

Designing Calibration and Expressivity-Efficient Instruction Sets for Quantum Computing

Lingling Lao*
University College London

Prakash Murali*
Princeton University

Margaret Martonosi
Princeton University

Dan Browne
University College London

Abstract—Near-term quantum computing (QC) systems have limited qubit counts, high gate (instruction) error rates, and typically support a minimal instruction set having one type of two-qubit gate (2Q). To reduce program instruction counts and improve application expressivity, vendors have proposed, and shown proof-of-concept demonstrations of richer instruction sets such as XY gates (Rigetti) and fSim gates (Google). These instruction sets comprise of families of 2Q gate types parameterized by continuous qubit rotation angles. That is, it allows a large set of different physical operations to be realized on the qubits, based on the input angles. However, having such a large number of gate types is problematic because each gate type has to be calibrated periodically, across the full system, to obtain high fidelity implementations. This results in substantial recurring calibration overheads even on current systems which use only a few gate types. Our work aims to navigate this tradeoff between application expressivity and calibration overhead, and identify what instructions vendors should implement to get the best expressivity with acceptable calibration time.

Studying this tradeoff is challenging because of the diversity in QC application requirements, the need to optimize applications for widely different hardware gate types and noise variations across gate types. Therefore, our work develops NuOp, a flexible compilation pass based on numerical optimization, to efficiently decompose application operations into arbitrary hardware gate types. Using NuOp and four important quantum applications, we study the instruction set proposals of Rigetti and Google, with realistic noise simulations and a calibration model. Our experiments show that implementing 4-8 types of 2Q gates is sufficient to attain nearly the same expressivity as a full continuous gate family, while reducing the calibration overhead by two orders of magnitude. With several vendors proposing rich gate families as means to higher fidelity, our work has potential to provide valuable instruction set design guidance for near-term QC systems.

Index Terms—instruction set architecture, compilation, quantum computing

I. INTRODUCTION

Quantum computing (QC) is an emerging paradigm that uses quantum mechanical principles to manipulate information. QC systems store information using *qubits* (quantum bits) and manipulate information using *gates* (operations). These operations allow the system to exploit quantum effects such as superposition, entanglement and interference to explore large state spaces efficiently, sometimes faster than classical (non-quantum) systems. In the near-future, QC is poised to drive research in domains such as machine learning [5], quantum

chemistry [6] and material science [7]. Several prototype QC systems have been built using hardware technologies such as superconducting and trapped-ion qubits. These systems are referred to as Noisy Intermediate-Scale Quantum (NISQ) systems [8] and have small qubit counts, restricted connectivity and high gate error rates. Although resource-constrained, NISQ systems are useful for demonstrating near-term applications [9] and milestones such as quantum supremacy [2].

Typically, NISQ systems support a *universal* gate set composed of continuous single-qubit (1Q) rotations and a few two-qubit (2Q) gate types that can express any application operation. Among them, 2Q gates generate entanglement between pairs of qubits and constitute a key building block for quantum algorithms. However, 2Q gates are more challenging to implement than 1Q gates and have higher error rates. For example, on IBM systems, 2Q gate error rates are 1-5% and 1Q gate error rates are less than 0.1%. Since a single entangling 2Q gate is sufficient for universality, most vendors use instruction sets having only one 2Q gate type. For example, in Rigetti's and Google's early QC systems, Controlled Z (CZ) gate was the only supported 2Q gate type [10]–[12].

The type and error rate of the hardware 2Q gate determines *application expressivity*, that is, the number of instructions required to implement an application and the overall execution fidelity. When a system has a single 2Q gate type, the compiler is forced to express all application operations using it, worsening executable instruction counts and duration. Importantly, the error rate of the available 2Q gate constrains the number of program instructions that can be reliably executed. Although vendors have been improving gate implementations, progress has been slow because of challenges in accurate qubit control, qubit defects from lithographic manufacturing and emergent sources of noise. Therefore, if an application cannot be expressed concisely, it is unlikely to execute successfully on near-term hardware.

To improve expressivity, and thereby improve execution success rate, industry and academic vendors [3], [4], [13] have proposed instruction sets which have continuous 2Q gate families. Table I shows the proposed XY gate family from Rigetti and the fSim gate family from Google. Each gate family is parameterized by rotation angles (e.g., θ and ϕ for Google's fSim gates) and forms a continuous set of two-qubit gate types (different unitaries). This means, in theory, there are an infinite number of hardware 2Q gate types that can be used to express applications succinctly. In practice, however,

*Lingling Lao and Prakash Murali are joint first authors. Corresponding author: pmurali@cs.princeton.edu

Rigetti		Google	
Current	Anticipated	Current	Anticipated
$CZ = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}$ $XY(\pi)$	$XY(\theta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta/2) & i\sin(\theta/2) & 0 \\ 0 & i\sin(\theta/2) & \cos(\theta/2) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$ CZ	$SYC = fSim(\pi/2, \pi/6)$ $\sqrt{i}SWAP = fSim(\pi/4, 0)$	$fSim(\theta, \phi) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta) & -i\sin(\theta) & 0 \\ 0 & -i\sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 0 & e^{-i\phi} \end{pmatrix}$
Fidelity $\sim 95\%$	95-99%	$\sim 99.6\%$	$\sim 99.6\%$

TABLE I: Current and anticipated two-qubit gate types in Rigetti and Google systems. Each two-qubit gate type corresponds to a 4x4 unitary matrix. Current gate sets are based on Rigetti’s Aspen-8 [1] and Google’s Sycamore [2] devices. Anticipated gate sets are based on [3] and [4]. Controlled Z (CZ) is an example of a fixed gate type. In contrast, $fSim(\theta, \phi)$ is parameterized by θ, ϕ . Varying these parameters allows a broad range of gate types to be realized in hardware. $XY(\theta)$ is a subset of $fSim$ gate family up to single-qubit rotations. (The actual instruction sets also include single-qubit operations for universality, they are not shown here.)

quantum gates are implemented using analog pulse sequences which must be calibrated to obtain high fidelity. Calibration routines require a large number of test executions to tune control parameters and characterize the hardware operations. Worse, calibration is not a one-time step (parameters drift over time) and must be performed periodically to maintain high fidelity. Even with a single type of 2Q gate, Google’s 54-qubit device requires up to four hours of calibration per day [2]. Therefore, across the full parameter space of a continuous gate set, calibration overheads are likely enormous, especially as devices scale up.

Our work studies the tradeoff between application expressivity and calibration overhead to determine what gate sets and gate parameter combinations vendors should implement to offer high-fidelity executions without tremendous calibration time. To accurately compare instruction set options, we require a toolflow to automatically decompose and optimize applications for different candidate gate sets. This is challenging because of the variety of gate options, variations in fidelity across gate parameter values, and the need to use different optimization strategies based on the hardware characteristics.

Figure 1 shows our simulation framework. Our work addresses the toolflow challenge by building NuOp, a flexible compilation pass based on numerical optimization techniques. NuOp accepts as input a set of target hardware gate types and uses noise-adaptive compilation across gate parameter values and approximate gate decomposition to heavily optimize applications, enabling accurate comparisons across gate choices. Using NuOp, we perform a simulation-driven study to identify the best instruction design choices across two vendor gate families and several NISQ applications.

Our contributions include:

- Current QC instruction sets belong to one of two camps: minimal sets having a single two-qubit gate type [14] (poor expressivity, easy to calibrate) or families of continuous gate sets [3], [4], [13] (rich expressivity, hard to calibrate). These two choices and their calibration tradeoffs are reminiscent of the tradeoffs in classical RISC vs. CISC ISA choices. Our work is the first to study the expressivity vs. calibration tradeoff in QC, and

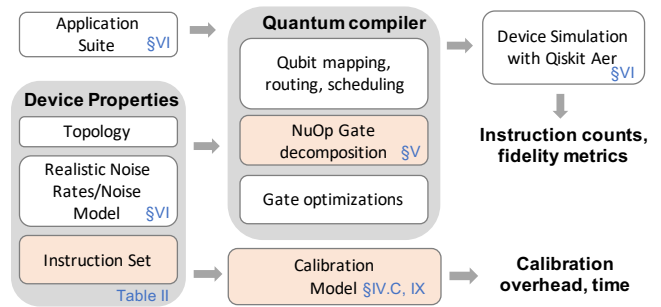


Fig. 1: Simulation framework. To study different instruction set choices, our work performs an application-driven simulation study with realistic device and calibration models.

design an instruction set that has rich expressivity and is easy to calibrate.

- To design this instruction set, we use NuOp to characterize instruction counts for four representative QC applications across the continuous parameter space for XY and $fSim$ gate families. Across the parameter space, one to six hardware 2Q gates are required to implement an application 2Q operation. Based on this characterization, we identify a small subset (4-8 discrete combinations) of these gate families that offers near-optimal expressivity i.e., instruction counts and fidelity that are better than instruction sets with a single gate type and close to the optimal values that can be obtained using the entire continuous parameter space of XY and $fSim$ families.
- Using a realistic calibration model for $fSim$ gates, we show that the proposed gate sets having 4-8 gate types, save two orders of magnitude calibration overhead, compared to using the entire parameter space. Since calibration is not a one-time overhead, this large reduction makes our instruction set practically viable for real-system implementations.
- Finally, QC systems are known to have highly variable error rates across qubits and time [15]–[18]. We demonstrate that exposing multiple 2Q gate types is a viable method to mitigate such noise variations. When

a noise-adaptive compiler pass such as NuOp is used in conjunction with multiple gate types, the compiler can choose different gate types on different qubits to obtain significant improvements in fidelity.

Our work makes the case for an instruction set having a small number of expressive two-qubit gate types. With major QC vendors aiming to offer more instruction types in the near-term [4], [13], we expect our work to significantly impact industry and academic instruction sets and provide useful guidance to hardware designers.

II. BACKGROUND ON QUANTUM GATES

Gate family vs. type: Quantum algorithms encode information using a set of qubits and manipulate information by applying different quantum gates (instructions). Each n -qubit quantum gate can be uniquely defined by a unitary matrix $U_{2^n \times 2^n}$. The unitary matrix of a gate may be a *fixed unitary* or a *parameterized unitary*. An example of a fixed unitary gate is the CZ gate shown in Table I. In contrast, $\text{fSim}(\theta, \phi)$ is parameterized by θ and ϕ . For such gates, we use the term *gate family* since it allows an infinite number of unitaries to be realized based on the parameter values. We use the term *gate type* to refer to fixed parameter values in this family e.g., $\text{fSim}(\pi, \pi/2)$ and $\text{fSim}(\pi/6, \pi/8)$ and $\text{CZ} = \text{fSim}(0, \pi)$ are three distinct gate types.

Gate Implementation in NISQ systems: Gate implementation depends on the qubit hardware technology and gate type. For example, in Google’s Sycamore system (transmon qubits), single-qubit gates are implemented using microwave pulses that are resonant with the qubit frequency; two-qubit gates (SYC and $\sqrt{\text{iSWAP}}$) are implemented by bringing neighboring qubits in resonance through a tunable coupler that mediates the interaction [2]. Gate implementations in all NISQ systems are error-prone. The quality of a gate implementation is characterized by its fidelity (1–error rate). Table I shows average gate fidelities for Rigetti and Google systems.

III. RELATED WORK

Most prior works [19]–[24] consider instruction set design issues assuming that systems have only one two-qubit gate type. For example, Rigetti Quil [19] and IBM OpenQASM [20] target only the CZ and CNOT gates, respectively. Murali et al. [16] compare native and software-visible instruction sets and recommend exposing all available gate types to the compiler to improve application fidelity. In contrast to these minimal instruction sets, we focus on instruction sets with multiple gate types to improve application expressivity.

Google’s Sycamore device has preliminary support for a continuous fSim gate [4]. Rigetti’s XY gates were demonstrated in [3]. Lacroix et al. [13] experimentally demonstrated that using continuous Controlled-Phase gates ($\text{CZ}(\theta)$) is beneficial for quantum optimization (QAOA). These works represent one extreme of the design space, i.e., the full parameter space is available to programs, which optimizes for application expressivity but requires demanding calibration schemes. Our work uses realistic gate errors and calibration models from

Google and Rigetti devices to study what subset of the gate parameter space should be exposed to obtain high application reliability with low calibration overhead.

Our work requires an automated method to decompose and optimize applications for arbitrary vendor gate sets. Several works have investigated compilation and program optimization techniques for QC devices. Early works focused on minimizing metrics such as gate count and circuit depth [25]–[29]. Recently noise-adaptive compilation passes [15], [17], [18], [30] have been developed to adapt program executions to spatio-temporal fidelity variations in real systems. However, these works and compilers such as IBM Qiskit [31] and TriQ [16] assume that the hardware instruction set has a single type of two-qubit gate. Hence, they are not directly applicable for studying instruction set design issues for multiple two-qubit gate types.

IV. INSTRUCTION SET DESIGN CHOICES AND TRADEOFFS

A. Background on ISAs with Continuous Two-Qubit Gates

Rigetti has proposed using $\text{XY}(\theta)$ gates along with the Controlled-Z gate (CZ) as the instruction set in their future systems. An early demonstration of this gate set on a two-qubit system in [3] shows executions of broad range of gate types with fidelities ranging from 95% to 99%. Rigetti’s current system (Aspen-8), a 30-qubit device, includes one gate type from this family, the $\text{XY}(\pi)$ gate, in addition to the CZ gate.

Similarly, Google has proposed a continuous $\text{fSim}(\theta, \phi)$ gate family for their devices. This gate family can be viewed as a composite of two families i.e., the CPHASE family, $\text{CZ}(\phi)$, obtained by varying ϕ , fixing $\theta = 0$, and the iSWAP -like family, $\text{iSWAP}(\theta)$, obtained by varying θ , fixing $\phi = 0$. In [4], Google demonstrated 525 instances of the fSim family with different θ and ϕ , on two qubits. Recent experiments from their 54-qubit Sycamore processor use two gate types from this family: SYC gate = $\text{fSim}(\pi/2, \pi/6)$ [2] and $\sqrt{\text{iSWAP}}$ gate = $\text{fSim}(\pi/4, 0)$ [34], with fidelities as high as 99.4%

In both Rigetti and Google instruction sets, multiple types of single-qubit gates (qubit rotations) are included for obtaining a universal gate set. However, these gates typically have one order of magnitude higher fidelity than two-qubit gates and have been optimized heavily using both hardware [35] and software techniques [31]. Hence, our focus is on designing efficient two-qubit instruction sets.

B. Application-level Benefits of Multiple Gate Types

Different gate types are advantageous to different QC applications. For example, the CZ gates are useful in quantum error correction. Excitation-preserving operations like the ones in XY and fSim families allow natural representations for quantum chemistry [6], [36] and combinatorial optimization [37]. Therefore, by offering more than one gate type, vendors aim to cater to the different application operation needs, and reduce the number of instructions required to express them.

Figure 2 shows decompositions of two application unitaries using two types of hardware gates. When only one of these gate types is available, the compiler has to decompose all

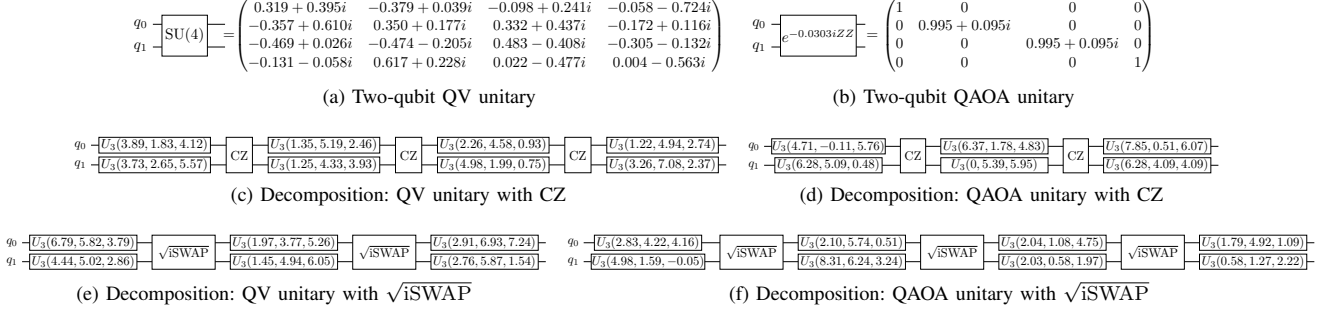


Fig. 2: Decomposition examples using NuOp. (a) shows a two-qubit gate with unitary from the $SU(4)$ group (Quantum volume circuits randomly sample gates from this group [32]). (b) shows a two-qubit Pauli interaction ($e^{-i\beta(Z\otimes Z)}$), which is used in QAOA circuits. (c-f) show exact implementation (decomposition error $\approx 10^{-8}$) of these unitaries using different two-qubit hardware gates from Rigetti (CZ) and Google (\sqrt{i} SWAP). The CZ gate is more expressive for QAOA unitaries and the \sqrt{i} SWAP gate is more expressive for QV unitaries. In (c-f), U3 gates are arbitrary single-qubit rotation operations¹.




Fig. 3: Layout of first 8 qubits in Rigetti Aspen-8 system. Nodes are qubits and edges are two-qubit gates labeled by the measured fidelity for XY(π) and CZ gates [33]. The best gate varies across qubit-pairs.

application unitaries using it, potentially increasing the number of instructions. Since NISQ systems have high two-qubit gate error rates, an increase in instruction count translates to a significant reduction in fidelity. Continuous gate families proposed by Rigetti and Google leverage this concept to an extreme. That is, they aim to offer a very large number of two-qubit gate types to applications using one or more controllable gate parameters. For Rigetti’s XY family, Peterson et al. [38] theoretically proves that offering applications access to the full continuous parameter space, can give $\sim 30\%$ two-qubit instruction count reduction for generic quantum circuits. Similarly, Foxen et al. [4] demonstrates that a continuous set of fSim gates can provide a 3X gate reduction compared to decomposition with one gate type. .

C. Experimental Difficulties of Multiple Gate Types

Calibration overhead: Implementing multiple two-qubit basis gates in the same device with high fidelity is experimentally

¹An arbitrary single-qubit gate can be described by a unitary matrix with three parameters

$$U3(\alpha, \beta, \lambda) = \begin{pmatrix} \cos(\frac{\alpha}{2}) & -e^{i\lambda} \sin(\frac{\alpha}{2}) \\ e^{i\beta} \sin(\frac{\alpha}{2}) & e^{i(\beta+\lambda)} \cos(\frac{\alpha}{2}) \end{pmatrix}.$$

difficult and requires demanding calibration strategies. First, each two-qubit gate on each qubit pair needs to be calibrated independently to find its optimal control parameters, and electronics need to be carefully calibrated to avoid any bleeding between control pulses of different gates. These are known challenges for continuous gate sets [4]. Second, on real-systems, calibration of different combinations of single-qubit and multi-qubit gates, and parallel operations are also required to obtain high fidelity [39]. Calibration complexity increases at least linearly with the number of two-qubit basis gates and becomes enormous for a continuous gate set. Since calibration parameters drift over time (causing gate error rate fluctuations of up to 10X [4]), frequent re-calibration is also required in practice. In Section IX, we model fSim gate calibration.

Fidelity variations across gate types: Continuous gate families also have variable fidelities across the parameter space, qubits in the system and time (similar to variability seen in systems with one type of two-qubit gate [16], [17]). Figure 3 shows that the best gate type varies across qubits in Rigetti’s Aspen-8 system. For XY gates, gate error rates vary up to 4X depending on the angle θ (see Figure 6 in [3]). For fSim gates, the error rates vary up to 5X based on θ and ϕ (see Figure 4 in [4]). If applications use a single gate type across the system, it may result in severe fidelity losses. Instead, by using multiple gate types with noise-aware compilation, applications can benefit from the best gates on each part of the device.

D. Our Work

Our work studies key instruction design questions considering application-level benefits, calibration overheads and noise variability. *Is it possible to design an instruction set having a small set of discrete gate types, but achieving the expressivity and fidelity of a continuous gate family? What are the calibration overheads of such an instruction set compared to a continuous gate family? Is it beneficial to expose multiple gate types and use noise adaptivity across gate types to mitigate noise variations?*

Table II shows the gate set options considered: instruction sets with a single type of two-qubit gate (S1-S7), our proposed instruction sets for Rigetti (R1-R5) and Google (G1-G7) and fully continuous Full_{XY} and Full_{fSim} gate families. To study these options, we use the framework shown in Figure 1.

V. TOOLFLOW FOR ARCHITECTURAL EXPLORATION

QC applications are typically written in terms of high-level device and gate set-independent operations. During compilation, program qubits are mapped onto the target device, communication operations such as SWAPs are inserted to comply with limited connectivity, and operations are translated or decomposed into the target instruction set. Figure 2 shows gate decomposition examples. Gate decomposition methods aim to represent the unitary of a desired application operation within a tolerable inaccuracy, using a circuit with short depth (to minimize decoherence), and small number of two-qubit gates (to minimize gate errors).

To probe instruction set design issues across different gate types and their combinations we require a flexible decomposition method that can support any application operation and hardware gate type. Current industry compilers focus on gate decomposition routines for specific gate types of interest to a particular hardware platform. These compilers typically use linear algebra-based methods such as KAK decomposition [40]–[45] or hard coded gate identities [46]. These methods are difficult to extend to continuous gate sets and arbitrary parameter combinations within these gate families. For example, IBM’s Qiskit compiler includes routines for decomposing arbitrary two-qubit unitaries to CNOT gates. Rigetti’s Quilc [47] includes support for CZ and XY(π) gates, but does not target arbitrary gate types and hence, was not suitable for our study. Google’s Cirq can support a variety of application unitaries and hardware gate types. However, Cirq only provides optimized decompositions for limited gate types such as SYC and CZ gates.

A. Numerical Optimization for Gate Decomposition

Our work develops NuOp, a numerical optimization technique to efficiently decompose any application unitary into any hardware gate. NuOp is based on recent prior work [48]–[50]. We describe the method for a two-qubit application unitary and then discuss how to apply it for a program circuit. NuOp takes an application unitary and a target hardware gate type as input and produces an optimized decomposition as output. At a high-level, NuOp constructs a series of parameterized template circuits and optimizes each template to achieve the desired unitary. Among the optimized templates, it selects the one which satisfies an input error threshold as the final decomposition.

Figure 4 shows an example template with i layers. Each template circuit has alternating layers of arbitrary single qubit gates and the target hardware two-qubit gate type (similar templates are used in the ADAPT-VQE algorithm [51]). For example, to decompose an application operation into a fSim($\pi/6, \pi/2$) hardware gate, the template has layers of

fSim($\pi/6, \pi/2$) gates sandwiched between arbitrary single-qubit operations. For a given template, the optimization variables are the rotation angles for the single qubit operations. By optimizing these angles and computing the accuracy w.r.t the desired unitary, NuOp obtains a highly accurate representation of the application operation. To compute accuracy, NuOp uses *decomposition fidelity* F_d , i.e., the accuracy of the unitary decomposition (without device/hardware noise), measured using the Hilbert-Schmidt inner product,

$$F_d = \frac{\langle U_d, U_t \rangle_{HS}}{\dim(U_t)} = \frac{\text{Tr}(U_d^\dagger U_t)}{\dim(U_t)}. \quad (1)$$

Here, U_t is the matrix of the application unitary and U_d is the matrix represented by the template. Each choice of single-qubit angles in the template gives rise to a different U_d . When U_d is far away from the desired unitary, F_d is close to 0 and when it matches U_t , F_d is close to 1, so NuOp maximizes F_d . To perform the optimization, NuOp uses a BFGS [52], a well-known numerical optimization method.

NuOp starts with template circuits having one layer and grows the number of layers one at a time. For each template size i , it uses the optimization to obtain the best F_d possible with i layers. Then NuOp takes a fidelity threshold as input (e.g., a decomposition fidelity of 99.999%) and selects the smallest number of layers that meets the threshold.

NuOp can be flexibly used for any input gate type, as well as any application unitary by changing the input template. To evaluate Full_{XY} and Full_{fSim} cases, i.e., the entire gate family is available, we allow NuOp to use templates where the two-qubit gate angles are also optimization variables. Our experiments in Section VII confirm that the decompositions produced by NuOp are comparable or better than optimized KAK-based decompositions in Google Cirq.

B. Approximation and Noise-Awareness Optimizations

Current NISQ systems have significantly higher hardware error rates (0.5-10%) than the decomposition error thresholds commonly used for gate decomposition (e.g., 10^{-8} for KAK-based approaches). Therefore, recent works have observed that a less accurate decomposition that requires fewer operations can be more beneficial than an exact decomposition [32]. We incorporate this optimization into NuOp by considering two kinds of errors: decomposition errors and hardware errors. Decomposition error is a result of inexact representation of the target unitary. Hardware error is a result of imperfect gate implementation on a real device. For a target gate type g_j and number of layers i , the decomposition errors are quantified using the decomposition fidelity, $F_d^{(g_j, i)}$ and hardware errors using hardware fidelity, $F_h^{(g_j, i)}$. Overall, the quality of a decomposition is the product of these two terms. By maximizing the product across gate types and the number of layers we can obtain a decomposition that has higher reliability overall, than maximizing F_d alone:

$$F_u^{\text{APD}} = \max_{g_j, i} F_d^{(g_j, i)} F_h^{(g_j, i)}. \quad (2)$$

Type	Instruction Sets (Up to single-qubit rotations, $XY(\theta) = \text{iSWAP}(\theta/2) = \text{fSim}(\theta/2, 0)$, $\text{CZ}(\phi) = \text{fSim}(0, \phi)$.)
Single two-qubit gate type	S1 = SYC = $\text{fSim}(\pi/2, \pi/6)$, S2 = $\sqrt{\text{iSWAP}} = \text{fSim}(\pi/4, 0)$, S3 = CZ = $\text{fSim}(0, \pi)$, S4 = $\text{iSWAP} = \text{fSim}(\pi/2, 0)$, S5 = $\text{fSim}(\pi/3, 0)$, S6 = $\text{fSim}(3\pi/8, 0)$, S7 = $\text{fSim}(\pi/6, \pi)$
Multiple gate types (Google)	G1 = {S1,S2}, G2 = {S1,S2,S3}, G3 = {S1,S2,S3,S4}, G4 = {S1,S2,S3,S4,S5}, G5 = {S1,S2,S3,S4,S5,S6}, G6 = {S1,S2,S3,S4,S5,S6,S7}, G7 = {S1,S2,S3,S4,S5,S6,S7,SWAP},
Multiple gate types (Rigetti)	R1 = {S3,S4}, R2 = {S2,S3,S4}, R3 = {S2,S3,S4,S5}, R4 = {S2,S3,S4,S5,S6}, R5 = {S2,S3,S4,S5,S6,SWAP}
Full continuous gate set	Full _{XY} = {XY(θ): $\theta \in [0, \pi]$ }, Full _{fSim} = {fSim(θ, ϕ): $\theta, \phi \in [0, \pi]$ }

TABLE II: Instruction sets studied in this work. Each set also includes arbitrary single qubit operations (not shown) to make it universal. S1-S7 were selected based on our application characterization experiments in Section VIII. See Figure 8 for a graphical view of S1-S7 on the fSim parameter space. For Google, G1-G7 are combinations of S1-S7, e.g., G2 has three two-qubit gate types. For Rigetti, R1-R5 are selected based on subsets which can be supported using the XY gate family. (For ease of presentation, the table expresses Rigetti gates with fSim notation; they can be translated to XY(θ) notation using the identities on the top of the table.)




Fig. 4: The template circuit used in the NuOp decomposition. It consists of i layers, alternating parameterized single-qubit gates (three parameters α, β, λ for each) and fSim gates. To compile for a fixed fSim gate, θ and ϕ are set to the desired values in the template. To compile for Full_{fSim}, θ and ϕ are also considered as optimization variables.




Fig. 5: Illustration of noise-adaptive approximate decomposition. (a) 3-qubit circuit with two gates, both having the matrix in Figure 2a. (b) decomposition circuit targeting qubits [2,3,4] on Rigetti Aspen-8 in Figure 3 (We assume the $\sqrt{\text{iSWAP}} = \text{XY}(\pi/2)$ gate has the same fidelity as the $\text{XY}(\pi)$ gate). The noise-aware gate compilation pass chooses the decomposition with CZ gates for qubit pair (2,3) because its high hardware fidelity leads to higher overall fidelity $F_u = 88.36\%$ (which is much higher than $F_h = 0.7$ of $\sqrt{\text{iSWAP}}$). Similarly, the decomposition with $\sqrt{\text{iSWAP}}$ is selected for qubit pair (3,4). Furthermore, the approximate decomposition maximizes the overall implementation fidelity by decomposing the unitary into only two CZ gates, compared to 3 CZ gates in the exact decomposition in Figure 2c.

$F_h^{(g_j, i)}$ is estimated by the product of the hardware fidelity (1-error rate) of every gate in the decomposition (this model has been shown to work well in real systems [2], [53]–[55]). Figure 5 shows an example decomposition with a CZ gate that has hardware fidelity 94%. When the above approximate optimization is used, only 2 applications of CZ are required to obtain a decomposition with overall fidelity of 88.36% ($F_d = 98.63\%$). If NuOp had used one more CZ gate to achieves an exact decomposition (decomposition infidelity 10^{-7}), the overall fidelity will be reduced to 83.06% because of its high hardware error rate from three CZ applications (0.94^3).

We can further improve decompositions by exploiting fidelity variations across gate types. In Figure 3, on qubit pair (2,3), the CZ gate has the highest fidelity, while on qubits (3,4), the $\text{XY}(\pi)$ gate has the highest fidelity. When multiple gates types (g_j) are available in a given device, we can choose the best gate type for each application operation, depending on the decomposition fidelity using that gate type and its hardware

fidelity. Figure 5 shows a noise-adaptive decomposition of an application circuit using two gate types. We can easily implement this idea within NuOp by computing $F_h^{(g_j, i)}$ using the hardware fidelity of each gate type, obtained from device calibration data. To our knowledge, our work is the first to adapt executables to noise variations across gate types.

To decompose a full application for a given instruction set having multiple two-qubit gate types, NuOp finds the best decomposition for each application unitary and each hardware gate type, using approximation if necessary. Then, for each unitary, NuOp selects the gate type which offers the highest overall fidelity F_u . NuOp creates an output circuit where application unitaries are replaced by their best decompositions.

VI. EXPERIMENTAL SETUP

Benchmarks: We use four applications for evaluation: Quantum Volume (QV) [32], Quantum Approximate Optimization Algorithm (QAOA) [37], Fermi-Hubbard model (FH), and Quantum Fourier Transform (QFT). They cover the main types

of circuits studied for QC systems i.e, random circuits (QV), quantum optimization circuits (QAOA), quantum simulation circuits (FH), and classic QC or longer-term circuits (QFT).

QV is a full-stack benchmark proposed by IBM to compare different QC devices and compiler stacks. Each random QV circuit with n qubits has n parallel layers of random two-qubit gates acting on a random pair of qubits.

QAOA is a promising NISQ application for solving combinatorial optimization problems. QAOA circuits have also been used to study real-system performance [34]. Our experiments use QAOA circuits with one layer of the MaxCut ansatz [56]. Each n -qubit QAOA circuit has $\sim n^3/4$ random two-qubit ZZ interactions, interleaved with single-qubit X rotations.

FH is a model of simulating interacting particles in a lattice, which becomes difficult to solve on classical computers for large problems. Our evaluation uses the 1D FH model circuits with one Trotter step [57]. Each n -qubit FH circuit consists of $2n$ ZZ interactions and $\sim 4n \frac{1}{2}(XX+YY)$ interactions.

QFT is a key subroutine in conventional quantum algorithms such as quantum phase estimation and Shor’s factoring [58]. An n -qubit QFT circuit consists of n Hadamard gates and $n(n-1)/2$ CZ($\pi/2^t$) gates ($t \in [1, \dots, n-1]$) [59].

QV and QAOA circuits are generated from Qiskit [31] and ReCirq [60], respectively. *Our toolflow is scalable for large circuit sizes, but current highly noisy systems can only achieve high-fidelity computation for small-scale circuits.* Therefore, our evaluation experiments uses QV, QAOA, and QFT benchmarks with 3-6 qubits and FH circuits with 10 and 20 qubits. For each n -qubit QV benchmark, we use 100 random circuits with different circuit structures and unitaries to test instructions sets against a wide range of application requirements. Similarly, for each n -qubit QAOA benchmark we use 100 random circuits with different unitaries.

Decomposition Algorithms: We implemented NuOp as a IBM Qiskit compiler pass using version 0.20. Numerical optimization was performed using the BFGS implementation [52] in scipy version 1.4.1. In NuOp we used templates with a maximum of ten two-qubit gate layers, but for most cases less than four layers were sufficient for obtaining high quality decompositions. To validate NuOp we compare it to the decomposition passes in Google Cirq v0.8.2.

QC Systems and Simulations: We use realistic noise simulations of two systems: Rigetti Aspen-8 and Google Sycamore. Aspen-8 has 30 qubits and has 4 connected rings with 8 qubits each (two qubits are not functional). Figure 3 shows the first ring of qubits. Sycamore is a 54-qubit device having grid connectivity. For Aspen-8 we obtain real calibration data from [33], including error rates for each gate type (XY(π), CZ, single-qubit), coherence times and readout errors. For experiments with arbitrary XY(θ) gates, we model the gate errors rates as a uniform distribution in the range 95-99% (based on experiments in [3]). For Sycamore we obtain qubit coherence times and readout errors from [2] and acquire error rates of simultaneous SYC gates and single-qubit gates via Google’s quantum computing service. We model the error rates of other types of two-qubit gates as a normal distribution with

$\mu = 0.62\%$, $\sigma = 0.24\%$ (based on the error rate distribution of SYC gates). For both systems, we use the Qiskit Aer simulator [31] to perform a noise simulation using the calibration data. Specifically, it applies single-qubit and two-qubit depolarizing noises based on single-qubit and two-qubit gate error rates. It implements amplitude damping and dephasing noise based on T1 and T2 times as well as gate duration.

Metrics: For QV, the standard metric is the average **probability of heavy output generation (HOP)** obtained using 100 random circuits [32], [61]. With n qubits, HOP greater than $2/3$ denotes that the set of qubits has quantum volume 2^n . For QAOA benchmarks, we measure the average **cross-entropy difference (XED)** [12] metric which tests the quality of the generated output with respect to the ideal distribution (noiseless simulation) and the uniform distribution. For FH benchmarks, we use the linear cross-entropy benchmarking **fidelity** [62] metric to compare the measured probabilities with the ideal probabilities. For QFT benchmarks, we measure the **success rate**, i.e., the probability of an execution getting correct quantum states. Higher values are better for all three metrics. For all circuits, we perform simulations for 10000 shots. We also report the **two-qubit instruction counts** required to express an application using a particular gate set.

Compilation setup and performance: Our experiments were performed on two systems: Intel Xeon Gold 6230 CPU 2.10GHz (using up to 32 threads, 128GB RAM) and Intel Xeon Gold 6140 CPU 2.30GHz (using up to 36 threads, 192GB RAM). NuOp’s compilation time scales as $O(GH)$ where G is the number of 2-qubit gates in the application and H is the number of available hardware 2-qubit gates. The compile time is independent of the number of qubits. Since decomposing an application gate to a hardware gate is independent of other gates, we use a parallel implementation of NuOp. With 32 threads, decomposing a circuit with 1000 2-qubit gates to one hardware gate type requires around 220 seconds.

VII. TOOLFLOW VALIDATION AND TUNING

A. Comparison of Cirq and exact decomposition with NuOp

Figure 6 compares the average hardware gate counts and decomposition error rates for 100 random QV, QAOA, and QFT unitaries using the gate types from Table I. Compared to the KAK-based decomposition in Cirq, exact decomposition with NuOp i.e., NuOp-100% finds decompositions with similar or less amount of applications of two-qubit hardware gates. For example, NuOp uses 3 CZ, 3 SYC, or 3 iSWAP gates for one QV unitary while Cirq requires 3 CZ, 6 SYC, or 4 iSWAP gates. Across applications, NuOp-100% reduces gate counts by 1.26X on average, compared to Cirq. Furthermore, NuOp can decompose any target unitary with any native gate while KAK-based decomposition like Cirq is typically implemented for decomposing unitaries into specific native gates. For example, Cirq does not support decompositions of QV unitaries with \sqrt{i} SWAP. Due to NuOp’s flexibility in hardware gate sets and gate reductions over Cirq, we use it for all other experiments in this work.




Fig. 6: Comparison with Cirq. Hardware gate counts (lower is better) averaged across applications. Legend shows the target two-qubit gate types. Cirq does not support decompositions for QV with $\sqrt{i\text{SWAP}}$. NuOp-99.9% denotes that 99.9% is the target hardware fidelity for the decomposition. On average, NuOp uses 1.3X-2.3X less gates than Cirq, depending on the hardware gate fidelity and gate type. For NuOp, the average decomposition errors are within a specified and tunable threshold based on the device requirements, e.g., for NuOp-99%, the average decomposition error is 0.12%.




Fig. 7: Comparison of exact and approximate decomposition on application performance with different average error rates of the SYC gate. The left and right y axis denote the average HOP of 100 5-qubit QV benchmarks and the average XED of 100 4-qubit QAOA benchmarks respectively (higher values are better). In the low noise regime, applications that are decomposed by the approximate approach have similar performance as the ones using exact decomposition. Approximate decomposition starts outperforming the exact approach when the device has average error rate around Sycamore’s (dashed vertical line).

B. Comparison of exact and approximate decomposition

From Figure 6, across gate sets, NuOp offers 1.33-1.68X average reduction in gate counts compared to Cirq when approximation is applied. Compared to NuOp-100%, the reductions with approximation versions of NuOp are average 1.05-1.33X. For example, NuOp-95% obtains QV decompositions that have an average of 1.8 CZ gates per unitary (40% lower than NuOp-100%) with decomposition infidelity around 3% (not shown). Figure 7 shows that for both QV and QAOA benchmarks, decomposition with the approximate approach achieves similar application performance as exact decomposition in less noisy systems and outperforms exact decomposition for systems with higher error rates (around 0.62%). Given that the average two-qubit gate error rates in

the Sycamore processor and the Aspen-8 processor are 0.62% and 95%, we use the approximate decomposition approach for evaluating different instruction sets on these two devices.

VIII. BENEFITS OF MULTIPLE GATE TYPES

A. Optimizing Instruction Count

Figure 8 shows the expressivity for application unitaries across the parameter space of XY and fSim gates. We uniformly discretized the continuous space into a 19×19 grid of θ and ϕ values for this experiment. The value range of θ and ϕ are set to be $[0, \pi/2]$ and $[0, \pi]$ (instead of $[0, 2\pi]$) because of the unitary symmetries under single-qubit phase rotations. For each parameter combination, we calculated the average two-qubit instruction counts for 1000 random QV unitaries, 1000 random QAOA unitaries, 10 QFT unitaries, 60 FH unitaries, and the SWAP unitary. The heatmap shows these instruction counts normalized by the number of instructions in the application i.e., the number of two-qubit instructions per application operation. Overall, the average instruction counts are strongly influenced by the application and gate parameter choices. For each application, instruction counts vary between 1 to 6 depending on the gate parameters. Across gate types and applications, these large differences in instruction counts imply that choosing the right hardware gate types, based on the application requirement is important for NISQ systems.

For QV unitaries, gates that are close to $\text{fSim}(5\pi/12, 0)$ and $\text{fSim}(\pi/6, \pi)$ are most expressive and near to the optimal value with a continuous set, i.e., only around 2 instructions are required per unitary (compared to 3 applications of CZ, SYC, or iSWAP). In contrast, the most expressive gates for QAOA unitaries are around iSWAP and CZ (2 applications per unitary). For QFT unitaries, most $\text{fSim}(\theta, \pi(0))$ and $\text{fSim}(\pi/2(0), \phi)$ gates are expressivity-efficient (2 applications per unitary). For FH unitaries, gates close to $\sqrt{i\text{SWAP}}$ are more expressive (2 applications per unitary). We also studied the instruction counts required for the SWAP unitary, since SWAPs are the fundamental means of qubit movement in superconducting systems. For a large fraction of fSim gates, 3 gate applications are required to construct one SWAP unitary. With $\text{fSim}(\pi/4, \pi/2)$, we require 2 gate applications and with $\text{fSim}(\pi/2, \pi)$, we require only 1 gate application since this gate is equivalent to the SWAP gate up to single-qubit rotations.

From these heatmaps, certain gate types perform well for multiple applications. For example, $\text{fSim}(\pi/6, \pi)$ gives low instruction counts for all cases. Selecting gates in this manner, we arrived at a set of seven two-qubit gate types that offer low instruction counts. These types are shown in Table II. This list also includes CZ, $\text{XY}(\pi)$, SYC, and $\sqrt{i\text{SWAP}}$ gates which are used in current Rigetti and Google systems. We use these gate types and their combinations for the subsequent experiments.

B. Optimizing Application Reliability

Figures 9 and 10 show the application reliability comparison among instruction sets with a single gate type (S1-S7), instruction sets having multiple gate types (R1-R5 for Rigetti and G1-G7 for Google), and instruction sets which

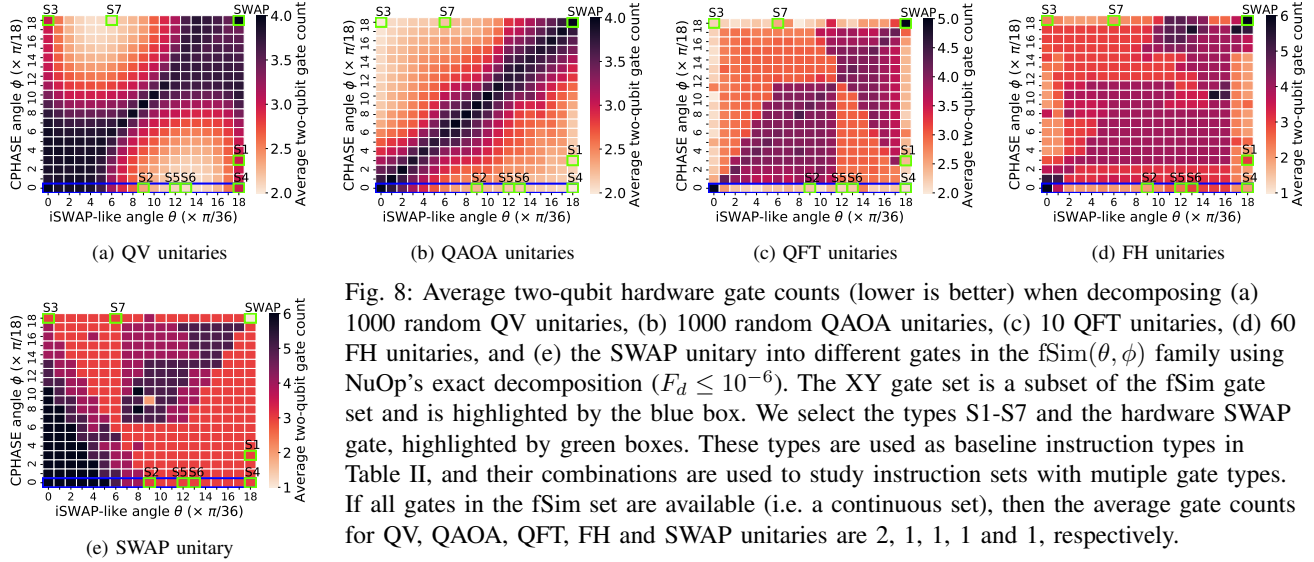



Fig. 8: Average two-qubit hardware gate counts (lower is better) when decomposing (a) 1000 random QV unitaries, (b) 1000 random QAOA unitaries, (c) 10 QFT unitaries, (d) 60 FH unitaries, and (e) the SWAP unitary into different gates in the $f\text{Sim}(\theta, \phi)$ family using NuOp’s exact decomposition ($F_d \leq 10^{-6}$). The XY gate set is a subset of the $f\text{Sim}$ gate set and is highlighted by the blue box. We select the types S1-S7 and the hardware SWAP gate, highlighted by green boxes. These types are used as baseline instruction types in Table II, and their combinations are used to study instruction sets with multiple gate types. If all gates in the $f\text{Sim}$ set are available (i.e. a continuous set), then the average gate counts for QV, QAOA, QFT, FH and SWAP unitaries are 2, 1, 1, 1 and 1, respectively.



(a) 100 3-qubit QV circuits (higher is better) (b) 100 4-qubit QAOA circuits (higher is better) (c) The 3-qubit QFT circuit (higher is better)

Fig. 9: Simulation results on Rigetti Aspen-8 with noise variations across gate types. The QV threshold is marked with a horizontal line at $2/3$. Instruction sets with multiple gate types (R1-R5) improve application performance compared to single-gate instruction sets (S2-S6). R5 that includes hardware SWAP gate, nearly achieves the same maximum reliability as Full_{XY} .

use the full continuous parameter space (Full_{XY} for Rigetti and $\text{Full}_{f\text{Sim}}$ for Google). In Figure 9a and 10a, instruction sets with multiple types increase the heavy output probability (HOP) by up to 6.6% (on average 2.3%), compared to the best instructions sets with one type. HOP is averaged across a large number of circuits and even small improvements are very significant for applications [32]. In particular, for Rigetti, none of the single gate type instruction sets reach the HOP threshold of $2/3$, while R3, R4 and R5 cross the threshold. That is, having access to these 4-6 types of two-qubit gates can double the quantum volume of the system. When the instruction set has multiple gate types, the cross-entropy differences of QAOA benchmarks improve by up to 12.9% (5.9% on average) in Google and 143.7% (43.1% on average) in Rigetti. Similarly for QFT, the success rates increase by up to 26% (11.6% on average) in Google and 56.5% (15.2% on average) in Rigetti. Finally, the fidelity of the FH circuit increases by up to 56.8% (29.8% on average) in Google (The fidelity for Rigetti is very low owing to high gate error rates, and is not reported in this paper).

We further compare the application reliability using multiple-gate sets with the values achieved by using full continuous sets. In Figure 9 and 10, across benchmarks and quantum systems, the applications decomposed with Full_{XY} and $\text{Full}_{f\text{Sim}}$ have much higher reliability than the decomposition with R1-R4 and G1-G6 instruction sets. This is because many routing operations (i.e., application SWAP gates in Qiskit compiler) are inserted to comply with the connectivity limitation in Sycamore and Aspen-8. This leads to a significant increase in the total instruction counts. These SWAP operations are difficult to express succinctly using the gates in R1-R4 and G1-G6. When the SWAP gate is added as a native type, i.e., in R5 and G7, the instruction set offers reliability close to the optimal values achieved by $\text{Full}_{f\text{Sim}}$. *Therefore, our work shows that implementing multiple types of two-qubit gates is beneficial and only a few application-expressive types are required to obtain reliability comparable to using the entire continuous set.*

Figure 10f compares the fidelity for FH circuits with 10 and 20 qubits. At high noise levels, having multiple gate types

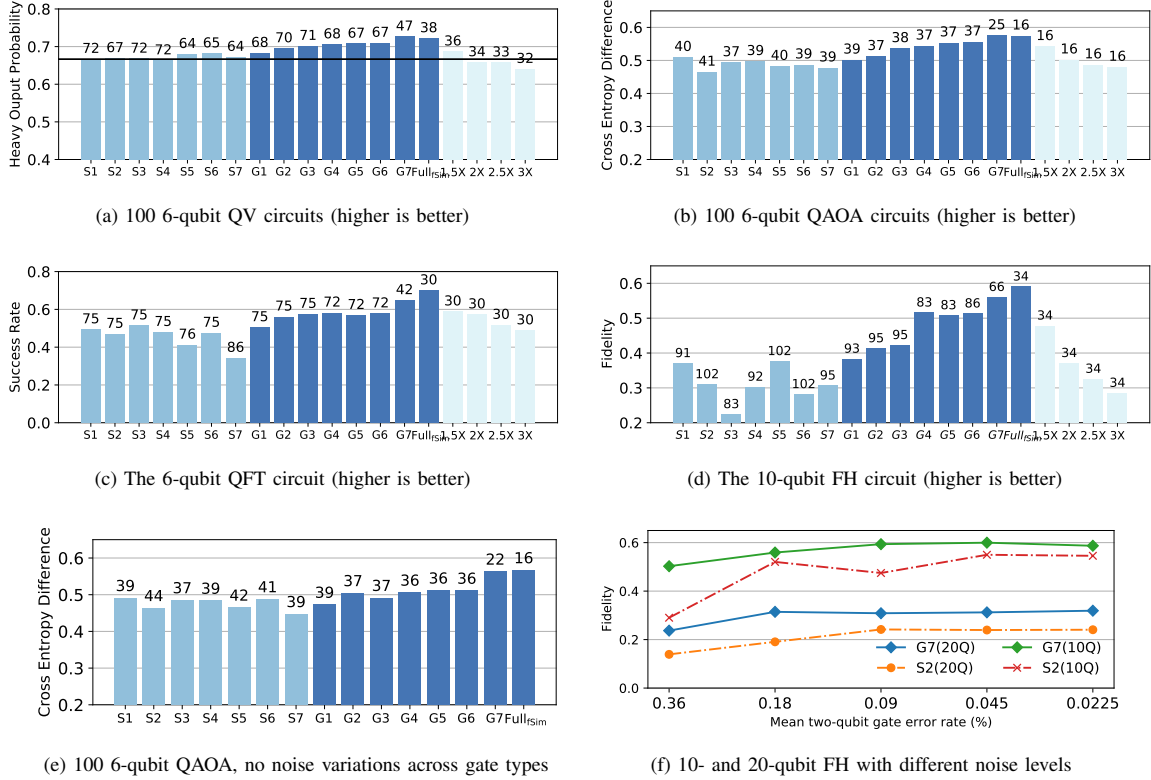


Fig. 10: Simulation results on Google Sycamore (a-d) with and (e) without noise variation across gate types. The QV threshold is marked with a horizontal line at $2/3$. Multi-gate instruction sets (G1-G7) improve application performance compared to single-gate instruction sets (S1-S7). Compared to (b), the improvements of G1-G7 in (e) are less if no noise variation across gate types. The instruction set with SWAP gate (G7) achieves comparable reliability as Full_{fSim}. The advantages of the continuous set reduces or even disappear when the average gate error rates in continuous set increase to 2X (the lightest blue bars). (f) Fidelity for Fermi-Hubbard model circuits with 10 and 20 qubits. Average two-qubit hardware error rates are varied from 0.02% (low noise) to 0.36% (high noise, current error rate). Across circuit sizes and noise ranges, multiple gate types (G7) outperforms the single gate type.


offers significantly better executions, with average 1.45X, up to 1.7X improvement in fidelity compared to a single gate type, for the 20-qubit case. Across circuit sizes, the instruction set with multiple gate types (G7) outperforms the set with only single gate type (S2) consistently, but the benefits reduce as the hardware error rate improves. Multiple instruction types are most beneficial at hardware noise levels that are expected in the next 5-10 years.

To understand why reliability improves using multiple gate types, we studied two sources of improvements: 1) reduction in instruction counts and 2) noise adaptivity across instruction types. In Figures 9 and 10, the bars are annotated on top with the respective two-qubit instruction counts. Compared to S1-S7, the number of instructions required to implement applications reduces when multiple instruction types are available in R1-R5 and G1-G7.² For Rigetti, multiple-gate sets R1-R4


²Note that these reductions are less prominent than the heatmaps in Figure 8 because the most expressive gate types were selected as the S1-S7 baselines.

reduce the instruction count by up to 14% compared to the single-instruction sets. When the SWAP gate is also added in R5, the reduction is on average 1.5X for QV, 2X for QAOA and 1.5X for QFT. Similarly, for Google, sets with multiple instruction types reduce the instruction count and G7 obtains average 1.9X reduction for QV, 1.6X for QAOA, 1.8X for QFT, and 1.3X for FH. In particular, the instruction counts of R5 match Full_{XY} for all the three applications and the counts of G7 are very close to Full_{fSim}. Hence, having a small number of instruction types is sufficient to obtain the same expressive power of a full continuous gate set for most applications.

To isolate the reliability benefits from gate count reductions from reductions obtained through noise adaptivity, in Figure 10e, we repeated the experiment in Figure 10b but with the assumption that there is no noise variation across gate types. Compared to Figure 10b, G7 in Figure 10e still reduces instruction counts and obtains higher cross entropy difference than the other discrete instruction sets. Importantly, it almost



(a)



(b)

Fig. 11: Tradeoffs between calibration overhead and application performance. (a) The number of calibration and benchmarking circuits grows with the number of hardware gate types and the device size (# qubits). (b) The calibration time cost (left y axis) and application performance (right y axis) increase with the number of hardware gate types.

matches the performance of $\text{Full}_{\text{fSim}}$. In contrast, even though the gate count reductions of G1-G6 remain similar ($\sim 10\%$), the application reliability improvements are lesser than the improvements in Figure 10b. Similar observations can be found for other benchmarks in Sycamore. This indicates that small improvements in gate counts are likely to have a low impact on application reliability in future systems where better gates and lower noise variability. However, in current and near-term systems, different gate types have significantly different error rates. NuOp exploits these noise differences to choose the best gate types across qubits and therefore achieves more improvements with G1-G6 in Figure 10b.

To study the importance of noise adaptivity, we compare the single-instruction sets S1-S4 with their combination G3 in Figure 10c. The two-qubit gates in S1-S4 have the same expressivity and same average gate error rates. However, S1-S4 have different performance because of fluctuations in error rates across the operating qubits. In comparison, G3 achieves much higher success rate than S1-S4 because exposing multiple instruction types allows noise-adaptivity across gate types and helps in mitigating the noise variations that arise when only type of gate is used. *Therefore, the fidelity improvements with multiple instruction types arise because of both improved expressivity and noise-adaptivity across gate types.*

IX. CALIBRATION OVERHEAD OF MULTIPLE GATE TYPES

In this section, we estimate the calibration overhead with multiple gate types by adopting the calibration model of the continuous fSim set in [4], where 525 gate types are calibrated on real hardware. As mentioned previously, a fSim gate can be seen as the composite of a CPHASE gate and an iSWAP-like gate. Therefore, the calibration of one single fSim gate type ($\text{fSim}(\theta, \phi)$) on one qubit pair can be divided into several steps. The first step is to calibrate the CPHASE gates with rotation angle ϕ and angle π . The second step calibrates iSWAP-like gates with angle 0 and angle $\pi/2$ while the CPHASE angle is tuned to ϕ . Then one tunes up the iSWAP-like gate angle θ with CPHASE angle π . Afterwards, the calibration data of these CPHASE and iSWAP-like gates are used to construct a pulse for the target $\text{fSim}(\theta, \phi)$ gate and unitary tomography is performed to further adjust its pulse parameters. Finally, the fidelity of this fSim gate is characterized by running a large number of cross entropy benchmarking circuits (1000 rounds). This is a conservative calibration model in which each gate type is calibrated individually on isolated qubits. In practice, more calibrations are required for minimizing errors due to pulse overlaps [4] and crosstalk.

Figure 11a shows the number of calibration circuits for different number of gate types and system sizes. The number of circuits increases linearly with system size and the number of gate types. For a 54-qubit device, we require $\sim 10^7$ circuits to calibrate 10 gate types. For a 1000-qubit system, nearly a billion circuits are required even for a small number of gate types. Comparing with current systems, calibrating two qubits in [63] takes up to four hours, including the calibration of electronics, qubit frequencies, single-qubit rotations, and a single two-qubit gate type. Even if we conservatively assume that a single two-qubit gate type takes ~ 2 hours, the linear scaling of calibration time with gate types and system sizes makes it practically intractable to calibrate a full gate family.

Figure 11b shows the tradeoffs between calibration time and application reliability (data from Figures 9 and 10). When more gates are available in the hardware, the average application reliability improves, but with diminishing returns after five gate types. Interestingly, the gate sets that have SWAP gates (R5 and G7) offer reliability comparable to the continuous sets. In practice, the continuous instruction set will likely have lower average gate fidelity than the discrete sets due to calibration difficulties and pulse overlaps [13]. The reliability improvements resulting from gate reduction may be compromised by the increased average gate error rates. To study this, we performed simulations using $\text{Full}_{\text{fSim}}$ with different increasing factors in the average error rates as shown in Figure 10a-10c. Compared to discrete multi-gate sets G1-G7 with average error rates 0.62%, the advantages of $\text{Full}_{\text{fSim}}$ disappear when its average error rates increases to 1.5X, 2X, 2.5X for QV, QAOA, and QFT benchmarks respectively.

Since calibration and benchmarking overheads increase with the number of gate types, having a small number of expressive two-qubit gates is beneficial and sufficient to obtain high

reliability. *To summarize, we recommend that QC systems implement 4-8 different two-qubit hardware gates that are expressive for most applications to obtain a sweet spot of high-fidelity application executions and low device calibration overheads. In addition, having hardware SWAP gates for connectivity-constrained devices or other application-specific gates can greatly improve reliability.*

X. CONCLUSIONS

Most QC instruction sets have a single type of two-qubit gate. This design choice sacrifices application expressivity, but simplifies calibration. On the other hand, vendors have proposed continuous gate families which maximize application expressivity, but are very difficult to calibrate especially as systems scale up. These tradeoffs are reminiscent of the instruction count vs. hardware complexity tradeoffs of RISC and CISC designs in classical computer architecture. Our work used an application-based simulation analysis of instruction sets, with realistic device and calibration models, to identify a small instruction set that is both *efficient for application decompositions and easy to calibrate*. Our proposed instruction sets R5 for Rigetti and G7 for Google have only 4-8 two-qubit gate types, but offer near-optimal application-expressivity and fidelity compared to a continuous instruction set, while reducing calibration overhead by two orders of magnitude. NuOp is available open-source at <https://github.com/prakashmurali/NuOp>.

With several vendors aiming to improve QC reliability by offering better gate types, our work is poised to guide instruction set design choices through the NISQ-regime.

XI. ACKNOWLEDGEMENTS

We thank Eric Peterson (Rigetti), Nathan Lacroix (ETH Zurich), Zhang Jiang (Google) and Tom O'Brien (Google) for discussions on compilation and instruction set design. This work is funded in part by Enabling Practical-scale Quantum Computation (EPiQC), an NSF Expedition in Computing, grant 1730082. P.M. was also supported by an IBM PhD Fellowship. D.B. and L.L. acknowledge support from the EPSRC Prosperity Partnership in Quantum Software for Modelling and Simulation, grant EP/S005021/1.

REFERENCES

- [1] Rigetti, "Rigetti Aspen-8," <https://medium.com/rigetti/rigetti-aspen-8-on-aws-236d9dc11613>, 2020.
- [2] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, A. Dunsworth, E. Farhi, B. Foxen, A. Fowler, C. Gidney, M. Giustina, R. Graff, K. Guerin, S. Habegger, M. P. Harrigan, M. J. Hartmann, A. Ho, M. Hoffmann, T. Huang, T. S. Humble, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, J. Kelly, P. V. Klimov, S. Knysh, A. Korotkov, F. Kostritsa, D. Landhuis, M. Lindmark, E. Lucero, D. Lyakh, S. Mandrà, J. R. McClean, M. McEwen, A. Megrant, X. Mi, K. Michielsen, M. Mohseni, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Y. Niu, E. Ostby, A. Petukhov, J. C. Platt, C. Quintana, E. G. Rieffel, P. Roushan, N. C. Rubin, D. Sank, K. J. Satzinger, V. Smelyanskiy, K. J. Sung, M. D. Trevithick, A. Vainsencher, B. Villalonga, T. White, Z. J. Yao, P. Yeh, A. Zalcman, H. Neven, and J. M. Martinis, "Quantum supremacy using a programmable superconducting processor," *Nature*, vol. 574, no. 7779, pp. 505–510, 2019.
- [3] D. M. Abrams, N. Didier, B. R. Johnson, M. P. da Silva, and C. A. Ryan, "Implementation of the XY interaction family with calibration of a single pulse," *arXiv:1912.04424*, 2019.
- [4] B. Foxen, C. Neill, A. Dunsworth, P. Roushan, B. Chiaro, A. Megrant, J. Kelly, Z. Chen, K. Satzinger, R. Barends, F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, S. Boixo, D. Buell, B. Burkett, Y. Chen, R. Collins, E. Farhi, A. Fowler, C. Gidney, M. Giustina, R. Graff, M. Harrigan, T. Huang, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, P. Klimov, A. Korotkov, F. Kostritsa, D. Landhuis, E. Lucero, J. McClean, M. McEwen, X. Mi, M. Mohseni, J. Y. Mutus, O. Naaman, M. Neeley, M. Niu, A. Petukhov, C. Quintana, N. Rubin, D. Sank, V. Smelyanskiy, A. Vainsencher, T. C. White, Z. Yao, P. Yeh, A. Zalcman, H. Neven, and J. M. Martinis, "Demonstrating a Continuous Set of Two-Qubit Gates for Near-Term Quantum Algorithms," *Phys. Rev. Lett.*, vol. 125, p. 120504, Sep 2020.
- [5] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, "Quantum machine learning," *Nature*, vol. 549, Sep 2017.
- [6] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, "A variational eigenvalue solver on a photonic quantum processor," *Nature Communications*, vol. 5, Jul 2014.
- [7] B. Bauer, S. Bravyi, M. Motta, and G. K.-L. Chan, "Quantum algorithms for quantum chemistry and quantum materials science," *arXiv:2001.03685*, 2020.
- [8] J. Preskill, "Quantum Computing in the NISQ era and beyond," *Quantum*, vol. 2, p. 79, Aug. 2018.
- [9] Y. Nam, J.-S. Chen, N. C. Pisenti, K. Wright, C. Delaney, D. Maslov, K. R. Brown, S. Allen, J. M. Amini, J. Apisdorf, K. M. Beck, A. Blinov, V. Chaplin, M. Chmielewski, C. Collins, S. Debnath, K. M. Hudek, A. M. Ducore, M. Keesan, S. M. Kreikemeier, J. Mizrahi, P. Solomon, M. Williams, J. D. Wong-Campos, D. Moehring, C. Monroe, and J. Kim, "Ground-state energy estimation of the water molecule on a trapped-ion quantum computer," *npj Quantum Information*, vol. 6, no. 1, p. 33, Apr 2020.
- [10] J. S. Otterbach, R. Manenti, N. Alidoust, A. Bestwick, M. Block, B. Bloom, S. Caldwell, N. Didier, E. S. Fried, S. Hong, P. Karalekas, C. B. Osborn, A. Papageorge, E. C. Peterson, G. Prawiroatmodjo, N. Rubin, C. A. Ryan, D. Scarabelli, M. Scheer, E. A. Sete, P. Sivarajah, R. S. Smith, A. Staley, N. Tezak, W. J. Zeng, A. Hudson, B. R. Johnson, M. Reagor, M. P. da Silva, and C. Rigetti, "Unsupervised machine learning on a hybrid quantum computer," *arXiv:1712.05771*, 2017.
- [11] J. Kelly, R. Barends, A. G. Fowler, A. Megrant, E. Jeffrey, T. C. White, D. Sank, J. Y. Mutus, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, I.-C. Hoi, C. Neill, P. J. J. O'Malley, C. Quintana, P. Roushan, A. Vainsencher, J. Wenner, A. N. Cleland, and J. M. Martinis, "State preservation by repetitive error detection in a superconducting quantum circuit," *Nature*, vol. 519, no. 7541, pp. 66–69, Mar 2015.
- [12] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, "Characterizing quantum supremacy in near-term devices," *Nature Physics*, vol. 14, no. 6, pp. 595–600, Jun 2018.
- [13] N. Lacroix, C. Hellings, C. K. Andersen, A. Di Paolo, A. Remm, S. Lazar, S. Krinner, G. J. Norris, M. Gabureac, J. Heinsoo, A. Blais, C. Eichler, and A. Wallraff, "Improving the performance of deep quantum optimization algorithms with continuous gate sets," *PRX Quantum*, vol. 1, no. 2, p. 110304, 2020.
- [14] IBM, "IBM Quantum Experience Devices," <https://quantum-computing.ibm.com/>, 2020.
- [15] P. Murali, J. M. Baker, A. Javadi-Abhari, F. T. Chong, and M. Martonosi, "Noise-Adaptive Compiler Mappings for Noisy Intermediate-Scale Quantum Computers," in *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, 2019, pp. 1015–1029.
- [16] P. Murali, N. M. Linke, M. Martonosi, A. J. Abhari, N. H. Nguyen, and C. H. Alderete, "Full-Stack, Real-System Quantum Computer Studies: Architectural Comparisons and Design Insights," in *Proceedings of the 46th International Symposium on Computer Architecture*, ser. ISCA '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 527–540. [Online]. Available: <https://doi.org/10.1145/3307650.3322273>
- [17] S. S. Tannu and M. K. Qureshi, "Not All Qubits Are Created Equal: A Case for Variability-Aware Policies for NISQ-Era Quantum Computers," in *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*. ACM, 2019, pp. 987–999.

- [18] S. Nishio, Y. Pan, T. Satoh, H. Amano, and R. V. Meter, "Extracting Success from IBM's 20-Qubit Machines Using Error-Aware Compilation," *ACM Journal on Emerging Technologies in Computing Systems (JETC)*, vol. 16, no. 3, pp. 1–25, 2020.
- [19] R. S. Smith, M. J. Curtis, and W. J. Zeng, "A practical quantum instruction set architecture," *arXiv:1608.03355*, 2016.
- [20] A. W. Cross, L. S. Bishop, J. A. Smolin, and J. M. Gambetta, "Open quantum assembly language," *arXiv:1707.03429*, 2017.
- [21] X. Fu, L. Riesebo, M. A. Rol, J. van Straten, J. van Someren, N. Khammassi, I. Ashraf, R. F. L. Vermeulen, V. Newsum, K. K. L. Loh, J. C. de Sterke, W. J. Vlothuizen, R. N. Schouten, C. G. Almudever, L. DiCarlo, and K. Bertels, "eQASM: An Executable Quantum Instruction Set Architecture," in *2019 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, 2019, pp. 224–237.
- [22] K. Bertels, A. Sarkar, T. Hubregtsen, M. Serrao, A. A. Mouedenne, A. Yadav, A. Krol, I. Ashraf, and C. G. Almudever, "Quantum computer architecture toward full-stack quantum accelerators," *IEEE Transactions on Quantum Engineering*, vol. 1, pp. 1–17, 2020.
- [23] M. Kjaergaard, M. E. Schwartz, A. Greene, G. O. Samach, A. Bengtsson, M. O'Keefe, C. M. McNally, J. Braumüller, D. K. Kim, P. Krantz, M. Marvian, A. Melville, B. M. Niedzielski, Y. Sung, R. Winik, J. Yoder, D. Rosenberg, K. Obenland, S. Lloyd, T. P. Orlando, I. Marvian, S. Gustavsson, and W. D. Oliver, "A quantum instruction set implemented on a superconducting quantum processor," 2020, *arXiv:2001.08838*.
- [24] X. Zou, S. P. Premaratne, M. A. Rol, S. Johri, V. Ostroukh, D. J. Michalak, R. Caudillo, J. S. Clarke, L. DiCarlo, and A. Y. Matsuura, "Enhancing a near-term quantum accelerator's instruction set architecture for materials science applications," *IEEE Transactions on Quantum Engineering*, vol. 1, pp. 1–7, 2020.
- [25] D. Maslov, S. M. Falconer, and M. Mosca, "Quantum circuit placement," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 27, no. 4, pp. 752–763, 2008.
- [26] M. Y. Siraichi, V. F. d. Santos, S. Collange, and F. M. Q. Pereira, "Qubit allocation," in *Proceedings of the 2018 International Symposium on Code Generation and Optimization*. ACM, 2018, pp. 113–125.
- [27] D. Venturelli, M. Do, E. Rieffel, and J. Frank, "Compiling quantum circuits to realistic hardware architectures using temporal planners," *Quantum Science and Technology*, vol. 3, no. 2, p. 025004, 2018.
- [28] A. Zulehner and R. Wille, "Compiling SU(4) quantum circuits to IBM QX architectures," in *Proceedings of the 24th Asia and South Pacific Design Automation Conference*. ACM, 2019, pp. 185–190.
- [29] G. Li, Y. Ding, and Y. Xie, "Tackling the qubit mapping problem for NISQ-era quantum devices," in *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*. ACM, 2019, pp. 1001–1014.
- [30] P. Murali, D. C. McKay, M. Martonosi, and A. Javadi-Abhari, "Software Mitigation of Crosstalk on Noisy Intermediate-Scale Quantum Computers," in *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, 2020, pp. 1001–1016.
- [31] H. Abraham, AduOffei, R. Agarwal, I. Y. Akhalwaya, G. Aleksandrowicz, T. Alexander, E. Arbel, A. Asfaw, C. Azaustre, AzizNgoueya, A. Bansal, P. Barkoutsos, G. Barron, L. Bello, Y. Ben-Haim, D. Bevenius, L. S. Bishop, S. Bolos, S. Bosch, S. Bravyi, D. Bucher, A. Burov, F. Cabrera, P. Calpin, L. Capelluto, J. Carballo, G. Carrascal, A. Chen, C.-F. Chen, E. Chen, J. C. Chen, R. Chen, J. M. Chow, S. Churchill, C. Claus, C. Clauss, R. Cocking, A. J. Cross, A. W. Cross, S. Cross, J. Cruz-Benito, C. Culver, A. D. Córcoles-Gonzales, S. Dague, T. E. Dandachi, M. Daniels, M. Dartiaill, DavideFrr, A. R. Davila, A. Dekusar, D. Ding, J. Doi, E. Drechsler, Drew, E. Dumitrescu, K. Dumon, I. Duran, K. EL-Safy, E. Eastman, P. Eendebak, D. Egger, M. Everitt, P. M. Fernández, A. H. Ferrera, R. Fouilland, FranckChevalier, A. Frisch, A. Fuhrer, M. GEORGE, J. Gacon, B. G. Gago, C. Gambella, J. M. Gambetta, A. Gammanpila, L. Garcia, S. Garion, A. Gilliam, A. Giridharan, J. Gomez-Mosquera, S. de la Puente González, J. Gorzinski, I. Gould, D. Greenberg, D. Grinko, W. Guan, J. A. Gunnels, M. Haglund, I. Haide, I. Hamamura, O. C. Hamido, V. Havlicek, J. Hellmers, L. Herok, S. Hillmich, H. Horii, C. Howington, S. Hu, W. Hu, R. Huisman, H. Imai, T. Imamichi, K. Ishizaki, R. Iten, T. Itoko, JamesSeaward, A. Javadi, A. Javadi-Abhari, Jessica, M. Jivrajani, K. Johns, Jonathan-Shoemaker, T. Kachmann, N. Kanazawa, Kang-Bae, A. Karazeev, P. Kassebaum, S. King, Knabberjoe, Y. Kobayashi, A. Kovyshin, R. Krishnakumar, V. Krishnan, K. Krsulich, G. Kus, R. LaRose, E. Lacal, R. Lambert, J. Lapeyre, J. Latone, S. Lawrence, C. Lee, G. Li, D. Liu, P. Liu, Y. Maeng, A. Malyshev, J. Manela, J. Marecek, M. Marques, D. Maslov, D. Mathews, A. Matsuo, D. T. McClure, C. McGarry, D. McKay, D. McPherson, S. Meesala, T. Metcalfe, M. Mevissen, A. Mezzacapo, R. Midha, Z. Mineev, A. Mitchell, N. Moll, M. D. Mooring, R. Morales, N. Moran, MrF, P. Murali, J. Müggengberg, D. Nadlinger, K. Nakanishi, G. Nannicini, P. Nation, E. Navarro, Y. Naveh, S. W. Neagle, P. Neuweiler, P. Niroula, H. Norlen, L. J. O'Riordan, O. Ogunbayo, P. Ollitrault, S. Oud, D. Padilha, H. Paik, Y. Pang, S. Perriello, A. Phan, F. Piro, M. Pistoia, C. Piveteau, A. Pozasi-Kerstjens, V. Prutyayov, D. Puzzuoli, J. Pérez, Quintiii, R. I. Rahman, A. Raja, N. Ramagiri, A. Rao, R. Raymond, R. M.-C. Redondo, M. Reuter, J. Rice, D. M. Rodríguez, RohithKarur, M. Rossmannek, M. Ryu, T. SAPV, SamFerracin, M. Sandberg, R. Sapra, H. Sargsyan, A. Sarkar, N. Sathaye, B. Schmitt, C. Schnabel, Z. Schoenfeld, T. L. Scholten, E. Schoute, J. Schwarm, I. F. Sertage, K. Setia, N. Shammah, Y. Shi, A. Silva, A. Simonetto, N. Singstock, Y. Siraichi, I. Sitdikov, S. Sivarajah, M. B. Sletfjerding, J. A. Smolin, M. Soeken, I. O. Sokolov, SooluThomas, Starfish, D. Steenken, M. Stypulkoski, S. Sun, K. J. Sung, H. Takahashi, I. Tavernelli, C. Taylor, P. Taylour, S. Thomas, M. Tillet, M. Tod, M. Tomasiak, E. de la Torre, K. Trabing, M. Treinish, TrishaPe, W. Turner, Y. Vaknin, C. R. Valcarce, F. Varchon, A. C. Vazquez, V. Villar, D. Vogt-Lee, C. Vuillot, J. Weaver, R. Wiecezorek, J. A. Wildstrom, E. Winston, J. J. Woehr, S. Woerner, R. Woo, C. J. Wood, R. Wood, S. Wood, S. Wood, J. Wootton, D. Yeralin, D. Yongemallo, R. Young, J. Yu, C. Zachow, L. Zdanski, H. Zhang, C. Zoufal, and Zoufal, "Qiskit: An open-source framework for quantum computing," 2019.
- [32] A. W. Cross, L. S. Bishop, S. Sheldon, P. D. Nation, and J. M. Gambetta, "Validating quantum computers using randomized model circuits," *Phys. Rev. A*, vol. 100, no. 3, p. 032328, 2019.
- [33] Rigetti, "Rigetti quantum computing service," <https://www.qcs.rigetti.com/>, 2020.
- [34] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, S. Boixo, M. Broughton, B. B. Buckley, D. A. Buell, B. Burkett, N. Bushnell, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, S. Demura, A. Dunswoth, D. Eppens, E. Farhi, A. Fowler, B. Foxen, C. Gidney, M. Giustina, R. Graff, S. Habegger, M. P. Harrigan, A. Ho, S. Hong, T. Huang, L. B. Ioffe, S. V. Isakov, E. Jeffrey, Z. Jiang, C. Jones, D. Kafri, K. Kechedzhi, J. Kelly, S. Kim, P. V. Klimov, A. N. Korotkov, F. Kostritsa, D. Landhuis, P. Laptev, M. Lindmark, M. Leib, E. Lucero, O. Martin, J. M. Martinis, J. R. McClean, M. McEwen, A. Megrant, X. Mi, M. Mohseni, W. Mroczkiewicz, J. Mutus, O. Naaman, M. Neeley, C. Neill, F. Neukart, H. Neven, M. Y. Niu, T. E. O'Brien, B. O'Gorman, E. Ostby, A. Petukhov, H. Putterman, C. Quintana, P. Roushan, N. C. Rubin, D. Sank, K. J. Satzinger, A. Skolik, V. Smelyanskiy, D. Strain, M. Streif, K. J. Sung, M. Szalay, A. Vainsencher, T. White, Z. J. Yao, P. Yeh, A. Zalcman, and L. Zhou, "Quantum approximate optimization of non-planar graph problems on a planar superconducting processor," *arXiv:2004.04197*, 2020.
- [35] D. C. McKay, C. J. Wood, S. Sheldon, J. M. Chow, and J. M. Gambetta, "Efficient Z gates for quantum computing," *Phys. Rev. A*, vol. 96, p. 022330, Aug 2017.
- [36] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, "Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets," *Nature*, vol. 549, Sep 2017.
- [37] E. Farhi, J. Goldstone, and S. Gutmann, "A quantum approximate optimization algorithm," *arXiv:1411.4028*, 2014.
- [38] E. C. Peterson, G. E. Crooks, and R. S. Smith, "Two-qubit circuit depth and the monodromy polytope," *Quantum*, vol. 4, p. 247, 2020.
- [39] P. V. Klimov, J. Kelly, J. M. Martinis, and H. Neven, "The snake optimizer for learning quantum processor control parameters," *arXiv:2006.04594*, 2020.
- [40] R. R. Tucci, "A rudimentary quantum compiler (2nd ed.)," *arXiv:9902062*, 1999.
- [41] N. Khaneja, R. Brockett, and S. J. Glaser, "Time optimal control in spin systems," *Phys. Rev. A*, vol. 63, no. 3, p. 032308, 2001.
- [42] B. Kraus and J. Cirac, "Optimal creation of entanglement using a two-qubit gate," *Phys. Rev. A*, vol. 63, no. 6, p. 062309, 2001.
- [43] F. Vatan and C. Williams, "Optimal quantum circuits for general two-qubit gates," *Phys. Rev. A*, vol. 69, no. 3, p. 032315, 2004.
- [44] M. Möttönen, J. J. Vartiainen, V. Bergholm, and M. M. Salomaa,

- “Quantum circuits for general multiqubit gates,” *Phys. Rev. Lett.*, vol. 93, no. 13, p. 130502, 2004.
- [45] R. Iten, R. Colbeck, and M. Christandl, “Quantum circuits for quantum channels,” *Phys. Rev. A*, vol. 95, no. 5, p. 052316, 2017.
- [46] D. Maslov, C. Young, D. M. Miller, and G. W. Dueck, “Quantum circuit simplification using templates,” in *Design, Automation and Test in Europe*, 2005, pp. 1208–1213 Vol. 2.
- [47] R. S. Smith, E. C. Peterson, M. Skilbeck, and E. Davis, “An open-source, industrial-strength optimizing compiler for quantum programs,” *Quantum Science and Technology*, 2020.
- [48] T. Jones and S. C. Benjamin, “Quantum compilation and circuit optimisation via energy dissipation,” *arXiv:1811.03147*, 2018.
- [49] S. Khatri, R. LaRose, A. Poremba, L. Cincio, A. T. Sornborger, and P. J. Coles, “Quantum-assisted quantum compiling,” *Quantum*, vol. 3, p. 140, 2019.
- [50] M. G. Davis, E. Smith, A. Tudor, K. Sen, I. Siddiqi, and C. Iancu, “Heuristics for quantum compiling with a continuous gate set,” *arXiv:1912.02727*, 2019.
- [51] H. R. Grimsley, S. E. Economou, E. Barnes, and N. J. Mayhall, “An adaptive variational algorithm for exact molecular simulations on a quantum computer,” *Nature communications*, vol. 10, no. 1, pp. 1–9, 2019.
- [52] R. Fletcher, *Practical Methods of Optimization*, 2nd ed. New York, NY, USA: John Wiley & Sons, 1987.
- [53] J. Liu and H. Zhou, “Reliability Modeling of NISQ-Era Quantum Computers,” in *2020 IEEE International Symposium on Workload Characterization*, ser. IISWC, 2020.
- [54] N. M. Linke, D. Maslov, M. Roetteler, S. Debnath, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe, “Experimental comparison of two quantum computing architectures,” *Proceedings of the National Academy of Sciences*, vol. 114, no. 13, pp. 3305–3310, 2017. [Online]. Available: <https://www.pnas.org/content/114/13/3305>
- [55] K. Wright, K. M. Beck, S. Debnath, J. M. Amini, Y. Nam, N. Grzesiak, J.-S. Chen, N. C. Pistenti, M. Chmielewski, C. Collins, K. M. Hudek, J. Mizrahi, J. D. Wong-Campos, S. Allen, J. Apisdorf, P. Solomon, M. Williams, A. M. Ducore, A. Blinov, S. M. Kreikemeier, V. Chaplin, M. Keesan, C. Monroe, and J. Kim, “Benchmarking an 11-qubit quantum computer,” *Nature Communications*, vol. 10, no. 1, p. 5464, Nov 2019.
- [56] Z. Wang, S. Hadfield, Z. Jiang, and E. G. Rieffel, “Quantum approximate optimization algorithm for MaxCut: A fermionic view,” *Phys. Rev. A*, vol. 97, no. 2, p. 022304, 2018.
- [57] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, A. Bengtsson, S. Boixo, M. Broughton, B. Buckley, D. A. Buell, B. Burkett, N. Bushnell, Y. Chen, Z. Chen, B. Chiaro, R. Collins, S. Cotton, W. Courtney, S. Demura, A. Derk, A. Dunsworth, D. Eppens, T. Eckl, C. Erickson, E. Farhi, A. Fowler, B. Foxen, C. Gidney, M. Giustina, R. Graff, J. A. Gross, S. Habegger, M. P. Harrigan, A. Ho, S. Hong, T. Huang, W. Huggins, L. Ioffe, S. V. Isakov, E. Jeffrey, Z. Jiang, C. Jones, D. Kafri, K. Kechedzhi, J. Kelly, S. Kim, P. Klimov, A. Korotkov, F. Kostritsa, D. Landhuis, P. Laptev, M. Lindmark, E. Lucero, M. Marthaler, O. Martin, J. M. Martinis, A. Marusczyk, S. McArdle, J. McClean, T. McCourt, M. McEwen, A. Megrant, C. Mejuto-Zaera, X. Mi, M. Mohseni, W. Mroczkiewicz, J. Mutus, O. Naaman, M. Neeley, C. Neill, H. Neven, M. Newman, M. Y. Niu, T. E. O’Brien, E. Ostby, B. Pat’o, A. Petukhov, H. Putterman, C. Quintana, J. Reiner, P. Roushan, N. Rubin, D. Sank, K. Satzinger, V. Smelyanskiy, D. Strain, K. J. Sung, P. Schmitteckert, M. Szalay, N. M. Tubman, A. Vainsencher, T. White, N. Vogt, Z. Yao, P. Yeh, A. Zalcman, and S. Zanker, “Observation of separated dynamics of charge and spin in the fermi-hubbard model,” *arXiv:2010.07965*, 2020.
- [58] P. W. Shor, “Algorithms for quantum computation: Discrete logarithms and factoring,” in *Proceedings 35th annual symposium on foundations of computer science*. Ieee, 1994, pp. 124–134.
- [59] M. A. Nielsen and I. Chuang, “Quantum computation and quantum information,” 2002.
- [60] Google, “Recirq,” <https://github.com/quantumlib/ReCirq>, 2020.
- [61] S. Aaronson and L. Chen, “Complexity-theoretic foundations of quantum supremacy experiments,” *arXiv:1612.05903*, 2016.
- [62] C. Neill, P. Roushan, K. Kechedzhi, S. Boixo, S. V. Isakov, V. Smelyanskiy, A. Megrant, B. Chiaro, A. Dunsworth, K. Arya, R. Barends, B. Burkett, Y. Chen, Z. Chen, A. Fowler, B. Foxen, M. Giustina, R. Graff, E. Jeffrey, T. Huang, J. Kelly, P. Klimov, E. Lucero, J. Mutus, M. Neeley, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, H. Neven, and J. M. Martinis, “A blueprint for demonstrating quantum supremacy with superconducting qubits,” *Science*, vol. 360, no. 6385, pp. 195–199, 2018. [Online]. Available: <https://science.sciencemag.org/content/360/6385/195>
- [63] C. Rigetti, J. M. Gambetta, S. Poletto, B. L. T. Plourde, J. M. Chow, A. D. Córcoles, J. A. Smolin, S. T. Merkel, J. R. Rozen, G. A. Keefe, M. B. Rothwell, M. B. Ketchen, and M. Steffen, “Superconducting qubit in a waveguide cavity with a coherence time approaching 0.1 ms,” *Phys. Rev. B*, vol. 86, p. 100506, Sep 2012. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevB.86.100506>