

Let Each Quantum Bit Choose Its Basis Gates

Sophia Fuhui Lin*, Sara Sussman†, Casey Duckering*, Pranav S. Mundada‡,
 Jonathan M. Baker*, Rohan S. Kumar*, Andrew A. Houck‡, Frederic T. Chong*

*Department of Computer Science, University of Chicago, Chicago, USA

Email: {sophialin1, cduck, jmbaker, rohansk}@uchicago.edu, chong@cs.uchicago.edu

†Department of Physics, Princeton University, Princeton, USA

Email: sarafs@princeton.edu

‡Department of Electrical Engineering, Princeton University, Princeton, USA

Email: pmundada@alumni.princeton.edu, aahouck@princeton.edu

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. 2Q entangling gates, coupled with some 1Q gates, allow for universal quantum computation. In superconducting technologies, the current state of the art is to implement the same 2Q gate between every pair of qubits (typically an XX- or XY-type gate). This strict hardware uniformity requirement for 2Q gates in a large quantum computer has made scaling up a time and resource-intensive endeavor in the lab.

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. This work aims to give quantum scientists the ability to run meaningful algorithms with qubit systems that are not perfectly uniform. Scientists will also be able to use a much broader variety of novel 2Q gates for quantum computing. We develop a theoretical framework for identifying good 2Q basis gates on “nonstandard” Cartan trajectories that deviate from “standard” trajectories like XX. We then introduce practical methods for calibration and compilation with nonstandard 2Q gates, and discuss possible ways to improve the compilation. To demonstrate our methods in a case study, we simulated both standard XY-type trajectories and faster, nonstandard trajectories using an entangling gate architecture with far-detuned transmon qubits. We identify efficient 2Q basis gates on these nonstandard trajectories and use them to compile a number of standard benchmark circuits such as QFT and QAOA. Our results demonstrate an 8x improvement over the baseline 2Q gates with respect to speed and coherence-limited gate fidelity.

Keywords-quantum computing; two-qubit gates;

I. INTRODUCTION

Quantum computers have the potential to solve problems currently intractable for conventional computers [1], but current computations are limited by errors [2], particularly when interacting two qubits to perform a quantum gate operation. This is not surprising, as qubits are engineered to preserve quantum state and be isolated from the environment, but a quantum operation is a moment in time where external control is applied from the environment to deliberately alter

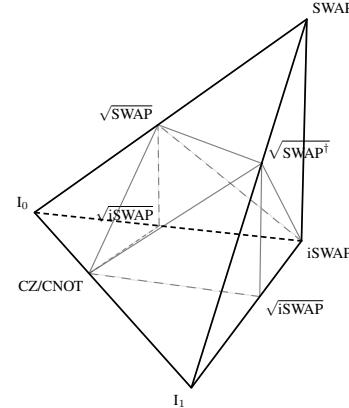


Figure 1. The Weyl chamber of 2Q quantum gates, explained in Sec II-B. The non-local part of a 2Q gate is fully described by its position in the Weyl chamber. As the duration of an entangling gate pulse increases, the 2Q gate evolves, traversing a Cartan trajectory in the Weyl chamber. CNOT and CZ are 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 .

a qubit’s state. To accomplish low-error gates, the control mechanisms are carefully designed and the control signals are 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 two-qubit (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 point that does not coincide with SWAP or Identity has nonzero entangling power. Any of these 2Q entangling gates can achieve universal computation when added to a continuous set of 1Q gates [3].

Using the minimal set of gates needed for universal computing is rarely a desirable thing to do. For example, while

the NAND is universal in classical computing, building circuits from it alone is less efficient than using a larger set of logic gates. However, the intensive calibrations necessary for high fidelity 2Q gates between qubits in a large quantum computer make it impractical to support a large set of 2Q basis gates. All logical 2Q gates scheduled to run on a quantum computer have to be decomposed by its compiler into alternating layers of pre-calibrated 1Q and 2Q basis gates. Thus, the choice of which 2Q gates to directly support is critical to enabling high-performance quantum computing. On the hardware side, the 2Q basis gates must have high-fidelity hardware implementations. On the software side, they must enable the low-depth decomposition of other 2Q gates.

Superconducting qubits support XX- and XY-type 2Q interactions [4], [5]. The strength of each of these interactions depends on the type of coupling, the coupling strength, and the frequency detuning between the qubits [5]. The Weyl chamber space of 2Q gates (Fig. 1) is a useful way to visualize these interactions: the coordinates of a gate correspond to its non-local part in Cartan’s KAK decomposition (see Section II-B). 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. Cartan trajectories are generated by increasing the duration of an entangling gate pulse, which evolves the 2Q gate.

Section IV describes the various difficulties associated with reliably performing standard 2Q gates like CZ and iSWAP on today’s superconducting quantum computers. Whether a deviation from a standard Cartan trajectory is a 2Q error depends on what the target 2Q gate is. If the target 2Q gate has to be a certain standard gate (e.g. iSWAP), even a small amount of coherent crosstalk between the two qubits will cause the gate to have a small CZ component, and so the gate will not be identically iSWAP. However, if that coherent crosstalk is a stable systematic that does not add noise or cause decoherence, the gate could still be an effective, high fidelity entangler that is useful for computing; the target 2Q gate would just have to be a nonstandard unitary rather than the standard iSWAP. In Section IV we also show an example of an experimentally measured Cartan trajectory of 2Q gates on a superconducting qubit device. The measured Cartan trajectory includes very fast nonstandard 2Q gates with high entangling power, but it does not pass through traditional basis gates like iSWAP and CZ. This example motivates us to develop methods that enable the use of nonstandard 2Q gates for quantum computing. If we allow for deviations from standard 2Q gates, we can use high fidelity, non-standard quantum hardware for practical quantum computing.

Using nonstandard 2Q basis gates requires methods for identifying good basis gates on a general 2Q gate trajectory, calibrating a nonstandard gate, and compiling with non-

standard basis gates. The primary focus of our work is to construct and demonstrate a framework for efficiently identifying a “good” set of 2Q basis gates from a nonstandard trajectory. But we also propose solutions for calibration and compilation.

What are our standards for a good set of 2Q basis gates? Following the principle of Amdahl’s Law, we pay most attention to the SWAP gate as a target for synthesis because of its importance for communication within programs executing on superconducting devices. To mitigate crosstalk and satisfy 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, our framework also allows one to prioritize other target gates, including but not limited to CNOT, iSWAP, and the B gate. It also enables the simultaneous prioritisation of multiple target gates.

The calibration of a non-standard 2Q basis gate requires identifying a gate duration that gives an ideal basis gate and then accurately characterizing the corresponding gate so that we use the right unitary for compiling. Our proposed calibration protocol address both without causing a long downtime on a quantum device. However, we point out that in order to *precisely* characterize a non-standard gate, one should consider using gate set tomography (GST) as opposed to quantum process tomography (QPT). The data collected from GST experiments may require several hours of classical processing. Before that finishes, one would have to use the calibration results from the previous cycle. The speedup of GST’s classical processing, which is an active field of research [6], would help reduce the cost of calibrating non-standard gates. In addition, we observed that the systematic deviations are stable over days (Fig. 5). If the change in deviation is negligible, one may not need to apply GST in every calibration cycle.

Compiling with non-standard 2Q basis gates requires a conversion from arbitrary 2Q gates into the basis gates. There isn’t a general analytical formula that works for arbitrary target and basis gate, so a numerical search is needed. However, we can analytically obtain information on the minimum circuit depth needed for a perfect synthesis and use it to facilitate the numerical search. Besides, the circuits that synthesise common gates from the basis gates can be pre-computed after each calibration cycle, so that one wouldn’t need to re-compute them for every program.

Our contributions are summarized below.

- Our work is the first to consider using 2Q basis gates from general non-standard gate trajectories that are not parametrized by a simple function.

- We provide a theoretical framework for identifying and visualizing the set of good 2Q basis gates, given a set of target 2Q gates to prioritize. With an emphasis on SWAP, we characterize the sets of gates that enable the synthesis of SWAP in 1, 2, and 3 layers, respectively. As another example, we visualize the gates that are able to both synthesize SWAP in 3 layers and CNOT in 2 layers. After identifying the volume of desirable basis gates in the Weyl chamber, one can select the first intersection of the trajectory with the volume as the 2Q basis gate. (Section V)
- We propose a practical calibration protocol that is agnostic as to whether a 2Q gate is standard or non-standard. (Section VI)
- We discuss a practical approach to compiling with non-standard 2Q basis gates. (Section VII)
- We apply our methods to a case study entangling gate architecture with far-detuned transmon qubits [7]. First, we use our theoretical framework to select 2Q basis gates from simulated nonstandard Cartan trajectories that are realistic for this case study architecture. By increasing the entangling pulse drive amplitude we get a significant 2Q basis gate speedup but introduce a deviation into the Cartan trajectory. Then we use these 2Q basis gates to run a variety of benchmark circuits including BV[8], QAOA[9], the QFT adder[10], and the Cuccaro Adder[11], and compare to the results from using the \sqrt{iSWAP} gate on the standard XY-type trajectory. (Section VIII)

II. BACKGROUND

A. 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 $|1\rangle$; A general quantum state can be expressed as $\alpha|0\rangle + \beta|1\rangle$ where α, β are complex amplitudes that satisfy $|\alpha|^2 + |\beta|^2 = 1$. Thus, 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.

B. Geometric characterization of 2Q gates

Two 2Q quantum gates $U_1, U_2 \in SU(4)$ are *locally equivalent* if it is possible to obtain one from the other by adding 1Q operations. In other words, 2Q operations U_1 and U_2 are locally equivalent if there exist $k_1, k_2 \in SU(2) \otimes SU(2)$ such that $U_1 = k_1 U_2 k_2$. For example, CNOT and CZ are locally equivalent via Hadamard gates.

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

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

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 [12]. 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 represent a different equivalence class of 2Q gates. We refer the interested readers to [13] 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 .

C. Entangling power of 2Q gates

The entangling power [14] is a widely accepted quantitative measure of the capacity of a 2Q gate to entangle the qubits that it acts on. It is typically a good indicator of the ability of a specific 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 [14]. 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. Conversely, 2Q gate U is called a *perfect entangler* if it can produce a maximally entangled state from an unentangled one[12]. Perfect entanglers (PE) have entangling power no less than $\frac{1}{6}$. They constitute 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*[15]. 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[16]. However, there has been no proposal to directly implement the B gate in hardware.

III. RELATED WORK

To the best of our knowledge, no prior work involves using 2Q basis gates from arbitrary nonstandard gate trajectories. In parallel with this work, Lao et al. [17] propose to mitigate coherent parasitic errors in 2Q gates by software and present methods of compilation. Our work

is more general than [17], although we share the insight that coherent errors in 2Q gates can be treated as part of the gate for compilation. While our framework works for general irregular trajectories and select basis gates on them using the approach detailed in Section V, they focus on iSWAP-like (XY) gates with an unwanted CPHASE (XX) component (which belongs to the fSim gate set so is not truly non-standard) and always use CPHASE(ψ)iSWAP($\pi/4$) because it has similar expressivity as iSWAP($\pi/4$) for small deviation ψ . They do not discuss calibration. Their baseline for evaluation is similar to the baseline in our case study, which is to make the trajectory more standard by lengthening the gate duration.

Recent research from both the experimental [18], [19], [20], [21] and theory sides has utilized 2Q (and 3Q) basis gates from a continuous set of standard gates, as opposed to only building and compiling with the best-known gates like CNOT and iSWAP. The works that are most relevant to this project are those that look for a small set of 2Q basis gates (from a continuous standard gate set) that are the most valuable to calibrate. Lao et al. [22] use a numerical approach to test the performance of different gates from the fSim and XY gate sets on a range of application circuits, with the overall circuit success rate as the objective. Peterson et al. [23] from IBM use analytic techniques to find that the gate set $\{CX, CX^{1/2}, CX^{1/3}\}$ is almost as good as the entire continuous set of XX gates in implementing random 2Q gates. They try to minimize the expected (average) infidelity in implementing random 2Q gates under an experimentally motivated error model. Huang et al. [24] proposes using the \sqrt{iSWAP} as 2Q basis gate, instead of using iSWAP or CNOT, and implement it using a 2-fluxonium qubit device. Recent proposals for novel nonstandard 2Q gates in the superconducting qubit literature that are informed by the current experimental challenges in scaling up with standard 2Q gates include [25], [26].

IV. SYSTEMATIC DEVIATIONS IN 2Q GATES

The 2Q gate is a critical building block that must be well-engineered before it is used to construct a quantum computer with many qubits. In practice, engineering 2Q gates in the lab involves iterating prototypes of the devices to minimize any and all systematic errors that result from imperfect device design or control along with nonuniformities in device fabrication. Even if unwanted crosstalk between the qubits is successfully reduced and the 2Q gate is shown to be an effective entangler with a consistent identity, if the gate's identity is somehow nonstandard, one would normally assume it is not useful. The constraint of requiring 2Q gates to be standard is most burdensome for the superconducting qubit platform, where device Hamiltonians are engineered from scratch and there is no 2Q gate that is truly native to

the platform - unlike, for instance, the SWAP gates that are native to atomic qubits [27].

Today's multi-qubit superconducting devices are not able to perform perfectly identical 2Q gates between every pair of qubits because of device-level imperfections, tradeoffs and uncertainties. Experimentalists model the expected rate of information leakage between on-chip elements using microwave circuit design software [28], [29], but it is inevitable that irregularities arise during device fabrication and packaging. The devices are at least partially handmade and every fabrication tool has a finite precision. Also, the various materials that make up the layers of the superconducting device can host physical two-level systems that act as sources of noise and even can coherently interact with qubits [30], [31]; reducing the effect of these two-level systems is an active field of research [32]. Another active field of research is reducing irregularities in the fabrication of Josephson junctions, which are critical on-chip elements [33], [34]. For a given device, it can be difficult for the experimentalist to determine whether a systematic 2Q gate deviation is caused by an imperfection in the device design or in its control. For example, a common source of systematic 2Q gate deviation is the imperfect mitigation of the static ZZ crosstalk which is a dominant source of 2Q gate error for transmon qubits [35], [36], [37], [38], [39], [40]. Devices can be designed to suppress the static ZZ crosstalk but unless the device is properly fabricated, packaged, biased and controlled there will be nonzero static ZZ crosstalk which will cause the 2Q gate to deviate from the target unitary.

Superconducting devices can also have higher order Hamiltonian terms that result in the experimentally measured Cartan trajectory of 2Q gates deviating from the expected Cartan trajectory. This deviation is particularly significant for fast gates enabled by large coupling or large drive strength [7], [41], [42]. Experimentalists have historically tried to suppress these deviations by reducing the 2Q gate drive strength, which has the negative consequence of slowing the 2Q gate down. It is in general difficult to accurately model the effect of the strong drives that perform fast 2Q gates on the Hamiltonian level, and this is an active field of research [7].

Plotting measured 2Q gates in Cartan coordinates is a valuable tool experimentalists can use to easily visualize and study any deviations their gates may have from the expected Cartan trajectory. For example, Figure 2 shows a measured Cartan trajectory that is nonstandard. This experimental data was collected from one of the first iterations of a superconducting device [43] that was designed to implement a recently proposed entangling gate architecture [7]. The data includes a very fast (13 ns) perfect entangler. Since the measured trajectory was systematically offset from the predicted one (XY), the experimentalists investigated potential sources of that systematic offset. Since this source of deviation could be eliminated with better device and

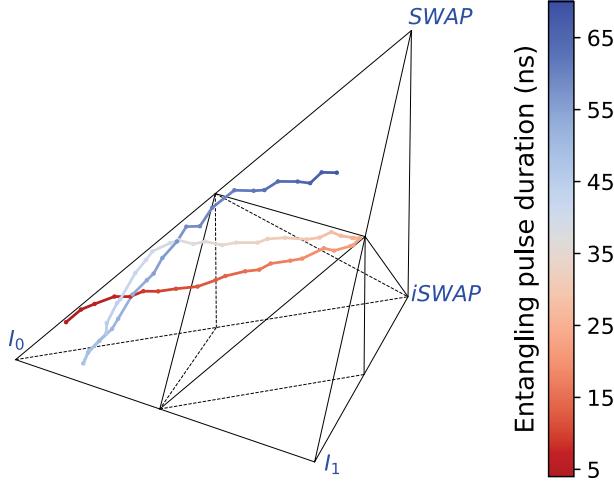


Figure 2. Experimental data showing a nonstandard Cartan coordinate trajectory. An experimental implementation [43] of the iSWAP gate with the entangler architecture proposed in [7] yielded a nonstandard Cartan coordinate trajectory close to the plane of I_0 , SWAP, and iSWAP. The first instance of a perfect entangler was at an entangler duration of 13 ns. In this nonstandard trajectory, the 13 ns entangler is offset from the Cartan coordinate for the square root of iSWAP and the 26 ns entangler is likewise offset from the Cartan coordinate for iSWAP. Note that due to an experimental hardware constraint the shortest possible entangling pulse duration was 4 ns, so the measured Cartan trajectory begins there.

control engineering, the experimentalists began to optimize their next device iteration accordingly. But in this work we suggest that there is nothing inherently unusable about measured Cartan trajectories that are nonstandard due to this kind of coherent systematic offset, and that the 13 ns nonstandard perfect entangler identified in Figure 2 could be treated as a native 2Q basis gate by the compiler.

Our work seeks to enable the use of the nonstandard 2Q gates that can be native to superconducting devices. If 2Q gate calibration and compiling protocols became more flexible, usable superconducting 2Q gate yield would increase considerably, enabling more rapid and effective prototyping of 2Q gates which could be scaled to a computer. Furthermore, any number of novel superconducting devices with very fast 2Q gates that happen to be nonstandard could be effectively utilized for computing.

V. IDENTIFYING GOOD 2Q BASIS GATES

A. 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 Figure 3 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.

B. 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. [44].

Without going into the technical details, here we summarize a key result from [44] that we adapt and apply in Section V-C and V-D.

Theorem 5.1: There exists a 2-layer decomposition of 2Q gate A into B , C , and 1Q gates as in Figure 3(a), 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 [44] or the implementation of the function in our code¹. Note that Reference [44] 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: $\text{LogSpec}(U) = (a, b, c, d)$ and $\rho(\text{LogSpec}(U)) = (c + \frac{1}{2}, d + \frac{1}{2}, a - \frac{1}{2}, b - \frac{1}{2})$. If $\text{LogSpec}(U) = \rho(\text{LogSpec}(U))$ for all A , B , and C , we only need to check one set of inequalities. If $\text{LogSpec}(U) \neq \rho(\text{LogSpec}(U))$ 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.

¹Our code can be found at https://github.com/SophLin/nonstandard_2qbasis_gates

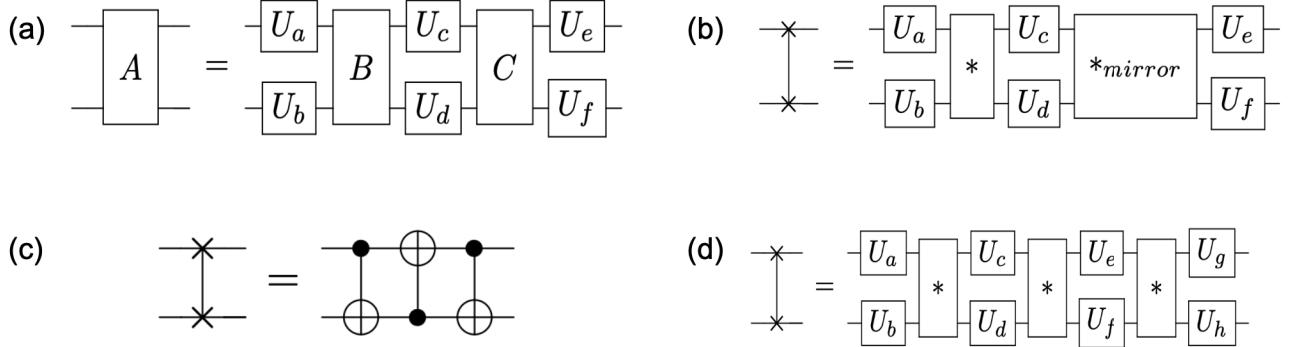


Figure 3. (a) Gate A, decomposed into 2 layers with 2Q gates B, C and 1Q gates $U_a, U_b, U_c, U_d, U_e, U_f$. (b) A general 2-layer decomposition of the SWAP gate. Here $*$, $*_{\text{mirror}}$ can be replaced by any pair of 2Q gates capable of synthesizing a SWAP in 2 layers. (c) The SWAP gate, decomposed into 3 CNOT gates. (d) 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.

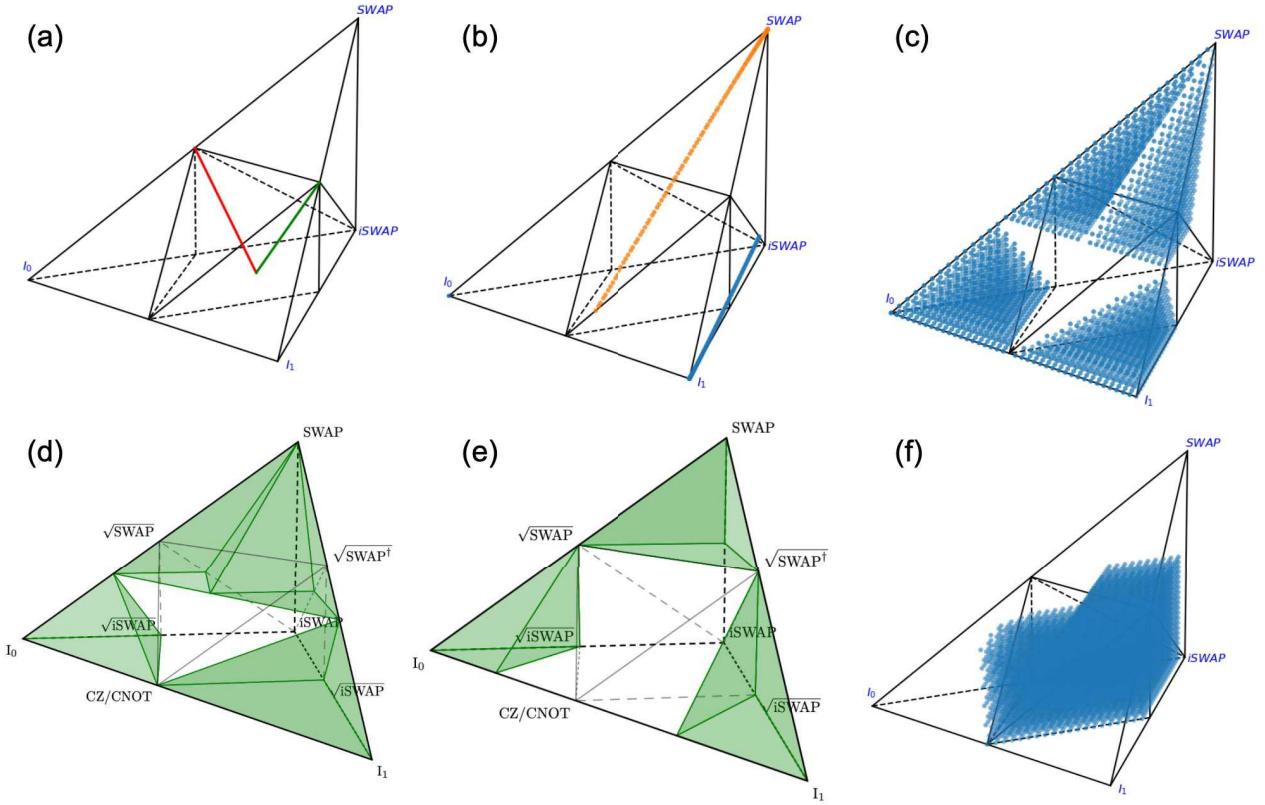


Figure 4. (a) Gates that are able to synthesize SWAP in 2 layers form 2 line segments in the Weyl chamber. The red one is from the B gate to \sqrt{SWAP} , and the green one is from the B gate to \sqrt{SWAP}^\dagger . (b) Pairs of gates that are able to synthesize a SWAP in 2 layers. In blue is an example 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. (c) Gates that are NOT able to synthesize a SWAP in 3 layers. (d) 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 $\{SWAP, (\frac{1}{2}, \frac{1}{6}, \frac{1}{6}), (\frac{5}{6}, \frac{1}{6}, \frac{1}{6}), (\frac{2}{3}, \frac{1}{3}, \frac{1}{6})\}$. (e) 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 $\{SWAP, \sqrt{SWAP}, \sqrt{SWAP}^\dagger, (\frac{1}{2}, \frac{1}{2}, \frac{1}{4})\}$. (f) Gates that are able to decompose SWAP in 3 layers and CNOT in 2 layers.

C. 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.

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. We discuss three 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 Figure 3(b). 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.

In the second case, for each point $*$ in the Weyl chamber, (as derived in Appendix B) there is exactly one point $*_{mirror}$ such that they together enable a 2-layer decomposition of SWAP (see Figure 3(b)). 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 (Figure 4(b)). 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.

SWAP in 3 layers: It is a well-known result that 3 invocations of CNOT are required to implement a SWAP [45]. We show the circuit in Figure 3(c). In fact, CNOT and iSWAP share the property that they can synthesize any arbitrary 2Q gate in 3 layers but only a 0-volume set of gates (in the Weyl chamber) in 2 layers [44].

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 Figure 3(d). 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 Figure 4(c). 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 Figure 4(d). We also learn that the volume of $S_{SWAP,3}$ is 68.5% the volume of the Weyl chamber.

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 Figure 4(d). 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 that the early part of the trajectory overlaps any of them.

D. 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 (denoted $S_{CNOT,2}$ here) takes up 75% of the volume in the Weyl chamber. The complement $\overline{S_{CNOT,2}}$ consists of 3 tetrahedra, as shown in Figure 4(e). 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 $S_{CNOT,2} \cap S_{SWAP,3}$ (See Figure 4(f)).

E. A strategy for locating good 2Q basis gates

Our framework allows one to prioritize different combinations of target 2Q gates. In Section VIII, we test the

following 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 V-C, 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 V-D, the gate that meets Criterion 2 can be found similarly. With this insight, we can locate a desired 2Q basis gate in an experimental setting using the methods in Section VI.

Our framework can be easily adapted to other criteria for selecting basis gates. For instance, we can select the fastest gate that can decompose another set of target gates within a certain number of layers. We can also incorporate other metrics like the entangling power into a criterion, e.g. we can locate the fastest gate on the trajectory that is both a PE and can synthesize SWAP in 3 layers.

VI. CALIBRATION OF NONSTANDARD 2Q GATES

We propose two stages for calibrating a 2Q basis gate on an unknown trajectory of 2Q gates: first, a more costly “initial tuneup” stage that does not assume any knowledge of the trajectory and then a less costly “retuning” stage that utilizes information from the last initial tuneup and the retunings after it. In a well-controlled industry setup we would imagine the initial tuneup being done once a month and retuning being done daily. In a less well-controlled environment (e.g. one prone to low frequency drift), the initial tuneup could be done more frequently, as needed.

Our proposed calibration approach uses two techniques for experimentally characterizing the unitary of a potentially non-standard 2Q gate: quantum process tomography (QPT) [46], [47] and gate set tomography (GST) [48], [49], [50], [51]. QPT is a simple way to estimate a unitary but it cannot separate state preparation and measurement (SPAM) errors from gate errors [52]. GST is a highly general and accurate tomography technique that characterizes all the operations in a gate set (including SPAM) simultaneously and self-consistently. GST is simple to run, taking minutes to acquire on a superconducting device. GST acquisition is followed by classical processing of the data that can be computed on a cluster in about two hours. Note that during the classical processing, the quantum device can still be used with gates from the previous calibration cycle. The speedup of GST’s classical processing is an active field of research and may be obtained by allowing physics to inform the dominant errors that are expected [6]. The most relevant returns for fine tuning the unitary are the error generators [53] for the gate set. The error generators are a basis for writing the transformation between the measured unitary and the

unitary that GST expects. It measures coherent differences and estimates stochastic noise levels. GST is thus a valuable tool for directly characterizing 2Q gates.

Here we list the steps in the initial tuneup stage.

- 1) Do preliminary coarse tuning experiments such as amplitude and frequency calibration of the entangling pulse drive to estimate the entangling pulse duration of interest. For example, a resonant *iSWAP*-like interaction may have an amplitude and a frequency to tune for optimal population swapping. (5 minutes per pulse)
- 2) Perform QPT for each 2Q gate in the Cartan trajectory leading up to the approximate 2Q gate of interest. The qubit controller resolution (typically ~ 1 ns) will determine the spacing between the trajectory points. Based on the findings in Step 1 the trajectory can be cropped around the entangling pulse duration of interest. The unitaries found will be the full list of candidate gates. (30-60 minutes per trajectory)
- 3) From the candidate gates in the previous steps, use Section V to identify which of them might be the fastest ones that also are good 2Q basis gates. In this step the list of candidate basis gates is narrowed down. We are not able to narrow down to one basis gate due to the imprecision of QPT.
- 4) Perform GST to obtain full information about each candidate 2Q gate, including an accurate gate unitary and a breakdown of error sources. Then the set of 2Q basis gates can be chosen. (~ 10 minutes for each 2Q gate, followed by classical processing)

The second calibration stage is the quick “retuning” of the 2Q basis gates that relies on the results of the initial tuneup. Once the precise unitary for each 2Q basis gate is found, the gates can be simply retuned using the coarse tuning procedures in Step 1 of the initial tuneup. The information gained in the initial tuneup would allow experimentalists to prescribe a different retuning protocol to each 2Q basis gate according to what it needs. In practice, retuning would most likely be a simple combination of amplitude calibration and frequency calibration of the elements involved in each 2Q basis gate, and it would take approximately 1-5 minutes per 2Q basis gate.

The extent to which previously gathered information can help reduce the cost of retuning depends on the stability of the gate trajectories over time. Figure 5 shows the nonstandard Cartan trajectories measured on two days, over two entangling pulse drive amplitudes. Over the five day period that Cartan trajectories were measured for this device, the trajectories were all found to look qualitatively similar, as in Figure 5. While limited, this experimental data suggests that the measured Cartan trajectories obtained in the initial tuneup stage could potentially be used for several days afterward to provide an initial guess for the duration of the

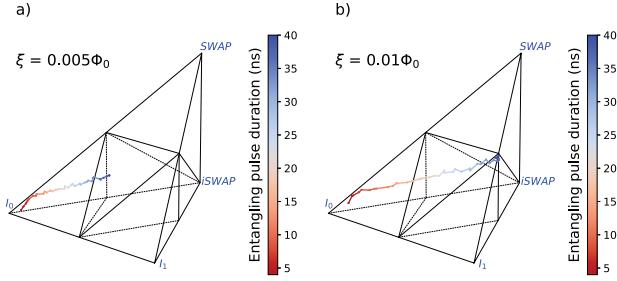


Figure 5. Stability over drive amplitude of the experimentally measured Cartan coordinate trajectories. In the same experimental implementation from Figure 2, as the entangling pulse drive amplitude ξ increased from $0.005\Phi_0$ to $0.01\Phi_0$, the Cartan coordinate trajectories were found to double in speed but still be qualitatively similar. The data was collected over a two day period. As in Figure 2, due to an experimental hardware constraint the shortest possible entangling pulse duration was 4 ns, so the measured Cartan trajectories begin there.

good 2Q basis gates.

Our calibration protocol does not include the use of randomized benchmarking (RB) [54], [55], [56]. RB is best suited for architectures with specific target gates that are members of the Clifford group. Furthermore, interleaved RB [57] will estimate the gate infidelity but will provide no information about an error budget. In our setting we do not have a fixed 2Q gate as the goal of implementation and understanding the gate unitaries themselves is a primary goal. We have thus decided GST and QPT are more suitable for precise gate characterization.

The scalability of our proposed calibration method is not different from traditional approaches. Calibration techniques like QPT, RB, and GST can be applied to multiple 2Q gates on the same device in parallel, as long as the gates do not act on the same qubits. One can use an edge-coloring of the device connectivity graph to determine which gates to calibrate simultaneously. An edge-coloring of the grid graph takes 4 colors, one for a sparser connectivity (e.g. heavy hexagonal) takes fewer colors. Thus, for a superconducting device with typical connectivity, the calibration overhead on the quantum device does not scale with the size of the device.

VII. COMPIILING WITH NON-STANDARD 2Q BASIS GATES

Most quantum programs and benchmarks are already specified at the 2 or 3 qubit gate level. Therefore, like previous works [24][22][23] that discuss choice of 2Q basis gate and how to use less conventional 2Q basis gates for compilation, we use a transpiler pass to convert other 2Q gates in a circuit into our own 2Q basis gates, instead of building an entirely new compiler.

Some of the prior works decompose 2Q gates from application circuits into 1Q gates and native 2Q gates using a numerical approach [22], while others take an analytical approach [24] [23]. Note that such a decomposition re-

quires finding the 1Q local unitaries, not just determining the required circuit depth. The analytical and numerical approaches each have their advantages. The numerical approach is more flexible. It can be applied to any 2Q basis and target gates. The analytic methods have limits on what gates they can be applied to, but are faster and some of them guarantee optimal results. There is currently no analytic formula that convert between arbitrary sets of 2Q gates. Huang et al. [24] and Peterson et al. [23] develop analytic algorithms that decomposes an arbitrary 2Q gates into $\sqrt{i}SWAP$ and discrete sets of XX-type gates, respectively. The `decompose_two_qubit_interaction_into_four_fsim_gates` function in Cirq [58] implements an analytic formula that decomposes an arbitrary 2Q gate into 4 layers of a given fSim gate, via the B gate.

In this project, we need to synthesize other 2Q gates from 2Q basis gates that are even less conventional than the ones considered in previous work. Therefore, we take a mostly numerical approach to gate synthesis and write our numerical search code based on NuOp from [22]. The difference is, we use knowledge about decomposition circuit depth computed analytically to inform and speedup the numerical search for 1Q local unitaries. NuOp first attempts to search for a 1-layer decomposition, and moves on to 1 more layer upon failure to find solution, until it meets the target decomposition error rate. Using the analytic techniques for determining circuit depth developed by [44] and extended by our work for SWAP, we are able to skip to the step in NuOp in which a perfect decomposition is guaranteed by theory. This significantly speeds up the numerical search and also guarantee that the solution has optimal depth.

Synthesizing all 2Q gates in the application programs directly into the basis gates might incur a compilation overhead. We avoid it by computing in advance and storing the decompositions of a few common 2Q gates into our basis gates. This only needs to be done once per calibration cycle (usually 1 day) and costs little time. In this work (see Section VIII) we only directly decompose SWAP and CNOT into our basis gates. But instead of taking this minimalist approach, one can alternatively prepare decompositions for a larger set of potential target gates into the basis gates. The cost would still quite small. We imagine that one can identify a set of potentially useful target gates using an approach similar to [22], except that [22] looks for a set of gates to calibrate instead of decompose. In addition, in the scenario where programs wait in long queues before execution, one might be able to afford directly decomposing all 2Q gates in the circuits into the basis gates.

VIII. CASE STUDY: ENTANGLING FIXED FREQUENCY FAR-DETUNED TRANSMONS WITH A TUNABLE

COUPLER

A. Introduction to the case study entangling gate architecture

Many efforts are being made in industry and academia to design a 2Q entangling gate architecture that can be used for scaling up to a general quantum computer [59], [18], [37]. The all-microwave cross-resonance gate was recently used by IBM to do a high fidelity CNOT gate in 90 ns [59], but to suppress the always-on ZZ crosstalk mentioned in Section IV, precise crosstalk cancellation pulses applied to both qubits during run time were required, adding complexity to the architecture. Google Quantum AI and MIT have each developed entangling gate architectures for high fidelity CZ and iSWAP gates, with Google's architecture supporting a continuous set of these standard gates [18], [37]. Google's architecture requires all qubits and couplers to be flux-tunable, which adds complexity and additional sources of leakage and noise to their architecture. Similarly, in order to suppress the always-on ZZ crosstalk, MIT's architecture requires one qubit per pair to be tunable as well as the coupler.

The unit cell of our case study entangling gate architecture is a pair of qubits and a coupler. This unit cell, first proposed in [7], was designed to perform a diverse set of 2Q gates, including iSWAP and CZ; the full list of 2Q gates can be found in Table 1 of [7]. The two qubits are fixed frequency transmon qubits; the benefits of fixed frequency transmons are that they are easy to fabricate and can be reliably engineered to have high coherence > 100 us [60]. The two qubits are also far detuned from each other so there is reduced single qubit control crosstalk. The coupler is a generalized flux qubit which has been designed to have several good properties for qubit control. Notably, because the coupler's positive anharmonicity has been designed to balance out the negative anharmonicity of the two qubits, the eigenspectrum of this architecture's unit cell can support a zero-ZZ crosstalk bias point. This architecture is relatively simple to implement because fixed frequency transmons have high coherence, there is only one flux-tunable element in the unit cell (the coupler), and it is easy to bias the unit cell to zero-ZZ crosstalk.

A model 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:

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

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 its first excited state, can be varied in time via the flux through its superconducting loop. Low-crosstalk 2Q gates are realized by AC modulating this coupler frequency after DC biasing it to the zero-ZZ crosstalk bias

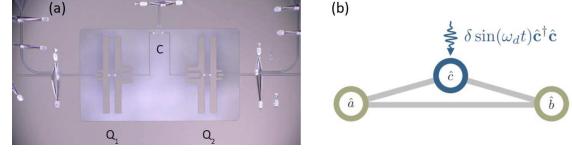


Figure 6. (a) Optical image of the device presented in [43] shows two fixed frequency transmons coupled via a tunable coupler. (b) Schematic for modelling the device adapted from [7].

point.

In [43] an early prototype device (shown in Fig. 6) for this case study architecture demonstrated a fast perfect entangler biased to zero-ZZ crosstalk. This device produced the nonstandard 2Q gate trajectory shown in Figure 2, which included a 13 ns perfect entangler. Figure 5 shows how the measured trajectories were similar over a range of entangling pulse drive amplitudes that did not exceed $\xi = 0.01\Phi_0$, the point at which strong drive effects would be expected to become non-negligible [7]. So in this early prototype device, the measured trajectories in Figures 2 and 5 were not nonstandard because of strong drive effects, but because of some other systematic in the experiment.

B. Our simulation approach

The case study entangling gate architecture natively supports strong parametrically activated interactions between the two qubits. Since the full Hamiltonian for this architecture is computationally intensive to model [7], for our simulation we use the simplified effective Hamiltonian from [7] that models the device using fewer parameters while still capturing all of the essential physics of the device (see Appendix A). Our general protocol for simulating Cartan trajectories is as follows:

- 1) We input the simulated device parameters into our Hamiltonian. These parameters include the qubit frequencies ω_a and ω_b , and the qubit coherence times. This generates the eigenspectrum of the simulated device.
- 2) 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.
- 3) We specify the drive amplitude ξ of our entangling pulse. In this case study we implement a iSWAP-like entangler, so the entangling pulse is driven at the frequency ω_d that generates maximal population swapping between the two qubits. For $\xi \leq 0.01\Phi_0$, the entangling pulse frequency ω_d is essentially identical to the difference frequency of the two qubits $|\omega_a - \omega_b|$. However, increasing $\xi > 0.01\Phi_0$ activates the two-photon process in Equation 2, causing population to enter the second excited state of the coupler and modify the entangling interaction. This in turn causes

ω_d to deviate from $|\omega_a - \omega_b|$. The entangling pulse is modulated by a rectangular envelope, as was done in experiment to obtain the measured trajectories; due to qubit controllers typically having a time resolution of 1 ns, short entangling gates ~ 10 ns have to be implemented using a pulse with a fast rise time. Experimentalists typically choose between a flat top Gaussian pulse with a short rise time, or a rectangular pulse for simplicity.

4) We evolve the time-dependent Hamiltonian and project the evolution propagator on the computational subspace to obtain the effective unitary operation with respect to the entangling pulse 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. We confirm that the leakage rates are much below the expected gate errors due to decoherence.

In this case study we simulate standard and nonstandard 2Q trajectories. The simplest and most consistent way to do this is to use the same simulated devices but to vary the drive power ξ . For $\xi \leq 0.01\Phi_0$ we expect the above protocol to result in a standard iSWAP interaction between the two qubits. But for $\xi > 0.01\Phi_0$, we expect strong drive effects to begin to emerge and cause the Cartan trajectory to deviate away from a standard iSWAP. We note that the simulated trajectories differ in several ways from the measured trajectories in Figures 2 and 5. Firstly, the measured trajectories are nonstandard even for $\xi \leq 0.01\Phi_0$ due to an additional systematic effect in the experiment. Secondly, the simulated trajectories are consistently slower than the measured trajectories; e.g. at $\xi = 0.01\Phi_0$, the simulated trajectories are slower by a factor of 3.5 than the measured trajectory, which included a 13 ns \sqrt{iSWAP} -like entangling gate. These discrepancies can both be explained by the simulation model Hamiltonian being significantly simpler than the true device Hamiltonian. Aside from these discrepancies, our simulations are realistic; our trajectories are generated using parameters and techniques that closely resemble those used in experiment and our method for generating standard and nonstandard trajectories using a single simulated device is physically intuitive.

Simulating Cartan trajectories over a range of entangling pulse amplitudes ξ we observe the correct intuitive behavior. The simulated trajectories deviate more and more from the standard iSWAP as the entangling pulse amplitude increases beyond $\xi = 0.01\Phi_0$. Secondly, the speed of the simulated trajectories scales linearly with ξ . This agrees with the experimental data shown in Figure 5 where the measured trajectory doubled in speed when ξ increased by a factor of two.

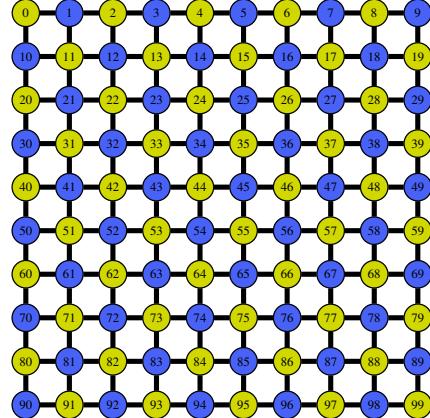


Figure 7. Device simulation. The high and low frequency qubits are shown in different colors. Each edge connects two qubits with different colors.

C. Methodology

We simulate a 10×10 device with grid connectivity (Fig. 7), where the qubit frequencies of each pair of neighbors are sampled from two normal distributions respectively with means that differ by 2 GHz. We use a 5% standard deviation for sampling the qubit frequencies. Improved fabrication techniques have reduced the smaller standard deviation to about 0.5% [34], but we use a larger standard deviation to show that our method is robust to variations in device fabrication.

Between each pair of neighboring qubits on the 10×10 grid, we simulate two types of 2Q trajectory by varying the entangling pulse amplitude ξ : 1) A baseline trajectory generated with a low entangling pulse amplitude of $\xi = 0.005\Phi_0$ and 2) a nonstandard trajectory due to strong drive effects resulting from a larger $\xi = 0.04\Phi_0$.

Then on each nonstandard trajectory, we select 2Q basis gates using Criterion 1 and 2 (respectively) introduced in Section V-E. We test these 3 sets of 2Q basis gates on common application circuits as benchmarks. We use the Qiskit[61] transpiler with the “SABRE”[62] layout and routing methods to map the benchmarks circuits to the 10×10 grid connectivity. With the nonstandard basis gates, we compile circuits using the methods from Section VII. With the \sqrt{iSWAP} from the standard trajectories, we use the analytic approach in [24]. Like the 2Q basis gates selected with Criterion 2, \sqrt{iSWAP} decomposes SWAP in 3 layers and CNOT in 2 layers, but we can also use it to directly decompose other 2Q gates (like the CRZ gates in the QFT benchmarks) analytically. For the 1Q gates in the gate and circuit synthesis, we use a duration of 20 ns, which is typical for fixed-frequency transmon qubit processors [42].

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 coherence time of the qubit. We

set T to a typical value of $80 \mu\text{s}$ 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 coherence-limited fidelity of a circuit is the product over the $e^{-t/T}$ term from each qubit. The decomposition errors in gate synthesis are negligible compared to the decoherence errors, and can be reduced to arbitrarily close to zero in theory. Thus we only show the coherence-limited fidelities in the results.

D. Results

Before discussing our results, as a disclaimer we note that while increasing the entangling pulse drive amplitude is one way to speed up 2Q gates, it is by no means an all-purpose solution that we generally advocate for. We chose to do this in our simulation case study only because it was a simple and intuitive way to compare standard and nonstandard simulated gates for the same case study entangling architecture. For this case study architecture, the drive amplitudes chosen were realistic in an experimental setting.

	Basis	SWAP	CNOT
Baseline	83.04 ns	329.1 ns	226.1 ns
	99.884%	99.541%	99.684%
Criterion 1	10.15 ns	110.5 ns	110.5 ns
	99.986%	99.845%	99.845%
Criterion 2	10.76 ns	112.3 ns	81.51 ns
	99.985%	99.843%	99.886%

Table I. Average duration (top) and coherence-limited gate fidelity (bottom) of the 2Q basis gates and the synthesized SWAP and CNOT gates, from baseline, Criterion 1, and Criterion 2.

The average durations and coherence limited fidelities (obtained using the Qiskit Ignis `coherence_limit` function [63]) of the synthesized SWAP and CNOT gates from the two approaches are summarized in Table I. In Table II, we show the coherence-limited circuit fidelities of 5 sets of benchmark circuits, when transpiled to different sets of 2Q basis gates. We first observe that the faster nonstandard 2Q basis gates have $\sim 8x$ lower coherence-limited infidelities than the baseline standard 2Q gates. We also observe that the synthesized SWAP (CNOT) gates from Criterion 1 and 2 are $3.0x$ and $2.9x$ ($2.0x$ and $2.8x$) faster than the baseline, respectively. Due to the relation between gate fidelity and circuit fidelity, fidelity improvements scale exponentially in benchmark size.

Next, we observe that Criterion 2 performs better than Criterion 1. This is not surprising since it has significantly faster CNOT gates and only slightly slower SWAP gates compared to Criterion 1.

For the baseline case, the 1Q gate duration is $4x$ shorter than the standard 2Q basis gate, and therefore $\sim 24\%$ of

the duration of the compiled SWAP/CNOT gate is spent performing 1Q gates. In contrast, for the nonstandard case, the 1Q gate duration is $2x$ longer than the nonstandard 2Q basis gate, and $\sim 72\%$ of the duration of the compiled SWAP/CNOT gate is spent performing 1Q gates. This puts us in the regime of today's fastest large superconducting processors such as Google's Sycamore device, where the optimal processor configuration that minimizes the overall effects of gate error has the 1Q gates being roughly twice as long as the 2Q gates [64].

IX. CONCLUSION

The idea of a uniform set of basis gates naturally arose from early notions of universal gate sets, which experimentalists then implemented on various qubit platforms. By looking at the theory of possible entanglers, we have found that there are many options for good 2Q basis gates, and that these gates behave differently on each pair of interacting qubits in a processor. This led us to a radically new idea, why be constrained to a single canonical gate (e.g. 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 instead adjust for these variations in software? If we do not treat all the coherent deviations in gate trajectories as errors, we will have more freedom in hardware design and achieve a higher gate fidelity.

In this paper, we examined the space of possible entanglers and developed a practical method for finding a high-fidelity entangler between every pair of qubits. In the case study, we find heterogeneous basis gates that are $\sim 8x$

Benchmark	Baseline	Criterion 1	Criterion 2
qft 10	58.2%	65.6%	70.8%
qft 20	1.33%	6.03%	9.94%
bv 9	88.7%	94.4%	95.3%
bv 19	79.3%	89.9%	91.0%
bv 29	44.5%	72.5%	74.3%
bv 39	26.8%	56.3%	59.7%
bv 49	27.7%	58.4%	62.4%
bv 59	12.5%	43.8%	47.4%
bv 69	9.15%	39.4%	43.2%
bv 79	0.428%	11.3%	14.2%
bv 89	2.44%	23.1%	26.3%
bv 99	0.06%	6.26%	7.97%
cuccaro 10	21.5%	46.3%	52.6%
cuccaro 20	0.800%	7.68%	11.8%
qaoa_0.1 10	97.2%	98.5%	98.8%
qaoa_0.1 20	84.4%	92.0%	93.6%
qaoa_0.1 30	14.4%	43.3%	49.0%
qaoa_0.1 40	0.00585%	5.59%	8.56%
qaoa_0.33 10	66.1%	81.0%	84.3%
qaoa_0.33 20	15.0%	42.2%	48.2%

Table II. The decoherence-limited fidelities of benchmark circuits, transpiled using the standard 2Q basis gates from baseline, and the nonstandard ones selected by Criterion 1 and 2. The QAOA benchmarks all have $p = 1$ where p is the number of times the protocol is repeated. The fractions 0.1 and 0.33 are the probabilities that an edge is created between a pair of nodes.

faster than the baseline, and use them to synthesize faster SWAP and CNOT gates than synthesized by the baseline \sqrt{iSWAP} gate from the standard XY-type trajectories. We then evaluate these heterogeneous basis gates on a number of benchmark circuits and find fidelity improvements that scale exponentially in benchmark size.

Our approach successfully uses software to overcome the limitations of today's hardware. Such types of adaptive basis-gate design will be essential to pioneering innovative future quantum systems.

ACKNOWLEDGMENTS

We would like to thank Charlie Guinn for discussions about calibration and tomography techniques.

This work is funded in part by EPiQC, an NSF Expedition in Computing, under grants CCF-1730082/1730449; in part by STAQ under grant NSF Phy-1818914; in part by NSF Grant No. 2110860; by the US Department of Energy Office of Advanced Scientific Computing Research, Accelerated Research for Quantum Computing Program; and in part by NSF OMA-2016136 and the Q-NEXT DOE NQI Center. SS is supported by the Department of Defense (DoD) through the National Defense Science & Engineering Graduate Fellowship (NDSEG) Program. FTC is Chief Scientist at Super.tech and an advisor to Quantum Circuits, Inc.

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 in [7]:

$$\hat{H}(t) = \hat{H}_a + \hat{H}_b + \hat{H}_c(t) + \hat{H}_g, \quad (3)$$

with

$$\begin{aligned} \hat{H}_a &= \omega_a \hat{a}^\dagger \hat{a} + \frac{\alpha_a}{2} \hat{a}^\dagger 2 \hat{a}^2, \\ \hat{H}_b &= \omega_b \hat{b}^\dagger \hat{b} + \frac{\alpha_b}{2} \hat{b}^\dagger 2 \hat{b}^2, \\ \hat{H}_c(t) &= \omega_c(t) \hat{c}^\dagger \hat{c} + \frac{\alpha_c}{2} \hat{c}^\dagger 2 \hat{c}^2. \\ \hat{H}_g &= -g_{ab} \hat{a}^\dagger \hat{b} - g_{bc} \hat{b}^\dagger \hat{c} - g_{ca} \hat{c}^\dagger \hat{a} \\ &\quad - g_{ab}^* \hat{a}^\dagger - g_{bc}^* \hat{b}^\dagger - g_{ca}^* \hat{c}^\dagger \end{aligned} \quad (4)$$

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 $\omega_c(t) = \omega_c^0 + \delta \sin(\omega_d t)$.

APPENDIX B.

SWAP SYNTHESIS IN 2 LAYERS

See the circuit in Fig. 3(a). 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. 3(a) 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] 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.
- [2] J. Preskill, "Quantum computing in the nisq era and beyond," *Quantum*, vol. 2, p. 79, 2018.
- [3] 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.
- [4] 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>
- [5] 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>
- [6] "pyGSTi," <https://www.pygsti.info>, 2022.
- [7] 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.
- [8] E. Bernstein and U. Vazirani, "Quantum complexity theory," *SIAM Journal on computing*, vol. 26, no. 5, pp. 1411–1473, 1997.

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

[10] 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.

[11] 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.

[12] 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.

[13] G. E. Crooks, “Gates, states, and circuits,” 2020.

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

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

[16] 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.

[17] L. Lao, A. Korotkov, Z. Jiang, W. Mruczkiewicz, T. E. O’Brien, and D. E. Browne, “Software mitigation of coherent two-qubit gate errors,” *Quantum Science and Technology*, vol. 7, no. 2, p. 025021, 2022.

[18] 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. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.125.120504>

[19] I. N. Moskalenko, I. A. Simakov, N. N. Abramov, A. A. Grigorev, D. O. Moskalev, A. A. Pishchimova, N. S. Smirnov, E. V. Zikiy, I. A. Rodionov, and I. S. Besedin, “High fidelity two-qubit gates on fluxoniums using a tunable coupler,” 2022. [Online]. Available: <https://arxiv.org/abs/2203.16302>

[20] H. Xiong, Q. Ficheux, A. Somoroff, L. B. Nguyen, E. Dogan, D. Rosenstock, C. Wang, K. N. Nesterov, M. G. Vavilov, and V. E. Manucharyan, “Arbitrary controlled-phase gate on fluxonium qubits using differential ac stark shifts,” *Phys. Rev. Research*, vol. 4, p. 023040, Apr 2022. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevResearch.4.023040>

[21] A. D. Hill, M. J. Hodson, N. Didier, and M. J. Reagor, “Realization of arbitrary doubly-controlled quantum phase gates,” 2021. [Online]. Available: <https://arxiv.org/abs/2108.01652>

[22] 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.

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

[24] 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.

[25] D. R. Perez, P. Varošy, Z. Li, T. Roy, E. Kapit, and D. Schuster, “Error-divisible two-qubit gates,” 2021. [Online]. Available: <https://arxiv.org/abs/2110.11537>

[26] X. Xu and M. H. Ansari, “Parasitic-free gate: A protected switch between idle and entangled states,” 2022. [Online]. Available: <https://arxiv.org/abs/2202.05208>

[27] U. Vool and M. Devoret, “Introduction to quantum electromagnetic circuits,” in *International Journal of Circuit Theory and Applications*, vol. 45, no. 7, 2017.

[28] “AWR Microwave Office,” <https://www.awr.com/>, 2022.

[29] “ANSYS Maxwell,” <https://www.ansys.com/>, 2022.

[30] S. Schlör, J. Lisenfeld, C. Müller, A. Bilmes, A. Schneider, D. P. Pappas, A. V. Ustinov, and M. Weides, “Correlating decoherence in transmon qubits: Low frequency noise by single fluctuators,” *Phys. Rev. Lett.*, vol. 123, p. 190502, Nov 2019. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.123.190502>

[31] M. Spiecker, P. Paluch, N. Drucker, S. Matityahu, D. Gusekova, N. Gosling, S. Günzler, D. Rieger, I. Takmakov, F. Valenti, P. Winkel, R. Gebauer, O. Sander, G. Catelani, A. Shnirman, A. V. Ustinov, W. Wernsdorfer, Y. Cohen, and I. M. Pop, “A quantum szilard engine for two-level systems coupled to a qubit,” 2022. [Online]. Available: <https://arxiv.org/abs/2204.00499>

[32] C. Wang, C. Axline, Y. Y. Gao, T. Brecht, Y. Chu, L. Frunzio, M. H. Devoret, and R. J. Schoelkopf, “Surface participation and dielectric loss in superconducting qubits,” *Applied Physics Letters*, vol. 107, 2015.

[33] J. M. e. a. Kreikebaum, “Improving wafer-scale josephson junction resistance variation in superconducting quantum coherent circuits,” *Supercond. Sci. Technol.*, vol. 33, p. 033201, 2020. [Online]. Available: <https://iopscience.iop.org/article/10.1088/1361-6668/ab8617/meta>

[34] 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.

[35] 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>

[36] 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>

[37] 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,” *Phys. Rev. X*, vol. 11, p. 021058, Jun 2021. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevX.11.021058>

[38] 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.

[39] 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.

[40] 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.

[41] 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>

[42] 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.

[43] C. Guinn, S. Sussman, P. S. Mundada, A. Vrajitoarea, C. Leroux, A. Place, C. L. Calonne, A. Di Paolo, A. Petrescu, A. Blais, and A. A. Houck, “Fast parametrically driven entangling gates in superconducting circuits using a tunable coupler,” *Bulletin of the American Physical Society*, 2022.

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

[45] 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.

[46] 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.

[47] 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.

[48] D. Greenbaum, “Introduction to quantum gate set tomography,” *arXiv preprint arXiv:1509.02921*, 2015.

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

[50] X. Xue, M. Russ, N. Samkharadze, B. Undseth, A. Sammak, G. Scappucci, and L. M. K. Vandersypen, “Quantum logic with spin qubits crossing the surface code threshold,” *Nature*, vol. 601, no. 7893, pp. 343–347, 2022. [Online]. Available: <https://doi.org/10.1038/s41586-021-04273-w>

[51] M. T. Mądzik, S. Asaad, A. Youssry, B. Joecker, K. M. Rudinger, E. Nielsen, K. C. Young, T. J. Proctor, A. D. Baczewski, A. Laucht, V. Schmitt, F. E. Hudson, K. M. Itoh, A. M. Jakob, B. C. Johnson, D. N. Jamieson, A. S. Dzurak, C. Ferrie, R. Blume-Kohout, and A. Morello, “Precision tomography of a three-qubit donor quantum processor in silicon,” *Nature*, vol. 601, no. 7893, pp. 348–353, 2022. [Online]. Available: <https://doi.org/10.1038/s41586-021-04292-7>

[52] S. T. Merkel, J. M. Gambetta, J. A. Smolin, S. Poletto, A. D. Córcoles, B. R. Johnson, C. A. Ryan, and M. Steffen, “Self-consistent quantum process tomography,” *Phys. Rev. A*, vol. 87, p. 062119, 2013.

[53] R. Blume-Kohout, M. P. da Silva, E. Nielsen, T. Proctor, K. Rudinger, M. Sarovar, and K. Young, “A taxonomy of small markovian errors,” *PRX Quantum*, vol. 3, no. 2, may 2022. [Online]. Available: <https://doi.org/10.1103/%2Fprxquantum.3.020335>

[54] 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.

[55] 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.

[56] 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.

[57] 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.

[58] C. Developers, “Cirq,” Aug 2021, see full list of authors on Github: <https://github.com/quantumlib/Cirq/graphs/contributors>.

[59] K. X. Wei, E. Magesan, I. Lauer, S. Srinivasan, D. F. Bogorin, S. Carnevale, G. A. Keefe, Y. Kim, D. Klaus, W. Landers, N. Sundaresan, C. Wang, E. J. Zhang, M. Steffen, O. E. Dial, D. C. McKay, and A. Kandala, "Hamiltonian engineering with multicolor drives for fast entangling gates and quantum crosstalk cancellation," *Phys. Rev. Lett.*, vol. 129, p. 060501, Aug 2022. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.129.060501>

[60] A. P. M. Place, L. V. H. Rodgers, P. Mundada, B. M. Smitham, M. Fitzpatrick, Z. Leng, A. Premkumar, J. Bryon, A. Vrajitoarea, S. Sussman, G. Cheng, T. Madhavan, H. K. Babla, X. H. Le, Y. Gang, B. Jäck, A. Gynis, N. Yao, R. J. Cava, N. P. de Leon, and A. A. Houck, "New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds," *Nature Communications*, vol. 12, no. 1, p. 1779, 2021. [Online]. Available: <https://doi.org/10.1038/s41467-021-22030-5>

[61] 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. Cholaraian, 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. Dartailh, 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. Foulland, 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, E. 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, Kovil1Day, 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. Moyer, 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. Nicanter, Nick-Singstock, P. Niroula, H. Norlen, NuoWenLei, L. J. O'Riordan, O. Ogundbayo, 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. Poreau, 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. Rossmanek, 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. Sิตdkov, S. Sivarajah, M. B. Sletfjerd, J. A. Smolin, M. Soeken, I. O. Sokolov, I. Sokolov, V. P. Soloviev, SooluThomas, Starfish, D. Steenken, M. Sty-pulkoski, A. Suau, S. Sun, K. J. Sung, M. Suwama, O. Słowiak, H. Takahashi, T. Takawale, I. Tavernelli, C. Taylor, P. Taylour, S. Thomas, K. Tian, M. Tillet, M. Tod, M. Tomasaki, C. Tornow, E. de la Torre, J. L. S. Toural, K. Trabing, M. Treinish, D. Trenev, TrishaPe, F. Truger, G. Tsilimigkounakis, D. Tului, 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, bcamorison, brandhsn, charmerDark, deeplokhande, dekel.meirom, dime10, dlasecki, ehchen, fanizzamarco, fs1132429, gadi, galeinston, georgezhou20, georgios ts, gruu, hhori, hykavitha, itoko, jeppevinkel, jessica angel7, jezerjojo14, jliu45, jscott2, klinvill, krutik2966, ma5x, michelle4654, msuwama, nico lgrs, ntgiwsvp, ordmoj, sagar pahwa, pritamsinha2304, ryan-cocuzzo, saktar unr, saswati qiskit, septembr, sethmerkel, sg495, shaashwat, smturro2, sternparky, strickroman, tiger-jack, tsura crisaldo, upsideon, vadebayo49, welien, willbang, wmurphy collabstar, yang.luh, and M. Čepulkovskis, "Qiskit: An open-source framework for quantum computing," 2021.

[62] 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.

[63] “Qiskit Ignis,” <https://github.com/Qiskit/qiskit-ignis>, 2019.

[64] 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. Kostitsa, D. Landhuis, M. Lindmark, E. Lucero, D. Lyakh, S. Mandrà, J. R. McClean, M. McEwen, A. Megrant, X. Mi, K. Michelsen, 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, pp. 505–510, 2019. [Online]. Available: <https://doi.org/10.1038/s41586-019-1666-5>