Compiling to Gate Sets Tailored to Qubit Variation

ABSTRACT

Near-term quantum computers are primarily limited by errors in quantum operations (or *gates*) between two quantum bits (or *qubits*). A physical machine typically provides a set of basis gates that include primitive 2-qubit (2Q) and 1-qubit (1Q) gates that can be implemented in a given technology. In superconducting technologies, the current state of the art is to implement the same 2-qubit *entangling* gate between every pair of qubits (typically an XX- or XY-type gate). Entangling gates, coupled with some 1Q gates, allow for universal quantum computation.

We propose a radical idea – allow the 2Q basis gate(s) to differ between every pair of qubits, selecting the best entangling gates that can be calibrated between given pairs of qubits. Then let the compiler synthesize the desired logical gates for quantum programs from the gates that the qubits prefer. We find that this approach can improve gate fidelities by more than 10X and improve program fidelities exponentially. We introduce a practical method for finding the best entangling gate between two qubits, as well as a method for synthesizing logical gate from these entanglers. Our proof of concept focuses on the efficient synthesis of the high-fidelity SWAP gate, the dominant gate in quantum programs on machines with near-neighbor interconnects, but we also ensure implementation of a universal gate set with good fidelity.

1. INTRODUCTION

Quantum computers have the potential to solve problems currently intractable for conventional computers [38], but current computations are limited by errors [33], particularly when interacting two qubits to perform a quantum gate operation. This is not surprising, as qubits are engineered to preserve quantum state and isolate from the environment, but a quantum operation is the moment in time where external control is applied from the environment to deliberately alter a qubit's state. To accomplish low-error gates, the control mechanisms are carefully designed and the control signals are carefully calibrated for each qubit or pair of qubits.

Similar to how classical computers use a small set of classical logic gates (AND, OR, NOT, XOR...) as building blocks for larger circuits, current superconducting quantum devices typically only directly support a universal gate set consisting of a few 2Q gates and a continuous set of single-qubit (1Q) gates. This paper will refer to the set of directly supported quantum gates as *basis gates*. In the space of 2Q gates (see Fig.1), any 2Q gate that is not locally equivalent to SWAP or Identity at least partially entangles the two qubits that they act on. Any of these 2Q entangling gates can achieve universal computation when added to 1Q gates [5].

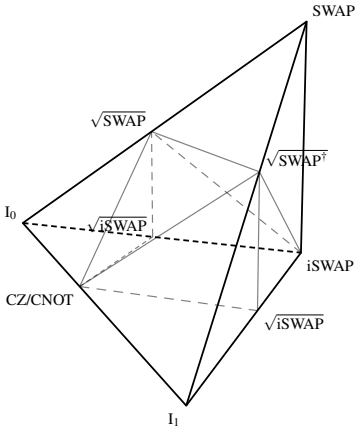

Figure 1: The Weyl chamber of 2Q quantum gates. The non-local part of a 2Q gate is fully described by its position in the Weyl chamber. CNOT and CZ are both represented by $(\frac{1}{2},0,0)$. The SWAP gate is at the top vertex $(\frac{1}{2},\frac{1}{2},\frac{1}{2})$. On the bottom surface, $(t_x,t_y,0)$ and $(1-t_x,t_y,0)$ represent the same equivalent class of gates. For example, the two points $I_0=(0,0,0)$ and $I_1=(1,0,0)$ both represent the 2Q identity gate I.

Universality is however not sufficient for a good set of basis gates. For example, in classical circuits while the NAND is universal, building circuits from it alone is less efficient than using a larger set of logic gates. On a quantum computer, all 2Q basis gates need to be routinely calibrated to maintain high fidelity in the presence of device variations over time. So it is impractical to support a large set of 2Q basis gate. The 2Q gates outside of a device's calibrated gate set have to be decomposed by the compiler into alternating layers of 1Q and 2Q basis gates. Thus, the choice of the few 2Q gates to directly support is important to device performance. On the hardware side, the 2Q basis gates must have high-fidelity hardware implementations. On the software side, they must enable low-depth decomposition of other 2Q gates.

The superconducting transmons support XX- and XY-type interactions [1, 18]. The strength of each of these interactions depend on the type of coupling, coupling strength and frequency detunings of the circuit elements [18]. A useful visualization is the Weyl chamber space of 2Q gates (Fig. 1), where the coordinates of a gate correspond to its non-local part in Cartan's KAK decomposition (see Section 2.2). In the Weyl chamber, gates in the XX family form a straight

trajectory from Identity to CNOT/CZ, while gates in the XY family form a trajectory from Identity to iSWAP. Traditional choices of 2Q basis gates like CNOT and iSWAP are already reasonably good at decomposing other 2Q gates. Three applications of either CNOT or iSWAP, along with 1Q gates, are sufficient to synthesize any 2Q gate. However, the current 2Q basis gates on fixed-frequency transmons tend to have much longer gate duration and up to about 10-50x higher error rates than 1Q gates [14]. This speed constraint on the two qubit gates typically arises because stronger drives (required for faster gates) leads to a combination of XX- and XY-type interactions that increase the error. Therefore, we see a larger room for improvement by relaxing the constraints on the hardware side while accounting for the deviations in the software.

There has been a recent proposal to implement faster and higher-fidelity 2Q gates by allowing an XX or XY gate trajectory to deviate from its standard path [25]. Quantum devices typically have higher order terms that result in the deviation of the experimentally realized native gate trajectory from the expected gate trajectory. This deviation is particularly significant for fast gates enabled by large coupling or large drive strength [14, 23, 31]. The standard approach is to suppress these deviations by slowing the gate speed which increases the susceptibility to the decoherence error. The alternative, however, is to let each pair of qubits "choose" its own 2Q gate trajectory. These deviated trajectories are more "native" to the device.

The truly native gate trajectories can benefit from extremely short gate times and hence higher fidelity. But they no longer pass through traditional basis gates like iSWAP and CZ. Furthermore, on the same device, the native 2Q gate trajectory on each pair of qubits is different. Thus, it is a challenge to efficiently identify a good set of 2Q basis gates from each native trajectory on a device. Our work addresses this problem, both in the scenario where we assume full information of the "native" trajectories, and in the less ideal case where characterizing each point on the trajectory requires nontrivial effort.

What are our requirements for a selected set of 2Q basis gates? Following the principle of Amdahl's Law, we pay most attention to optimizing the synthesis of the SWAP gate because of its importance to communication within programs executing on superconducting devices. Due to the need to mitigate crosstalk and other hardware constraints, superconducting devices usually have the sparse connectivity of a grid lattice or a hexagonal lattice. Therefore, the compiler has to schedule a series of SWAP gates before it can interact two qubits that are not adjacent to each other. Although clever mapping from logical to physical qubits can result in a smaller number of inserted SWAP gates, we still observe a high proportion of SWAP gates in post-mapping quantum circuits. Besides efficient synthesis of the SWAP gate, we also require the selected sets of 2Q basis gates to synthesize other 2Q gates more efficiently than the baseline.

Our contributions include:

- Current choices of 2Q basis gates are along the straight paths from Identity to standard gates like iSWAP and CZ. Our work is the first to consider selecting 2Q basis gates from gate trajectories that deviate from the standard paths.
- In the 3D space of 2Q gates, we characterize the sets of gates that enable the synthesis of SWAP in 1, 2, and 3 layers, respectively. Our approach can be applied with other 2Q gates as synthesis target as well. We also characterize the volumes of 2Q gates that enable the synthesis of CNOT and B gates in 2 layers, respectively. This provides guidance in selecting 2Q basis gates.
- We propose a method to identify good 2Q basis gate from a native gate trajectory in the setting where we have limited information about the trajectory. After obtaining the positions of no more than $\sim$10 points on the trajectory, we can identify, for example, the shortest gate on the trajectory that can synthesize SWAP in 3 layers, or the shortest gate on the trajectory that can both synthesize SWAP in 3 layers and CNOT in 2 layers.
- To test our method, we simulate gate trajectories from far-detuned transmon qubits, select 2Q basis gates from them, and then test on a variety of benchmark circuits including BV [4], QAOA [9], the QFT adder [35], and the Cuccaro Adder [8].

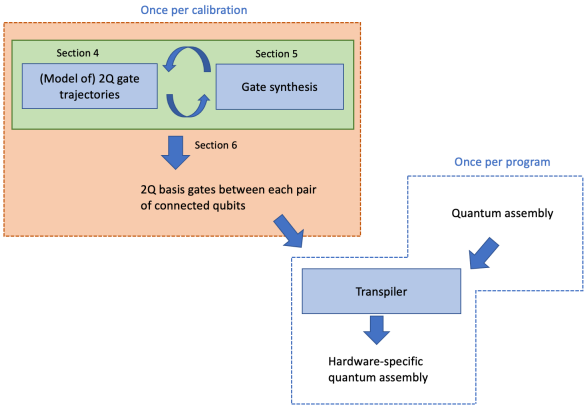

Figure 2: Overview of our work.

The paper is organized as follows (also see Fig. 2). Section 2 provides background on the theory of 2Q gates, including how they can be geometrically characterized as a Weyl chamber, and how their entangling power is quantified. Section 3 discusses related work. Section 4 describes the hardware that we simulate in this project to test our approach. Section 5 explains how we characterize the space of 2Q gates that we consider suitable for basis gates. Section 6 includes a discussion on calibration overhead, and our proposed method to identify good basis gates in the setting where we cannot afford to characterize too many points on each gate trajectory. Section 7 evaluates the performance gains. Finally, Section 8 and 9 discuss the implications and future directions.

2. BACKGROUND

2.1 Qubits and gates

Unlike a classical bit that is either 0 or 1, a quantum bit (qubit) can exist in a linear superposition of $|0\rangle$ and $|1\rangle$: $\alpha|0\rangle+\beta|1\rangle$ where α,β are complex amplitudes that satisfy $|\alpha|^2+|\beta|^2=1$. The state of one qubit can be represented by a 2-vector of the amplitudes α and β. A system of n qubits can exist in a superposition of up to 2^n basis states, and its state can be represented by a 2^n-vector of complex amplitudes. A quantum gate that acts on n qubits can be represented by a $2^n \times 2^n$ unitary matrix.

2.2 Geometric characterization of 2Q gates

Two $2Q$ quantum gates $U_1, U_2 \in SU(4)$ are *locally equivalent* if we can obtain one from the other by adding 1Q operations: $U_1 = k_1 U_2 k_2$ where $k_1, k_2 \in SU(2) \otimes SU(2)$. For example, CNOT and CZ are locally equivalent.

Any 2Q quantum gate $U \in SU(4)$ can be written in the form of

$$U = k_1 \exp(-i\frac{\pi}{2}(t_x X\otimes X + t_y Y\otimes Y + t_z Z\otimes Z))k_2 \qquad (1)$$

where X, Y, Z are the Pauli gates. This is called the Cartan decomposition.

The space of two-qubit quantum gates can be represented geometrically in a Weyl chamber (Fig. 1), where each point stands for a set of gates that are locally equivalent to each other [41]. The Cartan coordinates (t_x, t_y, t_z) in Eq. (1) are the coordinates of U in the Weyl chamber. They fully characterize the *non-local* part of a 2Q gate. On the bottom surface, $(t_x, t_y, 0)$ and $(1-t_x, t_y, 0)$ represent the same equivalent class of gates. The other points in the Weyl chamber each represents a different equivalence class of 2Q gates. We refer the interested readers to [7] for a more thorough introduction to the Weyl chamber. Note that other conventions of the Cartan coordinates are also common. They usually differ from ours by a constant factor of π or 2π.

In this paper, when we talk about some gate G in the Weyl chamber, we usually mean the local equivalence class of 2Q gates that includes G.

2.3 Entangling power of 2Q gates

The entangling power [40] is a widely accepted quantitative measure of the capacity of a 2Q gate to entangle the qubits that it acts on, and an indicator for the ability of a 2Q gate to synthesize arbitrary 2Q gates. For a unitary operator U, the entangling power $e_p(U) \in [0, \frac{2}{9}]$ is defined as the average linear entropy of the states produced by U acting on the manifold of all separable states [40]. It is solely based on the non-local part of U, which is characterized by the position of U in the Weyl chamber.

A 2Q gate has 0 entangling power if and only if it is locally equivalent to the Identity or the SWAP gate.

A 2Q gate U is called a *perfect entangler* if it can produce a maximally entangled state from an unentangled one [41]. Perfect entanglers have entangling power no less than $\frac{1}{6}$. They make up a polyhedron in the Weyl chamber that is exactly half of the total volume. The 6 vertices of the PE polyhedron are CZ(CNOT), iSWAP, $\sqrt{SWAP}$, $\sqrt{SWAP}^\dagger$, and the 2 points that both represent $\sqrt{iSWAP}$. The perfect entanglers with maximal entangling power of $\frac{2}{9}$ are also called *special perfect entanglers* [34]. In the Weyl chamber, they are on the line segment from CNOT to iSWAP. The B gate, which is at the midpoint of this line segment, has the property that it can synthesize any arbitrary 2Q gates within 2 layers [42]. However, there has been no proposal to directly implement the B gate in hardware.

3. RELATED WORK

To the best of our knowledge, no prior work studies the choice of 2Q basis gates from non-standard gate trajectories.

The most relevant to this project are recent works that analyze which small sets of 2Q basis gates from standard 2Q gate trajectories can provide the most performance gain. Murali et al. [19] test the performance of different points from the 2Q gate sets of Google's (fSim gates) and Rigetti's (XY gates) hardware on benchmark circuits by numerically decomposition, with the overall circuit success rate as the objective. Peterson el al. [29] use analytic techniques to study which small subset of the XX gates achieve the best expected fidelity when implementing random 2Q gates. They find that the gate set consisting CNOT, its square root, and its cubic root is almost as good as the entire set of XX gates in implementing random 2Q gates. Huang et al. [13] proposes using the square root of iSWAP as 2Q basis gate, instead of using iSWAP or CNOT.

4. SIMULATION OF A NATIVE GATE TRAJECTORY

A variety of artificial atoms can be constructed out of superconducting circuits where any two quantum levels can be used to realize a qubit. The key requirements for an ideal qubit are high coherence times, crosstalk-free coupling to other qubits and fast control. The ubiquitous qubit in the superconducting architecture called transmon can be modelled as an anharmonic oscillator. Capacitively coupling two transmons either directly or via another coupler element, typically leads to ZZ crosstalk. ZZ crosstalk is an always ON interaction which shifts the frequency of one qubit based on the state of the other qubit. This error is not corrected by most quantum error correction schemes including the surface code and leads to severe decrease in the circuit fidelity. Mitigation of this static ZZ crosstalk has been experimentally demonstrated by using carefully engineered couplers or microwave drives [15, 17, 24, 27, 39, 43]

Current superconducting NISQ devices either use tunable qubits for fast gates with poor coherence or fixed frequency qubits for slow gates but long coherence. Recently, there has been a new proposal [25] that demonstrated a 13ns perfect entangler between two fixed frequency transmons detuned by 2 GHz while simultaneously suppressing ZZ crosstalk. The high qubit detuning further addresses frequency crowding which reduces the yield of devices with increasing number of qubits. To exemplify our protocol, we consider the Hamiltonian for the new architecture in Ref. [25], but reiterate that the protocol can be applied to any quantum system.

The device image and schematic can be seen in Fig.3. The

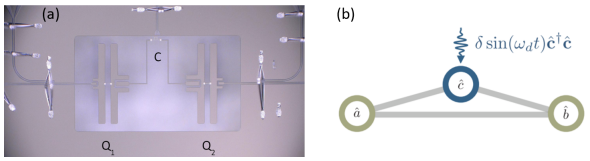

Figure 3: (a) Optical image of the device presented in [25] shows two fixed frequency transmons coupled via a tunable coupler. (b) Schematic for modelling the device adapted from [31]

full system Hamiltonian of the two qubits coupled with a tunable coupler is shown in Appendix A. Here, we highlight the time-dependent term, $\hat{H}_c(t)$, that describes the coupler dynamics-

$$\hat{H}_c(t) = \omega_c(t)\hat{c}^\dagger\hat{c} + \frac{\alpha_c}{2}\hat{c}^{\dagger 2}\hat{c}^2 \tag{2}$$

where α_c is the coupler anharmonicity, $\hat{c}$ is the annihilation operator and the coupler frequency $\omega_c(t)$, corresponding to the transition to it's first excited, can be varied in time via the flux through it's superconducting loop. We realize two qubit gates by either AC modulating or DC biasing of this coupler frequency.

Before generating two qubit gate trajectories, we need to choose the DC bias of the coupler frequency. We bias the coupler frequency (ω_c^0) between the two qubit frequencies (ω_a, ω_b) such that the static ZZ term (i.e. for $\delta(t) = 0$) between the two qubits is tuned to zero. This is typically achieved when α_c is positive.

This architecture natively supports strong parametrically activated interactions between the two qubits. To generate these native gate trajectories, we use a flux drive with carrier frequency $\omega_d = |\omega_a - \omega_b|$ and modulated by a flat-topped Gaussian envelope. We evolve the time-dependent Hamiltonian and project the evolution propagator on the computational subspace to obtain the effective unitary operation w.r.t the drive duration. This time ordered sequence of unitary operations can be represented as a trajectory in the Weyl space using Cartan coordinates. By examining the trace of the effective unitary propagator we can obtain the leakage outside the computational space. The observed leakage rates are much below the expected gate errors due to decoherence. Ideally the above parametric modulation should result in an iSWAP interaction between the two qubits. We note that using strong drives to realize fast entangling gates often leads to dynamical ZZ term that results in deviations from the ideal gate trajectories in the Weyl chamber. These deviations despite being completely coherent are traditionally categorized as errors. Our scheme unlocks the use of this coherent resource, thus enabling fast entangling native operations.

Cross-resonance gate, which is typically used for entangling fixed frequency transmons, is not feasible due to the high frequency detuning of 2GHz. So to provide a baseline for comparing to the native trajectory, we also simulate a direct CZ gate by applying a DC flux pulse to the coupler (i.e. $\omega_d = 0$). Note that because of it being a pure ZZ interaction, the trajectory is along the line connecting I_0 and CZ (or equivalently, from I_1 to CZ).

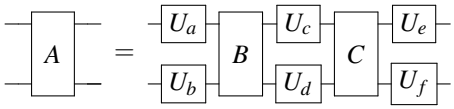

Figure 4: Gate A, decomposed into 2 layers with 2Q gates B, C and 1Q gates a,b,c,d,e,f

5. SYNTHESIS OF 2Q QUANTUM GATES

5.1 Fidelity of a synthesized gate

If a 2Q quantum gate is not directly supported on a device, it needs to be implemented by alternating layers of 1Q and 2Q gates from the set of basis gates that are directly supported. See Fig.4 and 12 for examples. We say that a decomposition is n-layer if it contains n layers of 2Q gates. Besides the errors that come from noises in the quantum hardware, a synthesized gate also suffers from the approximation error in gate decomposition. Thus the total fidelity of a gate should be the product of the hardware-limited fidelity and the decomposition fidelity. In this work, the decomposition errors are negligible compared to the hardware errors.

In our error model, decoherence is the dominant source of hardware error. So two factors determine whether a 2Q gate set is ideal for synthesizing a target gate: the duration of the basis gates, and the depth of the decomposition circuit. We need to take both into account when deciding on a strategy for selecting basis gates.

5.2 An analytic method for determining 2Q circuit depth

When deciding whether a potential basis gate is ideal for synthesizing a target gate, we consider the depth of the decomposition circuit as one of the factors. Given a 2Q target gate A, and a 2Q gate B (or a gate set S), how to determine the minimum circuit depth required for a decomposition of A into B (or S) and 1Q gates? One can take a practical, numerical approach to finding this decomposition. For a given number of layers, one can fix the 2Q gates and then numerically search for the 1Q gates that can minimize the discrepancy between the target unitary and the synthesized gate. One can start the numerical search from 1 layer, and increment the number of layers until the decomposition error gets below a threshold. But a more efficient and accurate way to determine the circuit depth is to apply the analytic method developed by Peterson et al. [30].

Without going into the technical details, here we summarize a key result from [34] that we adapt and apply in Section 5.3 and 5.4.

THEOREM 5.1. *There exists a 2-layer decomposition of 2Q gate A into B, C, and 1Q gates as in Fig. 4, if and only if any of the 1 to 8 sets of 72 inequalities that depend on the non-local parts of A, B, C is all satisfied.*

For details of the theorem, the readers can look at Theorem 23 of [30] or the implementation of the function in

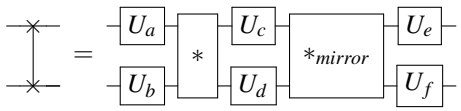

Figure 5: A general 2-layer decomposition of the SWAP gate. Here $*, *_{mirror}$ can be replaced by any pair of 2Q gates capable of synthesizing a SWAP in 2 layers.

our code [1]. Note that Reference [30] characterizes the space of 2q gates with LogSpec instead of the Cartan coordinates. Both are valid ways to represent the non-local part of a 2Q gate, but care must be taken when converting between the two. A gate U usually maps to 1 point in the Weyl chamber, but it usually maps to 2 points in the LogSpec space: $LogSpec(U) = (a,b,c,d)$ and $\rho(LogSpec(U)) = (c+\frac{1}{2}, d+\frac{1}{2}, a-\frac{1}{2}, b-\frac{1}{2})$. If $LogSpec(U) = \rho(LogSpec(U))$ for all A, B, and C, we only need to check one set of inequalities. If $LogSpec(U) \neq LogSpec(U)$ for 1, 2, or all 3 of A, B, and C, we need to plug in different versions of the LogSpec and check 2, 4, or 8 versions of the 72 inequalities, respectively.

5.3 Synthesis of the SWAP gate

On bounded connectivity architectures, SWAPs make up a significant portion of all two-qubit gates. A SWAP gate exchanges the quantum states of two neighboring qubits. A 2Q gate in a quantum program can be directly scheduled if it acts on two physical qubits that are connected to each other, but this is not the case in general. Superconducting devices are usually designed to have sparse connectivity, because otherwise crosstalk errors would be difficult to suppress. As a result, quantum programs usually contain a large proportion of SWAP gates after they are compiled to run on a superconducting device. In Section 7, we show the proportion of SWAP gates in the post-compilation benchmark circuits.

When we select the 2Q basis gate set for each pair of qubits, a top priority is to optimize the fidelity of the SWAP gate that is built from the gate set. In this section we address the following problem. Suppose we have full knowledge of $S \subset SU(4)$, a possibly infinite set of 2Q basis gates accessible to us, and the hardware errors of all the gates in S. Then which 2Q gate(s) from S will enable the optimal synthesis of the SWAP gate? This is an easier problem than what we face in practice. In a more realistic setting, we may not have full information of S and the gate errors.

We discuss 3 approaches towards synthesizing a SWAP gate: decompose it into 1, 2, or 3 layers of hardware 2Q gates.

SWAP in 1 layer: This requires a basis gate that is locally equivalent to SWAP. In other words, the trajectory of the available native gates needs to pass through the top vertex of the Weyl chamber.

SWAP in 2 layers: We consider 2 cases: 2-layer decomposition of SWAP using a single 2Q basis gate, and using two different 2Q basis gates.

In the first case, the set of 2Q gates that are capable of synthesizing SWAP in 2 layers are represented by 2 line segments in the Weyl chamber as shown in Fig.6. One is from the B gate to $\sqrt{SWAP}$ and the other is from B to $\sqrt{SWAP}^\dagger$. We denote them by L_0 and L_1, respectively.

[1] An anonymous link to our code: https://anonymous.4open.science/r/2q_gate_synthesis-42D4/

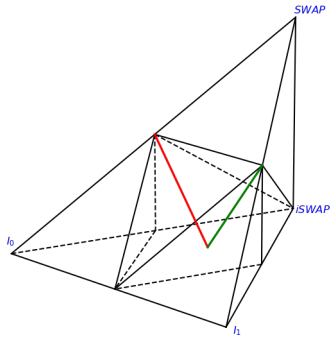

Figure 6: Gates that are able to synthesize SWAP in 2 layers form 2 line segments in the Weyl chamber. The red one is from the B gate to $\sqrt{SWAP}$, and the green one is from the B gate to $\sqrt{SWAP}^\dagger$.

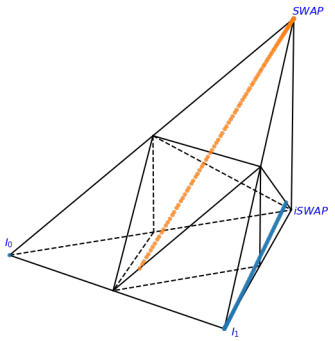

Figure 7: Pairs of gates that are able to synthesize a SWAP in 2 layers. In blue is an example trajectory that deviates from the standard XY interaction, in orange are the points that would complement the blue ones in synthesizing a SWAP in 2 layers.

In the second case, for each point $*$ in the Weyl chamber, we can find exactly one point $*_{mirror}$ such that they together enable a 2-layer decomposition of SWAP (see Fig. 5). The line segment from $*$ to $*_{mirror}$ always has one of L_0, L_1 as its perpendicular bisector. Thus, given $*$, we can locate $*_{mirror}$ by rotating $*$ by π around the closer one of L_0, L_1. One example pair of such points is CNOT and iSWAP. For a trajectory that deviates from the standard XY trajectory (goes from Identity to a point near iSWAP), its "mirror" is a trajectory from SWAP to a point near CNOT (Fig.7). Since there's no overlap between the example trajectory and the "mirror", we conclude that the trajectory does not contain any pair of points that is able to synthesize SWAP together in 2 layers.

In Appendix B we explain how to derive the results above.

SWAP in 3 layers: It is a well-known result that 3 invocations of CNOT are required to implement a SWAP [37]. We show the circuit in Fig. 8. In fact, CNOT and iSWAP share the property that they can synthesize any arbitrary 2Q

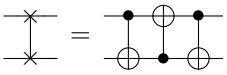

Figure 8: The SWAP gate, decomposed into 3 CNOT gates.

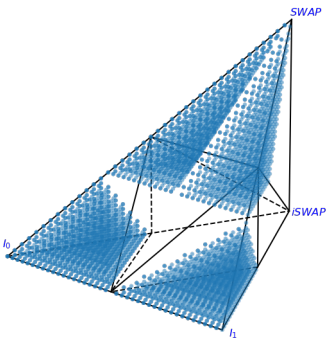

Figure 9: Gates that are NOT able to synthesize a SWAP in 3 layers.

gate in 3 layers but only a 0-volume set of gates (in the Weyl chamber) in 2 layers [30].

For our purpose, we need to know what other gates are capable of decomposing SWAP in 3 layers. We only consider 3-layer decomposition of SWAP using a single 2Q basis gate as in Fig. 12. Let $S_{SWAP,3}$ denote the set of gates that satisfy our requirement. To determine whether a 2Q basis gate G is in $S_{SWAP,3}$, we first locate the corresponding G_{mirror} such that G and G_{mirror} together can provide a 2-layer decomposition of SWAP. Then we apply Theorem 5.1 with G_{mirror} as target and G as basis gate to check if there exists a 2-layer decomposition of G_{mirror} into G.

We apply the method above to a sample of points in the Weyl chamber, and obtain the distribution of gates that are able to synthesize SWAP in 3 layers. Since the complement of the set has a simpler shape, here we show a plot of $\overline{S_{SWAP,3}}$, the points that are not able to synthesize SWAP in 3 layers, in Fig. 9. A visual inspection tells us $\overline{S_{SWAP,3}}$ consists of 4 tetrahedra in the Weyl chamber. After locating the vertices of the tetrahedra, we obtain Fig. 10. We also learn that the volume of $S_{SWAP,3}$ is 68.5% the volume of the Weyl chamber. In Fig. 11, we show how an example trajectory is divided into the part in $S_{SWAP,3}$ and the part in $\overline{S_{SWAP,3}}$.

A 2Q gate trajectory starts from either I_0 (or I_1) and goes out of the bottom left (or the bottom right) tetrahedron in Fig. 10. If the trajectory does not go directly to SWAP, it will enter $S_{SWAP,3}$ after leaving the bottom tetrahedron that it starts from. Thus, the fastest gate on the trajectory that synthesizes SWAP in 3 layers can be found by locating the intersection of the trajectory with the face $\{CZ, (\frac{1}{4},\frac{1}{4},0), (\frac{1}{6},\frac{1}{6},\frac{1}{6})\}$ or $\{CZ, (\frac{3}{4},\frac{1}{4},0), (\frac{5}{6},\frac{1}{6},\frac{1}{6})\}$.

Summary: Given a 2Q gate trajectory that deviates from XY or XX, the most suitable 2Q gate for SWAP synthesis is the fastest one on the trajectory that is capable of synthesizing SWAP in 3 layers. Although some gates in the Weyl chamber are able to synthesize SWAP in 1 or 2 layers, it is unlikely

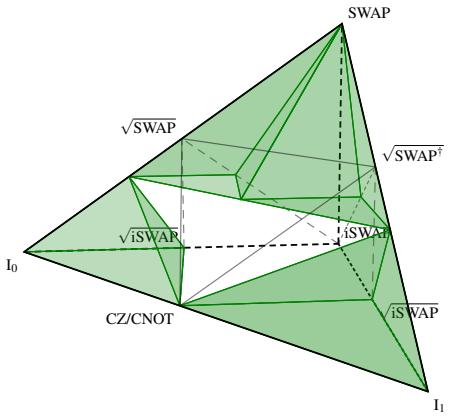

Figure 10: Gates that are NOT able to synthesize a SWAP in 3 layers. The 4 tetrahedra are defined by vertices $\{I_0, CZ, (\frac{1}{4},\frac{1}{4},0), (\frac{1}{6},\frac{1}{6},\frac{1}{6})\}$, $\{CZ, I_1, (\frac{3}{4},\frac{1}{4},0), (\frac{5}{6},\frac{1}{6},\frac{1}{6})\}$, $\{SWAP, (\frac{1}{2},\frac{1}{6},\frac{1}{6}), (\frac{1}{6},\frac{1}{6},\frac{1}{6}), (\frac{1}{3},\frac{1}{3},\frac{1}{6})\}$, and $\{SWAP, (\frac{1}{2},\frac{1}{6},\frac{1}{6}), (\frac{5}{6},\frac{1}{6},\frac{1}{6}), (\frac{2}{3},\frac{1}{3},\frac{1}{6})\}$.

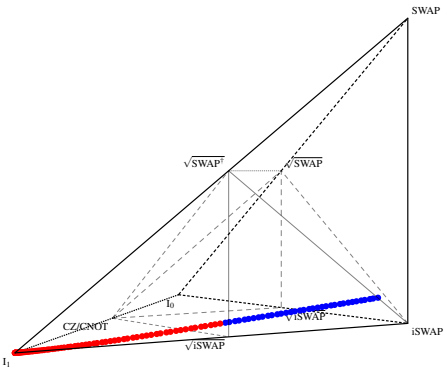

Figure 11: An example 2Q gate trajectory. In blue are the gates that are able to synthesize SWAP in 3 layers, in red are the gates that unable to do so.

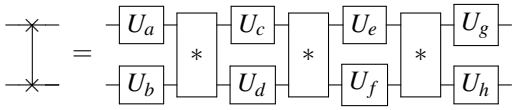

Figure 12: A general 3-layer decomposition of the SWAP gate. Here the $*$ can be replaced by any 2Q gate capable of synthesizing a SWAP in 3 layers.

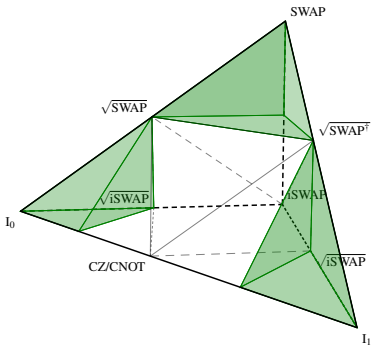

Figure 13: Gates that are NOT able to synthesize CNOT in 2 layers. The 3 tetrahedra in the plot are defined by vertices $\{I_0, (\frac{1}{4},0,0), (\frac{1}{4},\frac{1}{4},\frac{1}{4}), \sqrt{SWAP}\}$, $\{I_1, (\frac{3}{4},0,0), (\frac{3}{4},\frac{1}{4},0), \sqrt{SWAP}^\dagger\}$, and $\{SWAP, \sqrt{SWAP}, \sqrt{SWAP}^\dagger, (\frac{1}{2},\frac{1}{2},\frac{1}{4})\}$.

that the early part of the trajectory overlaps any of them.

5.4 Synthesis of other gates

The techniques that we use to study the synthesis of SWAP also applies to other 2Q gates. For example, by applying Theorem 5.1 to a sample of points in the Weyl chamber, with CNOT as target, we learn that the gates that are able to synthesize CNOT in 2 layers takes up 75% of the volume in the Weyl chamber. The complement $\overline{S_{CNOT,2}}$ consists of 3 tetrahedra, as shown in Fig. 13. Therefore, on a 2Q gate trajectory, we can locate the fastest gate that synthesizes CNOT in 2 layers by taking the intersection of the trajectory with the face $\{(\frac{1}{4},0,0), (\frac{1}{4},\frac{1}{4},\frac{1}{4}), \sqrt{SWAP}\}$ or $\{(\frac{3}{4},0,0), (\frac{3}{4},\frac{1}{4},0), \sqrt{SWAP}^\dagger\}$. We can also locate the fastest gate from the trajectory that can both synthesize CNOT in 2 layers and synthesize SWAP in 3 layers, by taking the first intersection of the trajectory with $S_{CNOT,2} \cap S_{SWAP,3}$.

6. SELECTING BASIS GATES IN A REALISTIC SETTING

In the previous section, we discussed the synthesis of the SWAP gate and other gates on a theoretical level. In an ideal setting where we have full information about a native 2Q trajectory, we will be able to immediately identify basis gate(s) that gives us the highest fidelity implementation of SWAP or other target gates. However, in practice, locating even 1 point on the trajectory requires a nontrivial amount of work so we can only afford to characterize a few points on each trajectory. In this section we review how to experimentally characterize an unknown 2Q gate, present an algorithm for selecting basis gate sets in the presence of calibration overhead, and discuss how the calibration overhead of our approach compares to the standard approach.

6.1 Characterization of an unknown 2Q gate

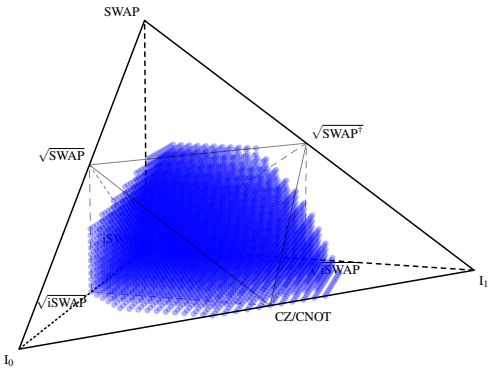

Figure 14: temp

Let t denotes the gate duration, and let $w(t)$ represent the Weyl coordinates of the gate on the trajectory with duration t. On a quantum device, given a value of t it is not hard to implement the gate corresponding to $w(t)$ if the pulse shape is a known function of t. However, characterizing the gate requires quantum process tomography (QPT). QPT [6,32] is a well-established technique that fully characterizes a quantum gate by computing the process matrix χ. For a 2Q gate G, $\chi(G)$ is a 16 by 16 matrix where each entry is defined as the probability

$$p_{ij} = \langle\langle E_j|G|\rho_i\rangle\rangle. \tag{3}$$

In Equation 3, E_j $(j = 1,\ldots,16)$ and $|\rho_i\rangle$ $(i = 1,\ldots,16)$ can be any complete bases for the Hilbert-Schmidt spaces they live in. Each probability can be estimated by taking the average of a set of measurement results.

Recently, gate set tomography (GST) [11] [26] has been proposed to isolate the state preparation and measurement (SPAM) errors in QPT, and selected quantum process tomography (SQPT) has been developed to compute selected elements in a process matrix [3,10,28].

6.2 Our strategy

In this work, we propose and test two criteria for selecting 2Q basis gates from native 2Q trajectories.

1. Select the fastest gate on the trajectory that can synthesize SWAP in 3 layers.

2. Select the fastest gate on the trajectory that can both synthesize SWAP in 3 layers and synthesize CNOT in 2 layers.

As explained in Section 5.3, the gate that meets Criterion 1 can be found at the intersection of the 2Q trajectory and one of the 2 faces $\{CZ,(\frac{1}{4},\frac{1}{4},0),(\frac{1}{6},\frac{1}{6},\frac{1}{6})\}$ and $\{CZ,(\frac{3}{4},\frac{1}{4},0),(\frac{5}{6},\frac{1}{6},\frac{1}{6})\}$. And as explained in Section 5.4, the gate that meets Criterion 2 can be found similarly. With this insight, we can locate the desired 2Q basis gate after characterizing a small number of points on the trajectory.

The algorithm for locating the desired gate under Criterion 1 is as follows. With a tolerance threshold t_{th}, we want to find t_{int} such that $w(t_{int})$ is the intersection of the trajectory and one of the two faces. Let $w(t_0)$ be the first point that we characterize on the trajectory. By taking the dot product of $w(t_0)$ and the normal vectors of the two faces, we can determine whether $w(t_0)$ is inside one of the bottom tetrahedra in Fig. 10. If $w(t_0)$ is inside one of the two tetrahedra, then $t_0 < t_{int}$. In this case we let $t_1 = 2t_0$ and characterize $w(t_1)$. We increase t until finding t_n such that $w(t_n)$ is outside both of the two tetrahedra. Then we know t_{int} is between t_{n-1} and t_n. We can perform a binary search in this interval to get within the tolerance threshold. If $w(t_0)$ is outside both of the two tetrahedra, the procedure is similar except we decrease t until $w(t_n)$ is inside one of the two tetrahedra.

The strategy for Criterion 2 is similar. After characterizing each point $w(t_i)$, instead of only checking whether it is inside one of the two bottom tetrahedra in Fig. 10, we also check whether it is inside one of the two bottom tetrahedra in Fig. 13. We want to locate the first point on the trajectory that is outside of all 4 tetrahedra.

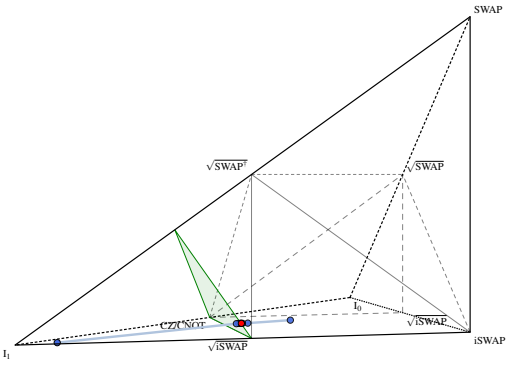

Figure 15: Locating a good 2Q basis gate on an example trajectory. The 1 red circle is the 2Q gate that we choose from this trajectory using Criterion 1. The 8 blue circles are the other gates we characterize in the process. The 7 circles (1 red and 6 blue) in the middle are close to each other in the figure.

In Fig. 15 we show an example of applying Criterion 1. In this example, 9 points are characterized in total, with $t = 10, 20, 30, 25, 22.5, 23.75, 23.125, 23.4375, 23.28125$ ns. The selected point is at $t_8 = 23.4375$ns.

Our strategy can be easily adapted to other criteria for selecting basis gates. For instance, we can locate the fastest perfect entangler on a 2Q trajectory, by finding the intersection of the trajectory with the faces $\{CZ,\sqrt{SWAP},(\frac{1}{4},\frac{1}{4},0)\}$ and $\{CZ,\sqrt{SWAP}^\dagger,(\frac{3}{4},\frac{1}{4},0)\}$. Combining the PE criterion with Criterion 1, we can locate the fastest gate on the trajectory that is both a PE and can synthesize SWAP in 3 layers. We can also try using the first local maximum of the entangling power on each trajectory as basis gate. These alternative criteria can be explored in future work.

6.3 Calibration overhead

Selecting a 2Q basis gate from a native trajectory requires more calibration time than calibrating a standard gate from a standard trajectory. The reason is two-fold. First, characterizing an unknown gate is harder than assessing the fidelity of a known gate, especially a standard gate like CNOT and iSWAP. Second, identifying the desired 2Q basis gate from a native trajectory requires characterizing multiple points.

Experimentally characterizing an unknown gate requires QPT, while the fidelity of a known gate can be assessed by randomized benchmarking (RB). Both QPT and RB [16,21, 36] are well-established techniques that vendors like IBM and Rigetti use to measure the fidelity of their gates. Either QPT or RB can be applied to multiple 2Q gates on the same device simultaneously, as long as they do not act on the same qubits. So the calibration overhead does not scale with the size of the device. Although QPT scales worse for multi-qubit systems, characterizing 2Q gates only involves systems of 2 qubits so

QPT is still viable. As explained in Section 6.1, applying QPT on a 2Q gate requires 16×16 sets of experiments. Each experiment circuit consists of a pre-rotation that prepares the input, applying the 2Q gate to the input, post-rotation into the desired basis, and measurement. Compared to QPT, RB requires fewer sets of experiments but each experiment has a longer circuit. RB measures the average gate fidelity in a system by running sequences of random Clifford gates that multiply to the Identity gate. The average gate fidelity is extracted from the plot of the success rates (of measuring the expected ground state) versus lengths of sequence. The fidelity of a Clifford gate like CNOT can be experimentally assessed by interleaved randomized benchmarking (iRB) [22], an extension of RB in which the gate of interest is interleaved throughout a sequence of random Clifford gates. To assess the fidelity of a non-Clifford gate like the square root of iSWAP, we need interleaved fully randomized benchmarking (iFRB) [13] where the random Clifford gates in iRB are replaced by Haar-random gates.

Our approach to selecting 2Q basis gates from a native trajectory involves characterizing multiple points, which also increases the calibration overhead. The number of points we need to characterize on each trajectory depends on the precision threshold we choose and the amount of information we learn from previous calibration cycles. On the 180 gate trajectories that we simulate (see Section 7 for details), we need to characterize an average of 8.9 (10.0) points on each trajectory when using Criterion 1 (Criterion 2) and a precision threshold of 0.2ns. One can reduce the calibration overhead by relaxing the precision threshold. Another way to lower the overhead is to estimate the duration of the desired basis gate on each trajectory using information from prior calibration cycles. With this information, the search of a desired basis gate can be simplified to conducting a binary search in a potentially narrow range. The extent to which we can reduce the calibration overhead is a topic for future work.

7. EVALUATION

7.1 Methodology

We simulate a 10 by 10 device with grid connectivity (Fig. 16), where the qubit frequencies are sampled from two normal distributions with means that differ by 2GHz. The high and low frequency qubits are shown in different colors. Each edge connects two qubits with different colors. We use a 5% standard deviation for sampling the qubit frequencies. Experimentalists have demonstrated a smaller standard deviation of about 0.5% [12], but we want to show that our method still works when there is a wider distribution in qubit frequency.

Between each pair of neighboring qubits on the 10×10 grid, we simulate a "native" 2Q trajectory and a direct CZ gate as baseline according to Section 4. Then on each "native" trajectory, we select 2Q basis gates using Criterion 1 and 2 (respectively) according to Section 6.2. We test these 3 sets of 2Q basis gates on common application circuits as benchmarks. We use the Qiskit [2] transpiler with the "SABRE" [20] layout and routing methods to map the benchmarks circuits to the 10×10 grid connectivity. Then the transpiler keeps the

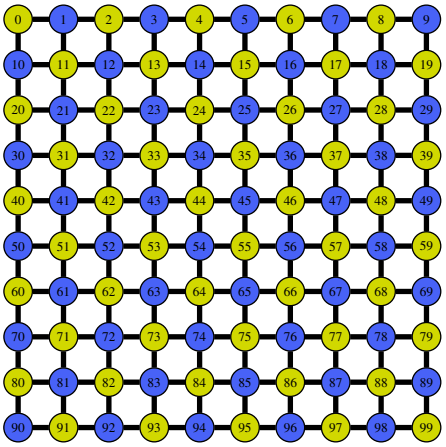

Figure 16: device simulation

SWAP gates and decompose all other 2Q gates (if any) to CNOT. Finally, we synthesize the SWAP and CNOT gates using each set of 2Q basis gates. This is a simplistic approach that favors the baseline since it directly implements a local equivalence of CNOT. The alternative of directly decomposing all gates into the basis gates takes a longer compilation time but would result in a larger gain for the "native" trajectories.

Decoherence is the dominant hardware noise in our noise model, because crosstalk is suppressed by the high detuning in the qubits. For each qubit, we model the decoherence error as $1 - e^{-t/T}$, where t is the *lifetime* of the qubit in a program and T is the coherence time of the qubit. We set T to a typical value of 8×10^{-5} s for all qubits. We compute t as $t_f - t_i$, where t_i is the start of the first gate on the qubit and t_f is the end of the last gate on the qubit. The total decoherence-limited fidelity of a circuit is the product over the $e^{-t/T}$ term from each qubit.

We also record the decomposition errors that come from approximations in circuit synthesis. However, they are negligible compared to the decoherence errors, and can be reduced to arbitrarily close to 0 in theory. Thus we do not include them in the results.

7.2 Results

First we synthesize SWAP and CNOT gates using the direct CZ gates from the baseline and the 2Q basis gates chosen using Criterion 1 and 2, respectively. The average durations of the synthesized SWAP and CNOT gates from the 3 approaches are summarized in Table 1. For the 1Q gates in the gate and circuit synthesis, we use 21.33 ns (a typical value on fixed-frequency transmons, taken from [14]) as the duration. The synthesized SWAP (CNOT) gates from Criterion 1 and 2 are 30.7x and 30.2x (10.3x and 14.3x) faster than the baseline, respectively. The differences in gate duration translate to 29.1x and 28.6x (10.1x and 14.1x) improvements in the decoherence errors of the individual gates.

In Fig.17,18,19,20,21, we show the success rates (circuit fidelities) of 5 sets of benchmark circuits, when transpiled to different sets of 2Q basis gates. The success rates are nor-

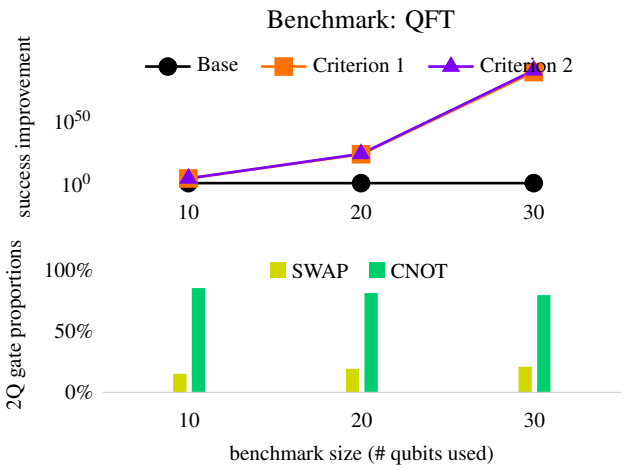

Figure 17: QFT Adder benchmarks

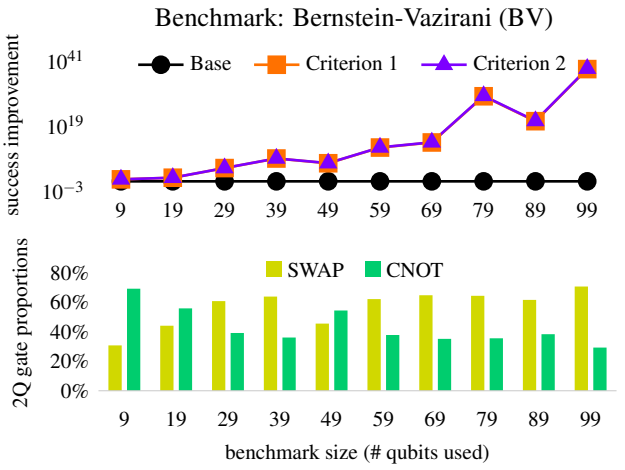

Figure 18: Bernstein-Vazirani (BV) benchmarks

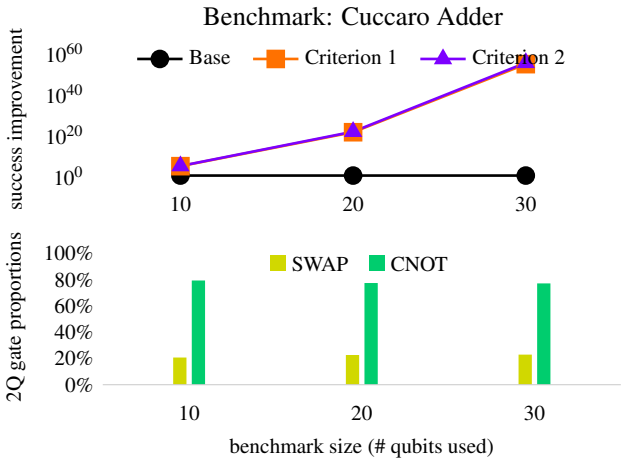

Figure 19: Cuccaro Adder benchmarks

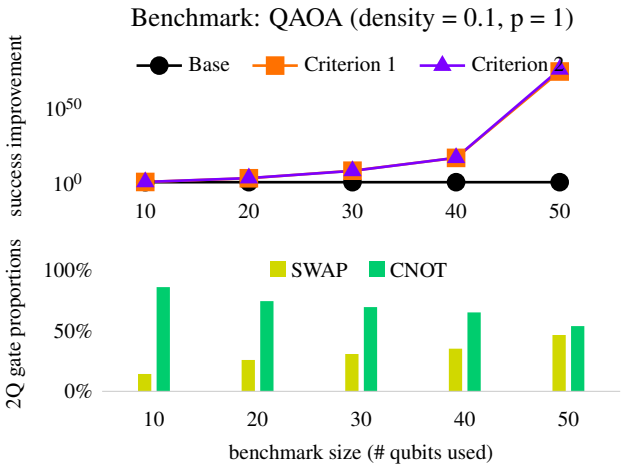

Figure 20: QAOA (density = 0.1, p = 1) benchmarks

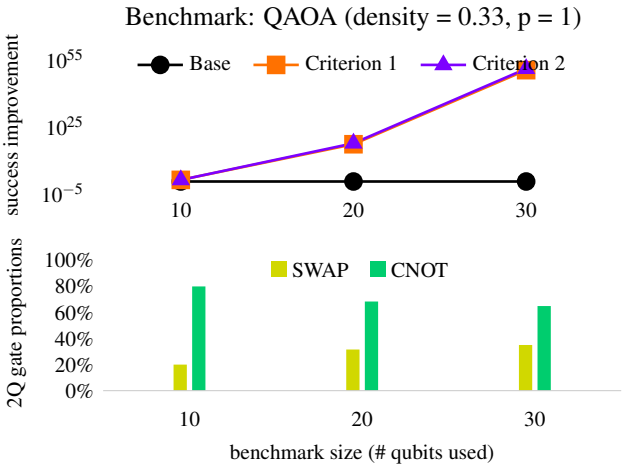

Figure 21: QAOA (density = 0.33, p = 1) benchmarks

	basis	SWAP	CNOT
Baseline	1607	4907	1650
Criterion 1	24.84	159.8	159.8
Criterion 2	25.66	162.3	115.3

Table 1: Average duration (in ns) of the 2Q basis gates and the synthesized SWAP and CNOT gates, from baseline, Criterion 1, and Criterion 2.

malized to the baseline and shown in log scale. Also shown in the graphs are the proportions of SWAP and CNOT gates in the post-compilation circuits. On each set of benchmarks, the improvement factors (of Criterion 1 and 2, relative to the baseline) grow exponentially in the size of the benchmark. On most of the benchmarks, Criterion 2 performs better since it has significantly faster CNOT gates and only sightly slower SWAP gates compared to Criterion 1.

8. DISCUSSION

It is not surprising that a 10x to 30x advantage in gate duration gets amplified to an exponential advantage in the overall circuit fidelity, given the relation between gate fidelity and circuit fidelity. Our results show that for fixed-frequency transmons, and perhaps also other technologies, we can improve the fidelity of 2Q gates if we do not treat the coherent deviations in gate trajectories as errors. A direction for future work is to experimentally demonstrate the selection of 2Q basis gates following the strategies in this work, and test the calibration overhead.

The direct CZ gates we use as baseline are slower than the 200ns direct CNOT gate recently demonstrated by [14], which is not supported on their devices yet. With a 200ns CNOT gate, one can synthesize a 600ns SWAP gate (Fig. 8). But the 200ns CNOT and 600ns SWAP gates are still slower than the SWAP and CNOT gates synthesized from the "native" trajectories. Also, note that while the gates we simulate do not suffer from leakage thanks to the high detuning and any coherent deviations (by the definition of native trajectory), the direct CNOT in [14] suffers from severe crosstalk in addition to the decoherence error. In order to keep the manuscript focused, we do not consider gates between tunable transmons.

9. CONCLUSION

The idea of a uniform set of basis gates naturally arose from early notions of universal gate sets, which experimentalists then adapted to their favorite technologies. By looking at the theory of possible entanglers, we have found that there are many options, and that these options behave differently on each pair of interacting qubits in a machine. This led us to a radically new idea, why be constrained to a single canonical gate (eg CX or CZ)? Why not tune up the gate that will have the highest fidelity between every pair of qubits, allowing each to differ and have adjust for these variations in software?

In this paper, we examined the space of possible entanglers and developed a method for practically finding a high-fidelity entangler between every pair of qubits. We find that these heterogeneous basis gates are more than 10X better in fidelity than traditional basis gates. We then evaluate these heterogeneous basis gates on a number of benchmark circuits and find fidelity improvements that scale exponentially in benchmark size. We believe that our approach is great example of letting the software make the hardware's job easier, and we hope that this kind of adaptive basis-gate design will lead to further innovation in future quantum systems.

APPENDIX

A. HAMILTONIAN OF 2 QUBITS COUPLED WITH A TUNABLE COUPLER

The system Hamiltonian of the two qubits coupled with a tunable coupler can be modelled as [31]-

$$\hat{H}(t) = \hat{H}_a + \hat{H}_b + \hat{H}_c(t) + \hat{H}_g, \tag{4}$$

with

$$\begin{aligned}
\hat{H}_a &= \omega_a \hat{a}^\dagger \hat{a} + \frac{\alpha_a}{2} \hat{a}^{\dagger 2} \hat{a}^2, \\
\hat{H}_b &= \omega_b \hat{b}^\dagger \hat{b} + \frac{\alpha_b}{2} \hat{b}^{\dagger 2} \hat{b}^2, \\
\hat{H}_c(t) &= \omega_c(t) \hat{c}^\dagger \hat{c} + \frac{\alpha_c}{2} \hat{c}^{\dagger 2} \hat{c}^2. \\
\hat{H}_g &= -g_{ab} \hat{a}^\dagger \hat{b} - g_{bc} \hat{b}^\dagger \hat{c} - g_{ca} \hat{c}^\dagger \hat{a} \\
&\quad - g_{ab}^* \hat{a} \hat{b}^\dagger - g_{bc}^* \hat{b} \hat{c}^\dagger - g_{ca}^* \hat{c} \hat{a}^\dagger
\end{aligned} \tag{5}$$

where $\omega_{a(b)}$ corresponds to the qubit a(b) frequency, g_{ij} represents capacitive coupling strength between elements i and j. The entangling interaction is realized by modulating the coupler frequency as $\omega_c(t) = \omega_c^0 + \delta \sin(\omega_d t)$.

B. SWAP SYNTHESIS IN 2 LAYERS

See the circuit in Fig. 4. Let $A = SWAP$ we get the equation

$$SWAP = (e \otimes f) C (c \otimes d) B (a \otimes b).$$

Move $e \otimes f$ and $a \otimes b$ to the other side and move $e \otimes f$ through SWAP,

$$\begin{aligned}
C(c \otimes d)B &= (e \otimes f)^\dagger SWAP (a \otimes b)^\dagger \\
&= SWAP (f \otimes e)^\dagger (a \otimes b)^\dagger \\
&= SWAP (fa \otimes eb)^\dagger.
\end{aligned}$$

Move $(fa \otimes eb)^\dagger$ to the LHS, and C to the RHS,

$$(c \otimes d) B (fa \otimes eb) = C^\dagger SWAP.$$

This equation tells us that, B and C can synthesize SWAP as in Fig. 4 if and only if the Cartan coordinates of B are equal to the Cartan coordinates of $C^\dagger SWAP$ up to canonicalization. Let $B \sim (x, y, z)$ and $C \sim (x', y', z')$, then we have $(x, y, z) \sim (-x', -y', -z') + (\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$. From this we can tell that for every local equivalence class $[B]$ of 2Q gates, there

is exactly one local equivalence class $[C]$ such that $[B]$ and $[C]$ together can synthesize SWAP in 2 layers. And since we know how to canonicalize Cartan coordinates into points within the Weyl chamber, given $[B]$ we will be able to find the corresponding $[C]$. Here we do not elaborate on how we identify the geometric relation between $[B]$ and $[C]$ inside the Weyl chamber, but the readers can check our claim by applying Theorem 5.1.

REFERENCES

[1] D. M. Abrams, N. Didier, B. R. Johnson, M. P. da Silva, and C. A. Ryan, "Implementation of XY entangling gates with a single calibrated pulse," *Nature Electronics*, vol. 3, no. 12, pp. 744–750, Nov. 2020. [Online]. Available: https://doi.org/10.1038/s41928-020-00498-1

[2] M. S. ANIS, Abby-Mitchell, H. Abraham, AduOffei, R. Agarwal, G. Agliardi, M. Aharoni, I. Y. Akhalwaya, G. Aleksandrowicz, T. Alexander, M. Amy, S. Anagolum, Anthony-Gandon, E. Arbel, A. Asfaw, A. Athalye, A. Avkhadiev, C. Azaustre, P. BHOLE, A. Banerjee, S. Banerjee, W. Bang, A. Bansal, P. Barkoutsos, A. Barnawal, G. Barron, G. S. Barron, L. Bello, Y. Ben-Haim, M. C. Bennett, D. Bevenius, D. Bhatnagar, A. Bhobe, P. Bianchini, L. S. Bishop, C. Blank, S. Bolos, S. Bopardikar, S. Bosch, S. Brandhofer, Brandon, S. Bravyi, N. Bronn, Bryce-Fuller, D. Bucher, A. Burov, F. Cabrera, P. Calpin, L. Capelluto, J. Carballo, G. Carrascal, A. Carriker, I. Carvalho, A. Chen, C.-F. Chen, E. Chen, J. C. Chen, R. Chen, F. Chevallier, K. Chinda, R. Cholarajan, J. M. Chow, S. Churchill, CisterMoke, C. Claus, C. Clauss, C. Clothier, R. Cocking, R. Cocuzzo, J. Connor, F. Correa, Z. Crockett, A. J. Cross, A. W. Cross, S. Cross, J. Cruz-Benito, C. Culver, A. D. Córcoles-Gonzales, N. D, S. Dague, T. E. Dandachi, A. N. Dangwal, J. Daniel, M. Daniels, M. Dartiailh, A. R. Davila, F. Debouni, A. Dekusar, A. Deshmukh, M. Deshpande, D. Ding, J. Doi, E. M. Dow, E. Drechsler, E. Dumitrescu, K. Dumon, I. Duran, K. EL-Safty, E. Eastman, G. Eberle, A. Ebrahimi, P. Eendebak, D. Egger, ElePT, Emilio, A. Espiricueta, M. Everitt, D. Facoetti, Farida, P. M. Fernández, S. Ferracin, D. Ferrari, A. H. Ferrera, R. Fouilland, A. Frisch, A. Fuhrer, B. Fuller, M. GEORGE, J. Gacon, B. G. Gago, C. Gambella, J. M. Gambetta, A. Gammanpila, L. Garcia, T. Garg, S. Garion, J. R. Garrison, J. Garrison, T. Gates, L. Gil, A. Gilliam, A. Giridharan, J. Gomez-Mosquera, Gonzalo, S. de la Puente González, J. Gorzinski, I. Gould, D. Greenberg, D. Grinko, W. Guan, D. Guijo, J. A. Gunnels, H. Gupta, N. Gupta, J. M. Günther, M. Haglund, I. Haide, I. Hamamura, O. C. Hamido, F. Harkins, K. Hartman, A. Hasan, V. Havlicek, J. Hellmers, Ł. Herok, S. Hillmich, H. Horii, C. Howington, S. Hu, W. Hu, J. Huang, R. Huisman, H. Imai, T. Imamichi, K. Ishizaki, Ishwor, R. Iten, T. Itoko, A. Ivrii, A. Javadi, A. Javadi-Abhari, W. Javed, Q. Jianhua, M. Jivrajani, K. Johns, S. Johnstun, Jonathan-Shoemaker, JosDenmark, JoshDumo, J. Judge, T. Kachmann, A. Kale, N. Kanazawa, J. Kane, Kang-Bae, A. Kapila, A. Karazeev, P. Kassebaum, T. Kehrer, J. Kelso, S. Kelso, V. Khanderao, S. King, Y. Kobayashi, Kovi11Day, A. Kovyrshin, R. Krishnakumar, V. Krishnan, K. Krsulich, P. Kumkar, G. Kus, R. LaRose, E. Lacal, R. Lambert, H. Landa, J. Lapeyre, J. Latone, S. Lawrence, C. Lee, G. Li, J. Lishman, D. Liu, P. Liu, Lolcroc, A. K. M, L. Madden, Y. Maeng, S. Maheshkar, K. Majmudar, A. Malyshev, M. E. Mandouh, J. Manela, Manjula, J. Marecek, M. Marques, K. Marwaha, D. Maslov, P. Maszota, D. Mathews, A. Matsuo, F. Mazhandu, D. McClure, M. McElaney, C. McGarry, D. McKay, D. McPherson, S. Meesala, D. Meirom, C. Mendell, T. Metcalfe, M. Mevissen, A. Meyer, A. Mezzacapo, R. Midha, D. Miller, Z. Minev, A. Mitchell, N. Moll, A. Montanez, G. Monteiro, M. D. Mooring, R. Morales, N. Moran, D. Morcuende, S. Mostafa, M. Motta, R. Moyard, P. Murali, J. Müggenburg, T. NEMOZ, D. Nadlinger, K. Nakanishi, G. Nannicini, P. Nation, E. Navarro, Y. Naveh, S. W. Neagle, P. Neuweiler, A. Ngoueya, T. Nguyen, J. Nicander, Nick-Singstock, P. Niroula, H. Norlen, NuoWenLei, L. J. O'Riordan, O. Ogunbayo, P. Ollitrault, T. Onodera, R. Otaolea, S. Oud, D. Padilha, H. Paik, S. Pal, Y. Pang, A. Panigrahi, V. R. Pascuzzi, S. Perriello, E. Peterson, A. Phan, K. Pilch, F. Piro, M. Pistoia, C. Piveteau, J. Plewa, P. Pocreau, A. Pozas-Kerstjens, R. Pracht, M. Prokop, V. Prutyanov, S. Puri, D. Puzzuoli, J. Pérez, Quant02, Quintiii, R. I. Rahman, A. Raja, R. Rajeev, I. Rajput, N. Ramagiri, A. Rao, R. Raymond, O. Reardon-Smith, R. M.-C. Redondo, M. Reuter, J. Rice, M. Riedemann, Rietesh, D. Risinger, M. L. Rocca, D. M. Rodríguez, RohithKarur, B. Rosand, M. Rossmannek, M. Ryu, T. SAPV, N. R. C. Sa, A. Saha, A. Ash-Saki, S. Sanand, M. Sandberg, H. Sandesara, R. Sapra, H. Sargsyan, A. Sarkar, N. Sathaye, B. Schmitt, C. Schnabel, Z. Schoenfeld, T. L. Scholten, E. Schoute, M. Schulterbrandt, J. Schwarm, J. Seaward, Sergi, I. F. Sertage, K. Setia, F. Shah, N. Shammah, R. Sharma, Y. Shi, J. Shoemaker, A. Silva, A. Simonetto, D. Singh, D. Singh, P. Singh, P. Singkanipa, Y. Siraichi, Siri, J. Sistos, I. Sitdikov, S. Sivarajah, M. B. Sletfjerding, J. A. Smolin, M. Soeken, I. O. Sokolov, I. Sokolov, V. P. Soloviev, SooluThomas, Starfish, D. Steenken, M. Stypulkoski, A. Suau, S. Sun, K. J. Sung, M. Suwama, O. Słowik, H. Takahashi, T. Takawale, I. Tavernelli, C. Taylor, P. Taylour, S. Thomas, K. Tian, M. Tillet, M. Tod, M. Tomasik, C. Tornow, E. de la Torre, J. L. S. Toural, K. Trabing, M. Treinish, D. Trenev, TrishaPe, F. Truger, G. Tsilimigkounakis, D. Tulsi, W. Turner, Y. Vaknin, C. R. Valcarce, F. Varchon, A. Vartak, A. C. Vazquez, P. Vijaywargiya, V. Villar, B. Vishnu, D. Vogt-Lee, C. Vuillot, J. Weaver, J. Weidenfeller, R. Wieczorek, J. A. Wildstrom, J. Wilson, E. Winston, WinterSoldier, J. J. Woehr, S. Woerner, R. Woo, C. J. Wood, R. Wood, S. Wood, J. Wootton, M. Wright, L. Xing, J. YU, B. Yang, U. Yang, J. Yao, D. Yeralin, R. Yonekura, D. Yonge-Mallo, R. Yoshida, R. Young, J. Yu, L. Yu, C. Zachow, L. Zdanski, H. Zhang, I. Zidaru, C. Zoufal, aeddins ibm, alexzhang13, b63, bartek bartlomiej, bcamorrison, brandhsn, charmerDark, deeplokhande, dekel.meirom, dime10, dlasecki, ehchen, fanizzamarco, fs1132429, gadial, galeinston, georgezhou20, georgios ts, gruu, hhorii, hykavitha, itoko, jeppevinkel, jessica angel7, jezerjojo14, jliu45, jscott2, klinvill, krutik2966, ma5x, michelle4654, msuwama, nico lgrs, ntgiwsvp, ordmoj, sagar pahwa, pritamsinha2304, ryancocuzzo, saktar unr, saswati qiskit, septembrr, sethmerkel, sg495, shaashwat, smturro2, sternparky, strickroman, tigerjack, tsura crisaldo, upsideon, vadebayo49, welien, willhbang, wmurphy collabstar, yang.luh, and M. Čepulkovskis, "Qiskit: An open-source framework for quantum computing," 2021.

[3] A. Bendersky, F. Pastawski, and J. P. Paz, "Selective and efficient estimation of parameters for quantum process tomography," *Physical review letters*, vol. 100, no. 19, p. 190403, 2008.

[4] E. Bernstein and U. Vazirani, "Quantum complexity theory," *SIAM Journal on computing*, vol. 26, no. 5, pp. 1411–1473, 1997.

[5] M. J. Bremner, C. M. Dawson, J. L. Dodd, A. Gilchrist, A. W. Harrow, D. Mortimer, M. A. Nielsen, and T. J. Osborne, "Practical scheme for quantum computation with any two-qubit entangling gate," *Physical review letters*, vol. 89, no. 24, p. 247902, 2002.

[6] I. L. Chuang and M. A. Nielsen, "Prescription for experimental determination of the dynamics of a quantum black box," *Journal of Modern Optics*, vol. 44, no. 11-12, pp. 2455–2467, 1997.

[7] G. E. Crooks, "Gates, states, and circuits," 2020.

[8] S. A. Cuccaro, T. G. Draper, S. A. Kutin, and D. P. Moulton, "A new quantum ripple-carry addition circuit," *arXiv preprint quant-ph/0410184*, 2004.

[9] E. Farhi, J. Goldstone, and S. Gutmann, "A quantum approximate optimization algorithm," *arXiv preprint arXiv:1411.4028*, 2014.

[10] A. Gaikwad, K. Shende, K. Dorai *et al.*, "Implementing efficient selective quantum process tomography of superconducting quantum gates on ibm quantum experience," *Scientific reports*, vol. 12, no. 1, pp. 1–11, 2022.

[11] D. Greenbaum, "Introduction to quantum gate set tomography," *arXiv preprint arXiv:1509.02921*, 2015.

[12] J. B. Hertzberg, E. J. Zhang, S. Rosenblatt, E. Magesan, J. A. Smolin, J.-B. Yau, V. P. Adiga, M. Sandberg, M. Brink, J. M. Chow *et al.*, "Laser-annealing josephson junctions for yielding scaled-up superconducting quantum processors," *npj Quantum Information*, vol. 7, no. 1, pp. 1–8, 2021.

[13] C. Huang, D. Ding, F. Wu, L. Kong, F. Zhang, X. Ni, Y. Shi, H.-h. Zhao, and J. Chen, "Towards ultra-high fidelity quantum operations: Sqisw gate as a native two-qubit gate," *arXiv preprint arXiv:2105.06074*, 2021.

[14] P. Jurcevic, A. Javadi-Abhari, L. S. Bishop, I. Lauer, D. F. Bogorin, M. Brink, L. Capelluto, O. Günlük, T. Itoko, N. Kanazawa *et al.*, "Demonstration of quantum volume 64 on a superconducting quantum computing system," *Quantum Science and Technology*, vol. 6, no. 2, p. 025020, 2021.

[15] A. Kandala, K. X. Wei, S. Srinivasan, E. Magesan, S. Carnevale, G. A. Keefe, D. Klaus, O. Dial, and D. C. McKay, "Demonstration of a high-fidelity cnot for fixed-frequency transmons with engineered zz suppression," 2020.

[16] 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," *Physical Review A*, vol. 77, no. 1, p. 012307, 2008.

[17] J. Ku, X. Xu, M. Brink, D. C. McKay, J. B. Hertzberg, M. H. Ansari, and B. L. T. Plourde, "Suppression of unwanted *zz* interactions in a hybrid two-qubit system," *Phys. Rev. Lett.*, vol. 125, p. 200504, Nov 2020. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.125.200504

[18] S. Kwon, A. Tomonaga, G. L. Bhai, S. J. Devitt, and J.-S. Tsai, "Gate-based superconducting quantum computing," *Journal of Applied Physics*, vol. 129, no. 4, p. 041102, Jan. 2021. [Online]. Available: https://doi.org/10.1063/5.0029735

[19] L. Lao, P. Murali, M. Martonosi, and D. Browne, "Designing calibration and expressivity-efficient instruction sets for quantum computing," in *2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 2021, pp. 846–859.

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

[21] E. Magesan, J. M. Gambetta, and J. Emerson, "Scalable and robust randomized benchmarking of quantum processes," *Physical review letters*, vol. 106, no. 18, p. 180504, 2011.

[22] 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 *et al.*, "Efficient measurement of quantum gate error by interleaved randomized benchmarking," *Physical review letters*, vol. 109, no. 8, p. 080505, 2012.

[23] D. C. McKay, S. Filipp, A. Mezzacapo, E. Magesan, J. M. Chow, and J. M. Gambetta, "Universal gate for fixed-frequency qubits via a tunable bus," *Phys. Rev. Applied*, vol. 6, p. 064007, Dec 2016. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevApplied.6.064007

[24] P. Mundada, G. Zhang, T. Hazard, and A. Houck, "Suppression of qubit crosstalk in a tunable coupling superconducting circuit," *Phys. Rev. Applied*, vol. 12, p. 054023, Nov 2019. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevApplied.12.054023

[25] P. S. Mundada, S. Sussman, A. Vrajitoarea, C. Guinn, C. L. Calonnec, C. Leroux, A. Di Paolo, A. Petrescu, A. Place, A. Blais, and A. A. Houck, "Fast crosstalk-free perfect entangler in a tunable coupling superconducting circuit," *Bulletin of the American Physical Society*, 2021, 2022.

[26] E. Nielsen, J. K. Gamble, K. Rudinger, T. Scholten, K. Young, and R. Blume-Kohout, "Gate set tomography," *Quantum*, vol. 5, p. 557, 2021.

[27] A. Noguchi, A. Osada, S. Masuda, S. Kono, K. Heya, S. P. Wolski, H. Takahashi, T. Sugiyama, D. Lachance-Quirion, and Y. Nakamura, "Fast parametric two-qubit gates with suppressed residual interaction using a parity-violated superconducting qubit," 2020.

[28] I. Perito, A. J. Roncaglia, and A. Bendersky, "Selective and efficient quantum process tomography in arbitrary finite dimension," *Physical Review A*, vol. 98, no. 6, p. 062303, 2018.

[29] E. C. Peterson, L. S. Bishop, and A. Javadi-Abhari, "Optimal synthesis into fixed xx interactions," *arXiv preprint arXiv:2111.02535*, 2021.

[30] E. C. Peterson, G. E. Crooks, and R. S. Smith, "Two-qubit circuit depth and the monodromy polytope," *Quantum*, vol. 4, p. 247, 2020.

[31] A. Petrescu, C. L. Calonnec, C. Leroux, A. Di Paolo, P. Mundada, S. Sussman, A. Vrajitoarea, A. A. Houck, and A. Blais, "Accurate methods for the analysis of strong-drive effects in parametric gates," *arXiv preprint arXiv:2107.02343*, 2021.

[32] J. Poyatos, J. I. Cirac, and P. Zoller, "Complete characterization of a quantum process: the two-bit quantum gate," *Physical Review Letters*, vol. 78, no. 2, p. 390, 1997.

[33] J. Preskill, "Quantum computing in the nisq era and beyond," *Quantum*, vol. 2, p. 79, 2018.

[34] A. Rezakhani, "Characterization of two-qubit perfect entanglers," *Physical Review A*, vol. 70, no. 5, p. 052313, 2004.

[35] L. Ruiz-Perez and J. C. Garcia-Escartin, "Quantum arithmetic with the quantum fourier transform," *Quantum Information Processing*, vol. 16, no. 6, pp. 1–14, 2017.

[36] C. Ryan, M. Laforest, and R. Laflamme, "Randomized benchmarking of single-and multi-qubit control in liquid-state nmr quantum information processing," *New Journal of Physics*, vol. 11, no. 1, p. 013034, 2009.

[37] V. V. Shende, S. S. Bullock, and I. L. Markov, "Recognizing small-circuit structure in two-qubit operators and timing hamiltonians to compute controlled-not gates," *arXiv preprint quant-ph/0308045*, 2003.

[38] P. W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer," *SIAM review*, vol. 41, no. 2, pp. 303–332, 1999.

[39] Y. Sung, L. Ding, J. Braumüller, A. Vepsäläinen, B. Kannan, M. Kjaergaard, A. Greene, G. O. Samach, C. McNally, D. Kim, A. Melville, B. M. Niedzielski, M. E. Schwartz, J. L. Yoder, T. P. Orlando, S. Gustavsson, and W. D. Oliver, "Realization of high-fidelity cz and zz-free iswap gates with a tunable coupler," 2020.

[40] P. Zanardi, C. Zalka, and L. Faoro, "Entangling power of quantum evolutions," *Physical Review A*, vol. 62, no. 3, p. 030301, 2000.

[41] J. Zhang, J. Vala, S. Sastry, and K. B. Whaley, "Geometric theory of nonlocal two-qubit operations," *Physical Review A*, vol. 67, no. 4, p. 042313, 2003.

[42] J. Zhang, J. Vala, S. Sastry, and K. B. Whaley, "Minimum construction of two-qubit quantum operations," *Physical review letters*, vol. 93, no. 2, p. 020502, 2004.

[43] P. Zhao, D. Lan, P. Xu, G. Xue, M. Blank, X. Tan, H. Yu, and Y. Yu, "Suppression of static zz interaction in an all-transmon quantum processor," 2020.