

Efficient Control of a Transmon Qudit Using Effective Spin-7/2 Rotations

Elizabeth Champion^{✉,*}, Zihao Wang^{✉,*}, Rayleigh W. Parker[✉], and Machiel S. Blok^{✉,†}

Department of Physics and Astronomy, *University of Rochester*, Rochester, New York 14627, USA

 (Received 13 May 2024; revised 18 February 2025; accepted 6 May 2025; published 18 June 2025)

Qudits hold great promise for efficient quantum computation and the simulation of high-dimensional quantum systems [1]. However, existing control and measurement schemes for qudit systems scale unfavorably with qudit dimension since they decompose $SU(d)$ operations into series of qubitlike rotations and perform measurements on a small number of states [2–6]. Here, we address these challenges by employing simultaneous multifrequency drives to generate rotations and projections in an effective spin-7/2 system mapped onto the energy eigenstates of a superconducting circuit. We implement single-shot readout of the eight states using a multitone dispersive readout ($F_{\text{assignment}} = 88.3\%$) and exploit the strong nonlinearity in a high- E_J/E_C transmon to simultaneously address each transition and realize a spin displacement operator. Combining this displacement operator with a virtual SNAP gate, we realize arbitrary single-qudit unitary operations in $\mathcal{O}(d)$ physical pulses and extract spin displacement gate fidelities ranging from 0.997 to 0.989 for virtual spins of size $j = 1$ to $j = 7/2$. We demonstrate the potential of our control scheme in three ways: the direct measurement of the spin qudit Wigner function, randomized benchmarking of a logical qubit encoded into the qudit state, and randomized benchmarking of the full qudit Clifford group. In the latter experiment, we implement the d -dimensional quantum Fourier transform with an average gate fidelity of 0.91(6) in $d = 8$. Our multifrequency approach to qudit control and measurement can be readily extended to other physical platforms that realize a multilevel system coupled to a cavity and can become a building block for efficient qudit-based quantum computation and simulation.

DOI: [10.1103/vbh4-lysv](https://doi.org/10.1103/vbh4-lysv)

Subject Areas: Quantum Physics, Quantum Information

I. INTRODUCTION

Quantum computing relies on the ability to encode, control, and measure quantum information with high precision in a large state space. The canonical approach uses qubits encoded in coupled (effective) spin-1/2 systems that can be controlled with high fidelity owing to their strong nonlinearity [7–10]. In contrast, continuous variable encodings use the infinite-dimensional state space of a linear bosonic mode, enabling, e.g., resource-efficient quantum error correction [11,12]. Qudits can combine the advantages of both approaches by encoding information in (effective) large spins: finite d -level systems ($d > 2$) with strong nonlinearity. Qudit algorithms are predicted to outperform their qubit counterparts at quantum information tasks including magic-state distillation for quantum error correction [13] and the synthesis of many-body gates [14].

Furthermore, qudits are promising candidates for the simulation of quantum fields [15–17] or large angular-momentum spins [18,19] owing to their natural correspondence with high-dimensional quantum systems.

Experimental qudit processors have recently been demonstrated using high-dimensional nuclear spins [2,3], hyperfine states of trapped ions [4], photonic circuits [5], and superconducting circuits [20]. So far, readout has been limited to sequential qubitlike readout or to up to four states simultaneously [21–26]. Universal qudit control in these processors is typically realized by sequentially applying Givens rotations between adjacent energy states requiring $\mathcal{O}(d^2)$ qubitlike rotations to execute arbitrary single-qudit gates [27,28]. This process scales poorly with qudit dimension, however, and the fact that most states are idling at any given time makes approaches like continuous dynamical decoupling difficult [29]. Thus far, experiments that explored simultaneous drives on multiple transitions, either with optimal control pulses [30] or with known decompositions [31,32], were limited to three- or four-level systems without a clear path towards high-fidelity qudit gates in larger dimensions. A platform capable of implementing native multilevel rotations and readout would unlock the power of qudits by making efficient use of the larger Hilbert space and allowing a direct mapping

*These authors contributed equally to this work.

†Contact author: machielblok@rochester.edu

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

between the platform-native operations and qudit algorithms or simulations.

In this work, we realize spin displacement operations in a superconducting transmon qudit with up to $d = 8$ levels. This process is facilitated by a modified transmon design with a deep cosine potential that encodes qudit states with coherence times close to the lifetime limit, enabling high-fidelity control (see Table III in Appendix A). Extending qubit dispersive cavity readout to a multitone measurement, we realize simultaneous single-shot readout of the eight energy states with average assignment fidelity of 88.3%. Unlike their bosonic counterpart, spin systems are inherently finite dimensional and have a spherical rather than planar phase space. The family of spin displacement unitaries likewise consists of rotations on this phase space. We leverage the strong nonlinearity of the transmon to realize spin displacement operators by simultaneously driving the individual transitions. Utilizing this family of displacement operations, we perform direct Wigner tomography of the qudit state and verify the Wigner function's expected behavior. Next, we prepare spin cat states by combining the displacement operator with a virtual selective number-dependent arbitrary phase (SNAP) gate, which is implemented with close-to-unity fidelity by updating the respective drive pulse phases. Finally, we demonstrate a gate set consisting of a single displacement pulse and the virtual SNAP gate to perform arbitrary qudit rotations in $\mathcal{O}(d)$ physical pulses. We benchmark the performance of the displacement pulse in two ways: randomized benchmarking of a logical qubit encoded in the qudit Hilbert space [33] and randomized benchmarking of the full qudit Clifford group [34]. We measure native gate fidelities of $F_{\hat{D}} = \{0.9971 \pm 0.0001, 0.9931 \pm 0.0001, 0.9895 \pm 0.0003\}$ for $d = 3$, $d = 5$, and $d = 8$, respectively, demonstrating the potential for our control scheme to realize high-fidelity qudit control.

II. RESULTS

A. Transmon qudits

The fixed-frequency superconducting transmon oscillator is often used as a physical qubit for quantum information processing [35]. It can be represented as a lumped-element circuit consisting of a capacitor with characteristic charging energy E_C and a nonlinear inductance in the form of a Josephson junction of characteristic energy E_J . The resulting circuit behaves as an anharmonic oscillator, with a Kerr nonlinearity as the lowest-order correction to the harmonic potential. The two lowest-lying energy eigenstates are typically used to encode a qubit in an effective spin-1/2 system. Typical ratios of $E_J/E_C \approx 50$ suppress charge noise-induced dephasing within this qubit subspace while producing an appreciable anharmonicity of around 300 MHz [35,36]. The qubit is controlled using a resonant microwave-frequency drive to address the




FIG. 1. Encoding large spins in high- E_J/E_C transmon qudits and mapping a large spin onto a transmon. We dispersively couple a transmon with $E_J/E_C \approx 250$ to a readout resonator. The resonator enables projective readout of the qudit state, while the large E_J/E_C increases the number of states confined in the cosine potential and suppresses charge noise in the highly excited states. We map the Hilbert space of a large spin onto the transmon qudit such that the spin's $\hat{J}_z$ eigenstates correspond to the transmon eigenstates. Under this mapping, our system resembles a large spin dispersively coupled to a cavity, with a small quadrupole term corresponding to the lowest-order (Kerr) nonlinearity of the transmon.

$|0\rangle \leftrightarrow |1\rangle$ transition. Readout is enabled by dispersively coupling the qubit to an LC oscillator realized as a coplanar waveguide resonator. The lumped-element representation of this circuit is shown in Fig. 1.

While the two lowest-energy eigenstates are well suited for qubit computation, one can include the higher excited states to realize a qudit [20,24,27,28]. The anharmonicity due to the Josephson junction ensures that each of the excited-state transitions is spectrally resolved. Typical E_J/E_C ratios permit the use of a few excited states beyond the qubit subspace; however, the number of states confined within the Josephson potential is relatively small, and the higher states are subject to appreciable charge noise [35]. We address both shortcomings with a modified device design having $E_J/E_C \approx 250$. The cosine potential and phase-basis energy-eigenstate wave functions are shown in the lower panel of Fig. 1. In the following experiments, we use up to seven excited states for qudit computation in $d = 8$ dimensions. The Hilbert space of an eight-level qudit is equivalent to an effective spin-7/2 system, and the spectral separation of our eigenstates can be seen as analogous to a high-dimensional spin with a quadrupole

term in its Hamiltonian. In this picture, the harmonic part of the transmon potential, $\hat{a}^\dagger \hat{a}$, plays the role of a Zeeman Hamiltonian $\gamma_n B_0 \hat{J}_z$ due to an applied magnetic field, while the Kerr term $\hat{a}^{\dagger 2} \hat{a}^2$ corresponds to the spin quadrupole term $Q \hat{J}_z^2$. In a transmon, the anharmonicity across all transitions is approximately given by $K/(2\pi) \approx -E_C$, where, in our device, we estimate $E_C = 108$ MHz. Including the readout resonator, one can imagine our system as a large spin dispersively coupled to a cavity, where the $\hat{J}_z$ eigenstates are mapped onto the transmon eigenstates (Fig. 1, right).

As in the case of a qubit, the dispersively coupled resonator facilitates readout of the transmon state. In the present case, however, we have optimized the device design to permit single-shot readout of multiple levels. We show the dispersive shift of the readout resonator spectrum in Fig. 2(b) for transmon states $|0\rangle$ through $|8\rangle$. A single-frequency readout tone is able to distinguish several of these states, limited by the linewidth of the resonator relative to the dispersive shift. We realize single-shot readout of the entire qudit by multiplexing three readout tones with frequencies chosen such that together they can distinguish each state from the others [23]. We demodulate the signal reflected from the resonator at each of these three frequencies and integrate each to produce three pairs of IQ values. We consider these pairs as a single point in a six-dimensional space, which is classified to determine the

readout result. Altogether, we find an assignment fidelity of $(88.3 \pm 0.2)\%$; see Appendix B for details. We note that these readout methods were demonstrated and characterized in much greater detail in Ref. [37], which was prepared concurrently with the present work.

The spectral separation of the qudit transitions allows us to selectively apply rotations between adjacent states $|n-1\rangle$ and $|n\rangle$, given that the drive strength Ω is considerably less than the anharmonicity, which, in our device, is $E_C = 108$ MHz. In combination with the qudit readout, this enables coherence time measurements of each transition. We measure the energy relaxation rate $T_1^{(n-1)(n)}$ for each given transition $|n\rangle \rightarrow |n-1\rangle$ and find $T_1^{01} = 46$ μ s down to $T_1^{67} = 13$ μ s, where the increased relaxation for the higher transitions is largely explained by the scaling of the matrix elements of the lowering operator, analogous to bosonic enhancement for harmonic modes. We perform similar measurements of T_2 and T_2^* using $\pi/2$ pulses to perform echo and Ramsey experiments within each subspace and find dephasing times (T_2) close to the lifetime limit; see Appendix A for the full set of coherence times.

B. Spin displacements

With our coherent, transition-selective control in mind, we now consider the implementation of a spin-qudit displacement operator, beginning with a brief overview of high-dimensional spins and their displacements.




FIG. 2. Implementing control and measurement with multitone drives. (a) Illustration of the high- E_J/E_C fixed-frequency transmon and dispersively coupled readout resonator. Transmon drives are applied through the readout line. (b) Readout resonator spectrum for each transmon eigenstate. We perform projective readout of the qudit state by multiplexing three readout tones at different frequencies such that each state can be distinguished from the others by at least one tone. (c) Qudit readout in the IQ plane of each readout tone. (d) Implementing the spin displacement operator by simultaneously driving multiple transitions. The pulse amplitudes are scaled such that the largest element of the displacement generator matrix is held roughly fixed, producing slower oscillations as the number of dimensions is increased. The solid lines show single-parameter fits to the experimental data accounting for an overall scale factor on the pulse amplitude, of order unity. We also compute the expectation value of the $\hat{J}_z$ angular-momentum operator for each case (bottom). The displacement drive induces sinusoidal oscillations in the angular momentum, realizing a complete inversion of the effective angular momentum when the displacement angle is $\theta = \pi$.

The Hilbert space of a spin- j system is spanned by a basis of $d = 2j + 1$ states, usually labeled $|j, m\rangle$ for $-j \leq m \leq j$. These states can be taken to be the eigenstates of the angular-momentum operator $\hat{J}_z$. The dynamical symmetry group of a spin of any dimension is $SU(2)$ with corresponding Lie algebra $\mathfrak{su}(2)$. The representation of this Lie algebra in a spin- j system consists of $d \times d$ Hermitian matrices, spanned by the basis $\{\hat{J}_x, \hat{J}_y, \hat{J}_z\}$. The displacement operator of a spin- j system is thus an element of the $SU(2)$ representation generated by this basis:

$$\hat{D}(\theta, \phi) = e^{-i\phi\hat{J}_z} e^{-i\theta\hat{J}_y} = e^{\alpha\hat{J}_+ - \alpha^*\hat{J}_-}, \quad (1)$$

where the spin ladder operators are $\hat{J}_\pm \equiv \hat{J}_x \pm i\hat{J}_y$ and $\alpha \equiv -(\theta/2)e^{-i\phi}$ [38,39]. The coherent states of a high-dimensional spin are those generated by applying the displacement operator to a fixed state in the Hilbert space, usually taken to be the ground state: $|\theta, \phi\rangle = \hat{D}(\theta, \phi)|j, j\rangle$. As in the case of a harmonic oscillator, applying a displacement operator to a coherent state yields another coherent state; furthermore, the coherent states form an overcomplete basis for the Hilbert space. The phase-space representation of a spin is closely related to these coherent states. The phase space itself is the surface of a sphere, and the displacement angles θ and ϕ specify polar and azimuthal angles on the phase-space sphere, respectively.

We translate this picture to our transmon qudits by mapping the spin state $|j, m\rangle$ onto the transmon oscillator state $|n\rangle = |j - m\rangle$. The ground state of the spin system, assuming a Zeeman Hamiltonian proportional to $-\hat{J}_z$, is $|j, j\rangle$, which, under this mapping, corresponds to the transmon's ground state, $|0\rangle$. In this notation, the action of a spin ladder operator is to increment the number of excitations in the transmon: $\hat{J}_+|n\rangle = \sqrt{n(d-n)}|n+1\rangle$ and $\hat{J}_-|n\rangle = \sqrt{(n+1)(d-n-1)}|n-1\rangle$. We note that, in the limit of $n \ll d$, these ladder operators approach the bosonic operators $\hat{a}$ and $\hat{a}^\dagger$, respectively, up to a multiplicative constant.

We implement high-dimensional spin displacements in transmons by exploiting the nonlinearity of the system to simultaneously drive multiple transitions [32]. We first note that, under the rotating-wave approximation (RWA), a single resonant drive at the $|n-1\rangle \leftrightarrow |n\rangle$ transition frequency produces an effective $\hat{\sigma}_x$ -like Hamiltonian within that subspace, thus inducing Rabi oscillations between those levels. The validity of the RWA is contingent on the applied drive strength being appreciably smaller than the anharmonicity of the qudit ($\Omega \ll K$). We calibrate the drive strength necessary to reach a given Rabi rate separately for each transition in our Hilbert space. Because the spin ladder operators $\hat{J}_\pm$ connect only adjacent states, we observe that we can drive displacement operations by simultaneously driving every transition. The drive strengths are scaled such that, if the drive on $|n-1\rangle \leftrightarrow |n\rangle$ were applied on its own,

it would yield Rabi oscillations of Rabi frequency $\Omega(t)\sqrt{n(d-n)}$ for some global time-dependent envelope $\Omega(t)$. The total displacement induced by these drives after a time T is $\hat{D}(\theta, \phi)$, where $\theta = \int_0^T \Omega(t)dt$ and ϕ is the phase applied to each drive. We use a flattop pulse with cosine ramps at the beginning and end, each having a duration of 1/4 of the total pulse length.

Applying the displacement drive to the ground state produces spin coherent states, shown in Fig. 2(d) for several different Hilbert space dimensions. We note that the dimensionality of the emulated spin is artificially imposed by restricting our drives to the appropriate subset of transitions. After applying a displacement pulse of a given duration, we perform qudit readout and plot the populations. We scale the drive amplitudes between dimensions such that the largest matrix element of the displacement generator is kept roughly fixed, producing slower oscillations as the number of dimensions is increased. Because we have chosen to map $\hat{J}_z$ eigenstates onto the transmon eigenstates, we can directly compute the expectation value $\langle \hat{J}_z \rangle$ from the populations, shown in the bottom panel. When the area of the displacement pulse is $\theta = \pi$, the population is transferred entirely to the highest excited state, $|d-1\rangle$, corresponding to a full inversion of the effective angular momentum. These results illustrate that we can drive high-fidelity $SU(2)$ rotations of emulated spins of varying dimension, limited only by the number of qudit states available in our system.

C. Wigner tomography

One application of this $SU(2)$ displacement operation is the direct measurement of the spin qudit Wigner function. Given a density matrix $\hat{\rho}$, the Wigner function at a phase-space coordinate α is given by $W(\alpha) = \text{Tr}[\hat{\rho} \hat{\Delta}(\alpha)]$, where $\hat{\Delta}(\alpha)$ is an operator-valued function over phase space known as the kernel. The appropriate kernel for a given system is not unique; it only needs to satisfy the Stratonovich-Weyl postulates [39–41]. In the case of a spin, it can be shown that one such kernel has the form of a displaced paritylike operator, in analogy with the harmonic oscillator: $\hat{\Delta}(\alpha) = 2\hat{D}(\alpha)\hat{\Pi}\hat{D}(-\alpha)$, where $\hat{\Pi}$ is diagonal in the $\hat{J}_z$ basis (Appendix C 3). We measure the Wigner function at a point α by applying a displacement $\hat{D}(-\alpha)$ and performing qudit readout to extract the expectation value of $\hat{\Pi}$. We demonstrate this process in Figs. 3(b) and 3(c) for energy eigenstates and superposition states, respectively. As in the case of a harmonic oscillator, energy eigenstates are completely delocalized in phase, producing rotationally symmetric phase-space distributions. The ground state and highest excited state are, themselves, examples of spin coherent states and therefore have nearly entirely positive Wigner functions, while the other energy eigenstates have large Wigner-negative regions indicative of nonclassicality.




FIG. 3. Representing transmon qudit states in phase space. We plot a polar projection of the $SU(2)$ phase-space sphere onto a disk such that the radial coordinate is θ and the angular coordinate is ϕ ; in this projection, the north pole of the $SU(2)$ sphere lies at the center of the disk, while the south pole is spread across its boundary. The color bar in (b) applies to all Wigner function plots. (a) Pulse sequence for Wigner tomography. Beginning in the ground state, the desired state is constructed using a series of π pulses for (b) and a $\pi/2$ pulse followed by π pulses for (c). Next, a displacement drive is applied, and the qudit state is measured. Populations estimated from many repetitions of this sequence are used to construct the Wigner function. (b) Measured Wigner functions for energy eigenstates in $d = 5$ dimensions. (c) Measured Wigner functions for equal superpositions of the ground state and each excited state. To the right, we show the real parts of the density matrices recovered from the Wigner functions. (d) SNAP-displacement sequence for encoding a cat state in $d = 8$. The sequence maps $|0\rangle$ onto the cat state $(|(\pi/2), 0\rangle + |-(\pi/2), 0\rangle)/\sqrt{2}$, which we use as the logical qubit state $|0_L\rangle$. We omit the initial SNAP gate layer since, in this case, it acts on $|0\rangle$ and thus has no physical effect.

D. Universal qudit control

In addition to enabling Wigner tomography, we now propose spin displacements as primitive gates for efficient qudit computation. Qudit gates are typically realized by decomposing the desired unitary into a series of Givens rotations within two-level subspaces of adjacent states $\{|n-1\rangle, |n\rangle\}$. Such decompositions require, in general, $d(d-1)/2$ such rotations [24,27,42], which can become costly for high-dimensional qudits. Inspired by the SNAP-displacement control scheme for bosonic systems [43–46], we find that interleaving N spin displacement operations with $N+1$ arbitrary phase rotations provides universal control over the qudit Hilbert space, given sufficiently large N (Appendix D):

$$\hat{U} = \hat{S}(\vec{\phi}^{(N)})\hat{D}(\theta^{(N)})\dots\hat{S}(\vec{\phi}^{(1)})\hat{D}(\theta^{(1)})\hat{S}(\vec{\phi}^{(0)}). \quad (2)$$

The required circuit depth for this scheme, N , is advantageous as compared to Givens rotations: Numerically, we find that only $\mathcal{O}(d)$ displacement layers are required to realize arbitrary operations in $SU(d)$. Moreover, unlike bosonic systems, we implement the SNAP layers simply by

updating subsequent drive phases, producing virtual rather than physical phase rotations [28,47]. Thus, any single-qudit unitary can, in principle, be realized by $\mathcal{O}(d)$ physical pulses. In principle, arbitrary $SU(2)$ rotations can be realized using two $\pi/2$ rotations alternating with $\hat{J}_z$ rotations (which are themselves examples of SNAP gates), so each general displacement $\hat{D}(\theta)$ could be replaced with two $\hat{D}(\pi/2)$ displacements with a virtual $\hat{J}_z$ rotation between them. In practice, however, we find this process unnecessary: Restricting each displacement in Eq. (2) to be $\hat{D}(\pi/2)$ has little effect on the number of layers required to implement an arbitrary unitary. We therefore focus our efforts on calibrating high-fidelity $\pi/2$ displacements and decompose arbitrary unitaries into sequences of this single pulse interleaved with virtual phase rotations. To implement a given target unitary, we then find the decomposition parameters $\vec{\phi}^{(i)}$ using numerical optimization.

We demonstrate the SNAP-displacement gate set in Fig. 3(d) by preparing a spin cat state, i.e., an equal superposition of opposite-phase spin coherent states, in $d = 8$ dimensions. In particular, we choose coherent states lying on the equator of the phase-space sphere, ideally

preparing the state $(|(\pi/2), 0\rangle + |-(\pi/2), 0\rangle)/\sqrt{2}$. In this case, we find that the cat state can be prepared with only two physical displacement pulses. We note that, much like the harmonic oscillator cat states, our spin cat state only contains even-parity eigenstates; likewise, the orthogonal cat state $(|(\pi/2), 0\rangle - |-(\pi/2), 0\rangle)/\sqrt{2}$ only contains odd-parity eigenstates. It is important to remember, however, that our cat state is a superposition of spin coherent states rather than harmonic oscillator coherent states, differing both in the amplitudes of the eigenstates and the overall dimensionality of the system.

1. Benchmarking the displacement fidelity

While the SNAP-displacement decomposition is universal over $SU(d)$, benchmarking its performance over the full Hilbert space presents a challenge. Below, we present results of qudit randomized benchmarking experiments, which yield estimates for the average fidelity with which we can perform qudit Clifford operations [34]. However, the size of the qudit Clifford group grows rapidly with dimension, as does the number of displacement gates required to implement each Clifford operation, making a direct analysis of the displacement fidelity itself challenging. Process tomography, meanwhile, is sensitive to SPAM errors, which poses a challenge when the gate errors themselves are smaller than SPAM errors. In this section, we instead perform randomized benchmarking of a logical qubit, namely, the spin cat encoding, where each logical Clifford gate is decomposed into a SNAP-displacement sequence [33]. We emphasize that, although this encoding has no beneficial error-correction properties given the noise channels in our system, it is a useful tool for benchmarking the displacement pulse. In particular, we find that all logical Clifford gates can be implemented in at most two displacement pulses, regardless of the underlying qudit dimension. This close relationship between the logical Clifford gates and the primitive pulse we wish to benchmark permits an accurate and reliable estimate of the displacement pulse fidelity.

We begin our randomized benchmarking procedure with the SNAP-displacement sequence discussed above, which maps the ground state $|0\rangle$ onto the logical state $|0_L\rangle$. We proceed in the typical fashion by sampling random sequences of Clifford gates of varying depth and appending the inverse Clifford operations to the end [48]. We compile each Clifford gate down to a sequence of alternating SNAP and displacement gates. As noted above, we find that all logical Clifford unitaries involving logical $\hat{\sigma}_x$ or $\hat{\sigma}_y$ rotations can be implemented with two displacement pulses, while those corresponding to pure logical $\hat{\sigma}_z$ rotations require only a single phase update per drive since the logical states are formed from disjoint sets of eigenstates. The final step is to invert the encoding sequence such that the population of $|0_L\rangle$ is mapped back to $|0\rangle$ prior to

readout. The average Clifford fidelity we recover is taken as a proxy for the overall fidelity of the displacement.

We show selected randomized benchmarking data in Fig. 4(b) for $d = 3$, $d = 5$, and $d = 8$. The resulting average logical Clifford gate fidelities are $F_{C_L} = \{0.9951 \pm 0.0002, 0.9886 \pm 0.0002, 0.9825 \pm 0.0006\}$ for $d = 3$, $d = 5$, and $d = 8$, respectively, obtained by fitting the exponential decay of the randomized benchmarking data. While in the idealized case our control pulses implement the generator of spin displacements, we find that off-resonant drive terms neglected in the RWA can cause coherent errors such as phase shifts and leakage, which can have a significant impact on the displacement fidelity for short pulse durations. We address this finding in several complementary ways. To lowest order, the phase offsets accumulated during the pulse can be compensated for by instantaneous phase corrections before and after the pulse. We compute these phase corrections from numerical simulations of the unitary evolution of the system under the displacement drive. We find that additional improvements to the fidelity are made by including small drive detunings and DRAG-like corrections [49], which correct the noncommuting phase errors not addressed by the instantaneous phase corrections and leakage, respectively. We calculate these corrections numerically using an optimal control code (see Appendix E and Ref. [50]). The optimal fidelities are achieved by utilizing all three types of corrections and by choosing displacement pulse lengths of 48 ns for $d = 3$, 100 ns for $d = 5$, and 140 ns for $d = 8$ to balance coherent errors, which increase with drive strength, and incoherent errors, which increase with pulse duration. We further demonstrate the effects of our pulse corrections and the overall duration of the displacement pulse in Fig. 4(e) for $d = 8$ by performing randomized benchmarking with displacement pulses of different durations and different combinations of pulse corrections. As a guide to the eye, we fit a phenomenological model in which the total infidelity is the sum of coherent and incoherent error terms: This model assumes that the coherent part decreases with the square of the pulse duration while the incoherent part increases linearly with pulse duration; it also assumes that each pulse correction type is subject to the same rate of incoherent error. We see that the pulse corrections indeed improve the fidelities, particularly for short pulse durations. For longer pulses, the incoherent error is dominant, so the fidelity curves converge to one another. We note that, because the dominant coherent errors are due to off-resonant driving (e.g., ac Stark shifts), the error is largely independent of the relative phases of the drives. The numerically optimized corrections are therefore independent of the drive phases, and our experiments make use of this single optimized pulse. We further test the stability of our control scheme and the reliability with which our pulse corrections improve the fidelity by repeating the same randomized benchmarking experiment




FIG. 4. Randomized benchmarking of a logical qubit controlled with SNAP-displacement decompositions. We form a two-dimensional logical basis of spin cat states embedded into the larger qudit Hilbert space. (a) Logical Bloch sphere with theoretical Wigner functions in $d = 8$. (b) Logical qubit randomized benchmarking in $d = 3$, $d = 5$, and $d = 8$. We extract average qubit Clifford gate fidelities of $F_{C_L} = \{0.9951 \pm 0.0002, 0.9886 \pm 0.0002, 0.9825 \pm 0.0006\}$, respectively, yielding estimated displacement fidelities of $F_{\hat{D}} = \{0.9971 \pm 0.0001, 0.9931 \pm 0.0001, 0.9895 \pm 0.0003\}$. (c) Validating displacement fidelities extracted from randomized benchmarking. We numerically simulate our randomized benchmarking procedure for displacement operations with varying (known) infidelities such that we can relate the underlying displacement fidelities to the Clifford gate fidelities extracted from exponential fits. We find that, regardless of the number of qudit dimensions being used, the results are well modeled by the relation $F_{\hat{D}} = (F_{C_L})^{1/N}$, with $N = 5/3$ being the average number of displacement pulses per Clifford operation (black line). (d) Cosine-square pulse shape used for the displacement pulses. The in-phase component of the drive is shown in red, and the out-of-phase component is in blue. DRAG-like corrections are calculated numerically, producing some amplitude in the out-of-phase part; detunings, meanwhile, can be viewed as oscillations in the envelopes. These corrections are applied per transition, i.e., each transition's drive has independent DRAG weights and detunings. Note that we have exaggerated the DRAG weight and detuning to better illustrate their effects on the pulse. (e) Randomized benchmarking fidelities in $d = 8$ as a function of displacement pulse duration for each level of correction. The fit model assumes that the total infidelity is the sum of coherent and incoherent parts, with the coherent error decreasing with the square of the pulse duration, reflecting phase errors and leakage due to off-resonant driving, and the incoherent error increasing linearly with duration, due to, e.g., energy relaxation. Our model further assumes that all correction settings are subject to the same rate of incoherent error. (f) Repeated randomized benchmarking in $d = 8$. We demonstrate the stability of our displacement pulse implementation as well as the significance of the pulse corrections by repeating the same randomized benchmarking experiment many times over the course of 82 hours, without recalibrating, showing the overall distribution of Clifford errors in the right panel.

many times, as shown in Fig. 4(f). The measured fidelities remain stable over roughly 82 hours, with the distributions of fidelities under each type of pulse correction remaining well separated.

We note that these randomized benchmarking results require careful interpretation for several reasons. First, the fidelity of the displacement over the full qudit Hilbert space is not necessarily equivalent to the fidelity within the logical subspace: It is conceivable that the displacement pulse is subject to errors affecting states orthogonal to the logical subspace, which would be suppressed in our results. Second, randomized benchmarking can potentially fail to account for errors resulting in leakage out of the logical subspace because such errors can, in general, produce

non-Markovian effects [51]. Finally, the logical Clifford operations require differing numbers of displacement pulses. We address these complications using a numerical model of our randomized benchmarking procedure in which randomly sampled, completely positive trace-preserving (CPTP) maps are applied after each displacement [33]. The average gate fidelity of this perturbed displacement is computed directly, which we take to be the underlying displacement fidelity $F_{\hat{D}}$ under a particular realization of the error channel. Meanwhile, the perturbed displacement is used to construct the Clifford set, from which we sample sequences of varying length. A fit to the resulting exponential decay is used to compute the average logical qubit Clifford fidelity F_{C_L} . Repeating this process

for many randomizations and weights of the error map ultimately allows us to estimate the relationship between these two fidelities, shown in Fig. 4(c) (Appendix F). We find that the simple parameter-free model $F_{\hat{D}} = (F_{C_L})^{1/N}$ represents the numerical data well, with $N = 5/3$ being the average number of displacement pulses per logical qubit Clifford gate. Using this relation and our Clifford fidelities from randomized benchmarking, we obtain estimated displacement pulse fidelities of $F_{\hat{D}} = \{0.9971 \pm 0.0001, 0.9931 \pm 0.0001, 0.9895 \pm 0.0003\}$ for $d = 3$, $d = 5$, and $d = 8$, respectively.

2. SNAP-displacement decomposition of qudit gates

In addition to the logical qubit gates discussed in the previous section, we also experimentally demonstrate qudit control by implementing the single-qudit Clifford group C_d . For any dimension d , the single-qudit Clifford group is generated by the set $\{\hat{1}, \hat{H}_d, \hat{Z}_d, \hat{S}_d\}$ [34]. Here, $\hat{H}_d$ is the qudit Hadamard gate, i.e., the single-qudit realization of the quantum Fourier transform (QFT):

$$\hat{H}_d = \frac{1}{\sqrt{d}} \sum_{n,m=0}^{d-1} \omega_d^{nm} |n\rangle\langle m|, \quad (3)$$

where $\omega_d = e^{i\frac{2\pi}{d}}$. This gate is a key component of qudit algorithms such as phase estimation and Shor's algorithm [1,52,53]. The remaining gates are

$$\hat{Z}_d = \sum_{n=0}^{d-1} \omega_d^n |n\rangle\langle n| \quad (4)$$

and

$$\hat{S}_d = \sum_{n=0}^{d-1} \omega_d^{-\frac{n(n+d)}{2}} |n\rangle\langle n|. \quad (5)$$

Because both of these gates are diagonal, they are implemented by virtual phase rotations. The qudit Hadamard gate, meanwhile, is realized by a SNAP-displacement decomposition.

We begin by constructing the full qudit Clifford group in terms of its generators for $d = 3$, $d = 5$, and $d = 8$ (see Appendix H for a description of the algorithm used). The groups contain 216, 3000, and 24,576 elements, respectively. The result is a decomposition of each Clifford operator into its generators, as well as the Cayley table representing the group's structure under multiplication. We note that, because $\hat{H}_d^4 = \hat{1}$ for any d , each Clifford operator's decomposition will contain at most three adjacent $\hat{H}_d$ gates. In light of this fact, we condense each Clifford operator's sequence such that it contains single instances of $\hat{H}_d$, $\hat{H}_d^2$, and $\hat{H}_d^3$ interleaved with arbitrary

numbers of $\hat{Z}_d$ and $\hat{S}_d$ gates. By precomputing SNAP-displacement decompositions for $\hat{H}_d$, $\hat{H}_d^2$, and $\hat{H}_d^3$, we can therefore generate the SNAP-displacement decomposition for any Clifford operator on the fly.

We find that, in general, $\hat{H}_d$ and $\hat{H}_d^3$ can each be realized by a SNAP-displacement decomposition of length $N = d - 1$, while $\hat{H}_d^2$ requires $N = d + 1$ displacement gates. For $d = 3$ and $d = 5$, we use these decomposition depths. For $d = 8$, however, we use a decomposition depth of $N = 4$ for each. At this depth, there is a decomposition error of roughly 5%, which would be decreased by adding more layers. However, the optimal experimental fidelity is achieved by balancing these decomposition errors with the pulse errors themselves, which were measured in the previous section to be on the order of 1% per displacement pulse for $d = 8$. This approach illustrates a key advantage of our control scheme as compared to Givens rotations: One can choose a number of layers that maximizes the fidelity of the desired gate under the error rates of the system at hand, balancing the detrimental effect of imperfect displacement pulses with that of an imperfect decomposition. This method is in contrast to Givens rotations, where there is no clear way to “prune” gates from the decomposition. We note that a similar conclusion can be drawn for bosonic systems from Ref. [54], strengthening the analogy between these methods.

Figure 5 summarizes the results of our experiments demonstrating qudit Clifford control. In Figs. 5(a)–5(c), we show density matrices recovered by performing state tomography after applying $\hat{H}_d$ to $|0\rangle$, which should ideally produce an equal superposition of the basis states. We recover state fidelities of 0.9996(5), 0.970(4), and 0.88(1) for $d = 3$, $d = 5$, and $d = 8$. The state tomography procedure and data analysis are discussed in detail in Appendix G.

While these state fidelities suggest that the qudit Hadamard gates behave as expected when applied to the ground state, they are not a reliable characterization of the gate as a whole or of the rest of the qudit Clifford group. We address this issue by performing qudit randomized benchmarking [28,34], showing the results in Figs. 5(d)–5(f). Like qubit randomized benchmarking, qudit randomized benchmarking involves sampling random sequences of Clifford gates and, for each, appending the inverse of the sequence to the end. Applying these sequences to the ground state and measuring the resulting population in $|0\rangle$ yields an exponential decay as a function of sequence length, from which we can estimate the average gate fidelity of the Clifford operations. We furthermore perform interleaved randomized benchmarking [55] to individually extract the average gate fidelities of $\hat{H}_d$, $\hat{H}_d^2$, and $\hat{H}_d^3$. The average gate fidelities extracted from all qudit randomized benchmarking experiments are summarized in Fig. 5(g), along with those resulting from the logical qubit randomized benchmarking experiments of the previous section.




FIG. 5. SNAP-displacement decomposition of the qudit Clifford group C_d . We first perform state tomography following the application of the qudit Hadamard gate, realized as a SNAP-displacement decomposition, in $d = 3$ (a), $d = 5$ (b), and $d = 8$ (c). The color indicates phase, with the color scale given to the right. Next, we measure the average qudit Clifford gate fidelity using qudit randomized benchmarking in $d = 3$ (d), $d = 5$ (e), and $d = 8$ (f), yielding average Clifford fidelities of $F_{C_d} = \{0.9850(6), 0.874(6), 0.57(3)\}$. We also show interleaved randomized benchmarking results, which independently measure the average gate fidelities of $\hat{H}_d$, $\hat{H}_d^2$, and $\hat{H}_d^3$. For instance, we find the Hadamard fidelities to be $F_{\hat{H}_d} = \{0.992(1), 0.956(4), 0.91(6)\}$. The average gate fidelities for all qudit randomized benchmarking experiments are summarized in panel (g) for $d = 3$ (red), $d = 5$ (blue), and $d = 8$ (green). For comparison, we also show the fidelities measured in the logical qubit randomized benchmarking experiments of the previous section, where C_L denotes the average over the logical qubit Clifford group and $\hat{D}$ is the displacement pulse itself.

In particular, we find average qudit Clifford fidelities of $F_{C_d} = \{0.9850(6), 0.874(6), 0.57(3)\}$ and Hadamard gate fidelities of $F_{\hat{H}_d} = \{0.992(1), 0.956(4), 0.91(6)\}$, for $d = 3$, $d = 5$, and $d = 8$, respectively. We note, however, that the qudit Clifford fidelities could be improved dramatically in the following way. In these experiments, we compile each Clifford gate into a sequence of generators, each of which is itself compiled into a SNAP-displacement sequence. However, one could instead directly decompose each qudit Clifford gate into a SNAP-displacement sequence, significantly decreasing the number of pulses required per Clifford gate and thereby increasing the average Clifford fidelity. We expect that the measured fidelities of $\hat{H}_d$, $\hat{H}_d^2$, and $\hat{H}_d^3$ are more representative of the fidelity with which an arbitrary qudit gate could be decomposed, given our error rates, than the measured Clifford fidelities. Our reason for approaching the Clifford gate decomposition in this way is that the large size of the qudit Clifford group makes the SNAP-displacement decomposition of each of its constituent gates impractical, given our current method for calculating the decomposition

parameters. We expect that this shortcoming will be alleviated by future work aimed at improving the numerical methods used to calculate these parameters.

A useful point of comparison for the performance of our $d = 8$ Hadamard is the decomposition of the equivalent three-qubit gate. In Table I, we show the number of primitive gates required to implement the equivalent three-qubit quantum Fourier transform as determined by the Qiskit transpiler [56] assuming either all-to-all or linear couplings. We assume a primitive gate set consisting of $\hat{X}(\pi/2)$, $\hat{Z}(\theta)$, and CNOT gates. From these gate counts, we estimate the achievable three-qubit QFT fidelities from

TABLE I. Number of primitive gates required to perform a three-qubit quantum Fourier transform given all-to-all and linear qubit couplings.

Coupling	$\hat{X}(\pi/2)$	$\hat{Z}(\theta)$	CNOT
All-to-all	36	42	16
Linear	50	62	22




FIG. 6. Estimated three-qubit quantum Fourier transform infidelities over a range of $\hat{X}(\pi/2)$ and CNOT fidelities assuming all-to-all or linear qubit couplings. Infidelities are represented by color, with the color bar on the right corresponding to both panels. The solid lines show our measured $d = 8$ quantum Fourier transform infidelity, with its uncertainty indicated by the dashed lines.

the simple model $F_{\text{QFT}} = (F_X)^{N_X} (F_{\text{CNOT}})^{N_{\text{CNOT}}}$, where N_X and N_{CNOT} are the $\hat{X}(\pi/2)$ and CNOT counts, respectively. We calculate these fidelities across a range of experimentally relevant fidelities and show the results in Fig. 6. The solid and dashed lines in each panel show our measured $d = 8$ QFT infidelity and its uncertainty region, respectively. We see that very high primitive qubit gate fidelities are required to match the performance of our equivalent qudit gate.

Finally, we point out that arbitrary diagonal unitaries can be implemented with near-unity fidelity using virtual SNAP gates [28,57], making the extension of our scheme to non-Clifford gates straightforward. For instance, the addition of the single-qudit T gate [1,58] yields a universal gate set for dimensions d that are primes or powers of primes. Universal qudit gate sets are not well understood for other d , but we nevertheless expect that, because we implement perfect diagonal gates, our qudit Clifford fidelities are representative of those of a universal single-qudit gate set.

III. DISCUSSION

In this paper, we have demonstrated a new method for controlling superconducting transmon qudits that combines concepts from bosonic computing platforms and high-dimensional spin systems to realize efficient, high-fidelity unitary control. Our method establishes a single high-dimensional spin displacement operation as a primitive gate for qudit computation, which, in combination with virtual phase rotations, is universal over the single-qudit Hilbert space. In combination with our single-shot projective readout of the full qudit state, this method takes full advantage of the qudit Hilbert space. Our gate set is highly efficient as compared to typical decompositions of qudit unitaries, requiring $\mathcal{O}(d)$ physical pulses in our implementation to produce any desired unitary. Furthermore, in principle, it does not require optimal

control methods, instead making use of the spectral separation of our qudit transitions to directly drive the generator of the spin displacements; this property makes it immediately applicable to transmon qudits of any dimension and, indeed, to any qudit platform in which the transitions are well separated in frequency. We note that, while in this work we apply optimal control in calculating the DRAG, detuning, and phase corrections, we expect that further theoretical work could yield analytic predictions for these parameters. It may additionally be possible to find simple experimental methods for directly calibrating these corrections without the need for analytic predictions or accurate numerical simulations, which we leave for future work.

A simple model of the displacement pulse fidelity in terms of coherent and incoherent errors, presented in Appendix I, suggests that an increase in coherence time would have a considerable effect on the displacement fidelity. We estimate that increasing our coherence times to be on par with state-of-the-art transmon devices [59] would increase our fidelities to $F_{\hat{D}}(\text{predicted}) = \{0.9993, 0.998, 0.996\}$ for $d = 3$, $d = 5$, and $d = 8$, respectively, even in the absence of pulse corrections. Our experimental results suggest that including pulse corrections could further decrease the infidelity by roughly a factor of 2. Alternative device designs such as the inductively shunted transmon [60] could eliminate the trade-off between the number of usable qudit states and the anharmonicity, further suppressing coherent errors.

We note that the proof of universality presented in Appendix D does not require any particular relationship between the drive strengths, and one could construct a universal single-qudit gate set using a different set of simultaneous drives on each transition, realizing some family of d -dimensional unitaries other than the spin displacements. The $SU(2)$ structure of our spin displacement scheme is attractive because it leads to a phase-space representation and because any $SU(2)$ rotation can be constructed by interleaving two $\pi/2$ displacements with $\hat{J}_z$ rotations. Furthermore, spin rotations are useful in themselves for analog quantum simulation [18,19]. However, the investigation of a set of drives with different relative strengths—for example, those saturating the limit imposed by the anharmonicity—is a promising direction for future work.

Qudit computation has the potential to significantly improve the capacity of current quantum computing architectures due to its intrinsic hardware efficiency, speeding up certain types of computations. Our work represents a major step towards this goal, simultaneously utilizing a high- E_J/E_C transmon design for many-level qudit computation and a SNAP-displacement decomposition for high-fidelity control. We have demonstrated that logical qubits encoded into the transmon qudit Hilbert space can be controlled with high fidelity, and furthermore, we have shown that

TABLE II. Measured transition frequencies and calculated E_J , E_C , bare resonator frequency f_r , and coupling strength g . All values are expressed in GHz.

Transmon	f_{01}	f_{12}	f_{23}	f_{34}	f_{45}	f_{56}	f_{67}	E_J	E_C	f_r	g
Q_4	4.864	4.751	4.632	4.507	4.374	4.234	4.082	28.70	0.108	6.412	0.025

important qudit gates such as the quantum Fourier transform can be efficiently realized by the SNAP-displacement decomposition.

Our multifrequency control and readout techniques can be integrated with recent advances on qudit entangling gates [24,61] to explore the advantages of qudit technology. For instance, there has recently been much interest in bosonic error-correcting codes in linear resonator modes [11,12,62–64], as well as error-correcting codes in high-dimensional spins [65,66]. In both cases, a high-dimensional Hilbert space is used to implement a stabilizer code designed to correct the most physically relevant error processes for the system at hand. Our work motivates investigation into whether similar error-correcting codes can be implemented in a transmon qudit, taking advantage of the high-fidelity control and readout enabled by our strong nonlinearity and multifrequency drives while retaining the hardware efficiency inherent to bosonic and spin qudit codes. More generally, our work provides a prescription for decomposing primitive qudit gates in a way that scales favorably with qudit dimension, paving the way for qudit experiments in higher dimensions. In the context of quantum simulation, meanwhile, qudits have recently been proposed as natural building blocks for simulating lattice gauge theories [15,67,68], owing to their ability to naturally encode multiple degrees of freedom per site. We ultimately expect that our control and readout scheme will facilitate more complex qudit computations and simulations than would otherwise be possible given current coherence times while remaining highly extensible and transferable between platforms.

Note added. Recently, we became aware of two complementary works on distinct hardware platforms that also demonstrate large-spin $SU(2)$ dynamics and Schrodinger spin cats [69,70].

ACKNOWLEDGMENTS

We would like to thank Shruti Puri and Daniel Weiss for fruitful discussions during the course of this work as well as feedback on the manuscript. This material is based upon work supported by the Air Force Office of Scientific Research under Award No. FA9550-23-1-0121. Devices were fabricated and provided by the Superconducting Qubits at Lincoln Laboratory (SQUILL) Foundry at MIT Lincoln Laboratory, with funding from the Laboratory for Physical Sciences (LPS) Qubit Collaboratory.

APPENDIX A: HIGH- E_J/E_C TRANSMON DEVICE

Our implementation of spin displacements and qudit readout relies on having a physical device that can encode a many-level system with minimal decoherence while maintaining spectral separation of its transitions. In this work, we propose a superconducting transmon qudit with a large ratio of E_J/E_C . The transmon Hamiltonian is

$$\hat{H} = 4E_C(\hat{n} - n_g)^2 - E_J \cos \hat{\phi}, \quad (\text{A1})$$

where E_C is the capacitive energy, E_J is the Josephson energy, $\hat{n}$ is the transmon charge operator, $\hat{\phi}$ is the superconducting phase operator, and n_g is a gate charge [36]. The $|0\rangle \leftrightarrow |1\rangle$ transition frequency is approximately given by $\omega_{01} = \sqrt{8E_J E_C} - E_C$, while the anharmonicity is $-E_C$. A rough estimate of the number of levels contained in the cosine potential can be obtained by dividing the depth of the potential, $2E_J$, by the plasma frequency $\sqrt{8E_J E_C}$ [71]:

$$n_{\text{levels}} \approx \sqrt{\frac{E_J}{2E_C}}. \quad (\text{A2})$$

Thus, we see that we should expect more excited states to be confined in the cosine potential as E_J/E_C is increased.

We must also consider the effects of charge-noise-induced dephasing in the higher excited states. The charge dispersion of the transmon state $|n\rangle$ is given by [35]

$$\epsilon_n = (-1)^n E_C \frac{2^{4n+5}}{n!} \sqrt{\frac{2}{\pi}} \left(\frac{E_J}{2E_C} \right)^{\frac{n+3}{4}} e^{-\sqrt{8E_J E_C}}. \quad (\text{A3})$$

TABLE III. Coherence times for each qudit transition. Each coherence time was measured 32 times over the course of 14 hours, and here we report the mean values and standard deviations.

Transition	Q_4		
	T_1 (μs)	T_2^* (μs)	T_2 (μs)
$ 0\rangle \leftrightarrow 1\rangle$	46 ± 7	52 ± 11	51 ± 7
$ 1\rangle \leftrightarrow 2\rangle$	25 ± 4	24 ± 6	37 ± 5
$ 2\rangle \leftrightarrow 3\rangle$	26 ± 3	15 ± 4	34 ± 3
$ 3\rangle \leftrightarrow 4\rangle$	14 ± 3	15 ± 5	25 ± 7
$ 4\rangle \leftrightarrow 5\rangle$	16 ± 2	48 ± 12	44 ± 6
$ 5\rangle \leftrightarrow 6\rangle$	14 ± 2	28 ± 10	27 ± 6
$ 6\rangle \leftrightarrow 7\rangle$	13 ± 2	20 ± 9	23 ± 6




FIG. 7. False-color image of the device design. All experiments in this work were carried out using the transmon labeled Q_4 .

The key point is that, while the charge dispersion increases exponentially with level, it decreases exponentially with $\sqrt{E_J/E_C}$. To realize a qudit with eight levels, we target a ratio of $E_J/E_C \approx 270$. We keep the transition frequency fixed to roughly 5 GHz, which necessitates a decrease in the capacitive energy E_C and therefore in the anharmonicity. However, we find that our anharmonicity is still large enough to permit high-fidelity control of the system. In Table II, we report the measured transition frequencies as well as calculated values for the underlying device parameters, and in Table III, we report the measured coherence times.

We show a false-color image of our device design in Fig. 7; the varying sizes of the transmon capacitor pads reflect the range of values of E_J/E_C chosen for this device. All experiments in this work were performed on the transmon labeled Q_4 .

APPENDIX B: SINGLE-SHOT DISPERSIVE QUDIT READOUT

Single-shot readout of the transmon qudit state is enabled by the dispersive coupling between the transmon and the readout resonator. In the lowest-order nonlinear




FIG. 8. The IQ readout data for each qudit state.

approximation to the transmon potential, the Hamiltonian of the transmon-resonator system is given by

$$\hat{H} = \omega_r \hat{a}^\dagger \hat{a} + \omega_q \hat{b}^\dagger \hat{b} - \frac{E_C}{2} \hat{b}^{\dagger 2} \hat{b}^2 + \sum_n (\Lambda_n + \chi_n \hat{a}^\dagger \hat{a}) |n\rangle \langle n|, \quad (\text{B1})$$

where $\hat{a}$ and $\hat{b}$ are the annihilation operators for the resonator and transmon, respectively; ω_r and ω_q are the bare resonator and transmon frequencies; Λ_n are Lamb shifts; and χ_n are dispersive shifts [35,36]. See Ref. [36] for the detailed forms of the Lamb and dispersive shifts; the salient point is that the resonator has an effective frequency shift that is different for each transmon state.

We optimize our device design to have dispersive shifts comparable in size to the linewidth of the readout resonator, ultimately producing the transmon-state-dependent spectrum depicted in Fig. 2(b) in the main text. In this regime, a single readout tone is able to distinguish several of the states, where there is appreciable overlap between the resonator spectra in those states. In order to read out the full qudit state, we apply three simultaneous readout tones with frequencies chosen such that, together, they are able to unambiguously distinguish between all qudit states. The signal transmitted through the readout line is demodulated at each of these frequencies and integrated to produce a six-dimensional vector $(I_1, Q_1, I_2, Q_2, I_3, Q_3)$.

We represent the resulting distributions in this six-dimensional space in Fig. 8, where points are colored according to the state that was prepared. We calibrate state classification by fitting a six-dimensional Gaussian mixture model to these data, with each Gaussian component corresponding to one of the qudit states. The resulting assignment probabilities are shown in Table IV, yielding an average assignment fidelity of $(88.3 \pm 0.2)\%$, where the uncertainty is calculated from the $N = 5000$ experimental shots and assuming binomial statistics. In practice, these probabilities are used to calculate a correction matrix, which is applied to the populations measured in each of the experiments in the main text.

TABLE IV. Probability that a prepared state $|n\rangle$ will be classified as $|m\rangle$.

		Prepared state							
		$ 0\rangle$	$ 1\rangle$	$ 2\rangle$	$ 3\rangle$	$ 4\rangle$	$ 5\rangle$	$ 6\rangle$	$ 7\rangle$
Assigned state	$ 0\rangle$	0.99	0.04	0.01	0.01	0.01	0.01	0.01	0.01
	$ 1\rangle$	0.01	0.96	0.07	0.01	0.01	0	0	0.01
	$ 2\rangle$	0	0	0.89	0.05	0.01	0	0	0
	$ 3\rangle$	0	0	0.03	0.92	0.18	0.02	0.01	0.01
	$ 4\rangle$	0	0	0.01	0.01	0.79	0.09	0.01	0.01
	$ 5\rangle$	0	0	0	0	0	0.86	0.10	0.08
	$ 6\rangle$	0	0	0	0	0	0.01	0.86	0.10
	$ 7\rangle$	0	0	0	0	0	0	0.01	0.79

APPENDIX C: THEORETICAL BACKGROUND

1. Lie groups and Lie algebras

We begin our discussion of qudit displacements, phase spaces, and Wigner functions with a pedagogical overview of Lie groups and their application to quantum systems with dynamical symmetries. Central to the idea of a phase space is the concept of a dynamical symmetry group. Recall that a group G is a collection of elements $g \in G$ and a binary operation between these group elements (which we denote by multiplication) having the following properties:

- (1) The group operation is associative: For any $g_a, g_b, g_c \in G$, we have $(g_a g_b) g_c = g_a (g_b g_c)$.
- (2) There exists a unique identity element $g_e \in G$ such that for any $g_a \in G$, $g_e g_a = g_a g_e = g_a$.
- (3) Every $g \in G$ has a unique inverse element, denoted by g^{-1} , such that $g^{-1} g = g g^{-1} = g_e$.

Any collection satisfying these properties is a group, irrespective of the underlying nature of its elements. If two groups share the same structure—in the sense that there is a bijective mapping between their elements under which they have the same multiplication table—then we say that they are isomorphic, meaning that they are, in some sense, the same group.

Groups frequently arise in physics when a system is symmetric under some set of transformations. We focus here on the case of a Lie group, which applies to continuous symmetries, such as rotational symmetry. A d -dimensional Lie group is a group that can be parametrized by d continuously varying, real-valued parameters $x_1, x_2, \dots, x_d$, which we collectively denote by a vector $\mathbf{x}$. In other words, it is a smooth manifold that also has a group structure. The group elements g are functions of the parameters: $g = G(\mathbf{x})$.

It is often convenient to construct a Lie group in terms of its Lie algebra, a set of algebraic objects that generate the Lie group through exponentiation. The generators are related to the group elements themselves in that they are the infinitesimal group elements. They represent a linearization of the action of group elements; concretely, they are given by

$$X_j = i \left(\frac{\partial g(\mathbf{x})}{\partial x_j} \right)_{\mathbf{x}=0}. \quad (\text{C1})$$

We therefore see that the number of generators is equal to d , the number of dimensions. The generators of a Lie group form a vector space—meaning that linear combinations of them are also generators—and therefore generate group elements. They obey the commutation relations

$$[X_i, X_j] = i f_{ijk} X_k, \quad (\text{C2})$$

where f_{ijk} are known as structure constants. When we think of a specific, concrete realization of a Lie group or

Lie algebra—for example, as operators acting on a Hilbert space—we are invoking the idea of a group representation. We denote the elements of a group's representation as D :

$$D[g(\mathbf{x})] \equiv D(\mathbf{x}). \quad (\text{C3})$$

The relationship between the representation of a group and the representation of its generators is given by an exponential map,

$$D[g(\mathbf{x})] = e^{-i\sum_j x_j X_j}. \quad (\text{C4})$$

2. Harmonic oscillator

It is helpful to begin with a description of the standard coherent states of a harmonic oscillator and its phase-space representation. We refer the reader to Refs. [38–41] for more comprehensive discussions of this and spin displacements. The harmonic oscillator is conveniently described by the bosonic creation and annihilation operators $\hat{a}^\dagger$ and $\hat{a}$, which, together with the identity operator $\hat{1}$, obey the Heisenberg commutation relations:

$$[\hat{a}, \hat{a}^\dagger] = \hat{1}, \quad (\text{C5})$$

$$[\hat{a}, \hat{1}] = [\hat{a}^\dagger, \hat{1}] = 0. \quad (\text{C6})$$

These operators therefore define a Lie algebra, having a general element of the form

$$t\hat{1} - i(\alpha\hat{a}^\dagger - \alpha^*\hat{a}) \quad (\text{C7})$$

for a real number t and complex number α (such a parametrization consists of three independent real values, reflecting the dimension of the algebra while enforcing Hermiticity). It follows that the operators

$$\hat{T}(t, \alpha) = e^{it} e^{\alpha\hat{a}^\dagger - \alpha^*\hat{a}} \quad (\text{C8})$$

are a representation of a Lie group, known as the Heisenberg-Weyl group H_3 . This Lie group represents the dynamical symmetry group of the harmonic oscillator when it is acted upon by an external drive. We therefore identify

$$\hat{D}(\alpha) \equiv e^{\alpha\hat{a}^\dagger - \alpha^*\hat{a}} \quad (\text{C9})$$

as the harmonic oscillator displacement operator; the multiplication of two such operators yields

$$\hat{D}(\alpha)\hat{D}(\beta) = e^{i\text{Im}(\alpha\beta^*)}\hat{D}(\alpha + \beta). \quad (\text{C10})$$

A set of coherent states can, in general, be constructed by applying these displacement operators to some fixed state

in the Hilbert space. Making the typical choice of $|0\rangle$ as our fixed state, we have

$$|\alpha\rangle \equiv \hat{D}(\alpha)|0\rangle. \quad (\text{C11})$$

This choice of fixed state yields a set of coherent states equivalent to those of the typical formulation, defined as eigenstates of the annihilation operator:

$$\hat{a}|\alpha\rangle = \alpha|\alpha\rangle. \quad (\text{C12})$$

We now turn our attention to the phase-space representation of a harmonic oscillator state—in particular, the Wigner function. For any quantum system, the Wigner function can be defined as

$$W(\alpha) = \text{Tr}[\hat{\rho} \hat{\Delta}(\alpha)], \quad (\text{C13})$$

where $\hat{\rho}$ is the density matrix and $\hat{\Delta}(\alpha)$ is a special operator-valued function over the phase space, called the kernel. For a harmonic oscillator, the kernel may be written in terms of a displacement in phase space and the parity operator as

$$\hat{\Delta}(\alpha) = 2\hat{D}(\alpha)\hat{\Pi}\hat{D}^\dagger(\alpha); \quad (\text{C14})$$

substituting this result into the definition of the Wigner function and exploiting the cyclic property of traces yields the familiar equation

$$W(\alpha) = 2\text{Tr}[\hat{D}^\dagger(\alpha)\hat{\rho}\hat{D}(\alpha)\hat{\Pi}] = 2\text{Tr}[\hat{D}(-\alpha)\hat{\rho}\hat{D}(\alpha)\hat{\Pi}]. \quad (\text{C15})$$

3. Spin systems

While the results in the previous section are specific to the case of a harmonic oscillator, the framework summarized here is far more general, allowing one to construct the Wigner function for a variety of systems. The mapping $W(\alpha) = \text{Tr}[\hat{\Delta}(\alpha)\hat{A}]$ from operators $\hat{A}$ to functions on a phase space is mediated by the self-dual kernel $\hat{\Delta}(\alpha)$, which can be derived for a specific system from a set of conditions known as the Stratonovich-Weyl postulates [39–41]. The process for finding such a kernel is intimately related to the displacement operation and therefore to the coherent states of the system. In its most general form, this calculation is very involved; thus, here we only discuss the construction of a spin displacement operator and then quote a result for the kernel in terms of the displacement and a generalization of a parity operator.

The Hilbert space of a spin- j system is spanned by a basis of $2j+1$ states, which we label $|j, m\rangle$ for $-j \leq m \leq j$. These states are the eigenstates of the angular-momentum operator $\hat{J}_z$. The dynamical symmetry group of a spin is

$SU(2)$, and the Lie algebra generating this group is spanned by the basis $\{\hat{J}_x, \hat{J}_y, \hat{J}_z\}$,

$$[\hat{J}_x, \hat{J}_y] = i\hat{J}_z, \quad [\hat{J}_y, \hat{J}_z] = i\hat{J}_x, \quad [\hat{J}_z, \hat{J}_x] = i\hat{J}_y. \quad (C16)$$

One can equivalently work in the basis $\{\hat{J}_+, \hat{J}_-, \hat{J}_z\}$, where

$$\hat{J}_\pm = \hat{J}_x \pm i\hat{J}_y, \quad (C17)$$

are ladder operators for the spin state:

$$\hat{J}_\pm |j, m\rangle = \sqrt{j(j+1) - m(m \pm 1)} |j, m \pm 1\rangle. \quad (C18)$$

In this basis, we have

$$[\hat{J}_z, \hat{J}_\pm] = \pm \hat{J}_\pm, \quad [\hat{J}_+, \hat{J}_-] = 2\hat{J}_z. \quad (C19)$$

A general group element can be parametrized as

$$\hat{g}(\psi, \theta, \phi) = e^{-i\phi\hat{J}_z} e^{-i\theta\hat{J}_y} e^{-i\psi\hat{J}_x}, \quad (C20)$$

where $\phi, \theta, \psi \in \mathbb{R}$ can be understood as Euler angles. We therefore see that the phase space associated with a spin is the surface of a sphere, S^2 .

Much like the case of a harmonic oscillator, the coherent states of a spin are the states generated by applying our group element representation $\hat{g}$ to a fixed state, which we

choose to be $|j, j\rangle$. We see that the effect of the $e^{-i\psi\hat{J}_x}$ term is to add a trivial phase factor to our states; thus, we take $\psi = 0$ and define our displacement operator:

$$\hat{D}_j(\theta, \phi) = e^{-i\phi\hat{J}_z} e^{-i\theta\hat{J}_y}. \quad (C21)$$

It can be shown [38] that this result is equivalent to

$$\hat{D}_j(\alpha) = e^{\alpha\hat{J}_+ - \alpha^*\hat{J}_-}, \quad (C22)$$

where

$$\alpha = -\frac{\theta}{2} e^{-i\phi}. \quad (C23)$$

Then, the coherent states of a spin are

$$|j, \alpha\rangle = \hat{D}_j(\alpha) |j, j\rangle \quad (C24)$$

for α lying in a disk of radius $\pi/2$ in the complex plane.

As stated above, the Wigner function for a given system is a mapping between operators on the Hilbert space and functions on the phase space. In the present case of a spin- j system, this mapping can be shown to be mediated by the kernel [39,41]

$$\hat{\Delta}(\alpha) = \sum_{m=-j}^j \sum_{l=0}^{2j} \frac{2l+1}{2j+1} \begin{pmatrix} j & l \\ m & 0 \end{pmatrix} \begin{pmatrix} j \\ m \end{pmatrix} \hat{D}(\alpha) |j, m\rangle \langle j, m| \hat{D}^\dagger(\alpha), \quad (C25)$$

where

$$\begin{pmatrix} j_1 & j_2 \\ m_1 & m_2 \end{pmatrix} \begin{pmatrix} J \\ M \end{pmatrix}$$

is a Clebsch-Gordan coefficient. Therefore, defining

$$2\hat{\Pi}_j = \sum_{m=-j}^j \sum_{l=0}^{2j} \frac{2l+1}{2j+1} \begin{pmatrix} j & l \\ m & 0 \end{pmatrix} \begin{pmatrix} j \\ m \end{pmatrix} |j, m\rangle \langle j, m|, \quad (C26)$$

we see that the kernel can be written as

$$\hat{\Delta}(\alpha) = 2\hat{D}(\alpha)\hat{\Pi}_j\hat{D}^\dagger(\alpha), \quad (C27)$$

in direct analogy to the harmonic oscillator. Note that $\hat{\Pi}_j$ is not exactly a parity operator in the typical sense, but it is diagonal in the basis of $\hat{J}_z$ eigenstates.

In Fig. 9 we show Wigner functions for $\hat{J}_z$ eigenstates for $j = 1/2$ through $j = 2$. Because these states are energy eigenstates and thus invariant under rotations about $\hat{J}_z$, the Wigner functions are rotationally symmetric, much like those of harmonic oscillator eigenstates. In each case, the ground state and highest excited state are examples of coherent states. As the size of the spin grows, the Wigner

functions for these states become almost entirely positive. This can be understood as the spin ladder operators approaching those of the harmonic oscillator in the high-dimensional limit.

APPENDIX D: UNITARY QUDIT CONTROL WITH DISPLACEMENT AND SNAP GATES

The only case for which the spin displacement unitary alone provides universal control over the qudit Hilbert space is $d = 2$. In this case, the displacement operator reduces to rotations on the Bloch sphere because the angular-momentum operators coincide with the Pauli operators. For $d > 2$, on the other hand, the angular-momentum operators alone do not form a basis for all $d \times d$ Hermitian matrices; therefore, there are some unitaries that cannot be realized by rotations of the angular momentum alone.

In bosonic systems, universal control is often realized by a decomposition of the target unitary into displacement gates and SNAP gates [43–46]. We show here that the same scheme can be utilized for transmon qudits with spin displacements, with the added benefit that our SNAP gates are implemented as instantaneous virtual rotations.




FIG. 9. Wigner functions for $\hat{J}_z$ eigenstates. The spherical phase space is flattened to a circle where θ and ϕ are the radial and angular coordinates, respectively. Then, the entire outer circumference of the plot represents the south pole of the phase space. The bottom row corresponds to Fig. 3(b) in the main text.

A SNAP gate is parametrized in terms of a d -dimensional vector of phases that are applied to each eigenstate:

$$\hat{S}(\vec{\phi}) = \prod_{n=0}^{d-1} e^{i\phi_n |n\rangle\langle n|}. \quad (\text{D1})$$

In this decomposition, an arbitrary unitary is decomposed into a sequence of $N + 1$ SNAP gates interleaved with N displacement gates:

$$\begin{aligned} \hat{U}(\vec{\phi}^{(0)}, \theta^{(1)}, \vec{\phi}^{(1)}, \dots, \theta^{(N)}, \vec{\phi}^{(N)}) \\ = \hat{S}(\vec{\phi}^{(N)}) \hat{D}(\theta^{(N)}) \dots \hat{S}(\vec{\phi}^{(1)}) \hat{D}(\theta^{(1)}) \hat{S}(\vec{\phi}^{(0)}). \end{aligned} \quad (\text{D2})$$

Note that we parametrize the displacement gates in terms of the amplitude θ , setting the phase of the displacement to zero. We use this approach because the displacement phase can be absorbed into the adjacent SNAP gates. In practice, we also find that fixing $\theta^{(i)} = \pi/2$ is sufficient to realize arbitrary unitaries. While this may come at the cost of moderately increasing the number of layers required for the decomposition, it comes with the significant advantage that we only need to calibrate a single physical pulse rather than a continuous set of them.

1. Proof of universality

Here, we show that the SNAP-displacement decomposition is universal over the qudit Hilbert space given large enough N , and we provide numerical evidence below that N scales as $\mathcal{O}(d)$. Consider a set of gates generated through the exponentiation of k different Hermitian operators $\hat{H}_1, \hat{H}_2, \dots, \hat{H}_k$ acting on a d -dimensional Hilbert space: $\hat{U}_m = e^{it\hat{H}_m}$, where t can be continuously varied. The set of unitary operations that this gate set can realize consists of all $\hat{U} = e^{i\hat{L}}$, where $\hat{L}$ is any element of the algebra of Hermitian operators generated by taking commutators and repeated commutators of the operators $\hat{H}_m$; for example, $i[\hat{H}_1, \hat{H}_2]$, $[\hat{H}_1, [\hat{H}_1, \hat{H}_3]]$, and so on [72]. If this algebra is a complete basis of $d \times d$ Hermitian matrices, then the gate set is capable of producing any rotation in $SU(d)$. In other words, it is universal.

With the phase of the displacement drive fixed to zero, the generator of displacements is the operator $\hat{J}_y$. To understand the effect of interleaving displacements with SNAP gates, we first need a basis for the generators of the SNAP gates. We are free to choose any complete basis of diagonal, real-valued $d \times d$ matrices, and we will see that the following operators are a particularly convenient choice:




FIG. 10. Numerical justification of the $\mathcal{O}(d)$ scaling of the SNAP-displacement decomposition. The circular markers correspond to each of the 100 Haar-random unitaries, while the stars correspond to the qudit Hadamard gate $\hat{H}_d$. Note that we have applied small horizontal offsets to the data points of differing decomposition depths to aid readability.

$$\hat{Q}_n = \sum_{q=0}^n |q\rangle\langle q| \quad (\text{D3})$$

for $0 \leq n \leq d-2$. Choosing one of these $\hat{Q}_n$ and taking its commutator with $\hat{J}_y$ yields

$$i[\hat{J}_y, \hat{Q}_n] = \frac{1}{2} \sqrt{(n+1)(d-(n+1))} (|n\rangle\langle n+1| + |n+1\rangle\langle n|). \quad (\text{D4})$$

We therefore see that, by choosing n , we can generate a $\hat{\sigma}_x$ -like operator in the $\{|n\rangle, |n+1\rangle\}$ subspace using a single commutator. Treating each of these subspaces as a qubit, we see that combining the $\hat{\sigma}_x$ rotations with virtual phase rotations produces $\hat{\sigma}_y$ rotations in the subspace in question. Finally, it is well known that this set of Pauli-like rotations between adjacent levels of a qudit is sufficient to generate any rotation in $SU(d)$ [24,27,42], so we conclude that our gate set is universal. We note that our proof is entirely analogous to that given in Ref. [44] for the case of a harmonic oscillator displacement.

2. Linear scaling of the SNAP-displacement decomposition depth with dimension

While the above proof shows that the SNAP-displacement decomposition is universal, it does not place any bounds on the number of layers required to decompose a given unitary with a satisfactory fidelity. We can place a lower bound on the required depth by noting that the number of independent parameters in a d -dimensional unitary is $d^2 - 1$. Assuming the displacement phase is fixed to zero, a displacement operator contains a single independent parameter, while a SNAP gate contains $d - 1$. A SNAP-displacement decomposition of depth N therefore has $N + (N + 1)(d - 1)$ parameters, so we expect to require a minimum of $N = d - 1$ SNAP-displacement layers to realize general unitaries.

We also note that it is possible to use fewer layers in cases where the desired unitary is closely related to the displacement operator itself, as is the case for the logical qubit Clifford rotations discussed in the main text.

The noncommuting nature of our SNAP and displacement gates makes a rigorous proof of universality for a fixed depth N difficult, and it is unclear whether such a proof can be found. We instead provide numerical evidence that the number of layers scales as $\mathcal{O}(d)$. For a given qudit dimension d , we sample 100 Haar-random $d \times d$ unitary matrices and attempt to find SNAP-displacement parameters producing each. We use SNAP-displacement depths of $N = d - 2$, $N = d - 1$, and $N = d$ for qudit dimensions between $d = 2$ and $d = 10$ and show the results in Fig. 10. The optimizer fails to find satisfactory decomposition parameters for $N = d - 2$, as expected from the number of independent parameters. For $N = d - 1$, meanwhile, the fidelities are very close to 1 in all but a few cases. It is unclear whether the points with unsatisfactory fidelities are due to the numerical optimizer becoming stuck in a local minimum or if the decomposition is fundamentally incapable of producing those unitaries. For $N = d$, on the other hand, the decomposition is able to faithfully produce the desired unitary in every case. We attribute the slight upwards trend in infidelity with dimension to the difficulty of performing numerical optimization in such a large parameter space. We also show decomposition fidelities for the case of the qudit Hadamard gate $\hat{H}_d$, implemented in Fig. 5 in the main text, which follow the same overall trends as the random unitaries.

3. Decomposing a target unitary

In this section, we describe the procedure for numerically optimizing the parameters of a SNAP-displacement decomposition to realize a target unitary. As discussed above, in practice, we fix all displacement angles $\theta^{(i)} = \pi/2$, so our decomposition is

$$\hat{U}(\vec{\phi}^{(0)}, \phi^{(1)}, \dots, \vec{\phi}^{(N)}) = \hat{S}(\vec{\phi}^{(N)})\hat{D}(\pi/2)\hat{S}(\vec{\phi}^{(N-1)})\dots\hat{S}(\vec{\phi}^{(1)})\hat{D}(\pi/2)\hat{S}(\vec{\phi}^{(0)}). \quad (\text{D5})$$

For some target unitary $\hat{U}_{\text{target}}$, we define a cost function

$$\mathcal{C}(\vec{\phi}^{(0)}, \phi^{(1)}, \dots, \vec{\phi}^{(N)}) = 1 - \frac{1}{d^2} \left| \text{Tr} \left[\hat{U}_{\text{target}}^\dagger \hat{U}(\vec{\phi}^{(0)}, \phi^{(1)}, \dots, \vec{\phi}^{(N)}) \right] \right|^2, \quad (\text{D6})$$

which reaches a minimum value of zero when the fidelity is 1.

We minimize this cost function over the decomposition parameters using the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm [73]. The gradient of this cost function is straightforward to evaluate analytically. The gradient of a SNAP gate $\hat{S}(\vec{\phi}^{(i)})$ with respect to a particular phase $\phi_n^{(i)}$ is

$$\frac{d}{d\phi_n^{(i)}} \hat{S}(\vec{\phi}^{(i)}) = ie^{i\phi_n^{(i)}} |n\rangle\langle n|. \quad (\text{D7})$$

Inserting this derivative at the appropriate location in a SNAP-displacement sequence therefore yields the overall derivative of the decomposed unitary with respect to each control parameter; for instance, if we want the derivative with respect to $\phi_n^{(1)}$, we evaluate

$$\frac{d}{d\phi_n^{(1)}} \hat{U} = \hat{S}(\vec{\phi}^{(N)})\hat{D}(\pi/2)\hat{S}(\vec{\phi}^{(N-1)})\dots\hat{D}(\pi/2)(ie^{i\phi_n^{(1)}} |n\rangle\langle n|)\hat{D}(\pi/2)\hat{S}(\vec{\phi}^{(0)}), \quad (\text{D8})$$

where, for brevity, we hereafter denote $\hat{U}(\vec{\phi}^{(0)}, \phi^{(1)}, \dots, \vec{\phi}^{(N)})$ simply by $\hat{U}$. The derivative with respect to any other SNAP phase is constructed similarly. The gradient of the cost function with respect to a parameter α (i.e., some SNAP phase) is obtained by simply propagating the derivative of the unitary through the cost function:

$$\frac{d\mathcal{C}}{d\alpha} = -\frac{2}{d^2} \text{Re} \left\{ \text{Tr} [\hat{U}_{\text{target}}^\dagger \hat{U}] \hat{U}_{\text{target}}^T \left(\frac{d}{d\alpha} \hat{U} \right)^* \right\}. \quad (\text{D9})$$

APPENDIX E: DRAG, DETUNING, AND PHASE CORRECTIONS

While our displacement drive ideally implements the generator of spin displacements, we find that off-resonant drives neglected in the RWA can cause phase errors and leakage. These effects can be mitigated by increasing the pulse duration, thereby decreasing the drive amplitudes. However, incoherent errors due to relaxation and dephasing limit the pulse duration.

We employ DRAG, detuning, and phase corrections to mitigate these coherent errors, permitting shorter pulses with smaller incoherent errors. We calculate these corrections from numerical simulations of our system under the displacement drive. We first compute the propagator of the system $\hat{U}(t)$ given the system Hamiltonian $\hat{H}(t) = \hat{H}_0 + \hat{H}_d(t)$ for transmon Hamiltonian $\hat{H}_0$ and time-dependent drive Hamiltonian $\hat{H}_d(t)$:

$$\frac{d}{dt} \hat{U}(t) = -i\hat{H}(t)\hat{U}(t). \quad (\text{E1})$$

This simulation is carried out in the lab frame without making the rotating-wave approximation, using the transmon charge operator as the drive operator, starting at $t = 0$ and ending at the pulse duration $t = T$. Following the simulation, this propagator is transformed into the transmon frame: $\tilde{U}(T) = e^{i\hat{H}_0 T} \hat{U}(T)$. Its fidelity with respect to the ideal displacement is then

$$F = \frac{1}{d^2} |\text{Tr} [\hat{D}^\dagger(\pi/2) \tilde{U}(T)]|^2. \quad (\text{E2})$$

Instantaneous phase corrections take the form of SNAP gates applied immediately before and after the pulse,

$$\hat{U}_{\text{corr}}(T) = \hat{S}(\vec{\phi}_{\text{post}}) \hat{U}(T) \hat{S}(\vec{\phi}_{\text{pre}}). \quad (\text{E3})$$

We compute the propagator once and then use numerical optimization to find the SNAP gate phases that maximize the fidelity with respect to the ideal displacement.

DRAG and detuning corrections, meanwhile, require optimizing the properties of the pulse waveform itself. To this end, we employ gradient-based optimal control. Given the time-dependent propagator $\hat{U}(t)$, its derivative with respect to a control parameter x is given by [50]

$$\frac{d}{dt} \hat{U}'_x(t) = -i(\hat{H}'_x(t)\hat{U}(t) + \hat{H}(t)\hat{U}'_x(t)). \quad (\text{E4})$$

This differential equation is solved for each control parameter (the DRAG weights and drive detunings); then, the gradient of the fidelity is computed using the chain rule. The fidelity and its gradient are passed to a numerical

optimizer, which terminates when the gradient becomes sufficiently small. We note that our transmon Hamiltonian is modeled using scqubits [74,75], and the simulations are, in part, carried out using QuTiP [76].

APPENDIX F: VALIDATING LOGICAL QUBIT RANDOMIZED BENCHMARKING FIDELITIES

The randomized benchmarking protocol utilized in the main text is generally only applicable under certain assumptions regarding the error processes affecting the logical state. In particular, leakage out of the logical subspace must be handled with care. Furthermore, even in the absence of leakage, the fidelity estimated by randomized benchmarking is only valid within the logical subspace itself, making generalization to the overall displacement fidelity difficult.

We address this challenge here using a numerical model for the effect of noise on the displacement unitary and the resulting Clifford operations. We begin with a general quantum channel describing a CPTP perturbation to the displacement:

$$\mathcal{D}(\hat{\rho}) = (1 - q)\hat{D}\hat{\rho}\hat{D}^\dagger + q\mathcal{A}(\hat{D}\hat{\rho}\hat{D}^\dagger). \quad (\text{F1})$$

The error weight q , $0 \leq q \leq 1$, describes the strength of the error channel, with $q = 0$ corresponding to a perfect displacement operation. The superoperator $\mathcal{A}$ is taken to be a random CPTP map; thus, in the case where $q = 1$, the system undergoes a completely random evolution instead of the displacement. We draw $\mathcal{A}$ from an ensemble of CPTP maps using QuTiP [76], which follows the protocol described in Ref. [77]. One can view the channel $\mathcal{D}$ as a Kraus decomposition in which we have the Kraus operator $\sqrt{1 - q}\hat{D}$ corresponding to a perfect displacement operation and Kraus operators $\sqrt{q}\hat{A}_i\hat{D}$ for each operator $\hat{A}_i$ in the Kraus decomposition of $\mathcal{A}$ corresponding to the error channel's effect on the ideal displacement.

For a particular realization of the perturbed displacement map, we begin by computing the average gate fidelity of this map with respect to the ideal displacement operation. We take this quantity to be the underlying fidelity of the displacement. To relate this number to a randomized benchmarking result, we compute the map corresponding to each logical Clifford operation using the perturbed displacement. We then simulate randomized benchmarking of these Clifford operations, extracting a Clifford fidelity from a fit to the exponential decay. By repeating this procedure for many values of q and randomizations of the error channel, we extract an approximate relationship between the measured Clifford fidelity and the underlying displacement fidelity, shown in Fig. 4(c) in the main text.

APPENDIX G: QUDIT TOMOGRAPHY

1. Bayesian state tomography

Figure 5 in the main text contains the results of qudit state tomography measurements, which we describe here. In order to uniquely characterize an arbitrary $d \times d$ density matrix $\hat{\rho}$, we measure the projection of $\hat{\rho}$ onto a set of pure states consisting of the energy eigenstates $|n\rangle$ as well as superposition states

$$\frac{|n\rangle + |m\rangle}{\sqrt{2}} \quad (\text{G1})$$

and

$$\frac{|n\rangle + i|m\rangle}{\sqrt{2}} \quad (\text{G2})$$

for $0 \leq n < m \leq d - 1$. The projections onto the eigenstates $|n\rangle$ are performed simultaneously by taking an ensemble of qudit measurements with no pulses, thereby extracting the populations of each eigenstate. The projections onto superposition states, meanwhile, are performed by first applying a series of π pulses to map $|m\rangle$ onto $|n + 1\rangle$ and then either a $\hat{Y}(-\pi/2)$ or $\hat{X}(\pi/2)$ rotation in the $\{|n\rangle, |n + 1\rangle\}$ subspace to rotate the amplitudes initially in $(|n\rangle + |m\rangle)/\sqrt{2}$ and $(|n\rangle + i|m\rangle)/\sqrt{2}$ onto $|n\rangle$, respectively.

If the tomography pulses and projective measurements were perfect, one could, in principle, construct the density matrix directly from linear combinations of the measured populations. However, imperfections in the pulses and readout, as well as shot noise, can, in general, lead to unphysical density matrices. Therefore, we turn to Bayesian parameter estimation. The goal of Bayesian parameter estimation is to estimate the posterior distribution $p(\hat{\rho}|\{N_{|n\rangle}^{(k)}\})$, where $N_{|n\rangle}^{(k)}$ is the number of experimental shots in which we measure the qudit in state $|n\rangle$ after performing tomography sequence k . Bayes' theorem relates this posterior probability to a likelihood function $\mathcal{L}(\{N_{|n\rangle}^{(k)}\}|\hat{\rho})$ and a prior distribution $p(\hat{\rho})$:

$$p(\hat{\rho}|\{N_{|n\rangle}^{(k)}\}) \propto \mathcal{L}(\{N_{|n\rangle}^{(k)}\}|\hat{\rho})p(\hat{\rho}). \quad (\text{G3})$$

a. Likelihood function

The likelihood function is constructed from the multinomial distribution as follows. We describe each tomography pulse sequence k by a unitary operator $\hat{U}^{(k)}$, such that the density matrix we measure is $\hat{U}^{(k)}\hat{\rho}(\hat{U}^{(k)})^\dagger$. The population measured in $|n\rangle$ following this is given by

$$P_{|n\rangle}^{(k)} = \text{Tr}[\hat{U}^{(k)} \hat{\rho} (\hat{U}^{(k)})^\dagger |n\rangle\langle n|] = \text{Tr}[\hat{\rho} (\hat{U}^{(k)})^\dagger |n\rangle\langle n| \hat{U}^{(k)}]. \quad (\text{G4})$$

For a fixed k , the probability of obtaining measurement counts $N_{|n\rangle}^{(k)}$ is given by the multinomial distribution

$$P(N_{|0\rangle}^{(k)}, N_{|1\rangle}^{(k)}, \dots, N_{|d-1\rangle}^{(k)} | \hat{\rho}) \propto \prod_{n=0}^{d-1} (P_{|n\rangle}^{(k)})^{N_{|n\rangle}^{(k)}}, \quad (\text{G5})$$

where the proportionality factor is constant with respect to $\hat{\rho}$ and can therefore be neglected. The likelihood function, then, is simply the product of these probabilities over all measurement sequences k :

$$\mathcal{L}(\{N_{|n\rangle}^{(k)}\} | \hat{\rho}) \propto \prod_k \prod_{n=0}^{d-1} (P_{|n\rangle}^{(k)})^{N_{|n\rangle}^{(k)}}. \quad (\text{G6})$$

In practice, we compute the log-likelihood function $\ln \mathcal{L}(\{N_{|n\rangle}^{(k)}\} | \hat{\rho})$, which, up to an additive constant, is given by

$$\ln \mathcal{L}(\{N_{|n\rangle}^{(k)}\} | \hat{\rho}) = \sum_k \sum_{n=0}^{d-1} N_{|n\rangle}^{(k)} \ln(P_{|n\rangle}^{(k)}). \quad (\text{G7})$$

Next, we describe a fast, vectorized method for evaluating this likelihood function, which is essential for Bayesian estimation of a large number of parameters. We begin by column-vectorizing the density matrix:

$$\vec{\rho} = \begin{pmatrix} \rho_{00} \\ \rho_{10} \\ \vdots \\ \rho_{d-1,0} \\ \vdots \\ \rho_{d-1,d-1} \end{pmatrix}. \quad (\text{G8})$$

Next, we form a set of measurement operators from the tomography unitaries described above:

$$\hat{E}^{(kn)} = (\hat{U}^{(k)})^\dagger |n\rangle\langle n| \hat{U}^{(k)}. \quad (\text{G9})$$

which are then row vectorized,

$$\vec{E}^{(kn)} = \begin{pmatrix} E_{00}^{(kn)} & E_{01}^{(kn)} & \dots & E_{0,d-1}^{(kn)} & \dots & E_{d-1,d-1}^{(kn)} \end{pmatrix}, \quad (\text{G10})$$

and we form a matrix $\mathbf{A}$ by stacking these rows:

$$\mathbf{A} = \begin{pmatrix} \vec{E}^{(00)} \\ \vec{E}^{(01)} \\ \vdots \\ \vec{E}^{(0,d-1)} \\ \vec{E}^{(10)} \\ \vdots \end{pmatrix}. \quad (\text{G11})$$

Then, it is straightforward to show that the matrix-vector product $\mathbf{A}\vec{\rho}$ yields the vectorized probabilities:

$$\vec{P} = \begin{pmatrix} P_{|0\rangle}^{(0)} \\ P_{|1\rangle}^{(0)} \\ \vdots \\ P_{|d-1\rangle}^{(0)} \\ P_{|0\rangle}^{(1)} \\ \vdots \end{pmatrix}. \quad (\text{G12})$$

The likelihood function evaluation is completed by first evaluating $\vec{P} = \mathbf{A}\vec{\rho}$, taking its elementwise logarithm, multiplying the result elementwise with the measured counts $\vec{N}$ (vectorized in the same order as $\vec{P}$), and then computing the sum of the entire resulting vector.

b. Density matrix parametrization and prior distribution

Bayesian parameter estimation requires that we impose some prior probability distribution $p(\hat{\rho})$ over the space of density matrices, which encodes our prior expectations for the state. Furthermore, any physically valid density matrix must be Hermitian, positive semidefinite, and trace normalized. Here, we follow the parametrization and prior distribution described in Ref. [78]. For d non-negative real numbers γ_i and d complex-valued column vectors $\mathbf{z}_i$, each of size d , the density matrix

$$\hat{\rho} = \sum_i \left(\frac{\gamma_i}{\sum_j \gamma_j} \right) \frac{\mathbf{z}_i \mathbf{z}_i^\dagger}{|\mathbf{z}_i|^2} \quad (\text{G13})$$

is guaranteed to be physically valid. Moreover, this parametrization permits a straightforward prior distribution:

$$p(\hat{\rho}) \propto \prod_i \gamma_i^{\alpha-1} e^{-\gamma_i} e^{-\frac{1}{2} \mathbf{z}_i^\dagger \mathbf{z}_i}. \quad (\text{G14})$$

As detailed in Ref. [78], α is a hyperparameter describing the expected purity of the state: For $\alpha = 1$, the prior is uniform over all physically realizable states, while for $\alpha < 1$, it favors lower-rank (i.e., more pure) density

matrices. However, because we expect state purities to be highly variable across experiments and qudit dimensions, here we adopt a hierarchical Bayesian model. We impose a uniform hyperprior on α over the range $0 \leq \alpha \leq 1$; hence, the prior over $\hat{\rho}$ alone is found by marginalizing over α :

$$p(\hat{\rho}) = \int_0^1 p(\hat{\rho}|\alpha)p(\alpha)d\alpha. \quad (\text{G15})$$

Noting that

$$p(\hat{\rho}) \propto \frac{\prod_i \gamma_i - 1}{\prod_i \gamma_i \ln(\prod_i \gamma_i)} \prod_i e^{-\gamma_i} e^{-\frac{1}{2} \mathbf{z}_i^\dagger \mathbf{z}_i}. \quad (\text{G17})$$

the α -independent prior is easily shown to be




FIG. 11. Posterior samples for the $d = 3$ Hadamard experiment [Fig. 5(a) in the main text]. The mean and 95% confidence intervals shown in the main text are computed directly from these samples. The red lines show the ideal values of the density matrix elements if the state were perfect.

c. Monte Carlo sampling

Given a posterior distribution $p(\hat{\rho}|\{N_{|n\rangle}^{(k)}\})$, the expectation value of any function $\phi(\hat{\rho})$ is given by

$$\langle\phi(\hat{\rho})\rangle = \int_{\hat{\rho}} p(\hat{\rho}|\{N_{|n\rangle}^{(k)}\}) \phi(\hat{\rho}) d\hat{\rho}. \quad (\text{G18})$$

This can be used to evaluate, for instance, the expectation value of $\hat{\rho}$ itself or the expectation value of the state fidelity. As the dimension of $\hat{\rho}$ increases, however, the direct evaluation of this integral quickly becomes intractable. The standard approach to this problem is to sample directly from the posterior itself. Given a set of M samples drawn from the posterior, any expectation value $\langle\phi(\hat{\rho})\rangle$ is easily estimated as its empirical mean:

$$\langle\phi(\hat{\rho})\rangle = \frac{1}{M} \sum_{i=1}^M \phi(\hat{\rho}_i). \quad (\text{G19})$$

In this work, we apply Markov chain Monte Carlo sampling to draw fair samples from the posterior using the emcee PYTHON package [79]. Figure 11 shows the posterior sample distribution over the density matrix elements for the $d = 3$ Hadamard experiment [Fig. 5(a) in the main text].

2. Wigner state tomography

As described above and in the main text, the spin Wigner function of a state $\hat{\rho}$ is given by

$$W(\theta, \phi) = 2\text{Tr}[\hat{D}^\dagger(\theta, \phi) \hat{\rho} \hat{D}(\theta, \phi) \hat{\Pi}], \quad (\text{G20})$$

where $\hat{\Pi}$ is a paritylike operator that is diagonal in the $\hat{J}_z$ eigenbasis. Here, we describe the data analysis used to extract the density matrices shown in Fig. 3(c) in the main text from experimentally measured Wigner functions.

Rather than the Bayesian methods used above, here we simply numerically optimize the density matrix to maximize the overlap between its Wigner function and the one we measure. For any two operators $\hat{A}$ and $\hat{B}$, the trace of their product can be expressed as an integral over the product of their Wigner functions [39,41]:

$$\text{Tr}[\hat{A} \hat{B}] = \frac{2j+1}{4\pi} \int_{\theta=0}^{\pi} \int_{\phi=0}^{2\pi} W_{\hat{A}}(\theta, \phi) W_{\hat{B}}(\theta, \phi) \sin(\theta) d\phi d\theta. \quad (\text{G21})$$

Because the Wigner function is measured at a discrete set of equally spaced θ and ϕ , we replace this integral with a sum and minimize the quantity

$$\mathcal{C}(\hat{\rho}) = -\sum_{i,j} W_{\hat{\rho}}(\theta_i, \phi_j) W_{\text{exp}}(\theta_i, \phi_j) \sin(\theta_i), \quad (\text{G22})$$

which, up to a constant factor, approximates the state overlap.

We begin by precomputing and storing the (complex-valued) Wigner functions of matrix elements $|n\rangle\langle m|$ at each discrete phase-space point measured in the experiment:

$$W_{nm}(\theta_i, \phi_j) = 2\text{Tr}[\hat{D}^\dagger(\theta_i, \phi_j) |n\rangle\langle m| \hat{D}(\theta_i, \phi_j) \hat{\Pi}]. \quad (\text{G23})$$

We parametrize the density matrix using the Cholesky decomposition,

$$\hat{\rho} = \frac{\hat{L} \hat{L}^\dagger}{\text{Tr}[\hat{L} \hat{L}^\dagger]}, \quad (\text{G24})$$

where $\hat{L}$ is a lower-triangular matrix with real, non-negative values on the diagonal. This parametrization guarantees that the density matrix is Hermitian, positive semidefinite, and trace normalized. The Wigner function of this density matrix is efficiently computed as

$$W_{\hat{\rho}}(\theta_i, \phi_j) = \sum_{n,m=0}^{d-1} \langle m|\hat{\rho}|n\rangle W_{nm}(\theta_i, \phi_j), \quad (\text{G25})$$

and we numerically optimize the elements of $\hat{L}$ to minimize $\mathcal{C}(\hat{\rho})$.

APPENDIX H: QUDIT CLIFFORD GROUP

In this appendix, we discuss the methods used in the qudit Clifford randomized benchmarking experiments shown in Fig. 5 in the main text—in particular, the construction of the qudit Clifford group itself. We begin with the generators discussed in the main text [34]: $\{\hat{1}, \hat{H}_d, \hat{Z}_d, \hat{S}_d\}$, where, for $\omega_d = e^{i\frac{2\pi}{d}}$,

$$\hat{H}_d = \frac{1}{\sqrt{d}} \sum_{n,m=0}^{d-1} \omega_d^{nm} |n\rangle\langle m|, \quad (\text{H1})$$

$$\hat{Z}_d = \sum_{n=0}^{d-1} \omega_d^n |n\rangle\langle n|, \quad (\text{H2})$$

and

$$\hat{S}_d = \sum_{n=0}^{d-1} \omega_d^{\frac{n(n+d)}{2}} |n\rangle\langle n|. \quad (\text{H3})$$

Our goal is to express every element of the Clifford group C_d as a product of these generators. We present the full algorithm here as PYTHON code. We assume we have a function `|equal_upto_phase(d, U1, U2)|`, which checks whether two d -dimensional unitaries are equal up to a global phase. The arguments to `|build_clifford_group|` are the dimension d and a dictionary mapping string names to the generators given above, stored as matrices.

```

1 def build_clifford_group(d, generators):
2     # This holds sequences of strings
3     clifford_sequences = [[g] for g in generators.keys()]
4     # This holds operators (matrices), in the same order as clifford_sequences
5     clifford_operators = [generators[g] for g in generators.keys()]
6     # Keep track of the current size of the group
7     current_cardinality = len(clifford_sequences)
8     # Stores a map of the form (g, C.i) → C.j, i.e., the action of a generator
9     # g on a group element C.i
10    generator_product_map = {}
11    index_start = 0
12    added_atleast_one = True
13    while added_atleast_one:
14        sequences_to_append = []
15        operators_to_append = []
16        added_atleast_one = False
17        for g in generators.keys():
18            G = generators[g]
19            for i in range(index_start, current_cardinality):
20                # Iterate over combinations of generators and Clifford group elements,
21                # taking their product and checking if the result is already in the
22                # group.
23                # We can speed this up considerably by only checking
24                # (generator, group element) combinations which have *not* been
25                # checked thus far.
26                sequence = clifford_sequences[i]
27                C = clifford_operators[i]
28                U = G @ C # Take their product
29                # To start, assume the product is *not* in the group.
30                already_in_group = False
31                for j in range(current_cardinality):
32                    # Check through the elements which were in the group prior to this
33                    # iteration.
34                    if equal_upto_phase(d, U, clifford_operators[j]):
35                        already_in_group = True
36                        generator_product_map[(g, i)] = j
37                        break
38                if not already_in_group:
39                    # Also check against those already added in this iteration, to avoid
40                    # double-adding them.
41                    for k in range(len(operators_to_append)):
42                        if equal_upto_phase(d, U, operators_to_append[k]):
43                            already_in_group = True
44                            generator_product_map[(g, i)] = current_cardinality + k
45                            break
46                if not already_in_group:
47                    sequences_to_append.append([g] + sequence)
48                    operators_to_append.append(U)
49                    generator_product_map[(g, i)] = current_cardinality + len(sequences_to_append) - 1
50                    added_atleast_one = True
51            index_start = current_cardinality
52            current_cardinality += len(sequences_to_append)
53            clifford_sequences = clifford_sequences + sequences_to_append
54            clifford_operators = clifford_operators + operators_to_append
55    return clifford_sequences, clifford_operators, generator_product_map

```

As seen from the return values, this function simultaneously does the following:

- (1) It decomposes each element of the Clifford group into a sequence of generators. Furthermore, we are guaranteed to have found the *shortest* sequence of generators producing each Clifford operator.
- (2) It produces the matrix representation of each group element.
- (3) It returns a map (PYTHON dictionary) indicating the action of each generator on each group element. In other words, we have a lookup table that stores the Clifford group element c_j obtained by taking the product of a generator g and any other Clifford group element c_i .

This final data structure permits the efficient construction of the group's Cayley table, which is essential for randomized benchmarking. For each pair of group elements (c_i, c_j) , the naive approach to finding the group element c_k

representing their product is to simply iterate over the full Clifford group, checking the product against each element. This approach has $\mathcal{O}(|C_d|^3 M(d))$ time complexity, where $|C_d|$ is the size of the group and $M(d)$ is the time complexity of checking whether two matrices are equal up to a global phase. We perform this check for unitaries $\hat{U}_1$ and $\hat{U}_2$ by checking whether $(1/d)\text{Tr}[\hat{U}_1^\dagger \hat{U}_2] - 1 = 0$; hence, $M(d) \sim d^4$. Therefore, the naive approach is prohibitively slow for even moderate values of d .

Instead, consider expressing c_i as a sequence of generators: $c_i = g_{(k)} g_{(k-1)} \dots g_{(0)}$. Then, the product $c_i \times c_j$ is

$$c_i \times c_j = (g_{(k)} g_{(k-1)} \dots g_{(0)}) \times c_j. \quad (\text{H4})$$

Rewriting this product as $c_i \times c_j = (g_{(k)} g_{(k-1)} \dots g_{(1)}) (g_{(0)} \times c_j)$, we refer to the lookup table returned by

`|build_clifford_group|` to find the Clifford group element $c'_j = g_{(0)} \times c_j$. Thus,

$$c_i \times c_j = (g_{(k)} g_{(k-1)} \dots g_{(1)}) \times c'_j. \quad (\text{H5})$$

We continue in this fashion, traversing the generator sequence right to left, ultimately yielding the desired $c_k = c_i \times c_j$. Importantly, this procedure does not involve any matrix multiplication: We instead need only refer to the label of each Clifford group element stored in the lookup table.

APPENDIX I: OUTLOOK

In this appendix, we propose a simplified model for the displacement pulse errors, and using this model, we present an estimate for realistically achievable fidelities given state-of-the-art transmon devices. We begin with the model presented in the main text, where the total error is a sum of coherent and incoherent parts that vary independently depending on the pulse properties and the device coherence times:

$$\mathcal{E} = \mathcal{E}_{\text{coh}} + \mathcal{E}_{\text{inc}}. \quad (\text{I1})$$

We assume that the incoherent part of the error is proportional to the length of the pulse, with the proportionality constant being the average decay rate of each qudit transition: For a pulse duration T ,

$$\mathcal{E}_{\text{inc}} = \frac{T}{\tau_{\text{inc}}}, \quad (\text{I2})$$

where for a $|0\rangle \leftrightarrow |1\rangle$ decay time T_{01} ,

$$\frac{1}{\tau_{\text{inc}}} = \frac{1}{d-1} \sum_{n=1}^{d-1} \frac{n}{T_{01}} = \frac{d}{2T_{01}}, \quad (\text{I3})$$

assuming a linear scaling of the decay rate with excitation number. Here, we consider transmons of varying qubit

frequency, in which case, the natural quantity to consider is the Q factor: $Q = T_{01} f_{01}$. We note that, while Q is not necessarily fixed as one varies qubit frequency, with its detailed behavior being dependent on the particular noise sources at hand, it is still the relevant quantity to use to relate the coherence properties of transmons across frequencies. With this in mind, we have

$$\mathcal{E}_{\text{inc}} = \frac{f_{01} d}{2Q} T. \quad (\text{I4})$$

Meanwhile, we primarily attribute the coherent errors to phase errors and leakage resulting from the off-resonant drive terms neglected in the RWA. This approach suggests that the coherent error should scale proportionally with the square of the drive strength and inversely with the square of the anharmonicity. In the following, we assume a fixed ratio $E_J/E_C = 270$, chosen to ensure a large enough number of confined states with sufficiently low charge dispersion to permit the qudit operations considered here. Using the approximate qubit frequency $f_{01} = \sqrt{8E_J E_C} - E_C$, we can write

$$E_C = \frac{f_{01}}{\sqrt{8 \frac{E_J}{E_C} - 1}}. \quad (\text{I5})$$

Noting that the anharmonicity is equal in magnitude to E_C ,

$$\mathcal{E}_{\text{coh}} = \frac{A(E_J, E_C, d)}{E_C^2 T^2}, \quad (\text{I6})$$

where $A(E_J, E_C, d)$ is a dimensionless phenomenological parameter capturing the time-independent scale of coherent errors; because E_J/E_C is fixed, we can equivalently parametrize it as $A(f_{01}, d)$. Substituting our equation for E_C , we have

$$\mathcal{E}_{\text{coh}} = \left(\sqrt{8 \frac{E_J}{E_C} - 1} \right)^2 \frac{A(f_{01}, d)}{f_{01}^2 T^2}. \quad (\text{I7})$$




FIG. 12. Modeling achievable displacement pulse fidelities.

Solving for the optimal pulse duration, we find

$$T_{\text{opt}} = \frac{1}{f_{01}} \left[4 \left(\sqrt{8 \frac{E_J}{E_C}} - 1 \right)^2 \frac{QA(f_{01}, d)}{d} \right]^{1/3}. \quad (\text{I8})$$

We note that upon substituting this solution into the equations for $\mathcal{E}_{\text{coh}}$ and $\mathcal{E}_{\text{inc}}$, the factor of f_{01} cancels out everywhere except in the argument to $A(f_{01}, d)$. In order to characterize this dependence, we perform a series of numerical simulations at differing transmon frequencies and extract $A(f_{01}, d)$ for each, showing the results in Fig. 12 (left). We find that, to a good approximation, $A(f_{01}, d)$ is independent of f_{01} , so we are justified in considering it to be a function of d only. Therefore, we ultimately find that the minimum infidelity is given by

$$\mathcal{E}_{\text{min}} = \frac{3}{2} \left(\frac{A(d)}{2} \right)^{1/3} \left(\frac{d}{Q} \right)^{2/3} \left(\sqrt{8 \frac{E_J}{E_C}} - 1 \right)^{2/3}. \quad (\text{I9})$$

With this model and the simulated $A(d)$ values in mind, we plot expected displacement fidelities as a function of Q in Fig. 12 (right). The vertical dashed lines represent our current device, with a frequency of roughly 5 GHz and a T_1 of roughly 50 μs , and a hypothetical device with an order of magnitude larger T_1 , representing the current state of the art [59]. The black dot is our experimentally measured fidelity for $d = 8$ in the absence of pulse corrections.

In this analysis, we have considered standard transmon devices only and have restricted E_J/E_C to a value appropriate for eight-dimensional qudits. In practice, one would ideally choose the device parameters that allow for the desired qudit dimension while also maximizing anharmonicity, which would further reduce coherent errors for $d < 8$. Furthermore, we have not considered the pulse corrections discussed in the main text, and in that sense, our model acts as an upper bound on the achievable error given a particular device. Finally, we note that alternative device designs such as the inductively shunted transmon [60] may be able to encode high-dimensional, long-coherence-time qudits without the restriction on anharmonicity inherent to high- E_J/E_C transmons.

-
- [1] Y. Wang, Z. Hu, B. C. Sanders, and S. Kais, *Qudits and high-dimensional quantum computing*, *Front. Phys.* **8**, 479 (2020).
 - [2] F. Dolde *et al.*, *High-fidelity spin entanglement using optimal control*, *Nat. Commun.* **5**, 3371 (2014).
 - [3] S. Asaad *et al.*, *Coherent electrical control of a single high-spin nucleus in silicon*, *Nature (London)* **579**, 205 (2020).
 - [4] M. Ringbauer *et al.*, *A universal qudit quantum processor with trapped ions*, *Nat. Phys.* **18**, 1053 (2022).
 - [5] Y. Chi *et al.*, *A programmable qudit-based quantum processor*, *Nat. Commun.* **13**, 1166 (2022).

- [6] R. Bianchetti, S. Filipp, M. Baur, J. M. Fink, C. Lang, L. Steffen, M. Boissonneault, A. Blais, and A. Wallraff, *Control and tomography of a three level superconducting artificial atom*, *Phys. Rev. Lett.* **105**, 223601 (2010).
- [7] S. J. Evered *et al.*, *High-fidelity parallel entangling gates on a neutral-atom quantum computer*, *Nature (London)* **622**, 268 (2023).
- [8] J.-S. Chen *et al.*, *Benchmarking a trapped-ion quantum computer with 30 qubits*, *Quantum* **8**, 1516 (2024).
- [9] R. Acharya *et al.*, *Suppressing quantum errors by scaling a surface code logical qubit*, *Nature (London)* **614**, 676 (2023).
- [10] X. Xue *et al.*, *Quantum logic with spin qubits crossing the surface code threshold*, *Nature (London)* **601**, 343 (2022).
- [11] Z. Ni *et al.*, *Beating the break-even point with a discrete-variable-encoded logical qubit*, *Nature (London)* **616**, 56 (2023).
- [12] V. V. Sivak *et al.*, *Real-time quantum error correction beyond break-even*, *Nature (London)* **616**, 50 (2023).
- [13] E. T. Campbell, H. Anwar, and D. E. Browne, *Magic-state distillation in all prime dimensions using quantum Reed-Muller codes*, *Phys. Rev. X* **2**, 041021 (2012).
- [14] P. Gokhale *et al.*, *Asymptotic improvements to quantum circuits via qutrits*, *Proceedings of the 46th International Symposium on Computer Architecture—ISCA '19* (2019), pp. 554–566, [arXiv:1905.10481](https://arxiv.org/abs/1905.10481).
- [15] A. Ciavarella, N. Klco, and M. J. Savage, *Trailhead for quantum simulation of $SU(3)$ Yang-Mills lattice gauge theory in the local multiplet basis*, *Phys. Rev. D* **103**, 094501 (2021).
- [16] D. González-Cuadra, T. V. Zache, J. Carrasco, B. Kraus, and P. Zoller, *Hardware efficient quantum simulation of non-Abelian gauge theories with qudits on Rydberg platforms*, *Phys. Rev. Lett.* **129**, 160501 (2022).
- [17] E. J. Gustafson, *Prospects for simulating a qudit-based model of $(1+1)D$ scalar QED*, *Phys. Rev. D* **103**, 114505 (2021).
- [18] G. Catarina and J. Fernández-Rossier, *Hubbard model for spin-1 Haldane chains*, *Phys. Rev. B* **105**, L081116 (2022).
- [19] Z.-X. Gong, M. F. Maghrebi, A. Hu, M. Foss-Feig, P. Richerme, C. Monroe, and A. V. Gorshkov, *Kaleidoscope of quantum phases in a long-range interacting spin-1 chain*, *Phys. Rev. B* **93**, 205115 (2016).
- [20] M. Blok *et al.*, *Quantum information scrambling on a superconducting qutrit processor*, *Phys. Rev. X* **11**, 021010 (2021).
- [21] M. J. Peterer *et al.*, *Coherence and decay of higher energy levels of a superconducting transmon qubit*, *Phys. Rev. Lett.* **114**, 010501 (2015).
- [22] X. Wu, S. L. Tomarken, N. A. Petersson, L. A. Martinez, Y. J. Rosen, and J. L. DuBois, *High-fidelity software-defined quantum logic on a superconducting qudit*, *Phys. Rev. Lett.* **125**, 170502 (2020).
- [23] L. Chen *et al.*, *Transmon qubit readout fidelity at the threshold for quantum error correction without a quantum-limited amplifier*, *npj Quantum Inf.* **9**, 26 (2023).
- [24] L. B. Nguyen *et al.*, *Empowering a qudit-based quantum processor by traversing the dual bosonic ladder*, *Nat. Commun.* **15**, 7117 (2024).

- [25] T. Kehrer, T. Nadolny, and C. Bruder, *Improving transmon qudit measurement on IBM quantum hardware*, *Phys. Rev. Res.* **6**, 013050 (2024).
- [26] S. Cao *et al.*, *Emulating two qubits with a four-level transmon qudit for variational quantum algorithms*, *Quantum Sci. Technol.* **9**, 035003 (2024).
- [27] P. Liu *et al.*, *Performing $SU(d)$ operations and rudimentary algorithms in a superconducting transmon qudit for $d = 3$ and $d = 4$* , *Phys. Rev. X* **13**, 021028 (2023).
- [28] A. Morvan *et al.*, *Qutrit randomized benchmarking*, *Phys. Rev. Lett.* **126**, 210504 (2021).
- [29] A. Stark *et al.*, *Clock transition by continuous dynamical decoupling of a three-level system*, *Sci. Rep.* **8**, 14807 (2018).
- [30] L. M. Seifert, Z. Li, T. Roy, D. I. Schuster, F. T. Chong, and J. M. Baker, *Exploring ququart computation on a transmon using optimal control*, *Phys. Rev. A* **108**, 062609 (2023).
- [31] M. A. Yurtalan, J. Shi, M. Kononenko, A. Lupascu, and S. Ashhab, *Implementation of a Walsh-Hadamard gate in a superconducting qutrit*, *Phys. Rev. Lett.* **125**, 180504 (2020).
- [32] M. Neeley *et al.*, *Emulation of a quantum spin with a superconducting phase qudit*, *Science* **325**, 722 (2009).
- [33] R. W. Heeres *et al.*, *Implementing a universal gate set on a logical qubit encoded in an oscillator*, *Nat. Commun.* **8**, 94 (2017).
- [34] M. Jafarzadeh, Y.-D. Wu, Y. R. Sanders, and B. C. Sanders, *Randomized benchmarking for qudit Clifford gates*, *New J. Phys.* **22**, 063014 (2020).
- [35] J. Koch *et al.*, *Charge-insensitive qubit design derived from the cooper pair box*, *Phys. Rev. A* **76**, 042319 (2007).
- [36] A. Blais, A. L. Grimsmo, S. M. Girvin, and A. Wallraff, *Circuit quantum electrodynamics*, *Rev. Mod. Phys.* **93**, 025005 (2021).
- [37] Z. Wang, R. W. Parker, E. Champion, and M. S. Blok, *High- E_J/E_C transmon qudits with up to 12 levels*, *Phys. Rev. Appl.* **23**, 034046 (2025).
- [38] A. M. Perelomov, *Generalized coherent states and some of their applications*, *Sov. Phys. Usp.* **20**, 703 (1977).
- [39] J. Davis, M. Kumari, R. B. Mann, and S. Ghose, *Wigner negativity in spin- j systems*, *Phys. Rev. Res.* **3**, 033134 (2021).
- [40] R. L. Stratonovich, *On distributions in representation space*, *J. Exp. Theor. Phys.* **4**, 891 (1957).
- [41] S. Heiss and S. Weigert, *Discrete Moyal-type representations for a spin*, *Phys. Rev. A* **63**, 012105 (2000).
- [42] H. de Guise, O. Di Matteo, and L. L. Sánchez-Soto, *Simple factorization of unitary transformations*, *Phys. Rev. A* **97**, 022328 (2018).
- [43] R. W. Heeres, B. Vlastakis, E. Holland, S. Krastanov, V. V. Albert, L. Frunzio, L. Jiang, and R. J. Schoelkopf, *Cavity state manipulation using photon-number selective phase gates*, *Phys. Rev. Lett.* **115**, 137002 (2015).
- [44] S. Krastanov, V. V. Albert, C. Shen, C. L. Zou, R. W. Heeres, B. Vlastakis, R. J. Schoelkopf, and L. Jiang, *Universal control of an oscillator with dispersive coupling to a qubit*, *Phys. Rev. A* **92**, 040303(R) (2015).
- [45] M. Kudra *et al.*, *Robust preparation of Wigner-negative states with optimized snap-displacement sequences*, *PRX Quantum* **3**, 030301 (2022).
- [46] W.-L. Ma *et al.*, *Quantum control of bosonic modes with superconducting circuits*, *Sci. Bull.* **66**, 1789 (2021).
- [47] D. C. McKay, C. J. Wood, S. Sheldon, J. M. Chow, and J. M. Gambetta, *Efficient Z gates for quantum computing*, *Phys. Rev. A* **96**, 022330 (2017).
- [48] E. Magesan, J. M. Gambetta, and J. Emerson, *Characterizing quantum gates via randomized benchmarking*, *Phys. Rev. A* **85**, 042311 (2012).
- [49] F. Motzoi, J. M. Gambetta, P. Rebentrost, and F. K. Wilhelm, *Simple pulses for elimination of leakage in weakly nonlinear qubits*, *Phys. Rev. Lett.* **103**, 110501 (2009).
- [50] S. Machnes, E. Assémat, D. Tannor, and F. K. Wilhelm, *Tunable, flexible, and efficient optimization of control pulses for practical qubits*, *Phys. Rev. Lett.* **120**, 150401 (2018).
- [51] J. M. Epstein, A. W. Cross, E. Magesan, and J. M. Gambetta, *Investigating the limits of randomized benchmarking protocols*, *Phys. Rev. A* **89**, 062321 (2014).
- [52] A. M. C. R. Stroud, *Quantum fast Fourier transform using multilevel atoms*, *J. Mod. Opt.* **49**, 2115 (2002).
- [53] H. S. Tonchev and N. V. Vitanov, *Quantum phase estimation and quantum counting with qudits*, *Phys. Rev. A* **94**, 042307 (2016).
- [54] T. Fösel, S. Krastanov, F. Marquardt, and L. Jiang, *Efficient cavity control with snap gates*, *arXiv:2004.14256*.
- [55] E. Magesan *et al.*, *Efficient measurement of quantum gate error by interleaved randomized benchmarking*, *Phys. Rev. Lett.* **109**, 080505 (2012).
- [56] A. Javadi-Abhari *et al.*, *Quantum computing with Qiskit*, *arXiv:2405.08810*.
- [57] D. C. McKay, C. J. Wood, S. Sheldon, J. M. Chow, and J. M. Gambetta, *Efficient z gates for quantum computing*, *Phys. Rev. A* **96**, 022330 (2017).
- [58] D. Amaro-Alcalá, B. C. Sanders, and H. de Guise, *Randomised benchmarking for universal qudit gates*, *New J. Phys.* **26**, 073052 (2024).
- [59] C. Wang *et al.*, *Towards practical quantum computers: Transmon qubit with a lifetime approaching 0.5 milliseconds*, *npj Quantum Inf.* **8**, 3 (2022).
- [60] F. Hassani *et al.*, *Inductively shunted transmons exhibit noise insensitive plasmon states and a fluxon decay exceeding 3 hours*, *Nat. Commun.* **14**, 3968 (2023).
- [61] N. Goss *et al.*, *High-fidelity qutrit entangling gates for superconducting circuits*, *Nat. Commun.* **13**, 7481 (2022).
- [62] D. Gottesman, A. Kitaev, and J. Preskill, *Encoding a qubit in an oscillator*, *Phys. Rev. A* **64**, 012310 (2001).
- [63] N. Ofek *et al.*, *Extending the lifetime of a quantum bit with error correction in superconducting circuits*, *Nature (London)* **536**, 441 (2016).
- [64] A. Grimm *et al.*, *Stabilization and operation of a Kerr-cat qubit*, *Nature (London)* **584**, 205 (2020).
- [65] J. A. Gross, *Designing codes around interactions: The case of a spin*, *Phys. Rev. Lett.* **127**, 010504 (2021).
- [66] S. Omanakuttan, V. Buchemavari, J. A. Gross, I. H. Deutsch, and M. Marvian, *Fault-tolerant quantum computation using large spin-cat codes*, *PRX Quantum* **5**, 020355 (2024).
- [67] M. Illa, C. E. P. Robin, and M. J. Savage, *Qu8its for quantum simulations of lattice quantum chromodynamics*, *Phys. Rev. D* **110**, 014507 (2024).

- [68] D. González-Cuadra, T. V. Zache, J. Carrasco, B. Kraus, and P. Zoller, *Hardware efficient quantum simulation of non-Abelian gauge theories with qudits on Rydberg platforms*, *Phys. Rev. Lett.* **129**, 160501 (2022).
- [69] S. Roy *et al.*, *Synthetic high angular momentum spin dynamics in microwave oscillator*, *Phys. Rev. X* **15**, 021009 (2025).
- [70] X. Yu *et al.*, *Schrödinger cat states of a nuclear spin qudit in silicon*, *Nat. Phys.* **21**, 362 (2025).
- [71] R. Shillito, A. Petrescu, J. Cohen, J. Beall, M. Hauru, M. Ganahl, A. G. M. Lewis, G. Vidal, and A. Blais, *Dynamics of transmon ionization*, *Phys. Rev. Appl.* **18**, 034031 (2022).
- [72] S. Lloyd, *Almost any quantum logic gate is universal*, *Phys. Rev. Lett.* **75**, 346 (1995).
- [73] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes 3rd Edition: The Art of Scientific Computing*, 3rd ed. (Cambridge University Press, Cambridge, England, 2007).
- [74] P. Groszkowski and J. Koch, *Scqubits: A PYTHON package for superconducting qubits*, *Quantum* **5**, 583 (2021).
- [75] S. P. Chitta, T. Zhao, Z. Huang, I. Mondragon-Shem, and J. Koch, *Computer-aided quantization and numerical analysis of superconducting circuits*, *New J. Phys.* **24**, 103020 (2022).
- [76] J. Johansson, P. Nation, and F. Nori, *QUTIP2: A PYTHON framework for the dynamics of open quantum systems*, *Comput. Phys. Commun.* **184**, 1234 (2013).
- [77] W. Bruzda, V. Cappellini, H.-J. Sommers, and K. Zyczkowski, *Random quantum operations*, *Phys. Lett. A* **373**, 320 (2009).
- [78] J. M. Lukens, K. J. H. Law, A. Jasra, and P. Lougovski, *A practical and efficient approach for bayesian quantum state estimation*, *New J. Phys.* **22**, 063038 (2020).
- [79] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, *emcee: The MCMC Hammer*, *Publ. Astron. Soc. Pac.* **125**, 306 (2013).