but very little intrinsic dephasing. The mechanical modes in ion traps tend to suffer from dephasing but have little amplitude damping.

[192] We conservatively define memory break-even as the logical qubit lifetime exceeding the lifetime of the best quantum component among the physical qubits comprising the logical qubit.

[193] L. Hu, Y. Ma, W. Cai, X. Mu, Y. Xu, W. Wang, Y. Wu, H. Wang, Y. P. Song, C.-L. Zou, S. M. Girvin, L.-M. Duan, and L. Sun, Quantum error correction and universal gate set operation on a binomial bosonic logical qubit, Nat. Phys. **15**, 503 (2019).

[194] P. Jordan and E. Wigner, Über das Paulische Äquivalenzverbot, Z. Phys. **47**, 631 (1928).

[195] G. Ortiz, J. E. Gubernatis, E. Knill, and R. Laflamme, Quantum algorithms for fermionic simulations, Phys. Rev. A **64**, 022319 (2001).

[196] G. Ortiz, J. E. Gubernatis, E. Knill, and R. Laflamme, Erratum: Quantum algorithms for fermionic simulations [Phys. Rev. A 64, 022319 2001)], Phys. Rev. A **65**, 029902(E) (2002).

[197] J. T. Seeley, M. J. Richard, and P. J. Love, The Bravyi-Kitaev transformation for quantum computation of electronic structure, J. Chem. Phys. **137**, 224109 (2012).

[198] V. c. v. Havlíček, M. Troyer, and J. D. Whitfield, Operator locality in the quantum simulation of fermionic models, Phys. Rev. A **95**, 032332 (2017).

[199] N. P. Sawaya and J. Huh, Quantum algorithm for calculating molecular vibronic spectra, J. Chem. Phys. Lett. **10**, 3586 (2019).

[200] A. F. Shaw, P. Lougovski, J. R. Stryker, and N. Wiebe, Quantum algorithms for simulating the lattice Schwinger model, Quantum **4**, 306 (2020).

[201] M. Soeken, M. Roetteler, N. Wiebe, and G. De Micheli, in *Proceedings of the 54th Annual Design Automation Conference 2017, DAC '17* (Association for Computing Machinery, New York, 2017).

[202] M. K. Bhaskar, S. Hadfield, A. Papageorgiou, and I. Petras, Quantum algorithms and circuits for scientific computing, Quant. Inf. Comput. **16**, 197 (2016).

[203] C. Sparrow, E. Martín-López, N. Maraviglia, A. Neville, C. Harrold, J. Carolan, Y. N. Joglekar, T. Hashimoto, N. Matsuda, J. L. O'Brien, D. P. Tew, and A. Laing, Simulating the vibrational quantum dynamics of molecules using photonics, Nature **557**, 660 (2018).

[204] Y. Shen, Y. Lu, K. Zhang, J. Zhang, S. Zhang, J. Huh, and K. Kim, Quantum optical emulation of molecular vibronic spectroscopy using a trapped-ion device, Chem. Sci. **9**, 836 (2018).

[205] L. Hu, Y.-C. Ma, Y. Xu, W.-T. Wang, Y.-W. Ma, K. Liu, H.-Y. Wang, Y.-P. Song, M.-H. Yung, and L.-Y. Sun, Simulation of molecular spectroscopy with circuit quantum electrodynamics, Sci. Bull. **63**, 293 (2018).

[206] C. S. Wang, N. E. Frattini, B. J. Chapman, S. Puri, S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf, Observation of wave-packet branching through an engineered conical intersection, Phys. Rev. X **13**, 011008 (2023).

[207] L. G. Lutterbach and L. Davidovich, Method for direct measurement of the Wigner function in cavity QED and ion traps, Phys. Rev. Lett. **78**, 2547 (1997).

[208] P. Bertet, A. Auffeves, P. Maioli, S. Osnaghi, T. Meunier, M. Brune, J. M. Raimond, and S. Haroche, Direct measurement of the Wigner function of a one-photon Fock state in a cavity, Phys. Rev. Lett. **89**, 200402 (2002).

[209] M. Hofheinz, H. Wang, M. Ansmann, R. C. Bialczak, E. Lucero, M. Neeley, A. O'Connell, D. Sank, J. Wenner, J. M. Martinis, *et al.*, Synthesizing arbitrary quantum states in a superconducting resonator, Nature **459**, 546 (2009).

[210] S. Haroche and J.-M. Raimond, *Exploring the Quantum* (Oxford University Press, Oxford, 2006).

[211] D. F. Walls and G. J. Milburn, *Quantum Optics* (Springer Science & Business Media, London, 2007).

[212] B. Vlastakis, A. Petrenko, N. Ofek, L. Sun, Z. Leghtas, K. Sliwa, Y. Liu, M. Hatridge, J. Blumoff, L. Frunzio, M. Mirrahimi, L. Jiang, M. H. Devoret, and R. J. Schoelkopf, Characterizing entanglement of an artificial atom and a cavity cat state with Bell's inequality, Nat. Commun. **6**, 8970 (2015).

[213] Azure quantum documentation, https://learn.microsoft.com/en-us/azure/quantum/, 2025.

[214] H. C. P. Kendell, G. Ferranti, and C. A. Weidner, Deterministic generation of highly squeezed GKP states in ultracold atoms, APL Quantum **1**, 026109 (2024).

[215] A. Krasnok, P. Dhakal, A. Fedorov, P. Frigola, M. Kelly, and S. Kutsaev, Superconducting microwave cavities and qubits for quantum information systems, Appl. Phys. Rev. **11**, 011302 (2024).

[216] P. Krantz, M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, and W. D. Oliver, A quantum engineer's guide to superconducting qubits, Appl. Phys. Rev. **6**, 021318 (2019).

[217] S. M. Girvin, in *Circuit QED: Superconducting Qubits Coupled to Microwave Photons, in Proceedings of the 2011 Les Houches Summer School on Quantum Machines*, edited by M. H. Devoret, R. Schoelkopf, B. Huard, and C. Leticia (Oxford University Press, Oxford, 2014).

[218] S. L. Braunstein and P. van Loock, Quantum information with continuous variables, Rev. Mod. Phys. **77**, 513 (2005).

[219] S. D. Bartlett, B. C. Sanders, S. L. Braunstein, and K. Nemoto, Efficient classical simulation of continuous variable quantum information processes, Phys. Rev. Lett. **88**, 097904 (2002).

[220] A. Gilchrist, K. Nemoto, W. J. Munro, T. C. Ralph, S. Glancy, S. L. Braunstein, and G. J. Milburn, Schrödinger cats and their power for quantum information processing, J. Opt. B: Quantum Semiclassical Opt. **6**, S828 (2004).

[221] T. Aoki, G. Takahashi, T. Kajiya, J.-i. Yoshikawa, S. L. Braunstein, P. van Loock, and A. Furusawa, Quantum error correction beyond qubits, Nat. Phys. **5**, 541 (2009).

[222] C. Weedbrook, S. Pirandola, R. García-Patrón, N. J. Cerf, T. C. Ralph, J. H. Shapiro, and S. Lloyd, Gaussian quantum information, Rev. Mod. Phys. **84**, 621 (2012).

[223] N. Liu, J. Thompson, C. Weedbrook, S. Lloyd, V. Vedral, M. Gu, and K. Modi, Power of one qumode for quantum computation, Phys. Rev. A **93**, 052304 (2016).

[224] V. V. Albert, C. Shu, S. Krastanov, C. Shen, R.-B. Liu, Z.-B. Yang, R. J. Schoelkopf, M. Mirrahimi, M. H. Devoret, and L. Jiang, Holonomic quantum control with continuous variable systems, Phys. Rev. Lett. **116**, 140502 (2016).

[225] D. D. Stancil and G. T. Byrd, *Principles of Superconducting Quantum Computers* (John Wiley & Sons, Hoboken, New Jersey, 2022).

[226] P. Kok, W. J. Munro, K. Nemoto, T. C. Ralph, J. P. Dowling, and G. J. Milburn, Linear optical quantum computing with photonic qubits, Rev. Mod. Phys. **79**, 135 (2007).

[227] E. Knill, R. Laflamme, and G. J. Milburn, A scheme for efficient quantum computation with linear optics, Nature **409**, 46 (2001).

[228] J. E. Bourassa, R. N. Alexander, M. Vasmer, A. Patil, I. Tzitrin, T. Matsuura, D. Su, B. Q. Baragiola, S. Guha, G. Dauphinais, K. K. Sabapathy, N. C. Menicucci, and I. Dhand, Blueprint for a scalable photonic fault-tolerant quantum computer, Quantum **5**, 392 (2021).

[229] S. Bartolucci, P. Birchall, H. Bombín, H. Cable, C. Dawson, M. Gimeno-Segovia, E. Johnston, K. Kieling, N. Nickerson, M. Pant, F. Pastawski, T. Rudolph, and C. Sparrow, Fusion-based quantum computation, Nat. Commun. **14**, 912 (2023).

[230] K. Sahay, J. Claes, and S. Puri, Tailoring fusion-based error correction for high thresholds to biased fusion failures, Phys. Rev. Lett. **131**, 120604 (2023).

[231] S. Takeda and A. Furusawa, Toward large-scale fault-tolerant universal photonic quantum computing, APL Photonics **4**, 060902 (2019).

[232] J. L. O'Brien, Optical quantum computing, Science **318**, 1567 (2007).

[233] J. D. Teoh, P. Winkel, H. K. Babla, B. J. Chapman, J. Claes, S. J. de Graaf, J. W. O. Garmon, W. D. Kalfus, Y. Lu, A. Maiti, K. Sahay, N. Thakur, T. Tsunoda, S. H. Xue, L. Frunzio, S. M. Girvin, S. Puri, and R. J. Schoelkopf, Dual-rail encoding with superconducting cavities, PNAS **120**, e2221736120 (2023).

[234] K. S. Chou, *et al.*, A superconducting dual-rail cavity qubit with erasure-detected logical measurements, Nat. Phys. **20**, 1454 (2024).

[235] T. Tsunoda, J. D. Teoh, W. D. Kalfus, S. J. de Graaf, B. J. Chapman, J. C. Curtis, N. Thakur, S. M. Girvin, and R.

J. Schoelkopf, Error-detectable bosonic entangling gates with a noisy ancilla, PRX Quantum **4**, 020354 (2023).

[236] S. J. de Graaf, S. H. Xue, B. J. Chapman, J. D. Teoh, T. Tsunoda, P. Winkel, J. W. Garmon, K. M. Chang, L. Frunzio, S. Puri, *et al.*, A mid-circuit erasure check on a dual-rail cavity qubit using the joint-photon number-splitting regime of circuit QED, npj Quantum Inf. **11**, 1 (2025).

[237] B. J. Chapman, S. J. de Graaf, S. H. Xue, Y. Zhang, J. Teoh, J. C. Curtis, T. Tsunoda, A. Eickbusch, A. P. Read, A. Koottandavida, S. O. Mundhada, L. Frunzio, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, High-on-off-ratio beam-splitter interaction for gates on bosonically encoded qubits, PRX Quantum **4**, 020355 (2023).

[238] D. I. Schuster, A. A. Houck, J. A. Schreier, A. Wallraff, J. M. Gambetta, A. Blais, L. Frunzio, J. Majer, B. Johnson, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Resolving photon number states in a superconducting circuit, Nature **445**, 515 (2007).

[239] S. S. Elder, C. S. Wang, P. Reinhold, C. T. Hann, K. S. Chou, B. J. Lester, S. Rosenblum, L. Frunzio, L. Jiang, and R. J. Schoelkopf, High-fidelity measurement of qubits encoded in multilevel superconducting circuits, Phys. Rev. X **10**, 011001 (2020).

[240] L. Sun, A. Petrenko, Z. Leghtas, B. Vlastakis, G. Kirchmair, K. M. Sliwa, A. Narla, M. Hatridge, S. Shankar, J. Blumoff, L. Frunzio, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf, Tracking photon jumps with repeated quantum non-demolition parity measurements, Nature **511**, 444 (2014).

[241] I. L. Chuang and Y. Yamamoto, Simple quantum computer, Phys. Rev. A **52**, 3489 (1995).

[242] A. Koottandavida, I. Tsioutsios, A. Kargioti, C. R. Smith, V. R. Joshi, W. Dai, J. D. Teoh, J. C. Curtis, L. Frunzio, R. J. Schoelkopf, and M. H. Devoret, Erasure detection of a dual-rail qubit encoded in a double-post superconducting cavity, Phys. Rev. Lett. **132**, 180601 (2024).

[243] A. Kubica, A. Haim, Y. Vaknin, H. Levine, F. Brandão, and A. Retzker, Erasure qubits: Overcoming the T_1 limit in superconducting circuits, Phys. Rev. X **13**, 041022 (2023).

[244] H. Levine, A. Haim, J. S. Hung, N. Alidoust, M. Kalaee, L. DeLorenzo, E. A. Wollack, P. A. Arriola, A. Khalajhedayati, Y. Vaknin, *et al.*, Demonstrating a long-coherence dual-rail erasure qubit using tunable transmons, Phys. Rev. X **14**, 011051 (2024).

[245] Y. Wu, S. Kolkowitz, S. Puri, and J. D. Thompson, Erasure conversion for fault-tolerant quantum computing in alkaline earth Rydberg atom arrays, Nat. Commun. **13**, 4657 (2022).

[246] M. Kang, W. C. Campbell, and K. R. Brown, Quantum error correction with metastable states of trapped ions using erasure conversion, PRX Quantum **4**, 020358 (2023).

[247] K. Sahay, J. Jin, J. Claes, J. D. Thompson, and S. Puri, High-threshold codes for neutral-atom qubits with biased erasure errors, Phys. Rev. X **13**, 041013 (2023).

[248] S. Ma, G. Liu, P. Peng, B. Zhang, S. Jandura, J. Claes, A. P. Burgers, G. Pupillo, S. Puri, and J. D. Thompson, High-fidelity gates and mid-circuit erasure conversion in an atomic qubit, Nature **622**, 279 (2023).

[249] The standard physics notation agrees that $|\uparrow\rangle = |e\rangle$ is the excited state and $|\downarrow\rangle = |g\rangle$ is the ground state. However, in traditional quantum physics the excited states are often labeled as $|e\rangle = |1\rangle$ and $|g\rangle = |0\rangle$, in contrast to the conventions of quantum information.

[250] A. Macridin, A. C. Y. Li, and P. Spentzouris, Qumode transfer between continuous and discrete variable devices by near-deterministic teleportation protocols, Phys. Rev. A **109**, 032419 (2024).

[251] A. Z. Goldberg, A. B. Klimov, G. Leuchs, and L. L. Sanchez-Soto, Covariant operator bases for continuous variables, Quantum **8**, 1363 (2024).

[252] A. A. Clerk, M. H. Devoret, S. M. Girvin, F. Marquardt, and R. J. Schoelkopf, Introduction to quantum noise, measurement, and amplification, Rev. Mod. Phys. **82**, 1155 (2010).

[253] E. Wigner, On the quantum correction for thermodynamic equilibrium, Phys. Rev. **40**, 749 (1932).

[254] Y.-S. Kim and M. E. Noz, *Phase Space Picture of Quantum Mechanics: Group Theoretical Approach* (World Scientific, 1991), Vol. 40.

[255] Here, we are using the standard definitions of position and momentum that have Poisson bracket equal to unity rather than choosing the value of one-half that would correspond in the quantum case to the choice of Wigner units.

[256] U. Chabaud, Ph.D. thesis, Sorbonne Université, https://doi.org/10.48550/arXiv.2102.05227, 2021.

[257] U. Chabaud, D. Markham, and F. Grosshans, Stellar representation of non-Gaussian quantum states, Phys. Rev. Lett. **124**, 063605 (2020).

[258] U. Chabaud and S. Mehraban, Holomorphic representation of quantum computations, Quantum **6**, 831 (2022).

[259] S. M. Girvin and T. Jach, Formalism for the quantum Hall effect: Hilbert space of analytic functions, Phys. Rev. B **29**, 5617 (1984).

[260] J. B. Conway, *Functions of One Complex Variable II* (Springer Science & Business Media, 2012), Vol. 159.

[261] Uniquely since the overall normalization of the state is fixed and the global phase is irrelevant.

[262] M. Walschaers, Non-Gaussian quantum states and where to find them, PRX Quantum **2**, 030204 (2021).

[263] A. Aralov, Émilie Gillet, V. Nguyen, A. Cosentino, M. Walschaers, and M. Frigerio, Photon catalysis for general multimode multi-photon quantum state preparation, ArXiv:2507.19397.

[264] F. Arzani, R. I. Booth, and U. Chabaud, Can effective descriptions of bosonic systems be considered complete? ArXiv:2501.13857.

[265] Y. Tong, V. V. Albert, J. R. McClean, J. Preskill, and Y. Su, Provably accurate simulation of gauge theories and bosonic systems, Quantum **6**, 816 (2022).

[266] R. W. Heeres, B. Vlastakis, E. Holland, S. Krastanov, V. V. Albert, L. Frunzio, L. Jiang, and R. J. Schoelkopf, Cavity state manipulation using photon-number selective phase gates, Phys. Rev. Lett. **115**, 137002 (2015).

[267] M. Kudra, M. Kervinen, I. Strandberg, S. Ahmed, M. Scigliuzzo, A. Osman, D. P. Lozano, M. O. Tholén, R. Borgani, D. B. Haviland, G. Ferrini, J. Bylander, A. F. Kockum, F. Quijandría, P. Delsing, and S. Gasparinetti,

Robust preparation of Wigner-negative states with optimized snap-displacement sequences, PRX Quantum **3**, 030301 (2022).
[268] H.-Y. Lo, D. Kienzler, L. de Clercq, M. Marinelli, V. Negnevitsky, B. C. Keitch, and J. P. Home, Spin-motion entanglement and state diagnosis with squeezed oscillator wavepackets, Nature **521**, 336 (2015).
[269] D. Kienzler, H.-Y. Lo, B. Keitch, L. de Clercq, F. Leupold, F. Lindenfelser, M. Marinelli, V. Negnevitsky, and J. P. Home, Quantum harmonic oscillator state synthesis by reservoir engineering, Science **347**, 53 (2015).
[270] D. Gottesman, in *Group22: Proceedings of the XXII International Colloquium on Group Theoretical Methods in Physics*, edited by S. P. Corney, R. Delbourgo, and P. D. Jarvis (International Press, Cambridge, MA, 1999), p. 32.
[271] S. Aaronson and D. Gottesman, Improved simulation of stabilizer circuits, Phys. Rev. A **70**, 052328 (2004).
[272] A. Mari and J. Eisert, Positive Wigner functions render classical simulation of quantum computation efficient, Phys. Rev. Lett. **109**, 230503 (2012).
[273] V. Veitch, N. Wiebe, C. Ferrie, and J. Emerson, Efficient simulation scheme for a class of quantum optics experiments with non-negative Wigner representation, New J. Phys. **15**, 013037 (2013).
[274] G. H. Low and I. L. Chuang, Optimal Hamiltonian simulation by quantum signal processing, Phys. Rev. Lett. **118**, 010501 (2017).
[275] Z. M. Rossi and I. L. Chuang, Multivariable quantum signal processing (M-QSP): Prophecies of the two-headed oracle, Quantum **6**, 811 (2022).
[276] J. M. Martyn, Y. Liu, Z. E. Chin, and I. L. Chuang, Efficient fully-coherent quantum signal processing algorithms for real-time dynamics simulation, J. Chem. Phys. **158**, 024106 (2023).
[277] Z. M. Rossi, V. M. Bastidas, W. J. Munro, and I. L. Chuang, Quantum signal processing with continuous variables, ArXiv:2304.14383.
[278] J. Sinanan-Singh, G. L. Mintzer, I. L. Chuang, and Y. Liu, Single-shot quantum signal processing interferometry, Quantum **8**, 1427 (2024).
[279] C. Shen, K. Noh, V. V. Albert, S. Krastanov, M. H. Devoret, R. J. Schoelkopf, S. M. Girvin, and L. Jiang, Quantum channel construction with circuit quantum electrodynamics, Phys. Rev. B **95**, 134501 (2017).
[280] B. Royer, S. Singh, and S. M. Girvin, Encoding qubits in multimode grid states, PRX Quantum **3**, 010335 (2022).
[281] J. Schwinger, Unitary transformations and the action principle, PNAS **46**, 883 (1960).
[282] D. Gottesman, A. Kitaev, and J. Preskill, Encoding a qubit in an oscillator, Phys. Rev. A **64**, 012310 (2001).
[283] C. Gerry and P. Knight, *Introductory Quantum Optics* (Cambridge University Press, 2004).
[284] R. Dassonneville, R. Assouly, T. Peronnin, A. A. Clerk, A. Bienfait, and B. Huard, Dissipative stabilization of squeezing beyond 3 dB in a microwave mode, PRX Quantum **2**, 020323 (2021).
[285] Y. Zhang, B. J. Lester, Y. Y. Gao, L. Jiang, R. J. Schoelkopf, and S. M. Girvin, Engineering bilinear mode coupling in circuit QED: Theory and experiment, Phys. Rev. A **99**, 012314 (2019).
[286] T.-C. Chien, O. Lanes, C. Liu, X. Cao, P. Lu, S. Motz, G. Liu, D. Pekker, and M. Hatridge, Multiparametric amplification and qubit measurement with a Kerr-free Josephson ring modulator, Phys. Rev. A **101**, 042336 (2020).
[287] J.-I. Yoshikawa, Y. Miwa, A. Huck, U. L. Andersen, P. van Loock, and A. Furusawa, Demonstration of a quantum nondemolition SUM gate, Phys. Rev. Lett. **101**, 250501 (2008).
[288] D. Marković, S. Jezouin, Q. Ficheux, S. Fedortchenko, S. Felicetti, T. Coudreau, P. Milman, Z. Leghtas, and B. Huard, Demonstration of an effective ultrastrong coupling between two oscillators, Phys. Rev. Lett. **121**, 040505 (2018).
[289] S. L. Braunstein and R. I. McLachlan, Generalized squeezing, Phys. Rev. A **35**, 1659 (1987).
[290] C. W. Sandbo Chang, C. Sabín, P. Forn-Díaz, F. Quijandría, A. M. Vadiraj, I. Nsanzineza, G. Johansson, and C. M. Wilson, Observation of three-photon spontaneous parametric down-conversion in a superconducting parametric cavity, Phys. Rev. X **10**, 011011 (2020).
[291] A. M. Eriksson, T. Sépulcre, M. Kervinen, T. Hillmann, M. Kudra, S. Dupouy, Y. Lu, M. Khanahmadi, J. Yang, C. Castillo-Moreno, P. Delsing, and S. Gasparinetti, Universal control of a bosonic mode via drive-activated native cubic interactions, Nat. Commun. **15**, 2512 (2024).
[292] D. Zhang, D. Barral, Y. Zhang, M. Xiao, and K. Bencheikh, Genuine tripartite non-Gaussian entanglement, Phys. Rev. Lett. **130**, 093602 (2023).
[293] A. Agustí, C. W. Sandbo Chang, F. Quijandría, G. Johansson, C. M. Wilson, and C. Sabín, Tripartite genuine non-Gaussian entanglement in three-mode spontaneous parametric down-conversion, Phys. Rev. Lett. **125**, 020502 (2020).
[294] K. Marshall, R. Pooser, G. Siopsis, and C. Weedbrook, Repeat-until-success cubic phase gate for universal continuous-variable quantum computation, Phys. Rev. A **91**, 032321 (2015).
[295] Y. Zheng, O. Hahn, P. Stadler, P. Holmvall, F. Quijandría, A. Ferraro, and G. Ferrini, Gaussian conversion protocols for cubic phase state generation, PRX Quantum **2**, 010327 (2021).
[296] G. Kirchmair, B. Vlastakis, Z. Leghtas, S. E. Nigg, H. Paik, E. Ginossar, M. Mirrahimi, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, Observation of quantum state collapse and revival due to the single-photon Kerr effect, Nature **495**, 205 (2013).
[297] H.-K. Lau and M. B. Plenio, Universal quantum computing with arbitrary continuous-variable encoding, Phys. Rev. Lett. **117**, 100501 (2016).
[298] Y. Y. Gao, B. J. Lester, K. Chou, L. Frunzio, M. H. Devoret, L. Jiang, S. M. Girvin, and R. J. Schoelkopf, Entanglement of bosonic modes through an engineered exchange interaction, Nature **566**, 509 (2019).
[299] D. C. McKay, C. J. Wood, S. Sheldon, J. M. Chow, and J. M. Gambetta, Efficient z gates for quantum computing, Phys. Rev. A **96**, 022330 (2017).
[300] S. L. Braunstein, Squeezing as an irreducible resource, Phys. Rev. A **71**, 055801 (2005).
[301] G. Cariolaro and G. Pierobon, Bloch-Messiah reduction of Gaussian unitaries by Takagi factorization, Phys. Rev. A **94**, 062109 (2016).

[302] A. L. Grimsmo and S. Puri, Quantum error correction with the Gottesman-Kitaev-Preskill code, PRX Quantum **2**, 020101 (2021).
[303] M. A. Castellanos-Beltran, K. D. Irwin, G. C. Hilton, L. R. Vale, and K. W. Lehnert, Amplification and squeezing of quantum noise with a tunable Josephson metamaterial, Nat. Phys. **4**, 929 (2008).
[304] X. Pan, J. Schwinger, N.-N. Huang, P. Song, W. Chua, F. Hanamura, A. Joshi, F. Valadares, R. Filip, and Y. Y. Gao, Protecting the quantum interference of cat states by phase-space compression, Phys. Rev. X **13**, 021004 (2023).
[305] C. Lang, C. Eichler, L. Steffen, J. M. Fink, M. J. Woolley, A. Blais, and A. Wallraff, Correlations, indistinguishability and entanglement in Hong-Ou-Mandel experiments at microwave frequencies, Nat. Phys. **9**, 345 (2013).
[306] Y. Y. Gao, B. J. Lester, Y. Zhang, C. Wang, S. Rosenblum, L. Frunzio, L. Jiang, S. M. Girvin, and R. J. Schoelkopf, Programmable interference between two microwave quantum memories, Phys. Rev. X **8**, 021073 (2018).
[307] E. Flurin, N. Roch, F. Mallet, M. H. Devoret, and B. Huard, Generating entangled microwave radiation over two transmission lines, Phys. Rev. Lett. **109**, 183901 (2012).
[308] G. Andersson, S. W. Jolin, M. Scigliuzzo, R. Borgani, M. O. Tholén, J. C. Rivera Hernández, V. Shumeiko, D. B. Haviland, and P. Delsing, Squeezing and multimode entanglement of surface acoustic wave phonons, PRX Quantum **3**, 010312 (2022).
[309] C. Eichler, D. Bozyigit, C. Lang, M. Baur, L. Steffen, J. M. Fink, S. Filipp, and A. Wallraff, Observation of two-mode squeezing in the microwave frequency domain, Phys. Rev. Lett. **107**, 113601 (2011).
[310] A. Roy and M. Devoret, Introduction to parametric amplification of quantum signals with Josephson circuits, C. R. Phys. **17**, 740 (2016).
[311] M. Esposito, A. Ranadive, L. Planat, S. Leger, D. Fraudet, V. Jouanny, O. Buisson, W. Guichard, C. Naud, J. Aumentado, F. Lecocq, and N. Roch, Observation of two-mode squeezing in a traveling wave parametric amplifier, Phys. Rev. Lett. **128**, 153603 (2022).
[312] I. Tzitrin, J. E. Bourassa, N. C. Menicucci, and K. K. Sabapathy, Progress towards practical qubit computation using approximate Gottesman-Kitaev-Preskill codes, Phys. Rev. A **101**, 032315 (2020).
[313] S. D. Bartlett and B. C. Sanders, Requirement for quantum computation, J. Mod. Opt. **50**, 2331 (2003).
[314] S. Singh, PhD Thesis, Yale University, 2025.
[315] G. C. Wick, The evaluation of the collision matrix, Phys. Rev. **80**, 268 (1950).
[316] M. Howard, J. Wallman, V. Veitch, and J. Emerson, Contextuality supplies the "magic" for quantum computation, Nature **510**, 351 (2014).
[317] H. Pashayan, J. J. Wallman, and S. D. Bartlett, Estimating outcome probabilities of quantum circuits using quasiprobabilities, Phys. Rev. Lett. **115**, 070501 (2015).
[318] J. R. Seddon and E. T. Campbell, Quantifying magic for multi-qubit operations, Proc. R. Soc. A **475**, 20190251 (2019).
[319] S. Bravyi, D. Browne, P. Calpin, E. Campbell, D. Gosset, and M. Howard, Simulation of quantum circuits by low-rank stabilizer decompositions, Quantum **3**, 181 (2019).
[320] T. Haug and M. S. Kim, Scalable measures of magic resource for quantum computers, PRX Quantum **4**, 010301 (2023).
[321] K. Bu, W. Gu, and A. Jaffe, Quantum entropy and central limit theorem, PNAS **120**, e2304589120 (2023).
[322] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, Randomized benchmarking of quantum gates, Phys. Rev. A **77**, 012307 (2008).
[323] J. Niset, J. Fiurášek, and N. J. Cerf, No-go theorem for Gaussian quantum error correction, Phys. Rev. Lett. **102**, 120501 (2009).
[324] C.-H. Wang, K. Noh, J. Lebreuilly, S. M. Girvin, and L. Jiang, Photon-number-dependent Hamiltonian engineering for cavities, Phys. Rev. Appl. **15**, 044026 (2021).
[325] K. Jacobs, Engineering quantum states of a nanoresonator via a simple auxiliary system, Phys. Rev. Lett. **99**, 117203 (2007).
[326] B. Royer, S. Singh, and S. M. Girvin, Stabilization of finite-energy Gottesman-Kitaev-Preskill states, Phys. Rev. Lett. **125**, 260509 (2020).
[327] Z. Leghtas, G. Kirchmair, B. Vlastakis, M. H. Devoret, R. J. Schoelkopf, and M. Mirrahimi, Deterministic protocol for mapping a qubit to coherent state superpositions in a cavity, Phys. Rev. A **87**, 042315 (2013).
[328] P. C. Haljan, K.-A. Brickman, L. Deslauriers, P. J. Lee, and C. Monroe, Spin-dependent forces on trapped ions for phase-stable quantum gates and entangled states of spin and motion, Phys. Rev. Lett. **94**, 153602 (2005).
[329] R. W. Heeres, P. Reinhold, N. Ofek, L. Frunzio, L. Jiang, M. H. Devoret, and R. J. Schoelkopf, Implementing a universal gate set on a logical qubit encoded in an oscillator, Nat. Commun. **8**, 94 (2017).
[330] J. Job, Efficient, direct compilation of SU(N) operations into SNAP & Displacement gates, ArXiv:2307.11900.
[331] K. Park, P. Marek, and R. Filip, Efficient quantum simulation of nonlinear interactions using SNAP and Rabi gates, Quantum Sci. Technol. **9**, 025004 (2024).
[332] P.-T. Fong and H.-K. Lau, Engineering non-Gaussian bosonic gates through quantum signal processing, ArXiv:2508.20261.
[333] N. E. Frattini, U. Vool, S. Shankar, A. Narla, K. M. Sliwa, and M. H. Devoret, 3-wave mixing Josephson dipole element, Appl. Phys. Lett. **110**, 222603 (2017).
[334] T. Hillmann, F. Quijandría, G. Johansson, A. Ferraro, S. Gasparinetti, and G. Ferrini, Universal gate set for continuous-variable quantum computation with microwave circuits, Phys. Rev. Lett. **125**, 160501 (2020).
[335] S. Sefi and P. van Loock, How to decompose arbitrary continuous-variable quantum operations, Phys. Rev. Lett. **107**, 170501 (2011).
[336] T. W. Borneman, C. E. Granade, and D. G. Cory, Parallel information transfer in a multinode quantum information processor, Phys. Rev. Lett. **108**, 140502 (2012).

[337] A. M. Childs and N. Wiebe, Product formulas for exponentials of commutators, J. Math. Phys. **54**, 062202 (2013).
[338] Y.-A. Chen, A. M. Childs, M. Hafezi, Z. Jiang, H. Kim, and Y. Xu, Efficient product formulas for commutators and applications to quantum simulation, Phys. Rev. Res. **4**, 013191 (2022).
[339] A. M. Childs, Y. Su, M. C. Tran, N. Wiebe, and S. Zhu, Theory of Trotter error with commutator scaling, Phys. Rev. X **11**, 011020 (2021).
[340] C. Kang, M. B. Soley, E. Crane, S. M. Girvin, and N. Wiebe, Leveraging Hamiltonian simulation techniques to compile operations on bosonic devices, J. Phys. A: Math. Theor. **58**, 175301 (2025).
[341] C.-F. Chen, H.-Y. Huang, R. Kueng, and J. A. Tropp, Concentration for random product formulas, PRX Quantum **2**, 040305 (2021).
[342] A. A. Diringer, E. Blumenthal, A. Grinberg, L. Jiang, and S. Hacohen-Gourgy, Conditional-NOT displacement: Fast multioscillator control with a single qubit, Phys. Rev. X **14**, 011055 (2024).
[343] M. Um, J. Zhang, D. Lv, Y. Lu, S. An, J.-N. Zhang, H. Nha, M. S. Kim, and K. Kim, Phonon arithmetic in a trapped ion system, Nat. Commun. **7**, 11410 (2016).
[344] S. An, J.-N. Zhang, M. Um, D. Lv, Y. Lu, J. Zhang, Z.-Q. Yin, H. T. Quan, and K. Kim, Experimental test of the quantum Jarzynski equality with a trapped-ion system, Nat. Phys. **11**, 193 (2015).
[345] F. Beaudoin, M. P. da Silva, Z. Dutton, and A. Blais, First-order sidebands in circuit QED using qubit frequency modulation, Phys. Rev. A **86**, 022305 (2012).
[346] P. J. Leek, S. Filipp, P. Maurer, M. Baur, R. Bianchetti, J. M. Fink, M. Göppl, L. Steffen, and A. Wallraff, Using sideband transitions for two-qubit operations in superconducting circuits, Phys. Rev. B **79**, 180511(R) (2009).
[347] A. Z. Goldberg and A. M. Steinberg, Transcoherent states: Optical states for maximal generation of atomic coherence, PRX Quantum **1**, 020306 (2020).
[348] A. Z. Goldberg, A. M. Steinberg, and K. Heshami, Beyond transcoherent states: Field states for effecting optimal coherent rotations on single or multiple qubits, Quantum **7**, 963 (2023).
[349] A. B. Özgüler and D. Venturelli, Numerical gate synthesis for quantum heuristics on bosonic quantum processors, Front. Phys. **10**, 900612 (2022).
[350] N. C. Menicucci, P. van Loock, M. Gu, C. Weedbrook, T. C. Ralph, and M. A. Nielsen, Universal quantum computation with continuous-variable cluster states, Phys. Rev. Lett. **97**, 110501 (2006).
[351] M. Gu, C. Weedbrook, N. C. Menicucci, T. C. Ralph, and P. van Loock, Quantum computing with continuous-variable clusters, Phys. Rev. A **79**, 062318 (2009).
[352] S. Ghose and B. C. Sanders, Non-Gaussian ancilla states for continuous variable quantum computation via Gaussian maps, J. Mod. Opt. **54**, 855 (2007).
[353] B. Q. Baragiola, G. Pantaleoni, R. N. Alexander, A. Karanjai, and N. C. Menicucci, All-Gaussian universality and fault tolerance with the Gottesman-Kitaev-Preskill code, Phys. Rev. Lett. **123**, 200502 (2019).
[354] C. Calcluth, N. Reichel, A. Ferraro, and G. Ferrini, Sufficient condition for universal quantum computation using bosonic circuits, PRX Quantum **5**, 020337 (2024).
[355] A. Tranter, P. J. Love, F. Mintert, and P. V. Coveney, A comparison of the Bravyi-Kitaev and Jordan-Wigner transformations for the quantum simulation of quantum chemistry, J. Chem. Theory Comput. **14**, 5617 (2018).
[356] E. Jaynes and F. Cummings, Comparison of quantum and semiclassical radiation theories with application to the beam maser, Proc. IEEE **51**, 89 (1963).
[357] H. A. Gersch and G. C. Knollman, Quantum cell model for bosons, Phys. Rev. **129**, 959 (1963).
[358] W. P. Su, J. R. Schrieffer, and A. J. Heeger, Solitons in polyacetylene, Phys. Rev. Lett. **42**, 1698 (1979).
[359] I. I. Rabi, Space quantization in a gyrating magnetic field, Phys. Rev. **51**, 652 (1937).
[360] K. Hepp and E. H. Lieb, On the superradiant phase transition for molecules in a quantized radiation field: The Dicke maser model, Ann. Phys. **76**, 360 (1973).
[361] T. Holstein, Studies of polaron motion: Part II. The "small" polaron, Ann. Phys. **8**, 343 (1959).
[362] E. A. Arsenault, A. J. Schile, D. T. Limmer, and G. R. Fleming, Vibronic coupling in energy transfer dynamics and two-dimensional electronic-vibrational spectra, J. Chem. Phys. **155**, 054201 (2021).
[363] C. Schweizer, F. Grusdt, M. Berngruber, L. Barbiero, E. Demler, N. Goldman, I. Bloch, and M. Aidelsburger, Floquet approach to $\mathbb{Z}_2$ lattice gauge theories with ultracold atoms in optical lattices, Nat. Phys. **15**, 1168 (2019).
[364] J. Murugan, D. Stanford, and E. Witten, More on supersymmetric and 2D analogs of the SYK model, J. High Energy Phys. **2017**, 146 (2017).
[365] W. Fu and S. Sachdev, Numerical study of fermion and boson models with infinite-range random interactions, Phys. Rev. B **94**, 035135 (2016).
[366] I. Affleck, T. Kennedy, E. H. Lieb, and H. Tasaki, Rigorous results on valence-bond ground states in antiferromagnets, Phys. Rev. Lett. **59**, 799 (1987).
[367] A. Szabo and N. S. Ostlund, *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory* (Courier Corporation, 2012).
[368] D. P. Arovas, E. Berg, S. A. Kivelson, and S. Raghu, The Hubbard model, Annu. Rev. Condens. Matter Phys. **13**, 239 (2022).
[369] S. Lloyd, Coherent quantum feedback, Phys. Rev. A **62**, 022108 (2000).
[370] R. J. Nelson, Y. Weinstein, D. Cory, and S. Lloyd, Experimental demonstration of fully coherent quantum feedback, Phys. Rev. Lett. **85**, 3045 (2000).
[371] J. Tilly, H. Chen, S. Cao, D. Picozzi, K. Setia, Y. Li, E. Grant, L. Wossnig, I. Rungger, G. H. Booth, *et al.*, The variational quantum eigensolver: A review of methods and best practices, Phys. Rep. **986**, 1 (2022).
[372] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, *et al.*, Variational quantum algorithms, Nat. Rev. Phys. **3**, 625 (2021).
[373] K. Bharti, A. Cervera-Lierta, T. H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J. S.

Kottmann, T. Menke, W.-K. Mok, S. Sim, L.-C. Kwek, and A. Aspuru-Guzik, Noisy intermediate-scale quantum algorithms, Rev. Mod. Phys. **94**, 015004 (2022).
[374] K. Temme, S. Bravyi, and J. M. Gambetta, Error mitigation for short-depth quantum circuits, Phys. Rev. Lett. **119**, 180509 (2017).
[375] Y. Suzuki, S. Endo, K. Fujii, and Y. Tokunaga, Quantum error mitigation as a universal error reduction technique: Applications from the NISQ to the fault-tolerant quantum computing eras, PRX Quantum **3**, 010345 (2022).
[376] S. Endo, S. C. Benjamin, and Y. Li, Practical quantum error mitigation for near-future applications, Phys. Rev. X **8**, 031027 (2018).
[377] Z. Cai, R. Babbush, S. C. Benjamin, S. Endo, W. J. Huggins, Y. Li, J. R. McClean, and T. E. O'Brien, Quantum error mitigation, Rev. Mod. Phys. **95**, 045005 (2023).
[378] T. Steckmann, D. Luo, Y.-X. Wang, S. R. Muleady, A. Seif, C. Monroe, M. J. Gullans, A. V. Gorshkov, O. Katz, and A. Schuckert, Error mitigation of shot-to-shot fluctuations in analog quantum simulators, ArXiv:2506.16509.
[379] D. Su, R. Israel, K. Sharma, H. Qi, I. Dhand, and K. Brádler, Error mitigation on a near-term quantum photonic device, Quantum **5**, 452 (2021).
[380] J. Ang, *et al.*, ARQUIN: Architectures for multinode superconducting quantum computers, ACM Trans. Quantum Comput. **5**, 3 (2024).
[381] S. P. Kumar and I. Dhand, Optimal modular architectures for universal linear optics, ArXiv:2001.02012.
[382] A. J. Friedman, C. Yin, Y. Hong, and A. Lucas, Locality and error correction in quantum dynamics with measurement, ArXiv:2206.09929.
[383] L. Piroli, G. Styliaris, and J. I. Cirac, Quantum circuits assisted by local operations and classical communication: Transformations and phases of matter, Phys. Rev. Lett. **127**, 220503 (2021).
[384] N. Tantivasadakarn, A. Vishwanath, and R. Verresen, Hierarchy of topological order from finite-depth unitaries, measurement, and feedforward, PRX Quantum **4**, 020339 (2023).
[385] K. C. Smith, E. Crane, N. Wiebe, and S. M. Girvin, Deterministic constant-depth preparation of the AKLT state on a quantum processor using fusion measurements, PRX Quantum **4**, 020315 (2023).
[386] K. C. Smith, A. Khan, B. K. Clark, S. M. Girvin, and T.-C. Wei, Constant-depth preparation of matrix product states with adaptive quantum circuits, PRX Quantum **5**, 030344 (2024).
[387] S. H. Choe and R. König, How to fault-tolerantly realize any quantum circuit with local operations, PRX Quantum **6**, 010357 (2025).
[388] M. Arienzo, D. Grinko, M. Kliesch, and M. Heinrich, Bosonic randomized benchmarking with passive transformations, PRX Quantum **6**, 020305 (2025).
[389] M. Kliesch and I. Roth, Theory of quantum system certification, PRX Quantum **2**, 010201 (2021).
[390] J. Eisert, D. Hangleiter, N. Walk, I. Roth, D. Markham, R. Parekh, U. Chabaud, and E. Kashefi, Quantum certification and benchmarking, Nat. Rev. Phys. **2**, 382 (2020).
[391] I. Šupić and J. Bowles, Self-testing of quantum systems: A review, Quantum **4**, 337 (2020).
[392] U. Chabaud, T. Douce, F. Grosshans, E. Kashefi, and D. Markham, in *15th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2020)*, Leibniz International Proceedings in Informatics (LIPIcs), edited by S. T. Flammia (Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2020), Vol. 158, p. 3:1.
[393] U. Chabaud, F. Grosshans, E. Kashefi, and D. Markham, Efficient verification of boson sampling, Quantum **5**, 578 (2021).
[394] C. Wang, Y. Y. Gao, P. Reinhold, R. W. Heeres, N. Ofek, K. Chou, C. Axline, M. Reagor, J. Blumoff, K. M. Sliwa, L. Frunzio, S. M. Girvin, L. Jiang, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf, A Schrödinger cat living in two boxes, Science **352**, 1087 (2016).
[395] H. Jeon, J. Kang, W. Choi, K. Kim, J. You, and T. Kim, Multimode bosonic state tomography with single-shot joint parity measurement of a trapped ion, ArXiv:2506.12628.
[396] D. Gottesman, Stabilizer codes and quantum error correction, ArXiv:quant-ph/9705052.
[397] E. Magesan, J. M. Gambetta, B. R. Johnson, C. A. Ryan, J. M. Chow, S. T. Merkel, M. P. da Silva, G. A. Keefe, M. B. Rothwell, T. A. Ohki, M. B. Ketchen, and M. Steffen, Efficient measurement of quantum gate error by interleaved randomized benchmarking, Phys. Rev. Lett. **109**, 080505 (2012).
[398] R. Harper and S. T. Flammia, Estimating the fidelity of T gates using standard interleaved randomized benchmarking, Quantum Sci. Technol. **2**, 015008 (2017).
[399] J. T. Iosue, K. Sharma, M. J. Gullans, and V. V. Albert, Continuous-variable quantum state designs: Theory and applications, Phys. Rev. X **14**, 011013 (2024).
[400] J. B. Altepeter, D. Branning, E. Jeffrey, T. C. Wei, P. G. Kwiat, R. T. Thew, J. L. O'Brien, M. A. Nielsen, and A. G. White, Ancilla-assisted quantum process tomography, Phys. Rev. Lett. **90**, 193601 (2003).
[401] U. Chabaud, G. Roeland, M. Walschaers, F. Grosshans, V. Parigi, D. Markham, and N. Treps, Certification of non-Gaussian states with operational measurements, PRX Quantum **2**, 020333 (2021).
[402] S. Becker, N. Datta, L. Lami, and C. Rouze, Classical shadow tomography for continuous variables quantum systems, IEEE Trans. Inf. Theory **70**, 3427 (2024).
[403] S. Gandhari, V. V. Albert, T. Gerrits, J. M. Taylor, and M. J. Gullans, Precision bounds on continuous-variable state tomography using classical shadows, PRX Quantum **5**, 010346 (2024).
[404] K. Banaszek, C. Radzewicz, K. Wódkiewicz, and J. S. Krasiński, Direct measurement of the Wigner function by photon counting, Phys. Rev. A **60**, 674 (1999).
[405] M.-D. Choi, Completely positive linear maps on complex matrices, Linear Algebra Appl. **10**, 285 (1975).
[406] A. Jamiołkowski, Linear transformations which preserve trace and positive semidefiniteness of operators, Rep. Math. Phys. **3**, 275 (1972).
[407] N. Johnston and D. W. Kribs, Quantum gate fidelity in terms of Choi matrices, J. Phys. A: Math. Theor. **44**, 495303 (2011).

[408] J. M. Chow, J. M. Gambetta, A. D. Córcoles, S. T. Merkel, J. A. Smolin, C. Rigetti, S. Poletto, G. A. Keefe, M. B. Rothwell, J. R. Rozen, M. B. Ketchen, and M. Steffen, Universal quantum gate set approaching fault-tolerant thresholds with superconducting qubits, Phys. Rev. Lett. **109**, 060501 (2012).

[409] E. Nielsen, J. K. Gamble, K. Rudinger, T. Scholten, K. Young, and R. Blume-Kohout, Gate set tomography, Quantum **5**, 557 (2021).

[410] In practice this is a poor approximation since low-frequency dephasing noise often has a $1/f$ -like spectrum.

[411] A. Joshi, K. Noh, and Y. Y. Gao, Quantum information processing with bosonic qubits in circuit QED, Quantum Sci. Technol. **6**, 033001 (2021).

[412] W. Cai, Y. Ma, W. Wang, C.-L. Zou, and L. Sun, Bosonic quantum error correction codes in superconducting quantum circuits, Fundam. Res. **1**, 50 (2021).

[413] M. Mazyar, L. Zaki, V. A. Victor, T. Steven, J. S. Robert, J. Liang, and H. D. Michel, Dynamically protected cat-qubits: A new paradigm for universal quantum computation, New J. Phys. **16**, 045014 (2014).

[414] V. V. Albert, K. Noh, K. Duivenvoorden, D. J. Young, R. T. Brierley, P. Reinhold, C. Vuillot, L. Li, C. Shen, S. M. Girvin, B. M. Terhal, and L. Jiang, Performance and structure of single-mode bosonic codes, Phys. Rev. A **97**, 032346 (2018).

[415] A. L. Grimsmo, J. Combes, and B. Q. Baragiola, Quantum computing with rotation-symmetric bosonic codes, Phys. Rev. X **10**, 011058 (2020).

[416] M. H. Michael, M. Silveri, R. T. Brierley, V. V. Albert, J. Salmilehto, L. Jiang, and S. M. Girvin, New class of quantum error-correcting codes for a bosonic mode, Phys. Rev. X **6**, 031006 (2016).

[417] For a discussion of why one cannot achieve hardware scalability by replacing a *large* number of physical qubits by a single oscillator mode, see Ref. [733].

[418] D. Lachance-Quirion, M.-A. Lemonde, J. O. Simoneau, L. St-Jean, P. Lemieux, S. Turcotte, W. Wright, A. Lacroix, J. Fréchette-Viens, R. Shillito, *et al.*, Autonomous quantum error correction of Gottesman-Kitaev-Preskill states, Phys. Rev. Lett. **132**, 150607 (2024).

[419] S. Rosenblum, P. Reinhold, M. Mirrahimi, L. Jiang, L. Frunzio, and R. J. Schoelkopf, Fault-tolerant detection of a quantum error, Science **361**, 266 (2018).

[420] O. Vy, X. Wang, and K. Jacobs, Error-transparent evolution: The ability of multi-body interactions to bypass decoherence, New J. Phys. **15**, 053002 (2013).

[421] E. Kapit, Error-transparent quantum gates for small logical qubit architectures, Phys. Rev. Lett. **120**, 050503 (2018).

[422] W.-L. Ma, M. Zhang, Y. Wong, K. Noh, S. Rosenblum, P. Reinhold, R. J. Schoelkopf, and L. Jiang, Path-independent quantum gates with noisy ancilla, Phys. Rev. Lett. **125**, 110503 (2020).

[423] W.-L. Ma, S.-S. Li, and L. Jiang, Algebraic structure of path-independent quantum control, Phys. Rev. Res. **4**, 023102 (2022).

[424] E. Knill, R. Laflamme, R. Martinez, and C. Negrevergne, Benchmarking quantum computers: The five-qubit error correcting code, Phys. Rev. Lett. **86**, 5811 (2001).

[425] D. W. Leung, M. A. Nielsen, I. L. Chuang, and Y. Yamamoto, Approximate quantum error correction can lead to better codes, Phys. Rev. A **56**, 2567 (1997).

[426] C. Gidney and M. Ekerå, How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits, Quantum **5**, 433 (2021).

[427] E. Gouzien, D. Ruiz, F.-M. Le Régent, J. Guillaud, and N. Sangouard, Performance analysis of a repetition cat code architecture: Computing 256-bit elliptic curve logarithm in 9 hours with 126 133 cat qubits, Phys. Rev. Lett. **131**, 040602 (2023).

[428] N. Raveendran, N. Rengaswamy, F. Rozpędek, A. Raina, L. Jiang, and B. Vasić, Finite rate QLDPC-GKP coding scheme that surpasses the CSS Hamming bound, Quantum **6**, 767 (2022).

[429] I. L. Chuang, D. W. Leung, and Y. Yamamoto, Bosonic quantum codes for amplitude damping, Phys. Rev. A **56**, 1114 (1997).

[430] P. T. Cochrane, G. J. Milburn, and W. J. Munro, Macroscopically distinct quantum-superposition states as a bosonic code for amplitude damping, Phys. Rev. A **59**, 2631 (1999).

[431] K. Noh, V. V. Albert, and L. Jiang, Quantum capacity bounds of Gaussian thermal loss channels and achievable rates with Gottesman-Kitaev-Preskill codes, IEEE Trans. Inf. Theory **65**, 2563 (2019).

[432] Q. Xu, P. Zeng, D. Xu, and L. Jiang, Fault-tolerant operation of bosonic qubits with discrete-variable ancillae, Phys. Rev. X **14**, 031016 (2024).

[433] M.-J. Liang, T. Chen, and Z.-Y. Xue, Error-mitigated geometric quantum control over an oscillator, Phys. Rev. Appl. **23**, 024033 (2025).

[434] D. S. Schlegel, F. Minganti, and V. Savona, Quantum error correction using squeezed Schrödinger cat states, Phys. Rev. A **106**, 022431 (2022).

[435] P. David Rodríguez and K. Eliot, Improved autonomous error correction using variable dissipation in small logical qubit architectures, Quantum Sci. Technol. **6**, 015006 (2020).

[436] J. Lebreuilly, K. Noh, C.-H. Wang, S. M. Girvin, and L. Jiang, Autonomous quantum error correction and quantum computation, ArXiv:2103.05007.

[437] Z. Li, T. Roy, D. Rodríguez Pérez, K.-H. Lee, E. Kapit, and D. I. Schuster, Autonomous error correction of a single logical qubit using two transmons, Nat. Commun. **15**, 1681 (2024).

[438] Z. Li, T. Roy, D. R. Pérez, D. I. Schuster, and E. Kapit, Hardware-efficient autonomous error correction with linear couplers in superconducting circuits, Phys. Rev. Res. **6**, 013171 (2024).

[439] Q. Xu, G. Zheng, Y.-X. Wang, P. Zoller, A. A. Clerk, and L. Jiang, Autonomous quantum error correction and fault-tolerant quantum computation with squeezed cat qubits, npj Quantum Inf. **9**, 78 (2023).

[440] S. Touzard, A. Grimm, Z. Leghtas, S. O. Mundhada, P. Reinhold, C. Axline, M. Reagor, K. Chou, J. Blumoff, K. M. Sliwa, *et al.*, Coherent oscillations inside a quantum manifold stabilized by dissipation, Phys. Rev. X **8**, 021005 (2018).

[441] C. Berdou, *et al.*, One hundred second bit-flip time in a two-photon dissipative oscillator, PRX Quantum **4**, 020350 (2023).
[442] U. Réglade, A. Bocquet, R. Gautier, J. Cohen, A. Marquet, E. Albertinale, N. Pankratova, M. Hallén, F. Rautschke, L.-A. Sellem, P. Rouchon, A. Sarlette, M. Mirrahimi, P. Campagne-Ibarcq, R. Lescanne, S. Jezouin, and Z. Leghtas, Quantum control of a cat qubit with bit-flip times exceeding ten seconds, Nature **629**, 778 (2024).
[443] S. Puri, S. Boutin, and A. Blais, Engineering the quantum states of light in a Kerr-nonlinear resonator by two-photon driving, npj Quantum Inf. **3**, 18 (2017).
[444] S. Puri, L. St-Jean, J. A. Gross, A. Grimm, N. E. Frattini, P. S. Iyer, A. Krishna, S. Touzard, L. Jiang, A. Blais, S. T. Flammia, and S. M. Girvin, Bias-preserving gates with stabilized cat qubits, Sci. Adv. **6**, eaay5901 (2020).
[445] A. Grimm, N. E. Frattini, S. Puri, S. O. Mundhada, S. Touzard, M. Mirrahimi, S. M. Girvin, S. Shankar, and M. H. Devoret, Stabilization and operation of a Kerr-cat qubit, Nature **584**, 205 (2020).
[446] H. Putterman, J. Iverson, Q. Xu, L. Jiang, O. Painter, F. G. S. L. Brandão, and K. Noh, Stabilizing a bosonic qubit using colored dissipation, Phys. Rev. Lett. **128**, 110502 (2022).
[447] N. E. Frattini, R. G. Cortiñas, J. Venkatraman, X. Xiao, Q. Su, C. U. Lei, B. J. Chapman, V. R. Joshi, S. Girvin, R. J. Schoelkopf, S. Puri, and M. H. Devoret, Observation of pairwise level degeneracies and the quantum regime of the Arrhenius law in a double-well parametric oscillator, Phys. Rev. X **14**, 031040 (2024).
[448] J. Venkatraman, R. G. Cortinas, N. E. Frattini, X. Xiao, and M. H. Devoret, Quantum interference of tunneling paths under a double-well barrier, ArXiv:2211.04605.
[449] V. E. Manucharyan, J. Koch, L. I. Glazman, and M. H. Devoret, Fluxonium: Single Cooper-pair circuit free of charge offsets, Science **326**, 113 (2009).
[450] A. Somoroff, Q. Ficheux, R. A. Mencia, H. Xiong, R. Kuzmin, and V. E. Manucharyan, Millisecond coherence in a superconducting qubit, Phys. Rev. Lett. **130**, 267001 (2023).
[451] V. V. Albert, S. O. Mundhada, A. Grimm, S. Touzard, M. H. Devoret, and L. Jiang, Pair-cat codes: Autonomous error-correction with low-order nonlinearity, Quantum Sci. Technol. **4**, 035007 (2019).
[452] J. Harrington and J. Preskill, Achievable rates for the Gaussian quantum channel, Phys. Rev. A **64**, 062301 (2001).
[453] J. Conrad, J. Eisert, and F. Arzani, Gottesman-Kitaev-Preskill codes: A lattice perspective, Quantum **6**, 648 (2022).
[454] S. P. Jain, J. T. Iosue, A. Barg, and V. V. Albert, Quantum spherical codes, Nat. Phys. **20**, 1300 (2024).
[455] J. Conrad, J. Eisert, and J.-P. Seifert, Good Gottesman-Kitaev-Preskill codes from the NTRU cryptosystem, Quantum **8**, 1398 (2024).
[456] M. Lin, C. Chamberland, and K. Noh, Closest lattice point decoding for multimode Gottesman-Kitaev-Preskill codes, PRX Quantum **4**, 040334 (2023).
[457] J. Schwinger, On angular momentum, Technical Report (1952), https://doi.org/10.2172/4389568..
[458] M. Mathur, I. Raychowdhury, and R. Anishetty, SU(N) irreducible Schwinger bosons, J. Math. Phys. **51**, 093504 (2010).
[459] A. V. Gorshkov, M. Hermele, V. Gurarie, C. Xu, P. S. Julienne, J. Ye, P. Zoller, E. Demler, M. D. Lukin, and A. M. Rey, Two-orbital SU(N) magnetism with ultracold alkaline-earth atoms, Nat. Phys. **6**, 289 (2010).
[460] S. Chaudhury, S. Merkel, T. Herr, A. Silberfarb, I. H. Deutsch, and P. S. Jessen, Quantum control of the hyperfine spin of a Cs atom ensemble, Phys. Rev. Lett. **99**, 163002 (2007).
[461] C. Vuillot, H. Asasi, Y. Wang, L. P. Pryadko, and B. M. Terhal, Quantum error correction with the toric Gottesman-Kitaev-Preskill code, Phys. Rev. A **99**, 032344 (2019).
[462] K. Noh, S. M. Girvin, and L. Jiang, Encoding an oscillator into many oscillators, Phys. Rev. Lett. **125**, 080503 (2020).
[463] J. Wu and Q. Zhuang, Continuous-variable error correction for general Gaussian noises, Phys. Rev. Appl. **15**, 034073 (2021).
[464] J. Wu, A. J. Brady, and Q. Zhuang, Optimal encoding of oscillators into more oscillators, Quantum **7**, 1082 (2023).
[465] L. Hänggli and R. König, Oscillator-to-oscillator codes do not have a threshold, IEEE Trans. Inf. Theory **68**, 1068 (2022).
[466] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, Surface codes: Towards practical large-scale quantum computation, Phys. Rev. A **86**, 032324 (2012).
[467] K. Noh and C. Chamberland, Fault-tolerant bosonic quantum error correction with the surface–Gottesman-Kitaev-Preskill code, Phys. Rev. A **101**, 012316 (2020).
[468] K. Noh, C. Chamberland, and F. G. S. L. Brandão, Low-overhead fault-tolerant quantum error correction with the surface-GKP code, PRX Quantum **3**, 010315 (2022).
[469] B. M. Terhal, J. Conrad, and C. Vuillot, Towards scalable bosonic quantum error correction, Quantum Sci. Technol. **5**, 043001 (2020).
[470] A. S. Darmawan, B. J. Brown, A. L. Grimsmo, D. K. Tuckett, and S. Puri, Practical quantum error correction with the XZZX code and Kerr-cat qubits, PRX Quantum **2**, 030345 (2021).
[471] S. Singh, A. S. Darmawan, B. J. Brown, and S. Puri, High-fidelity magic-state preparation with a biased-noise architecture, Phys. Rev. A **105**, 052410 (2022).
[472] C. Chamberland, K. Noh, P. Arrangoiz-Arriola, E. T. Campbell, C. T. Hann, J. Iverson, H. Putterman, T. C. Bohdanowicz, S. T. Flammia, A. Keller, G. Refael, J. Preskill, L. Jiang, A. H. Safavi-Naeini, O. Painter, and F. G. S. L. Brandão, Building a fault-tolerant quantum computer using concatenated cat codes, PRX Quantum **3**, 010329 (2022).
[473] J. Guillaud and M. Mirrahimi, Repetition cat qubits for fault-tolerant quantum computation, Phys. Rev. X **9**, 041053 (2019).
[474] H. Putterman, *et al.*, Hardware-efficient quantum error correction via concatenated bosonic qubits, Nature **638**, 927 (2025).

[475] D. Ruiz, J. Guillaud, A. Leverrier, M. Mirrahimi, and C. Vuillot, LDPC-cat codes for low-overhead quantum computing in 2D, Nat. Commun. **16**, 1040 (2025).
[476] K. Fukui, A. Tomita, and A. Okamoto, Analog quantum error correction with encoding a qubit into an oscillator, Phys. Rev. Lett. **119**, 180507 (2017).
[477] Y. Xu, Y. Wang, E.-J. Kuo, and V. V. Albert, Qubit-oscillator concatenated codes: Decoding formalism and code comparison, PRX Quantum **4**, 020342 (2023).
[478] S. Singh, B. Royer, and S. M. Girvin, Towards non-Abelian quantum signal processing: Efficient control of hybrid continuous- and discrete-variable architectures, ArXiv:2504.19992.
[479] J. Hastrup, K. Park, J. B. Brask, R. Filip, and U. L. Andersen, Measurement-free preparation of grid states, npj Quantum Inf. **7**, 17 (2021).
[480] C. K. Law and J. H. Eberly, Arbitrary control of a quantum electromagnetic field, Phys. Rev. Lett. **76**, 1055 (1996).
[481] Disentangling the ancilla from cat states parallel to position axis in phase space using non-Abelian QSP entangles the ancilla with cat states parallel to the momentum axis.
[482] This is admissible since it can be tracked in software with a deterministic Pauli frame change after each stabilization step.
[483] X. Zhou, D. W. Leung, and I. L. Chuang, Methodology for quantum logic gate construction, Phys. Rev. A **62**, 052316 (2000).
[484] The sBs circuit rotates the ancilla state entangled with $|P_+\rangle_L$ ($|P_-\rangle_L$) by 4π (2π) Thus, the qubit is disentangled with an additional local phase of π on the logical state. Hence, this operation applies a deterministic logical Pauli P operation along with the logical single-qubit rotation. See Ref. [478] for details.
[485] I. Rojkov, P. M. Röggla, M. Wagener, M. Fontboté-Schmidt, S. Welte, J. Home, and F. Reiter, Two-qubit operations for finite-energy Gottesman-Kitaev-Preskill encodings, Phys. Rev. Lett. **133**, 100601 (2024).
[486] S. Chakraborty, Implementing any linear combination of unitaries on intermediate-term quantum computers, Quantum **8**, 1496 (2024).
[487] C. Yuan, C. McNally, and M. Carbin, TWIST: Sound reasoning for purity and entanglement in quantum programs, Proc. ACM Program. Lang. **6**, 1 (2022).
[488] S. Becker, N. Datta, L. Lami, and C. Rouzé, Energy-constrained discrimination of unitaries, quantum speed limits, and a Gaussian Solovay-Kitaev theorem, Phys. Rev. Lett. **126**, 190504 (2021).
[489] M. Gschwendtner and A. Winter, Infinite-dimensional programmable quantum processors, PRX Quantum **2**, 030308 (2021).
[490] K. Sharma, S. Khatri, M. Cerezo, and P. J. Coles, Noise resilience of variational quantum compiling, New J. Phys. **22**, 043006 (2020).
[491] H. Wang, Y. Ding, J. Gu, Y. Lin, D. Z. Pan, F. T. Chong, and S. Han, in *IEEE International Symposium on High-Performance Computer Architecture, HPCA '22* (2022), p. 692.
[492] P. Murali, J. M. Baker, A. Javadi-Abhari, F. T. Chong, and M. Martonosi, in *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS '19* (Association for Computing Machinery, New York, 2019), p. 1015.
[493] G. H. Low, T. J. Yoder, and I. L. Chuang, Methodology of resonant equiangular composite quantum gates, Phys. Rev. X **6**, 041067 (2016).
[494] D. J. Singh and L. Nordstrom, *Planewaves, Pseudopotentials, and the LAPW Method* (Springer Science & Business Media, 2006).
[495] G. J. Milburn, Simulating nonlinear spin models in an ion trap, ArXiv:quant-ph/9908037.
[496] Using vector identities, it can also be shown that $W_{\eta_2,\eta_1} = \frac{1}{4}[\hat{I}(1+\eta_1\eta_2\hat{c}\cdot\hat{b}) + \eta_2\hat{c}\cdot\vec{\sigma} + \eta_1\hat{b}\cdot\vec{\sigma} + i\eta_1\eta_2\hat{c}\times\hat{b}\cdot\vec{\sigma}]$.
[497] B. Dias and R. König, Classical simulation of non-Gaussian bosonic circuits, Phys. Rev. A **110**, 042402 (2024).
[498] A. W. Cross and J. M. Gambetta, Optimized pulse shapes for a resonator-induced phase gate, Phys. Rev. A **91**, 032325 (2015).
[499] H. Paik, A. Mezzacapo, M. Sandberg, D. T. McClure, B. Abdo, A. D. Corcoles, O. Dial, D. F. Bogorin, B. L. T. Plourde, M. Steffen, A. W. Cross, J. M. Gambetta, and J. M. Chow, Experimental demonstration of a resonator-induced phase gate in a multiqubit circuit-QED system, Phys. Rev. Lett. **117**, 250502 (2016).
[500] S. Puri and A. Blais, High-fidelity resonator-induced phase gate with single-mode squeezing, Phys. Rev. Lett. **116**, 180501 (2016).
[501] S. Sheldon, E. Magesan, J. M. Chow, and J. M. Gambetta, Procedure for systematically tuning up cross-talk in the cross-resonance gate, Phys. Rev. A **93**, 060302(R) (2016).
[502] A. M. Childs and N. Wiebe, Hamiltonian simulation using linear combinations of unitary operations, Quantum Inf. Comput. **12**, 901 (2012).
[503] J. Han, W. Cai, L. Hu, X. Mu, Y. Ma, Y. Xu, W. Wang, H. Wang, Y. P. Song, C.-L. Zou, and L. Sun, Experimental simulation of open quantum system dynamics via Trotterization, Phys. Rev. Lett. **127**, 020504 (2021).
[504] M. Suzuki, General theory of fractal path integrals with applications to many-body theories and statistical physics, J. Math. Phys. **32**, 400 (1991).
[505] J. Hastrup, K. Park, R. Filip, and U. L. Andersen, Unconditional preparation of squeezed vacuum from Rabi interactions, Phys. Rev. Lett. **126**, 153602 (2021).
[506] M. Ayyash, X. Xu, S. Ashhab, and M. Mariantoni, Driven multiphoton qubit-resonator interactions, Phys. Rev. A **110**, 053711 (2024).
[507] M. Ayyash, Multimode qubit-conditional operations via generalized cross-resonance, ArXiv:2503.15941.
[508] M. K. Hope, J. Lidal, and F. Massel, Preparation of conditionally-squeezed states in qubit-oscillator systems, ArXiv:2504.01664.
[509] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, A variational eigenvalue solver on a photonic quantum processor, Nat. Commun. **5**, 5213 (2014).
[510] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, ArXiv:1411.4028.
[511] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, Simulating Hamiltonian dynamics with

a truncated Taylor series, Phys. Rev. Lett. **114**, 090502 (2015).
[512] G. H. Low and I. L. Chuang, Hamiltonian simulation by uniform spectral amplification, ArXiv:1707.05391.
[513] L. Bell, Y. Wang, K. C. Smith, Y. Liu, E. Dumitrescu, and S. M. Girvin, Co-designing eigen- and singular-value transformation oracles: From algorithmic applications to hardware compilation, ArXiv:2502.16029.
[514] A. Gilyén, Ph.D. thesis, Universiteit van Amsterdam, https://hdl.handle.net/11245.1/20e9733e-6014-402d-afa9-20f3cc4a0568, 2019.
[515] N. Stamatopoulos and W. J. Zeng, Derivative pricing using quantum signal processing, Quantum **8**, 1322 (2024).
[516] S. Wimperis, Broadband, narrowband, and passband composite pulses for use in advanced NMR experiments, J. Magn. Reson., Ser. A **109**, 221 (1994).
[517] 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, PRX Quantum **3**, 040305 (2022).
[518] R. Chao, D. Ding, A. Gilyen, C. Huang, and M. Szegedy, Finding angles for quantum signal processing with machine precision, ArXiv:2003.02831.
[519] Z. M. Rossi, J. L. Ceroni, and I. L. Chuang, Modular quantum signal processing in many variables, Quantum **9**, 1776 (2025).
[520] B. Németh, B. Kövér, B. Kulcsár, R. B. Miklósi, and A. Gilyén, On variants of multivariate quantum signal processing and their characterizations, ArXiv:2312.09072.
[521] H. Le Jeannic, A. Cavaillès, K. Huang, R. Filip, and J. Laurat, Slowing quantum decoherence by squeezing in phase space, Phys. Rev. Lett. **120**, 073603 (2018).
[522] T. Hillmann and F. Quijandría, Quantum error correction with dissipatively stabilized squeezed-cat qubits, Phys. Rev. A **107**, 032423 (2023).
[523] S. Kan, M. Palma, Z. Du, S. A. Stein, C. Liu, J. Chen, A. Li, and Y. Mao, in *2024 IEEE International Conference on Quantum Computing and Engineering, QCE '24* (2024), Vol. 1, p. 1558.
[524] V. V. Sivak, A. Eickbusch, H. Liu, B. Royer, I. Tsioutsios, and M. H. Devoret, Model-free quantum control with reinforcement learning, Phys. Rev. X **12**, 011059 (2022).
[525] T. Schulte-Herbrüggen, A. Spörl, N. Khaneja, and S. J. Glaser, Optimal control for generating quantum gates in open dissipative systems, J. Phys. B: At. Mol. Opt. Phys. **44**, 154013 (2011).
[526] H. Ball, M. J. Biercuk, A. R. R. Carvalho, J. Chen, M. Hush, L. A. D. Castro, L. Li, P. J. Liebermann, H. J. Slatyer, C. Edmunds, V. Frey, C. Hempel, and A. Milne, Software tools for quantum control: Improving quantum computer performance through noise and error suppression, Quantum Sci. Technol. **6**, 044011 (2021).
[527] D. P. Kingma and J. Ba, in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7–9, 2015, Conference Track Proceedings*, edited by Y. Bengio and Y. LeCun (2015).
[528] We note through private communication that the fidelities presented here for points at circuit depth of $N = 2$ correct a minor simulation error that was later found for points at $N = 2$ in Fig. 2(a) of Ref. [60]. To the best of our knowledge, at $N = 2$ the results here represent the best known numerically optimized protocols.
[529] B. Zhang and Q. Zhuang, Energy-dependent barren plateau in bosonic variational quantum circuits, Quantum Sci. Technol. **10**, 015009 (2024).
[530] R. Porotti, V. Peano, and F. Marquardt, Gradient-ascent pulse engineering with feedback, PRX Quantum **4**, 030305 (2023).
[531] S. Khatri, R. LaRose, A. Poremba, L. Cincio, A. T. Sornborger, and P. J. Coles, Quantum-assisted quantum compiling, Quantum **3**, 140 (2019).
[532] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, Cost function dependent barren plateaus in shallow parametrized quantum circuits, Nat. Commun. **12**, 1791 (2021).
[533] A. K. Pati and S. L. Braunstein, in *Quantum Information with Continuous Variables*, edited by S. L. Braunstein and A. K. Pati (Springer Netherlands, Dordrecht, 2003), p. 31.
[534] A. K. Pati, S. L. Braunstein, and S. Lloyd, Quantum searching with continuous variables, ArXiv:quant-ph/0002082.
[535] S. J. Lomonaco Jr and L. H. Kauffman, A continuous variable Shor algorithm, ArXiv:quant-ph/0210141.
[536] S.-C. Wang, Y. Chi, L. Yu, Z. Sun, Q.-P. Su, and C.-P. Yang, Implementing a quantum search algorithm with nonorthogonal states, Phys. Rev. A **103**, 032413 (2021).
[537] A. P. Lund, A. Laing, S. Rahimi-Keshari, T. Rudolph, J. L. O'Brien, and T. C. Ralph, Boson sampling from a Gaussian state, Phys. Rev. Lett. **113**, 100502 (2014).
[538] C. S. Hamilton, R. Kruse, L. Sansoni, S. Barkhofen, C. Silberhorn, and I. Jex, Gaussian boson sampling, Phys. Rev. Lett. **119**, 170501 (2017).
[539] H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding, Y. Hu, *et al.*, Quantum computational advantage using photons, Science **370**, 1460 (2020).
[540] L. S. Madsen, F. Laudenbach, M. F. Askarani, F. Rortais, T. Vincent, J. F. Bulmer, F. M. Miatto, L. Neuhaus, L. G. Helt, M. J. Collins, *et al.*, Quantum computational advantage with a programmable photonic processor, Nature **606**, 75 (2022).
[541] Y. Liu, J. M. Martyn, J. Sinanan-Singh, K. C. Smith, S. M. Girvin, and I. L. Chuang, Toward mixed analog-digital quantum signal processing: Quantum AD/DA conversion and the Fourier transform, IEEE Trans. Signal Process. **73**, 3641 (2025).
[542] Here, we are taking $|x, \Delta\rangle_B$ to be an exact position eigenstate; we will remedy this assumption and evaluate performance on $|x, \Delta\rangle_B$ (which is not an exact position eigenstate) shortly.
[543] J. Hastrup, K. Park, J. B. Brask, R. Filip, and U. L. Andersen, Universal unitary transfer of continuous-variable quantum states into a few qubits, Phys. Rev. Lett. **128**, 110503 (2022).

[544] As mentioned in Ref. [543], this exact state is unphysical because it has infinite energy; however, it can be well approximated by a squeezed vacuum.

[545] P. Shor, in *Proceedings 35th Annual Symposium on Foundations of Computer Science, FOCS '94* (1994), p. 124.

[546] A. Gilyén, S. Arunachalam, and N. Wiebe, in *Proceedings of the 2019 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '19* (2019), p. 1425.

[547] National Science Technology Council, National quantum initiative supplement to the presidents FY 2022 budget (2022), https://www.quantum.gov/wp-content/uploads/2021/12/NQI-Annual-Report-FY2022.pdf.

[548] European Quantum Flagship, European quantum flagship strategic research agenda (2020), https://qt.eu//app/uploads/2020/04/Strategic_Research-_Agenda_d_FINAL.pdf.

[549] C. W. Bauer, *et al.*, Quantum simulation for high-energy physics, PRX Quantum **4**, 027001 (2023).

[550] Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, N. P. D. Sawaya, S. Sim, L. Veis, and A. Aspuru-Guzik, Quantum chemistry in the age of quantum computing, Chem. Rev. **119**, 10856 (2019).

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

[552] H.-D. Meyer, U. Manthe, and L. Cederbaum, The multi-configurational time-dependent Hartree approach, Chem. Phys. Lett. **165**, 73 (1990).

[553] M. Marques and E. Gross, Time-dependent density functional theory, Annu. Rev. Phys. Chem. **55**, 427 (2004).

[554] A. Polkovnikov, Phase space representation of quantum dynamics, Ann. Phys. **325**, 1790 (2010).

[555] R. J. Bartlett, Many-body perturbation theory and coupled cluster theory for electron correlation in molecules, Annu. Rev. Phys. Chem. **32**, 359 (1981).

[556] R. J. Bartlett and M. Musiał, Coupled-cluster theory in quantum chemistry, Rev. Mod. Phys. **79**, 291 (2007).

[557] M. C. Bañuls, Tensor network algorithms: A route map, Annu. Rev. Condens. Matter Phys. **14**, 173 (2023).

[558] W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, Quantum Monte Carlo simulations of solids, Rev. Mod. Phys. **73**, 33 (2001).

[559] J. Danzer, C. Gattringer, L. Liptak, and M. Marinkovic, A study of the sign problem for lattice QCD with chemical potential, Phys. Lett. B **682**, 240 (2009).

[560] L. Bongiovanni, Numerical methods for the sign problem in lattice field theory, ArXiv:1603.06458.

[561] R. P. Feynman, Quantum mechanical computers, Opt. News **11**, 11 (1985).

[562] It is important to distinguish between two cases: (a) simulation of the time evolution of a quantum system with a specified Hamiltonian and a specified starting state; and (b) finding the ground state (or other eigenstate) of a specified Hamiltonian. For most physical systems, the former is in the complexity class BQP and the latter is in QMA which is a quantum analog of NP and is believed to be exponentially more challenging than BQP [1].

[563] E. Altman, *et al.*, Quantum simulators: Architectures and opportunities, PRX Quantum **2**, 017003 (2021).

[564] A. J. Daley, I. Bloch, C. Kokail, S. Flannigan, N. Pearson, M. Troyer, and P. Zoller, Practical quantum advantage in quantum simulation, Nature **607**, 667 (2022).

[565] Y. Salathé, M. Mondal, M. Oppliger, J. Heinsoo, P. Kurpiers, A. Potočnik, A. Mezzacapo, U. Las Heras, L. Lamata, E. Solano, S. Filipp, and A. Wallraff, Digital quantum simulation of spin models with circuit quantum electrodynamics, Phys. Rev. X **5**, 021027 (2015).

[566] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. Van Den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme, *et al.*, Evidence for the utility of quantum computing before fault tolerance, Nature **618**, 500 (2023).

[567] R. Barends, L. Lamata, J. Kelly, L. García-Álvarez, A. G. Fowler, A. Megrant, E. Jeffrey, T. C. White, D. Sank, J. Y. Mutus, *et al.*, Digital quantum simulation of fermionic models with a superconducting circuit, Nat. Commun. **6**, 7654 (2015).

[568] T. V. Zache, D. González-Cuadra, and P. Zoller, Fermion-qudit quantum processors for simulating lattice gauge theories with matter, Quantum **7**, 1140 (2023).

[569] J. Simon, W. S. Bakr, R. Ma, M. E. Tai, P. M. Preiss, and M. Greiner, Quantum simulation of antiferromagnetic spin chains in an optical lattice, Nature **472**, 307 (2011).

[570] C. Monroe, W. C. Campbell, L.-M. Duan, Z.-X. Gong, A. V. Gorshkov, P. W. Hess, R. Islam, K. Kim, N. M. Linke, G. Pagano, P. Richerme, C. Senko, and N. Y. Yao, Programmable quantum simulations of spin systems with trapped ions, Rev. Mod. Phys. **93**, 025001 (2021).

[571] M. Greiner, O. Mandel, T. Esslinger, T. W. Hänsch, and I. Bloch, Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms, Nature **415**, 39 (2002).

[572] A. Mazurenko, C. S. Chiu, G. Ji, M. F. Parsons, M. Kanász-Nagy, R. Schmidt, F. Grusdt, E. Demler, D. Greif, and M. Greiner, A cold-atom Fermi-Hubbard antiferromagnet, Nature **545**, 462 (2017).

[573] X. Zhang, K. Zhang, Y. Shen, S. Zhang, J.-N. Zhang, M.-H. Yung, J. Casanova, J. S. Pedernales, L. Lamata, E. Solano, and K. Kim, Experimental quantum simulation of fermion-antifermion scattering via boson exchange in a trapped ion, Nat. Commun. **9**, 195 (2018).

[574] A. Mil, T. V. Zache, A. Hegde, A. Xia, R. P. Bhatt, M. K. Oberthaler, P. Hauke, J. Berges, and F. Jendrzejewski, A scalable realization of local U(1) gauge invariance in cold atomic mixtures, Science **367**, 1128 (2020).

[575] W. Domcke and D. R. Yarkony, Role of conical intersections in molecular spectroscopy and photoinduced chemical dynamics, Annu. Rev. Phys. Chem. **63**, 325 (2012).

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

[577] J.-P. Bureik, G. Hercé, M. Allemand, A. Tenart, T. Roscilde, and D. Clément, Suppression of Bogoliubov momentum pairing and emergence of non-Gaussian correlations in ultracold interacting Bose gases, Nat. Phys. **21**, 57 (2025).

[578] A. Kitaev and J. Preskill, Topological entanglement entropy, Phys. Rev. Lett. **96**, 110404 (2006).

[579] M. Levin and X.-G. Wen, Detecting topological order in a ground state wave function, Phys. Rev. Lett. **96**, 110405 (2006).
[580] S. Lloyd, Universal quantum simulators, Science **273**, 1073 (1996).
[581] G. H. Low, V. Kliuchnikov, and N. Wiebe, Well-conditioned multiproduct Hamiltonian simulation, ArXiv: 1907.11679.
[582] A. M. Childs, Universal computation by quantum walk, Phys. Rev. Lett. **102**, 180501 (2009).
[583] D. W. Berry and A. M. Childs, Black-box Hamiltonian simulation and unitary implementation, Quantum Inf. Comput. **12**, 29 (2012).
[584] D. W. Berry, A. M. Childs, and R. Kothari, in *2015 IEEE 56th Annual Symposium on Foundations of Computer Science (FOCS)* (2015).
[585] G. H. Low and I. L. Chuang, Hamiltonian simulation by qubitization, Quantum **3**, 163 (2019).
[586] E. Campbell, Random compiler for fast Hamiltonian simulation, Phys. Rev. Lett. **123**, 070503 (2019).
[587] P. K. Faehrmann, M. Steudtner, R. Kueng, M. Kieferová, and J. Eisert, Randomizing multi-product formulas for Hamiltonian simulation, Quantum **6**, 806 (2022).
[588] A. Y. Kitaev, Quantum measurements and the Abelian stabilizer problem, ArXiv:quant-ph/9511026.
[589] S. Bravyi, A. Chowdhury, D. Gosset, and P. Wocjan, Quantum Hamiltonian complexity in thermal equilibrium, Nat. Phys. **18**, 1367 (2022).
[590] S. Lu, M. C. Bañuls, and J. I. Cirac, Algorithms for quantum simulation at finite energies, PRX Quantum **2**, 020321 (2021).
[591] A. Schuckert, A. Bohrdt, E. Crane, and M. Knap, Probing finite-temperature observables in quantum simulators of spin systems with short-time dynamics, Phys. Rev. B **107**, L140410 (2023).
[592] Y. Zhang, S. Jahanbani, A. Riswadkar, S. Shankar, and A. C. Potter, Sequential quantum simulation of spin chains with a single circuit QED device, Phys. Rev. A **109**, 022606 (2024).
[593] S. Morita and H. Nishimori, Mathematical foundation of quantum annealing, J. Math. Phys. **49**, 125210 (2008).
[594] S. Lee, J. Lee, H. Zhai, Y. Tong, A. M. Dalzell, A. Kumar, P. Helms, J. Gray, Z.-H. Cui, W. Liu, M. Kastoryano, R. Babbush, J. Preskill, D. R. Reichman, E. T. Campbell, E. F. Valeev, L. Lin, and G. K.-L. Chan, Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry, Nat. Commun. **14**, 1952 (2023).
[595] F. Pollmann, E. Berg, A. M. Turner, and M. Oshikawa, Symmetry protection of topological phases in one-dimensional quantum spin systems, Phys. Rev. B **85**, 075125 (2012).
[596] A. Taylor, G. Bressanini, H. Kwon, and M. S. Kim, Quantum error cancellation in photonic systems: Undoing photon losses, Phys. Rev. A **110**, 022622 (2024).
[597] H.-Y. Huang, R. Kueng, and J. Preskill, Predicting many properties of a quantum system from very few measurements, Nat. Phys. **16**, 1050 (2020).
[598] C. Eichler, D. Bozyigit, C. Lang, L. Steffen, J. Fink, and A. Wallraff, Experimental state tomography of itinerant single microwave photons, Phys. Rev. Lett. **106**, 220503 (2011).
[599] D. Braak, Integrability of the Rabi model, Phys. Rev. Lett. **107**, 100401 (2011).
[600] A. F. Kockum, A. Miranowicz, S. D. Liberato, S. Savasta, and F. Nori, Ultrastrong coupling between light and matter, Nat. Rev. Phys. **1**, 19 (2019).
[601] P. Forn-Díaz, L. Lamata, E. Rico, J. Kono, and E. Solano, Ultrastrong coupling regimes of light-matter interaction, Rev. Mod. Phys. **91**, 025005 (2019).
[602] S. Aaronson and A. Arkhipov, The computational complexity of linear optics, Theory Comput. **9**, 143 (2013).
[603] S. M. Girvin, in *Topological Aspects of Low-dimensional Systems: Proceedings, Les Houches Summer School of Theoretical Physics, Session 69: Les Houches, France, July 7–31, 1998*, Les Houches—École d'Été de Physique Theorique, edited by A. Comtet, T. Jolicœur, S. Ouvry, and F. David (Springer, Berlin, 1999), Vol. 69.
[604] E. Kapit, Quantum simulation architecture for lattice bosons in arbitrary, tunable, external gauge fields, Phys. Rev. A **87**, 062336 (2013).
[605] M. Hafezi, P. Adhikari, and J. M. Taylor, Engineering three-body interaction and Pfaffian states in circuit QED systems, Phys. Rev. B **90**, 060503(R) (2014).
[606] C. Owens, A. LaChapelle, B. Saxberg, B. M. Anderson, R. Ma, J. Simon, and D. I. Schuster, Quarter-flux Hofstadter lattice in a qubit-compatible microwave cavity array, Phys. Rev. A **97**, 013818 (2018).
[607] J. C. Owens, M. G. Panetta, B. Saxberg, G. Roberts, S. Chakram, R. Ma, A. Vrajitoarea, J. Simon, and D. I. Schuster, Chiral cavity quantum electrodynamics, Nat. Phys. **18**, 1048 (2022).
[608] P. D. Kurilovich, V. D. Kurilovich, J. Lebreuilly, and S. M. Girvin, Stabilizing the Laughlin state of light: Dynamics of hole fractionalization, SciPost Phys. **13**, 107 (2022).
[609] M. Aidelsburger, *Artificial Gauge Fields with Ultracold Atoms in Optical Lattices* (Springer International Publishing, 2016).
[610] J. Léonard, S. Kim, J. Kwan, P. Segura, F. Grusdt, C. Repellin, N. Goldman, and M. Greiner, Realization of a fractional quantum Hall state with ultracold atoms, Nature **619**, 495 (2023).
[611] L. W. Clark, N. Schine, C. Baum, N. Jia, and J. Simon, Observation of Laughlin states made of light, Nature **582**, 41 (2020).
[612] C. Wang, F.-M. Liu, M.-C. Chen, H. Chen, X.-H. Zhao, C. Ying, Z.-X. Shang, J.-W. Wang, Y.-H. Huo, C.-Z. Peng, *et al.*, Realization of fractional quantum Hall state with interacting photons, Science **384**, 579 (2024).
[613] M. C. Tran, D. K. Mark, W. W. Ho, and S. Choi, Measuring arbitrary physical properties in analog quantum simulation, Phys. Rev. X **13**, 011049 (2023).
[614] A. Lukin, M. Rispoli, R. Schittko, M. E. Tai, A. M. Kaufman, S. Choi, V. Khemani, J. Léonard, and M. Greiner, Probing entanglement in a many-body–localized system, Science **364**, 256 (2019).
[615] R. N. Palmer and D. Jaksch, High-field fractional quantum Hall effect in optical lattices, Phys. Rev. Lett. **96**, 180407 (2006).
[616] S. G. Tabrizi, A. V. Arbuznikov, and M. Kaupp, Construction of giant-spin Hamiltonians from many-spin

Hamiltonians by third-order perturbation theory and application to an Fe_3Cr single-molecule magnet, Chem. Eur. J. **22**, 6853 (2016).
[617] J. Schwinger, in *Quantum Mechanics* (Springer, 2001), p. 149.
[618] U. Wiese, Ultracold quantum gases and lattice systems: Quantum simulation of lattice gauge theories, Ann. Phys. **525**, 777 (2013).
[619] S. Roy, A. Senanian, C. S. Wang, O. C. Wetherbee, L. Zhang, B. Cole, C. P. Larson, E. Yelton, K. Arora, P. L. McMahon, B. L. T. Plourde, B. Royer, and V. Fatemi, Synthetic high angular momentum spin dynamics in a microwave oscillator, Phys. Rev. X **15**, 021009 (2025).
[620] E. Champion, Z. Wang, R. W. Parker, and M. S. Blok, Efficient control of a transmon qudit using effective spin-7/2 rotations, Phys. Rev. X **15**, 021096 (2025).
[621] P. Giorda, P. Zanardi, and S. Lloyd, Universal quantum control in irreducible state-space sectors: Application to bosonic and spin-boson systems, Phys. Rev. A **68**, 062320 (2003).
[622] B. E. Anderson, H. Sosa-Martinez, C. A. Riofrío, I. H. Deutsch, and P. S. Jessen, Accurate and robust unitary transformations of a high-dimensional quantum system, Phys. Rev. Lett. **114**, 240401 (2015).
[623] R. G. Newton and B.-L. Young, Measurability of the spin density matrix, Ann. Phys. **49**, 393 (1968).
[624] M. A. Perlin, D. Barberena, and A. M. Rey, Spin qudit tomography and state reconstruction error, Phys. Rev. A **104**, 062413 (2021).
[625] S. M. Girvin and K. Yang, *Modern Condensed Matter Physics* (Cambridge University Press, 2019).
[626] R. F. Evans, W. J. Fan, P. Chureemart, T. A. Ostler, M. O. Ellis, and R. W. Chantrell, Atomistic spin model simulations of magnetic nanomaterials, J. Phys.: Condens. Matter **26**, 103202 (2014).
[627] J. J. Hopfield, Neural networks and physical systems with emergent collective computational abilities, PNAS **79**, 2554 (1982).
[628] M. P. Fisher, Quantum cognition: The possibility of processing with nuclear spins in the brain, Ann. Phys. **362**, 593 (2015).
[629] A. Winter and H. Rieger, Quantum phase transition and correlations in the multi-spin-boson model, Phys. Rev. B **90**, 224401 (2014).
[630] S. Schmidt and G. Blatter, Strong coupling theory for the Jaynes-Cummings-Hubbard model, Phys. Rev. Lett. **103**, 086403 (2009).
[631] K. C. Smith, A. Bhattacharya, and D. J. Masiello, Exact k-body representation of the Jaynes-Cummings interaction in the dressed basis: Insight into many-body phenomena with light, Phys. Rev. A **104**, 013707 (2021).
[632] S. Kundu and N. Makri, Intramolecular vibrations in excitation energy transfer: Insights from real-time path integral calculations, Annu. Rev. Phys. Chem. **73**, 349 (2022).
[633] M. R. Wasielewski, M. D. Forbes, N. L. Frank, K. Kowalski, G. D. Scholes, J. Yuen-Zhou, M. A. Baldo, D. E. Freedman, R. H. Goldsmith, T. Goodson III, *et al.*, Exploiting chemistry and molecular systems for quantum information science, Nat. Rev. Chem. **4**, 490 (2020).
[634] M. Kebric, L. Barbiero, C. Reinmoser, U. Schollwöck, and F. Grusdt, Confinement and Mott transitions of dynamical charges in one-dimensional lattice gauge theories, Phys. Rev. Lett. **127**, 167203 (2021).
[635] C. Ratti, Lattice QCD and heavy ion collisions: a review of recent progress, Rep. Prog. Phys. **81**, 084301 (2018).
[636] J. A. Flores-Livas, L. Boeri, A. Sanna, G. Profeta, R. Arita, and M. Eremets, A perspective on conventional high-temperature superconductors at high pressure: Methods and materials, Phys. Rep. **856**, 1 (2020).
[637] R. Dutta, *et al.*, Simulating chemistry on bosonic quantum devices, J. Chem. Theory Comput. **20**, 6426 (2024).
[638] R. Dutta, N. P. Vu, C. Xu, D. G. A. Cabral, N. Lyu, A. V. Soudackov, X. Dan, H. Li, C. Wang, and V. S. Batista, Simulating electronic structure on bosonic quantum computers, J. Chem. Theory Comput. **21**, 2281 (2025).
[639] F. Verstraete and J. I. Cirac, Mapping local Hamiltonians of fermions to local Hamiltonians of spins, J. Stat. Mech: Theory Exp. **2005**, P09012 (2005).
[640] C. Derby, J. Klassen, J. Bausch, and T. Cubitt, Compact fermion to qubit mappings, Phys. Rev. B **104**, 035118 (2021).
[641] I. D. Kivlichan, J. McClean, N. Wiebe, C. Gidney, A. Aspuru-Guzik, GarnetKin-Lic Chan, and R. Babbush, Quantum simulation of electronic structure with linear depth and connectivity, Phys. Rev. Lett. **120**, 110501 (2018).
[642] M. G. Algaba, P. V. Sriluckshmy, M. Leib, and F. Šimkovic IV, Low-depth simulations of fermionic systems on square-grid quantum hardware, Quantum **8**, 1327 (2024).
[643] Y. Aharonov, L. Davidovich, and N. Zagury, Quantum random walks, Phys. Rev. A **48**, 1687 (1993).
[644] E. Flurin, V. V. Ramasesh, S. Hacohen-Gourgy, L. S. Martin, N. Y. Yao, and I. Siddiqi, Observing topological invariants using quantum walks in superconducting circuits, Phys. Rev. X **7**, 031023 (2017).
[645] S. P. Jordan, K. S. Lee, and J. Preskill, Quantum computation of scattering in scalar quantum field theories, Quantum Inf. Comput. **14**, 1014 (2014).
[646] F. Gray, Pulse code communication, United States Patent Number 2,632,058 (1953), https://patentimages.storage.googleapis.com/a3/d7/f2/0343f5f2c0cf50/US2632058.pdf.
[647] N. P. D. Sawaya, T. Menke, T. H. Kyaw, S. Johri, A. Aspuru-Guzik, and G. G. Guerreschi, Resource-efficient digital quantum simulation of *d*-level systems for photonic, vibrational, and spin-*s* Hamiltonians, npj Quantum Inf. **6**, 49 (2020).
[648] Note, however, that L cannot be made too large because the maximum energy stored in the oscillator (and thus its decay rate) grows exponentially with L. See Ref. [733] for a general discussion of why the Hilbert space of quantum computers must be built up from tensor products of multiple subsystems.
[649] D. Janković, J.-G. Hartmann, M. Ruben, and P.-A. Hervieux, Noisy qudit vs multiple qubits: Conditions on gate efficiency for enhancing fidelity, npj Quantum Inf. **10**, 59 (2024).

[650] S. Rosenblum, Y. Y. Gao, P. Reinhold, C. Wang, C. J. Axline, L. Frunzio, S. M. Girvin, L. Jiang, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf, A CNOT gate between multiphoton qubits encoded in two cavities, Nat. Commun. **9**, 652 (2018).
[651] https://en.wikipedia.org/wiki/Gray_code.
[652] X. Li, G. Yang, C. M. Torres, D. Zheng, and K. L. Wang, A class of efficient quantum incrementer gates for quantum circuit synthesis, Int. J. Mod. Phys. B **28**, 1350191 (2014).
[653] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, Elementary gates for quantum computation, Phys. Rev. A **52**, 3457 (1995).
[654] T. Häner, M. Roetteler, and K. M. Svore, Optimizing quantum circuits for arithmetic, ArXiv:1805.12445.
[655] N. Wiebe, D. Berry, P. Høyer, and B. C. Sanders, Higher order decompositions of ordered operator exponentials, J. Phys. A: Math. Theor. **43**, 065203 (2010).
[656] A. M. Krol and Z. Al-Ars, Beyond quantum Shannon: Circuit construction for general *n*-qubit gates based on block ZXZ-decomposition, ArXiv:2403.13692.
[657] C. F. Kane, S. Hariprakash, N. S. Modi, M. Kreshchuk, and C. W. Bauer, Block encoding bosons by signal processing, Quantum **9**, 1747 (2025).
[658] M. Hanada, S. Matsuura, E. Mendicelli, and E. Rinaldi, Exponential improvement in quantum simulations of bosons, ArXiv:2505.02553.
[659] D. A. Patterson and J. L. Hennessy, *Computer Organization and Design RISC-V Edition: The Hardware Software Interface* (Morgan Kaufmann Publishers Inc., San Francisco, California, 2017), 1st ed.
[660] D. R. Weinberg, C. J. Gagliardi, J. F. Hull, C. F. Murphy, C. A. Kent, B. C. Westlake, A. Paul, D. H. Ess, D. G. McCafferty, and T. J. Meyer, Proton-coupled electron transfer, Chem. Rev. **112**, 4016 (2012).
[661] S. Hammes-Schiffer, Proton-coupled electron transfer: Moving together and charging forward, JACS **137**, 8860 (2015).
[662] J. Bloch, A. Cavalleri, V. Galitski, M. Hafezi, and A. Rubio, Strongly correlated electron-photon systems, Nature **606**, 41 (2022).
[663] R. F. Ribeiro, L. A. Martínez-Martínez, M. Du, J. Campos-Gonzalez-Angulo, and J. Yuen-Zhou, Polariton chemistry: Controlling molecular dynamics with optical cavities, Chem. Sci. **9**, 6325 (2018).
[664] A. V. Oppenheim and R. W. Schafer, *Discrete-Time Signal Processing* (3rd Ed.), (Pearson, Upper Saddle River, Boston, Columbus, San Francisco, New York, Indianapolis, London, Toronto, Sydney, Singapore. Tokyo, Montreal, Dubai, Madrid, Hong Kong, Mexico City, Munich, Paris, Amsterdam, Cape Town, 2009).
[665] C. M. Dawson and M. A. Nielsen, The Solovay-Kitaev algorithm, Quantum Inf. Comput. **6**, 81 (2006).
[666] V. Kliuchnikov, D. Maslov, and M. Mosca, Asymptotically optimal approximation of single qubit unitaries by Clifford and *T* circuits using a constant number of ancillary qubits, Phys. Rev. Lett. **110**, 190502 (2013).
[667] J. Haah, Product decomposition of periodic functions in quantum signal processing, Quantum **3**, 190 (2019).
[668] E. Gouzien and N. Sangouard, Factoring 2048-bit RSA integers in 177 days with 13 436 qubits and a multimode memory, Phys. Rev. Lett. **127**, 140503 (2021).
[669] C. Duckering, J. M. Baker, D. I. Schuster, and F. T. Chong, in *53rd Annual IEEE/ACM International Symposium on Microarchitecture, MICRO '20* (Association for Computing Machinery, New York, 2020), p. 173.
[670] I. Pietikäinen, O. Černotík, A. Eickbusch, A. Maiti, J. W. O. Garmon, R. Filip, and S. M. Girvin, Strategies and trade-offs for controllability and memory time of ultra-high-quality microwave cavities in circuit QED, PRX Quantum **5**, 040307 (2024).
[671] A. Krishna and J.-P. Tillich, Towards low overhead magic state distillation, Phys. Rev. Lett. **123**, 070507 (2019).
[672] E. T. Campbell, H. Anwar, and D. E. Browne, Magic-state distillation in all prime dimensions using quantum Reed-Muller codes, Phys. Rev. X **2**, 041021 (2012).
[673] Y. Yamamoto, K. Aihara, T. Leleu, K.-i. Kawarabayashi, S. Kako, M. Fejer, K. Inoue, and H. Takesue, Coherent Ising machines—optical neural networks operating at the quantum limit, npj Quantum Inf. **3**, 49 (2017).
[674] T. Wang, L. Wu, P. Nobel, and J. Roychowdhury, Solving combinatorial optimisation problems using oscillator based Ising machines, Nat. Comput. **20**, 287 (2021).
[675] M. Bonnin, F. L. Traversa, and F. Bonani, in *Advances in Artificial Intelligence-based Technologies: Selected Papers in Honour of Professor Nikolaos G. Bourbakis—Vol. 1*, edited by M. Virvou, G. A. Tsihrintzis, L. H. Tsoukalas, and L. C. Jain (Springer International Publishing, Cham, 2022), p. 179.
[676] B. Coecke and R. Duncan, Interacting quantum observables: Categorical algebra and diagrammatics, New J. Phys. **13**, 043016 (2011).
[677] J. van de Wetering, ZX-calculus for the working quantum computer scientist, ArXiv:2012.13966.
[678] H. Nagayoshi, W. Asavanant, R. Ide, K. Fukui, A. Sakaguchi, J. ichi Yoshikawa, N. C. Menicucci, and A. Furusawa, ZX graphical calculus for continuous-variable quantum processes, Phys. Rev. Res. **7**, 033141 (2025).
[679] C. Chamberland and M. E. Beverland, Flag fault-tolerant error correction with arbitrary distance codes, Quantum **2**, 53 (2018).
[680] E. van den Berg, S. Bravyi, J. M. Gambetta, P. Jurcevic, D. Maslov, and K. Temme, Single-shot error mitigation by coherent Pauli checks, Phys. Rev. Res. **5**, 033193 (2023).
[681] Y. Liu, S. Singh, K. C. Smith, E. Crane, J. M. Martyn, A. Eickbusch, A. Schuckert, R. D. Li, J. Sinanan-Singh, M. B. Soley, *et al.*, Hybrid Oscillator-Qubit Quantum Processors: Instruction Set Architectures, Abstract Machine Models, and Applications [GitHub], 2025, https://github.com/shraggy/ISA-simulations.
[682] J. Koch, T. M. Yu, J. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Charge-insensitive qubit design derived from the Cooper pair box, Phys. Rev. A **76**, 042319 (2007).
[683] J. Gambetta, A. Blais, D. I. Schuster, A. Wallraff, L. Frunzio, J. Majer, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Qubit-photon interactions in a cavity:

Measurement-induced dephasing and number splitting, Phys. Rev. A **74**, 042318 (2006).
[684] Y. Shi, P. Gokhale, P. Murali, J. M. Baker, C. Duckering, Y. Ding, N. C. Brown, C. Chamberland, A. Javadi-Abhari, A. W. Cross, *et al.*, Resource-efficient quantum computing by breaking abstractions, Proc. IEEE **108**, 1353 (2020).
[685] C. Arenz, D. I. Bondar, D. Burgarth, C. Cormick, and H. Rabitz, Amplification of quadratic Hamiltonians, Quantum **4**, 271 (2020).
[686] S. C. Burd, R. Srinivas, J. J. Bollinger, A. C. Wilson, D. J. Wineland, D. Leibfried, D. H. Slichter, and D. T. C. Allcock, Quantum amplification of mechanical oscillator motion, Science **364**, 1163 (2019).
[687] Y. Shapira, R. Shaniv, T. Manovitz, N. Akerman, L. Peleg, L. Gazit, R. Ozeri, and A. Stern, Theory of robust multiqubit nonadiabatic gates for trapped ions, Phys. Rev. A **101**, 032330 (2020).
[688] N. A. Gershenfeld and I. L. Chuang, Bulk spin-resonance quantum computation, Science **275**, 350 (1997).
[689] U. Warring, C. Ospelkaus, Y. Colombe, R. Jördens, D. Leibfried, and D. J. Wineland, Individual-ion addressing with microwave field gradients, Phys. Rev. Lett. **110**, 173002 (2013).
[690] S. J. Evered, D. Bluvstein, M. Kalinowski, S. Ebadi, T. Manovitz, H. Zhou, S. H. Li, A. A. Geim, T. T. Wang, N. Maskara, H. Levine, G. Semeghini, M. Greiner, V. Vuletić, and M. D. Lukin, High-fidelity parallel entangling gates on a neutral-atom quantum computer, Nature **622**, 268 (2023).
[691] M. O. Brown, S. R. Muleady, W. J. Dworschack, R. J. Lewis-Swan, A. M. Rey, O. Romero-Isart, and C. A. Regal, Time-of-flight quantum tomography of an atom in an optical tweezer, Nat. Phys. **19**, 569 (2023).
[692] Y. Wang, A. Kumar, T.-Y. Wu, and D. S. Weiss, Single-qubit gates based on targeted phase shifts in a 3D neutral atom array, Science **352**, 1562 (2016).
[693] L. F. Buchmann, K. Mølmer, and D. Petrosyan, Creation and transfer of nonclassical states of motion using Rydberg dressing of atoms in a lattice, Phys. Rev. A **95**, 013403 (2017).
[694] K. Eckert, J. Mompart, X. X. Yi, J. Schliemann, D. Bruß, G. Birkl, and M. Lewenstein, Quantum computing in optical microtraps based on the motional states of neutral atoms, Phys. Rev. A **66**, 042317 (2002).
[695] L. Lami and M. M. Wilde, Exact solution for the quantum and private capacities of bosonic dephasing channels, Nat. Photonics **17**, 525 (2023).
[696] M. Brownnutt, M. Kumph, P. Rabl, and R. Blatt, Ion-trap measurements of electric-field noise near surfaces, Rev. Mod. Phys. **87**, 1419 (2015).
[697] Q. A. Turchette, D. Kielpinski, B. E. King, D. Leibfried, D. M. Meekhof, C. J. Myatt, M. A. Rowe, C. A. Sackett, C. S. Wood, W. M. Itano, C. Monroe, and D. J. Wineland, Heating of trapped ions from the quantum ground state, Phys. Rev. A **61**, 063418 (2000).
[698] C. Noel, M. Berlin-Udi, C. Matthiesen, J. Yu, Y. Zhou, V. Lordi, and H. Häffner, Electric-field noise from thermally activated fluctuators in a surface ion trap, Phys. Rev. A **99**, 063427 (2019).
[699] R. Dubessy, T. Coudreau, and L. Guidoni, Electric field noise above surfaces: A model for heating-rate scaling law in ion traps, Phys. Rev. A **80**, 031402(R) (2009).
[700] G. H. Low, P. F. Herskind, and I. L. Chuang, Finite-geometry models of electric field noise from patch potentials in ion traps, Phys. Rev. A **84**, 053425 (2011).
[701] M. Kumph, C. Henkel, P. Rabl, M. Brownnutt, and R. Blatt, Electric-field noise above a thin dielectric layer on metal electrodes, New J. Phys. **18**, 023020 (2016).
[702] M. Teller, D. A. Fioretto, P. C. Holz, P. Schindler, V. Messerer, K. Schüppert, Y. Zou, R. Blatt, J. Chiaverini, J. Sage, and T. E. Northup, Heating of a trapped ion induced by dielectric materials, Phys. Rev. Lett. **126**, 230505 (2021).
[703] A. Safavi-Naini, P. Rabl, P. F. Weck, and H. R. Sadeghpour, Microscopic model of electric-field-noise heating in ion traps, Phys. Rev. A **84**, 023412 (2011).
[704] B. L. Foulon, K. G. Ray, C.-E. Kim, Y. Liu, B. M. Rubenstein, and V. Lordi, $1/\omega$ electric-field noise in surface ion traps from correlated adsorbate dynamics, Phys. Rev. A **105**, 013107 (2022).
[705] K. G. Ray, B. M. Rubenstein, W. Gu, and V. Lordi, van der Waals-corrected density functional study of electric field noise heating in ion traps caused by electrode surface adsorbates, New J. Phys. **21**, 053043 (2019).
[706] E. Kim, A. Safavi-Naini, D. A. Hite, K. S. McKay, D. P. Pappas, P. F. Weck, and H. R. Sadeghpour, Electric-field noise from carbon-adatom diffusion on a Au(110) surface: First-principles calculations and experiments, Phys. Rev. A **95**, 033407 (2017).
[707] C. D. Bruzewicz, J. M. Sage, and J. Chiaverini, Measurement of ion motional heating rates over a range of trap frequencies and temperatures, Phys. Rev. A **91**, 041402(R) (2015).
[708] A. Pagano, D. Jaschke, W. Weiss, and S. Montangero, Optimal control transport of neutral atoms in optical tweezers at finite temperature, Phys. Rev. Res. **6**, 033282 (2024).
[709] R. A. Bertlmann and P. Krammer, Bloch vectors for qudits, J. Phys. A: Math. Theor. **41**, 235303 (2008).
[710] F. Albertini and D. D'Alessandro, Notions of controllability for bilinear multilevel quantum systems, IEEE Trans. Autom. Control **48**, 1399 (2003).
[711] D. Dong and I. Petersen, Quantum control theory and applications: A survey, IET Control Theory Appl. **4**, 2651 (2010).
[712] M. Suzuki, Generalized Trotter's formula and systematic approximants of exponential operators and inner derivations with applications to many-body problems, Commun. Math. Phys. **51**, 183 (1976).
[713] D. Edward Bruschi, A. Xuereb, and R. Zeier, Deciding finiteness of bosonic dynamics with tunable interactions, J. Phys. A: Math. Theor. **58**, 025204 (2024).
[714] F. A. Mele, A. A. Mele, L. Bittel, J. Eisert, V. Giovannetti, L. Lami, L. Leone, and S. F. E. Oliviero, Learning quantum states of continuous variable systems, ArXiv:2405.01431.
[715] C. H. Valahu, T. Navickas, M. J. Biercuk, and T. R. Tan, Benchmarking bosonic modes for quantum information

with randomized displacements, PRX Quantum **5**, 040337 (2024).

[716] A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, J. Majer, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Approaching unit visibility for control of a superconducting qubit with dispersive readout, Phys. Rev. Lett. **95**, 060501 (2005).

[717] T. Walter, P. Kurpiers, S. Gasparinetti, P. Magnard, A. Potočnik, Y. Salathé, M. Pechal, M. Mondal, M. Oppliger, C. Eichler, and A. Wallraff, Rapid high-fidelity single-shot dispersive readout of superconducting qubits, Phys. Rev. Appl. **7**, 054020 (2017).

[718] P. D. Kurilovich, T. Connolly, C. G. L. Bøttcher, D. K. Weiss, S. Hazra, V. R. Joshi, A. Z. Ding, H. Nho, S. Diamond, V. D. Kurilovich, W. Dai, V. Fatemi, L. Frunzio, L. I. Glazman, and M. H. Devoret, High-frequency readout free from transmon multi-excitation resonances, ArXiv:2501.09161.

[719] K. Vogel and H. Risken, Determination of quasiprobability distributions in terms of probability distributions for the rotated quadrature phase, Phys. Rev. A **40**, 2847 (1989).

[720] M. Ahmad, S. Qamar, and M. S. Zubairy, Quantum-state tomography using phase-sensitive amplification, Phys. Rev. A **62**, 043814 (2000).

[721] M. Ahmad, S. Qamar, and M. S. Zubairy, Reconstruction of a multimode entangled state using a two-photon phase-sensitive linear amplifier, Phys. Rev. A **67**, 043815 (2003).

[722] A. Ourjoumtsev, R. Tualle-Brouri, and P. Grangier, Quantum homodyne tomography of a two-photon Fock state, Phys. Rev. Lett. **96**, 213601 (2006).

[723] K. He, M. Yuan, Y. Wong, S. Chakram, A. Seif, L. Jiang, and D. I. Schuster, Efficient multimode Wigner tomography, Nat. Commun. **15**, 4138 (2024).

[724] C. Shen, R. W. Heeres, P. Reinhold, L. Jiang, Y.-K. Liu, R. J. Schoelkopf, and L. Jiang, Optimized tomography of continuous variable systems using excitation counting, Phys. Rev. A **94**, 052327 (2016).

[725] T. Krisnanda, C. Y. Fontaine, A. Copetudo, P. Song, K. X. Lee, N.-N. Huang, F. Valadares, T. C. H. Liew, and Y. Y. Gao, Demonstrating efficient and robust bosonic state reconstruction via optimized excitation counting, PRX Quantum **6**, 010303 (2025).

[726] S. T. Flammia, D. Gross, Y.-K. Liu, and J. Eisert, Quantum tomography via compressed sensing: Error bounds, sample complexity and efficient estimators, New J. Phys. **14**, 095022 (2012).

[727] K. R. Brown, A. W. Harrow, and I. L. Chuang, Arbitrarily accurate composite pulse sequences, Phys. Rev. A **70**, 052318 (2004).

[728] K. R. Brown, A. W. Harrow, and I. L. Chuang, Erratum: Arbitrarily accurate composite pulse sequences [Phys. Rev. A 70, 052318 (2004)], Phys. Rev. A **72**, 039905(E) (2005).

[729] W. Gong and S. Choi, Robust and parallel control of many qubits, ArXiv:2312.08426.

[730] L. Ying, Stable factorization for phase factors of quantum signal processing, Quantum **6**, 842 (2022).

[731] Y. Dong, X. Meng, K. B. Whaley, and L. Lin, Efficient phase-factor evaluation in quantum signal processing, Phys. Rev. A **103**, 042419 (2021).

[732] J. Wang, Y. Dong, and L. Lin, On the energy landscape of symmetric quantum signal processing, Quantum **6**, 850 (2022).

[733] R. Blume-Kohout, C. M. Caves, and I. H. Deutsch, Climbing mount scalable: Physical resource requirements for a scalable quantum computer, Found. Phys. **32**, 1641 (2002).