Supplementary information:
Engineering direct nonlinear coupling between microwave photons

Adrian Copetudo,[1] Amon M. Kasper,[1] Tanjung Krisnanda,[1] Gregoire
Veyrac,[1] Shushen Qin,[1] Hui Khoon Ng,[2, 1] and Yvonne Y. Gao[1, 2, *]
[1] *Centre for Quantum Technologies, National University of Singapore, Singapore*
[2] *Department of Physics, National University of Singapore, Singapore*
(Dated: June 6, 2026)

I. COMPARISON BETWEEN ENTANGLING OPERATIONS IN CQED

Entanglement between bosonic modes can be generated either by applying linear operations to non-classical states or by engineering nonlinear interactions between the modes themselves. A paradigmatic example of the former is the beam-splitter (BS) operation [1], which is a purely linear (Gaussian) transformation that can generate entanglement when acting on non-classical states such as single-photon or squeezed states. However, the BS is not an intrinsically entangling operation: when applied to classical states, such as vacuum or coherent states, it does not generate entanglement. Moreover, the BS does not generally preserve the individual code spaces of the oscillators, as it redistributes excitations between the modes. For example, when each oscillator is encoded in the single-photon manifold, the logical basis state

$$|1,1\rangle \xrightarrow{\text{BS}} \frac{|2,0\rangle + |0,2\rangle}{\sqrt{2}}, \tag{1}$$

thereby leaving the encoded subspace. The final state in Eq. 1 is a NOON state with $N=2$, which belongs to the class of Gaussian-entanglable (GE) states, defined as those states that can be generated by Gaussian operations acting on separable states [2]. However, GE states are a limited subset of all possible entangled states. For example, higher-order NOON states with $N \geq 3$, essential for quantum metrology applications, are non-GE. Hence, intrinsically nonlinear operations are a necessary ingredient to generate arbitrary entangled states [3].

The standard approach to generating entanglement between bosonic modes in cQED exploits the nonlinear dispersive interaction between the oscillators and an ancillary superconducting qubit. Two prominent examples are the swap–wait–swap protocol [4, 5], in which a sideband transition transfers an excitation from one oscillator to the ancilla and the resulting dispersive interaction imprints a conditional phase on the second oscillator, and the geometric-phase protocol [6], where the ancilla is driven around a closed trajectory in phase space conditioned on the photon occupations of the cavities, thereby acquiring a conditional geometric phase. Although these protocols implement nonlinear entangling gates, they temporarily populate the ancillary element and do not preserve the individual code spaces of the oscillators throughout the operation. Consequently, they are susceptible to ancilla-induced errors and decoherence. In this work, we demonstrate an alternative approach to entangling linear bosonic modes by dynamically engineering a direct nonlinear interaction between them. Unlike previous methods, our protocol preserves the code space of each oscillator throughout the gate while keeping the ancillary nonlinear element only virtually populated. A comparison of these approaches is presented in Table I.

II. EXPERIMENTAL DEVICE AND SYSTEM PARAMETERS

The experimental device used in this work consists of two oscillators in the form of the standard three-dimensional stub cavities, Alice and Bob, machined out of high-purity (4N6) aluminum. They are coaxial $\lambda/4$-resonators with cut-off frequencies around $f_{\text{cut}} \sim 600\,\text{MHz}$ and stub lengths of 16.5 mm and 15.3 mm, which correspond to frequencies of 4.12 GHz and 4.41 GHz, respectively. The external surface layer ($\sim 0.15\,\text{mm}$) of the device has been chemically removed with aluminium etchant type A to reduce fabrication imperfections, achieving Q factors of $\sim 2{\cdot}10^7$ and $\sim 3\ {\cdot}10^7$, respectively. Each oscillator is capacitively coupled to a standard auxiliary transmon that is used for state preparation and readout. The electromagnetic field distributions were simulated using Ansys HFSS, and the Hamiltonian parameters were determined using the energy participation ratio (EPR) approach provided by the python package pyEPR [10]. The Hamiltonian of both cavities and their side chips has the form

$$\begin{aligned} H_i/\hbar = {} & \omega_i \hat{i}^\dagger \hat{i} + \omega_{q_i} \hat{q}_i^\dagger \hat{q}_i + \omega_{rr_i} \hat{r}_i^\dagger \hat{r}_i \\ & - \frac{\alpha_{q_i}}{2} \hat{q}_i^\dagger \hat{q}_i^\dagger \hat{q}_i \hat{q}_i - \frac{K_i}{2} \hat{i}^\dagger \hat{i}^\dagger \hat{i} \hat{i} \\ & - \chi_{i\, q_i}\, \hat{i}^\dagger \hat{i}\, \hat{q}_i^\dagger \hat{q}_i - \chi_{q_i\, rr_i}\, \hat{q}_i^\dagger \hat{q}_i\, \hat{r}_i^\dagger \hat{r}_i, \end{aligned} \tag{2}$$

where $\hat{i} = \hat{a}, \hat{b}$, for both oscillators. All the parameters of the Hamiltonian and their physical meaning are summarized in table II.

To allow for fast initial state preparation and fast tomography, both oscillators couple to the respective auxiliary transmons in the strong dispersive coupling regime.

---

[*] Corresponding author: yvonne.gao@nus.edu.sg

	Dispersive coupling	Beamsplitter	Direct cross-Kerr
Underlying mechanism	Swap - wait - swap Geometric phase	Four-wave mixing Three-wave mixing	Raman transition
Crucial dynamics	$\hat{a}^\dagger\hat{a}\,\hat{c}^\dagger\hat{c}\,\&\,\hat{b}^\dagger\hat{b}\,\hat{c}^\dagger\hat{c}$	$\hat{a}^\dagger\hat{b}+\hat{a}\,\hat{b}^\dagger$	$\hat{a}^\dagger\hat{a}\,\hat{b}^\dagger\hat{b}$
Code-space preserving	No	No	Yes
Ancilla excitation	Yes	No (0−2%)	No (<1%)
Nonlinear process	Yes	No	Yes
References	[4–7]	[1, 8, 9]	This work

TABLE I. **Entangling operations between bosonic modes in cQED.** Comparison of representative two-mode entangling protocols for bosonic oscillators, including their underlying implementation mechanism, the interaction dynamics they rely on, whether they preserve the code space of each oscillator throughout the gate, whether they involve ancillary excitations, and whether they implement an intrinsically nonlinear interaction.

Parameter	Description	Value
$\omega_a/2\pi$	Alice frequency	4.12 GHz
$\omega_{q_a}/2\pi$	Transmon A frequency	5.65 GHz
$\omega_{rr_a}/2\pi$	Resonator A frequency	7.20 GHz
$\alpha_{q_a}/2\pi$	Transmon A anharmonicity	118.3 MHz
$K_a/2\pi$	Alice self-Kerr	1.3 kHz
$\chi_{a\,q_a}/2\pi$	Transmon A - Alice disp. shift	0.77 MHz
$\chi_{q_a\,rr_a}/2\pi$	Transmon A - Resonator A disp. shift	0.6 MHz
$\omega_b/2\pi$	Bob frequency	4.41 GHz
$\omega_{q_b}/2\pi$	Transmon B frequency	5.18 GHz
$\omega_{rr_b}/2\pi$	Resonator B frequency	7.32 GHz
$\alpha_{q_b}/2\pi$	Transmon B anharmonicity	137.9 MHz
$K_b/2\pi$	Bob self-Kerr	2.4 kHz
$\chi_{b\,q_b}/2\pi$	Transmon B - Bob disp. shift	2.16 MHz
$\chi_{q_b\,rr_b}/2\pi$	Transmon B - Resonator B disp. shift	0.45 MHz
$\omega_c/2\pi$	Coupler frequency	5.22 GHz
$\alpha_c/2\pi$	Coupler anharmonicity	196 MHz
$\chi_a/2\pi$	Coupler - Alice disp. shift	0.8 MHz
$\chi_b/2\pi$	Coupler - Bob disp. shift	0.5 MHz
$\omega_{rr_c}/2\pi$	Coupler readout frequency	7.49 GHz
$\chi_{c\,rr_c}/2\pi$	Coupler - readout disp. shift	0.8 MHz

TABLE II. **Hamiltonian parameters of the device.** The parameters of the coupler and its readout resonator correspond to the operating flux-bias point used in Figures 3 and 4 of the main text.

The strength of this dispersive coupling is chosen such that the inherited self-Kerr non-linearity of the cavity modes, $K_{a/b}$, is in an acceptable range of 1-3 kHz. All the states used in the main text are prepared using numerically optimized Gradient Ascent Pulse Engineering (GRAPE) pulses [11] of $2\,\mu$s length. Tomography is performed either by sampling the Wigner function and measuring the parity of the cavity state at different points in phase space using a standard Ramsey technique, or by using the Optimized Reconstruction via Excitation Number Sampling (ORENS) technique [12]. See Section X for further details.

A. Coupler frequency arrangement

Both oscillators are capacitively coupled to a flux-tunable asymmetric Superconducting Quantum Interference Device (SQUID) [13] that acts as the coupler in these experiments. While flux-tunability is not strictly necessary for our protocol, it provides valuable flexibility by allowing us to avoid frequency configurations that may introduce unintended higher-order transitions and to explore different coupler frequencies within a single cooldown.

As the auxiliary transmons do not participate in the main dynamics of the engineered cross-Kerr coupling, we now focus only on the Hamiltonian of the coupler and both oscillators, given by:

$$\begin{aligned} H_c/\hbar = {}& \omega_a\hat{a}^\dagger\hat{a} - \frac{K_a}{2}\hat{a}^\dagger\hat{a}^\dagger\hat{a}\hat{a} \\ &+ \omega_b\hat{b}^\dagger\hat{b} - \frac{K_b}{2}\hat{b}^\dagger\hat{b}^\dagger\hat{b}\hat{b} \\ &+ \omega_c\hat{c}^\dagger\hat{c} - \frac{\alpha_c}{2}\hat{c}^\dagger\hat{c}^\dagger\hat{c}\hat{c} \\ &- \chi_a\hat{a}^\dagger\hat{a}\hat{c}^\dagger\hat{c} - \chi_b\hat{b}^\dagger\hat{b}\hat{c}^\dagger\hat{c} - \frac{K_{ab}}{2}\hat{a}^\dagger\hat{a}\hat{b}^\dagger\hat{b} \\ &+ \omega_{rr_C}\hat{r}^\dagger\hat{r} - \chi_{c\,rr_c}\hat{c}^\dagger\hat{c}\hat{r}^\dagger\hat{r}, \end{aligned} \tag{3}$$

where the experimental parameters associated with each mode used in Figures 3 and 4 in the main text are summarized in table II.

To engineer the Raman-assisted interaction described in the main text, we parametrically drive the coupler at a frequency $\omega_d \approx \omega_c \pm |\omega_b - \omega_a|$. Drive frequencies closer to the coupler frequency translate into a higher effective drive strength ξ. Thus, we design the oscillator frequencies ω_a and ω_b at close spectral proximity $(\omega_b - \omega_a)/2\pi \approx 284$ MHz. Since this detuning is comparable to the coupler anharmonicity, we chose to drive the $a^\dagger b|g\rangle\langle e| + h.c.$ rather than the $ab^\dagger|g\rangle\langle e| + h.c.$ transition, in order to avoid driving the coupler in its straddling regime. Both processes are suitable to generate the desired cross-Kerr dynamics.

Placing the coupler frequency above the cavities allows us to tune the coupling regime of the system from essentially decoupled to almost resonant, as we increase the biasing flux. Coupler frequencies close to the cavities result in an increased always-on natural cross-Kerr, decreasing the on/off ratio of the gate. On the other hand, when the coupler is close to its upper sweet spot, we observe a drop in both oscillators dephasing times. We attribute this to photon shot noise caused by thermal excitations in the coupler and readout mode. Since the strength of the coherent exchange interaction g_1 is proportional to the dispersive shift of both cavities, we must consider the tradeoff between gate speed and on-off ratio. Balancing these considerations, we typically operate at the coupler frequency range of 5.1 ± 0.2 GHz, where the natural cross-Kerr K_{ab} is below the kHz level.

B. Coupler flux delivery

Introducing magnetic flux while preserving the coherence of the cavities in a superconducting package is a challenging task [14]. Here, we follow the approach in [9], where an on-chip superconducting pick-up loop [15] transfers flux from a nearby magnetic coil to the SQUID. The loop works as a transformer, where the magnetic flux originating from a small superconducting coil threads the pick-up loop, inducing a screening current which then biases the SQUID loop.

The magnetic coil generating the initial magnetic field consists of ~ 500 turns of NbTi superconducting wire (SC-54S43-0.178 mm wire from Supercon) coiled around an OFHC copper spool of 1 cm length and 4.5 mm diameter. The terminals of the coil are soldered to the DC lines at the mixing chamber (MXC) stage of the dilution refrigerator. The spool is thermally anchored to the lid of the cavities, which is also made of OFHC copper, and is thermally anchored directly to the MXC stage using a copper braid. We observed a maximum MXC stage temperature of 26 mK when passing a current of 60 mA.

The design of the SQUID chip is shown in Fig. 1(a). The geometry of the pick-up loop was optimized using Ansys Maxwell to maximize the flux delivered to the device while minimizing the current required in the coil. Further, we also verify that the design preserves the Q factors of both cavities and does not introduce undesired modes in the system. The final configuration consists of 4 distinct sections: (i) a circular region optimized to only collect the magnetic flux piercing the chip; (ii) a constriction to avoid collecting the returning field; (iii) the long sections of the transformer, spaced out to reduce the loop self-inductance; and (iv) a narrowed section near the SQUID loop to avoid interfering with the junction fabrication during the lithography step.

The length of the pick-up loop is chosen to balance two effects. First, it needs to be long enough such that the copper coil and spool are spatially separated from the cavities, both for ease of assembly and to prevent the electric field of the cavity TEM modes from leaking out through the aperture of the coil, which would degrade the cavity lifetimes. Second, the parasitic resonant modes of the transformer, currently at 6.2 GHz and 8.1 GHz, decrease in frequency as its length increases, potentially interfering with other modes in the system. Finally, the loop formed by both junctions of the SQUID, which has an area of 800 μm^2, is placed at a distance of $\sim 50\,\mu$m from the flux transformer, see inset in Fig. 1(a). The transformer has a width of 10 μm at its narrowest point, ensuring that it remains superconducting even with large induced currents.

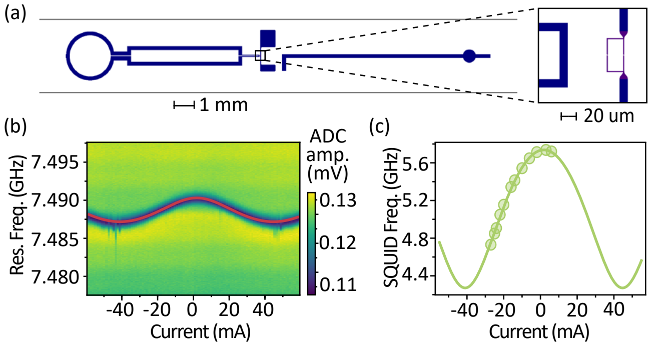

FIG. 1. **SQUID coupler tunability.** (a) Layout of the SQUID coupler chip, with the pick-up loop transformer (left), the SQUID (middle), and its readout resonator (right). (b) Readout frequency as a function of the current applied to bias the SQUID. Solid line shows the fit used to extract the current-flux relation. (c) SQUID frequency as a function of current. Solid line shows the fitted frequency.

To mitigate the dephasing of the flux-tunable coupler, we compromise on the frequency tunability range by making both junctions asymmetric. Given the junctions Josephson energies E_{J_1} and E_{J_2}, the effective Josephson energy of the SQUID is

$$E_J\left(\Phi_{\mathrm{ext}}\right) = E_{J,\Sigma}\cos\left(\pi\frac{\Phi_{\mathrm{ext}}}{\Phi_0}\right)\sqrt{1+\frac{E_{J,\Delta}}{E_{J,\Sigma}}\tan^2\left(\pi\frac{\Phi_{\mathrm{ext}}}{\Phi_0}\right)}, \tag{4}$$

where $E_{J,\Sigma} = E_{J_1} + E_{J_2}$ and $E_{J,\Delta} = E_{J_1} - E_{J_2}$. The measured frequency of the asymmetric SQUID as a function of the applied current on the coil is shown in Fig. 1(c). Due to the increased attenuation of the drive line at lower frequencies and the reduced dispersive coupling to its readout resonator, we cannot measure the coupler frequency for values below 4.7 GHz. However, by fitting the SQUID readout resonator frequency as a function of the applied current through the coil, we extract the relation between the current and the flux in the SQUID loop, with 42.89 mA corresponding to half a flux quantum, see Fig. 1(b). Finally, to extract the whole tunability range of the SQUID and the individual Josephson energies of both junctions, we consider the coupler-resonator Hamiltonian,

$$
\begin{aligned}
\hat{H}_c/\hbar = {} & 4E_C\hat{n}^2 - E_J\left(\Phi_{\text{ext}}\right)\cos\hat{\phi} + \omega_{rr_c}\hat{r}^\dagger\hat{r} \\
& + g(\hat{c} - \hat{c}^\dagger)(\hat{r} - \hat{r}^\dagger)
\end{aligned}
\tag{5}
$$

where E_C is the charging energy of the SQUID, $\hat{n}$ and $\hat{\phi}$ are the charge and phase operators, respectively, ω_{rr_c} is the readout resonator frequency and g is the SQUID-resonator coupling strength. Fitting the frequencies of these two modes as a function of flux, constrained by the measured anharmonicity and the current-flux relation, we obtain the parameters shown in Table III. This allows us to determine the full SQUID tunability, spanning 4.3 GHz to 5.7 GHz, see solid line in Fig. 1(c), as well as the SQUID asymmetry, $E_{J_1}/E_{J_2} \approx 3.7$, which closely matches the designed asymmetry during fabrication. These are the parameters used in Floquet simulation in Section VI.

Parameter	Value	Description
E_{J_1}	$19.0 \pm 0.1\,\text{GHz}$	Josephson energy
E_{J_2}	$5.2 \pm 0.1\,\text{GHz}$	Josephson energy
E_C	$181.4 \pm 0.1\,\text{MHz}$	Charging energy
g	$0.10 \pm 0.02\,\text{GHz}$	Coupling constant
$\omega_{r_0}/2\pi$	$7.484 \pm 0.002\,\text{GHz}$	Resonator frequency

TABLE III. **Hamiltonian parameters of the SQUID coupler**. Parameters were fitted to reproduce the resonator and coupler flux tunability as a function of the applied current, with the additional constraint of reproducing the experimental anharmonicity of the coupler.

C. System coherences

Each transmon and its readout resonator are addressed through a shared coupling port, whose position is optimized in Ansys HFSS to suppress radiative decay of the transmons into the measurement lines, thereby achieving intrinsic Purcell filtering [16]. This approach preserves transmon coherence without the need for on-chip Purcell filters, reducing chip size, mitigating susceptibility to vibrations, and enabling a more compact device. The coherence times of both side transmons and the SQUID coupler are shown in Table IV. The coupler coherence times depend strongly on the it frequency configuration, with T_1 times ranging from $50\,\mu$s at the lowest-measured frequency to $10\,\mu$s at its upper sweet spot, where the coupler hybridizes more strongly with its readout resonator. On the other hand, the dephasing times range from up to $20\,\mu$s at the flux sweet spot to $1\,\mu$s at the most sensitive location. The coupler coherence times in the configuration used to obtain the final data for Figures 3 and 4 in the main text are shown in Table IV. Importantly, although the coupler is operated near its most flux-sensitive point, its intrinsic decoherence does not significantly affect the quality of the gate, since the Raman-assisted interaction is mediated through virtual processes that avoid populating its excited states.

Param.	Value (μs)	Param.	Value (μs)	Param.	Value (μs)
T_{1q_a}	25	T_{1q_b}	50-60	T_{1c}	50
$T^*_{2q_a}$	18	$T^*_{2q_b}$	40-50	T^*_{2c}	2.5
$T^E_{2q_a}$	22	$T^E_{2q_b}$	50	T^E_{2c}	10
T_{1a}	700-900	T_{1b}	800-1000		
T_{2a}	500-600	T_{2b}	700-900		

TABLE IV. **Device coherence times.** Coherence times for the coupler are quoted at the flux point used in the main text. Transmon T_2^* were measured with a standard Ramsey technique, while T_2^E times were obtained via a Hahn-echo experiment. Oscillator T_1 times were measured by fitting the decay of both coherent states and single photons. Oscillator T_2 times were measured through a Ramsey experiment using SNAP pulses.

In contrast, cavity decoherence directly limits the achievable quality of the engineered cross-Kerr interaction. The cavities are designed to be internal Q limited, with coupling ports designed to have coupling quality factors on the order of $\sim 10^9$. Good care was also taken to minimize the cavity photon field from leaking into the coil used to bias the SQUID. We measured cavity T_1 times both by initializing a coherent state and measuring its characteristic decay, and by loading a photon in the cavity using SNAP pulses and fitting an exponential decay. Both methods give consistent results with average single-photon lifetimes of $800\,\mu$s and $900\,\mu$s for Alice and Bob, respectively. The cavity T_2 times were measured using a standard Ramsey technique, where they are initialized in $(|0\rangle + |1\rangle)/\sqrt{2}$ using SNAP pulses and, after a variable wait time, they are displaced by $D(\alpha = -0.8e^{i\omega t})$. This displacement maximizes the overlap with vacuum for $|0\rangle + |1\rangle)$ and minimizes it for $|0\rangle - |1\rangle)$, and ω is a fixed virtual detuning that facilitates extracting the T_2 dephasing time by fitting the data to an exponentially decaying sinusoid.

D. Auxiliary transmon readout

Both the auxiliary transmons and the SQUID coupler are measured using Cavity Level Excitation and Reset (CLEAR) readout pulses [17] applied to their respective readout resonators. In contrast to a conventional constant-amplitude readout pulse, the CLEAR protocol employs a sequence of five piecewise-constant segments: two ring-up segments, one steady-state segment, and two ring-down segments. This pulse shaping offers high-fidelity state discrimination while significantly reducing the residual photon population after the measurement pulse. Compared to standard square pulses, CLEAR pulses achieve faster resonator reset and therefore support reduced measurement-induced backaction and shorter experimental cycle times. The amplitudes and duration of the individual CLEAR drive segments are optimized running an algorithm that minimizes the population of the resonator after the CLEAR pulse given the decay time of the resonator $1/\kappa$ and the dispersive shift to the transmon χ.

In addition to CLEAR pulse shaping, our readout is typically optimized to differentiate between states $|g\rangle$ and $|f\rangle$, which improves the separation of the histogram counts for the same integration time compared to discriminating between $|g\rangle$ and $|e\rangle$, as $\chi_{gf} \approx 2\,\chi_{ge}$. After each experimental sequence we map the $|e\rangle$ population to the $|f\rangle$ state using a $\sigma = 24$ ns gaussian pulse. The resulting confusion matrices for both auxiliary transmons used in the tomography and state preparation steps are:

$$M_{\mathrm{RO},A} = \begin{pmatrix} P(g|g) = 97.5\% & P(e|g) = 2.5\% \\ P(g|e) = 2.8\% & P(e|e) = 97.2\% \end{pmatrix}$$

$$M_{\mathrm{RO},B} = \begin{pmatrix} P(g|g) = 95.9\% & P(e|g) = 4.1\% \\ P(g|e) = 6.4\% & P(e|e) = 93.6\% \end{pmatrix}$$

E. Chip fabrication

The ancillary transmon chips and the SQUID coupler chip are fabricated in aluminum on a sapphire substrate. A HEMEX sapphire wafer is cleaned in a 2:1 piranha solution for 20 minutes and rinsed in de-ionized (DI) water for another 20 minutes. It is then quickly rinsed in methanol and blown-dry with nitrogen. The wafer is coated with 700 nm of MMA and 200 nm of PMMA resist by spinning it for 100 seconds at 2000 rpm and baking it for 5 minutes at 200°C. A $\sim$10-nm discharge gold layer is sputtered in a Cressington sputterer at 30 mA for 30 seconds. The design is patterned using a Raith electron-beam lithography machine. The gold layer is removed in a KI solution and rinsed in DI water, prior to developing the resist in a 3:1 mixture of DI water and isopropanol at 6°C for 2 minutes. The wafer is then loaded into an Angstrom Engineering double-angle evaporator and pumped to 10^{-8} mbar, after which the resist is ion milled in a mixture of 85% O2 and 15% Argon, at 400 V for 15 s at -20° and 15 s at 20° to clean any resist residues. We deposit two aluminum layers of 20 nm and 30 nm thickness at -20° and +20°, respectively, separated by an oxidation step with a mixture of 85% O2 and 15% Argon at 20 mbar for 20 minutes. Before unloading the wafer, a final capping layer with the same oxidation conditions is performed to grow a more controlled surface oxide. The remaining resist is then lifted off in NMP at 90°C for 3 hours, then rinsed in acetone and methanol. Finally, a protective layer of AZ1512 photoresist is spun at 2000 rpm and baked at 80°C for 1 minute, and the individual chips are diced on an Accretech machine with a resin blade at 15000 rpm. The chips are finally cleaned in NMP, acetone and methanol, blown-dry and inserted into the waveguides using an aluminum clamps with indium wire to improve thermalization.

III. SNAP PULSES

We use Selective Number-dependent Arbitrary Phase (SNAP) pulses [18] to prepare the oscillators in Fock state $|1\rangle$ and the superposition state $(|0\rangle + |1\rangle)/\sqrt{2}$, to measure the single-photon decay rate and dephasing rate of the oscillators, respectively. The SNAP pulses consist of an initial displacement of the oscillators by α_1, a phase, θ_n, imparted on the transmon state conditioned on the oscillator having n photons, and a final displacement of the oscillators by α_2. All parameters used can be found in Table V. Oscillator displacements were performed using 100 ns constant pulses with 32 ns cosine ramps at the resonant frequency of the oscillators. The conditional $\theta_n = \pi$ phases were realized by driving the transmons through a full 2π-rotation around the Bloch sphere, using gaussian pulses of $\sigma_a = 960$ ns and $\sigma_b = 400$ ns, applied at the resonance frequency of each transmon conditioned on the oscillator containing n photons.

Transition	α_1	θ_0	θ_1	α_2
$\|0\rangle \to \|1\rangle$	1.14	π	0	-0.58
$\|0\rangle \to \frac{\|0\rangle+\|1\rangle}{\sqrt{2}}$	0.56	π	0	-0.24
$\frac{\|0\rangle-\|2\rangle}{\sqrt{2}} \to \|0\rangle$	-0.35	π	π	1.04

TABLE V. **Different SNAP parameters used in the main text**. First row corresponds to the creation of a single photon in the oscillators, used to measure their decay time. Second row corresponds to the creation of a superposition state in the oscillators, used to measure their dephasing time in a Ramsey experiment. Last row corresponds to the pulse used in the direct cavity-cavity parity mapping, shown in Fig. 4(c) of the main text.

In the direct parity mapping protocol described in Fig. 4(c) in the main text, we use a SNAP pulse to map the even and odd superposition states of Bob to

$$\frac{|0\rangle - |2\rangle}{\sqrt{2}} \to |0\rangle$$
$$\frac{|0\rangle + |2\rangle}{\sqrt{2}} \to 0.9\,|1\rangle + 0.4\,|2\rangle + \dots$$

A selective π-pulse to Bob's auxiliary transmon conditioned on Bob being in vacuum hence maps $(|0\rangle - |2\rangle)/\sqrt{2}$ and $(|0\rangle + |2\rangle)/\sqrt{2}$ to $|e\rangle$ and $|g\rangle$, respectively. The displacement amplitudes used to implement this transformation were obtained via a Nelder–Mead optimization that maximizes the probability of the final state being either in vacuum or not vacuum. The parameters of the SNAP sequence are listed in the last row of Table V.

IV. WIRING DIAGRAM

The microwave signals to drive the oscillators, auxiliary transmons, coupler, and readout resonators of our device are directly generated from digital waveforms via Direct-Digital-Synthesis (DDS) by a Digital-to-Analog Converter (DAC) of a Quantum Machines OPX1000 Field-Programmable Gate Array (FPGA). We also place band-pass filters at the output line of the OPX1000, which increased the transmon coherence times significantly. Additionally, a RF switch (HMC-C058) placed on the coupler line provides an additional layer of noise protection. Importantly, most of the coupler drive line attenuation at the MXC stage is replaced by a Mini-circuits 5500+ low-pass filter and Mini-circuits 5500+ high-pass filter, which provide -30 dB attenuation at 5.0 GHz but only -10 dB at 5.5 GHz to be able to drive the coupler strongly, while still providing at least 60 dB attenuation at the oscillator frequencies and 35 dB at the resonator frequency.

Readout is performed using the standard reflection technique. The reflected signal, amplified by a High-Electron mobility transistor (HEMT - LNF-LNC48C) as well as a ZVA-183-S+ room temperature amplifier, is directly sampled by a high-speed Analog-to-Digital Converter (ADC) of the OPX1000 FPGA. The digital signal is then downconverted and separated into I and Q components. The DC current sent to the coil is generated in a YOKOGAWA GS200 DC voltage/current source. To shield our device from external electromagnetic noise it is placed in a standard Cryoperm shield and anchored directly at the MXC plate. The schematic of the wiring can be seen below in Fig. 2

V. ANALYTICAL DERIVATION OF THE ENGINEERED DYNAMICS

The Hamiltonian of the two oscillators and the SQUID coupler is

$$\begin{aligned} H(t)/\hbar &= \tilde{\omega}_a\ \hat{a}^\dagger\hat{a} + \tilde{\omega}_b\ \hat{b}^\dagger\hat{b} \\ &+ \tilde{\omega}_c\ \hat{c}^\dagger\hat{c} - E_J(\Phi_{\mathrm{ext}})\left(\cos\hat{\phi} + \frac{1}{2!}\hat{\phi}^2\right) \\ &+ \epsilon(t)(\hat{c}+\hat{c}^\dagger), \end{aligned} \tag{6}$$

where $\tilde{\omega}_a$ and $\tilde{\omega}_b$ are the bare (uncoupled) oscillator frequencies, $\tilde{\omega}_c = \sqrt{8E_J(\Phi_{\mathrm{ext}})E_C}$ is the Josephson plasma frequency of the coupler, $\epsilon(t)$ is the time-dependent drive on the coupler and

$$\hat{\phi} = \phi_c\ (\hat{c}+\hat{c}^\dagger) + \phi_a\ (\hat{a}+\hat{a}^\dagger) + \phi_b\ (\hat{b}+\hat{b}^\dagger)$$

is the quantized phase of the coupler, that considers the participation of both oscillator fields in the junction.

The zero-point fluctuations of the coupler in the junction is defined as $\phi_c = \left(2E_C/E_J(\Phi_{\mathrm{ext}})\right)^{1/4}$, while the zero-point fluctuations of the oscillators in the junction are approximated by $\phi_a \approx \frac{g_a}{\Delta_a}\phi_c$ and $\phi_b \approx \frac{g_b}{\Delta_b}\phi_c$. Considering a drive of the form $\epsilon(t) = 2\varepsilon\cos(\omega_d\ t + \phi_d)$, we can go to the displaced frame $\hat{c} \rightarrow \hat{c} + \xi e^{-i\omega_d t}$ with $\xi = \varepsilon/(\omega_d - \tilde{\omega}_c)$, where the drive term is absorbed into the phase variable through

$$\hat{\phi} = \phi_c\ (\hat{c}+\hat{c}^\dagger+\xi e^{-i\omega_d t}+\xi^* e^{i\omega_d t}) + \phi_a\ (\hat{a}+\hat{a}^\dagger) + \phi_b\ (\hat{b}+\hat{b}^\dagger).$$

By expanding the cosine in Eq. 6 up to fourth order in $\hat{\phi}$, and keeping all the terms that are naturally on-resonance, we obtain the full Hamiltonian of the system

$$\begin{aligned} \hat{H}_0 &= \omega_a\ \hat{a}^\dagger\hat{a} - \frac{K_a}{2}\hat{a}^\dagger\hat{a}^\dagger\hat{a}\hat{a} \\ &+ \omega_b\ \hat{b}^\dagger\hat{b} - \frac{K_b}{2}\hat{b}^\dagger\hat{b}^\dagger\hat{b}\hat{b} \\ &+ \omega_c\ \hat{c}^\dagger\hat{c} - \frac{\alpha}{2}\hat{c}^\dagger\hat{c}^\dagger\hat{c}\hat{c} \\ &- \chi_a\ \hat{a}^\dagger\hat{a}\ \hat{c}^\dagger\hat{c} - \chi_b\ \hat{b}^\dagger\hat{b}\ \hat{c}^\dagger\hat{c} - K_{ab}\ \hat{a}^\dagger\hat{a}\ \hat{b}^\dagger\hat{b}, \end{aligned} \tag{7}$$

where K_a and K_b are the cavity self-Kerrs inherited from the coupler non-linearity, α is the coupler non-linearity (anharmonicity), χ_a and χ_b are the coupler-cavity cross-Kerrs (dispersive shifts), and K_{ab} is the residual always-on cavity-cavity cross-Kerr. All these terms produce Lamb and Stark shifts on the bare mode frequencies, which get renormalized as

$$\begin{aligned} \omega_a &= \tilde{\omega}_a - \frac{\chi_a}{2} - K_a - \frac{K_{ab}}{2} - \chi_a|\xi|^2 \\ \omega_b &= \tilde{\omega}_b - \frac{\chi_b}{2} - K_b - \frac{K_{ab}}{2} - \chi_b|\xi|^2 \\ \omega_c &= \tilde{\omega}_c - \alpha - \frac{\chi_a}{2} - \frac{\chi_b}{2} - 2\alpha|\xi|^2, \end{aligned}$$

where the last term of each equation corresponds to the drive-induced AC Stark shifts.

By moving to a frame rotating at $\omega_a\hat{a}^\dagger\hat{a} + \omega_b\hat{b}^\dagger\hat{b} + \omega_c\hat{c}^\dagger\hat{c}$ and restricting the coupler basis to the first two energy levels $\{|g\rangle, |e\rangle\}$, we find that the drive frequency

$$\omega_d = \omega_c + \omega_b - \chi_b - \omega_a + \Delta, \tag{8}$$

where Δ is the detuning to the resonance condition, successfully actives the dynamics given by

$$H_d(t) = g_1 e^{it\Delta}\hat{a}^\dagger\hat{b}|g\rangle\langle e| + g_1^* e^{-it\Delta}\hat{a}\hat{b}^\dagger\hat{}|e\rangle\langle g|, \tag{9}$$

which corresponds to Eq. 2 in the main text. The strength of this interaction is given by

$$g_1 = E_J\phi_c^2\phi_a\phi_b\xi^* \approx \sqrt{\chi_a\chi_b}\xi^*. \tag{10}$$

We measured the coherent exchange rate, g_1, for different effective drive strengths $|\xi|$ obtained from the coupler AC Stark shift, as shown in Fig. 3(a). The slope of the interaction rate is in very good agreement with the prediction of Eq. 10 using the experimentally measured dispersive shifts (solid line).

To find an analytical equation for the engineered cross-Kerr, we diagonalize the interaction term of Eq. 9, by applying a time-independent Schrieffer-Wolff transformation (SWT) [19, 20] in the rotating frame of the drive.

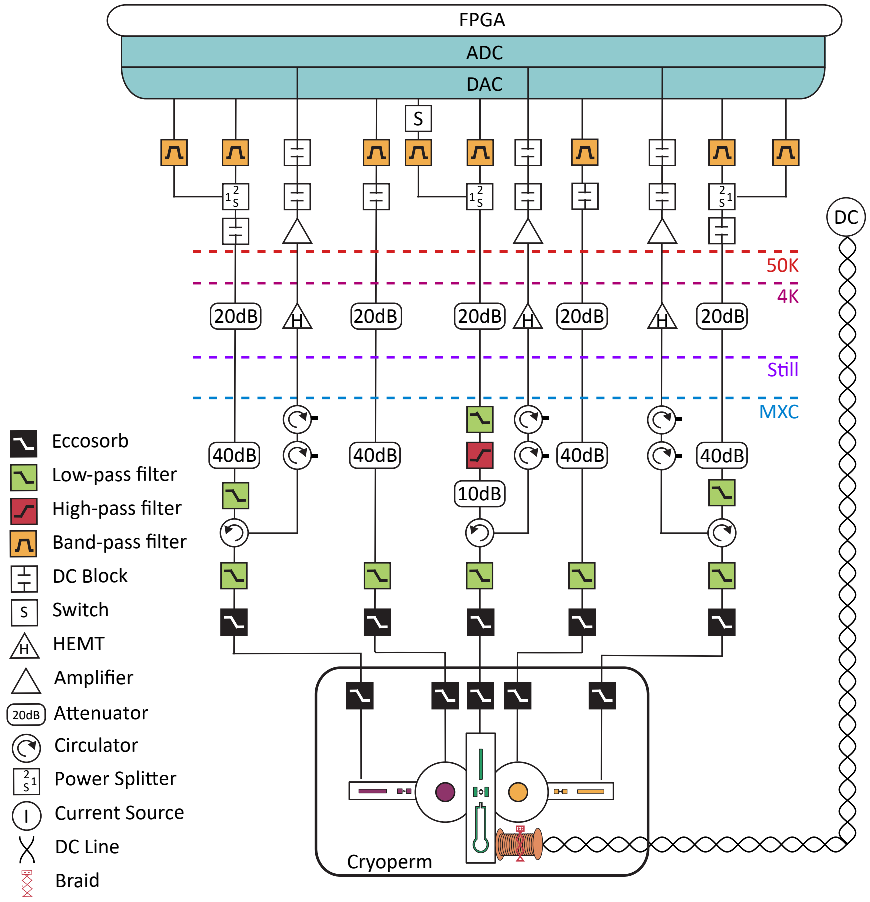

FIG. 2. **Experimental wiring setup**. Schematic of the RF components and connections at room temperature and inside the Bluefors dilution refrigerator. Most of the attenuation for the coupler drive line is replaced by a low-pass filter in series with a high-pass filter that act as a 5-7 GHz band-pass filter.

This is equivalent to applying a time-dependent SWT in the rotating frame of the coupler, which is also equivalent to performing a second-order RWA [21].

In the rotating frame of the drive, the diagonal part of the Hamiltonian reads

$$\begin{aligned}\hat{H}_0'/\hbar &= \omega_a\ \hat{a}^\dagger\hat{a} - \frac{K_a}{2}\hat{a}^\dagger\hat{a}^\dagger\hat{a}\hat{a}\\ &+ \omega_b\ \hat{b}^\dagger\hat{b} - \frac{K_b}{2}\hat{b}^\dagger\hat{b}^\dagger\hat{b}\hat{b}\\ &+ (\omega_c - \omega_d)\ \hat{c}^\dagger\hat{c} - \frac{\alpha}{2}\hat{c}^\dagger\hat{c}^\dagger\hat{c}\hat{c}\\ &- \chi_a\ \hat{a}^\dagger\hat{a}\ \hat{c}^\dagger\hat{c} - \chi_b\ \hat{b}^\dagger\hat{b}\ \hat{c}^\dagger\hat{c} - K_{ab}\ \hat{a}^\dagger\hat{a}\ \hat{b}^\dagger\hat{b},\end{aligned} \tag{11}$$

while the perturbation is given by

$$\hat{V}/\hbar = g_1\hat{a}^\dagger\hat{b}|g\rangle\langle e| + g_1^*\hat{a}\hat{b}^\dagger|e\rangle\langle g|. \tag{12}$$

The SWT consists of finding an anti-Hermitian generator $\hat{S} = -\hat{S}^\dagger$ that satisfies

$$\hat{V} = \left[\hat{H}_0', \hat{S}\right], \tag{13}$$

which transforms the Hamiltonian to

$$\hat{H}' = e^{\hat{S}}\left(\hat{H}_0' + \hat{V}\right)e^{-\hat{S}} = H_0' + \hat{V} + \left[\hat{S}, \hat{H}_0'\right] + \left[\hat{S}, \hat{V}\right] + \frac{1}{2!}\left[\hat{S}, \left[\hat{S}, \hat{H}_0'\right]\right] + \cdots \approx \hat{H}_0' + \frac{1}{2}\left[\hat{S}, \hat{V}\right] \tag{14}$$

By choosing a generator of the form

$$\hat{S} = \beta\left(\hat{a}^\dagger\hat{b}|g\rangle\langle e| - \hat{a}\hat{b}^\dagger|e\rangle\langle g|\right), \tag{15}$$

we can solve Eq. 13 to find the expression for β. Since the commutator

$$\left[\hat{H}_0', \hat{a}^\dagger\hat{b}|g\rangle\langle e|\right] = \left(\omega_d - \omega_c + \omega_a - \omega_b - K_a(\hat{n}_a - 1) + K_b\hat{n}_b + \chi_a(\hat{n}_a - 1) + \chi_b(\hat{n}_b + 1) - K_{ab}(\hat{n}_b - \hat{n}_a + 1)\right)\ \hat{a}^\dagger\hat{b}|g\rangle\langle e|$$

depends on the number of photons in both cavities, we expand the transformation generator as

$$\hat{S} = \sum_{n_a, n_b} \beta_{n_a n_b}\left(|n_a, n_b, g\rangle\langle n_a - 1, n_b + 1, e| - |n_a - 1, n_b + 1, e\rangle\langle n_a, n_b, g|\right), \tag{16}$$

where

$$\beta_{n_a n_b} = \frac{\sqrt{n_a(n_b+1)}\ g_1}{\omega_d - \omega_c + \omega_a - \omega_b - K_a(n_a - 1) + K_b n_b + \chi_a(n_a - 1) + \chi_b(n_b + 1) - K_{ab}(n_b - n_a + 1)}. \tag{17}$$

Finally, solving Eq. 14 leads to

$$\hat{H}' = \hat{H}_0' + \sum_{n_a, n_b} \beta_{n_a n_b}\sqrt{n_a(n_b+1)}\, g_1\left(|n_a, n_b, g\rangle\langle n_a, n_b, g| - |n_a - 1, n_b + 1, e\rangle\langle n_a - 1, n_b + 1, e|\right). \tag{18}$$

With these analytical results, we can see that when we parametrically drive this dynamics, the energies of the states $|1\rangle|0\rangle|g\rangle$ and $|1\rangle|1\rangle|g\rangle$ are shifted by

$$\Delta E_{10} = \frac{g_1^2}{\Delta}, \tag{19}$$

$$\Delta E_{11} = \frac{2\,g_1^2}{\Delta + \chi_b + K_b - K_{ab}} \approx \frac{2\,g_1^2}{\Delta + \chi_b}. \tag{20}$$

Eq. 19 shows that in presence of the parametric drive, despite Bob being in vacuum, the state $|+\rangle|0\rangle|g\rangle$ will acquire a phase at a rate ΔE_{10}. Hence, it is important to track the frame of the oscillators (on software) for any subsequent operation. For instance, the frame is tracked as a virtual detuning applied to the displacement pulses used for Wigner tomography. The phase acquired by the state $|+\rangle|0\rangle|g\rangle$ was also experimentally measured as a function of the detuning Δ, see green markers in Fig. 3(b), which agrees with the prediction of Eq. 19 (solid green line). Finally, the total cross-Kerr induced in the 0/1 photon subspace is thus

$$\begin{aligned} g_{ab} &= E_{11} - E_{10} - E_{01} + E_{00} \\ &= -K_{ab} + \Delta E_{11} - \Delta E_{10} \\ &= -K_{ab} + \frac{g_1^2}{\Delta}\frac{\Delta - (\chi_b + K_b - K_{ab})}{\Delta + (\chi_b + K_b - K_{ab})}, \end{aligned} \tag{21}$$

where E_{ij} refers to the energy of the state $|i\rangle|j\rangle|g\rangle$ and K_{ab} is the always-on residual cross-Kerr. We measured the cross-Kerr g_{ab} by tracking the oscillations of $|+\rangle|1\rangle|g\rangle$ in the frame where $|+\rangle|0\rangle|g\rangle$ does not oscillate, at different drive powers for a fixed detuning of $\Delta = -5$ MHz, see blue markers in Fig. 3(c). Notice that at different drive powers, the AC Stark shifted frequency of the coupler is different, which subsequently modifies the coherent exchange frequency ω_e. The solid blue line corresponds to the analytical behavior of Eq. 21, which underpredicts the measured interaction strength. We believe this discrepancy arises because the analytical model does not account for the enhanced hybridization between the modes under the strong drive, which renormalizes the dispersive shifts and the detunings. A full Floquet simulation (orange triangles) reproduces the experimental observations in very good agreement. For more details on the Floquet simulation, see Section VI.

The analytical model still qualitatively captures the asymmetric behavior of the engineered cross-Kerr between positive and negative detunings, see purple markers in Fig. 3(d). The green (yellow) dotted line corresponds to $-\Delta E_{10}$ $(+\Delta E_{11})$, which diverges at $\Delta = 0$ $(\Delta \approx -\chi_b)$. The solid purple line is the sum of both, which qualitatively reproduces the main features observed in the experimental data. For a quantitative agreement, it is desirable to model the system using Floquet theory, as shown in Fig. 2(b) of the main text.

VI. FLOQUET SIMULATIONS

In the presence of strong parametric drives, the system is governed by a time-periodic Hamiltonian in which the drive-induced shifts of the energy level give rise to unintended multi-photon transitions between the oscillators and higher levels of the coupler, beyond the coherent exchange interaction considered analytically. This scenario is appropriately addressed within the Floquet formalism [8, 22, 23]. We perform Floquet simulations using the QuTiP python package [24] to analyze the effects of the parametric drive on the subsystem composed

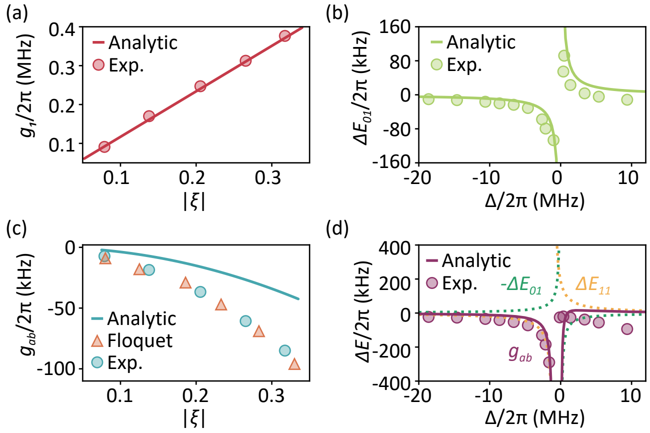

FIG. 3. **Analytical predictions based on the SWT.** (a) Coherent exchange rate, g_1, as a function of the effective drive strength, $|\xi|$ (red circles). The analytical prediction of Eq. 10 is in close agreement (solid line). (b) Phase rate acquired by the state $|+\rangle|0\rangle|g\rangle$ (green circles) as a function of the detuning Δ to the coherent exchange condition. Analytics correspond to the prediction from Eq. 19 (solid line). (c) Cross-Kerr strength (blue circles) measured at a fixed $\Delta = -5\,$MHz, as a function of drive power. Analytical prediction (solid line) based on Eq. 21 under-predicts the strength by a factor of 2. However, a full Floquet simulation (orange triangles) is in very good agreement with the data. (d) Cross-Kerr strength (purple circles) as a function of the detuning $\Delta = -5\,$MHz. The purple line corresponds to the sum of the yellow and green lines, which exhibit discontinuities at $\Delta = 0$ and $\Delta \approx -\chi_b$, respectively.

of the SQUID coupled to both cavities, ignoring the side auxiliary transmons. The Hilbert spaces are truncated to dimension $D = 12$ for the SQUID and $D = 4$ for the cavities. The full Hamiltonian considered in the simulation is

$$\begin{aligned} H/\hbar &= 4E_C\hat{n}^2 - E_J\left(\Phi_{\mathrm{ext}}\right)\cos\hat{\phi} \\ &+ \omega_a\hat{a}^\dagger\hat{a} + \omega_b\hat{b}^\dagger\hat{b} \\ &- g_{ac}(\hat{a}^\dagger - \hat{a})(\hat{c}^\dagger - \hat{c}) \\ &- g_{bc}(\hat{b}^\dagger - \hat{b})(\hat{c}^\dagger - \hat{c}) \\ &+ \epsilon\cos\left(\omega_d t\right)(\hat{c}^\dagger + \hat{c}), \end{aligned} \tag{22}$$

where the coupler charging and Josephson energies are those obtained by fitting the experimental data, and summarized in Table III; the oscillator frequencies are those measured experimentally through two-tone spectroscopy; the capacitive coupling strength between the oscillators and the SQUID are taken as $g_a/2\pi = 43.07\,$MHz and $g_b/2\pi = 22.56\,$MHz, to reproduce the experimentally measured dispersive shifts; and the drive amplitude ϵ is calibrated to reproduce the measured coupler AC Stark shift.

The 200 ns $\cos^2$ ramps for the parametric drive ensure that the undriven eigenstates adiabatically map onto Floquet modes. For a given drive frequency ω_d and amplitude ϵ, we label each Floquet mode by the eigenstate of the undriven system with which it has the highest overlap. The Floquet quasienergies, defined modulo ω_d, are unfolded into real energies by adding or subtracting integer multiples of ω_d such that they lie within $\pm 20\%$ of the energy of their corresponding undriven eigenstate. The cross-Kerr between the cavities is then obtained as the difference in the transition energy of Alice when Bob is in vacuum or in $|1\rangle$. The cross-Kerr strength obtained from Floquet simulations, as a function of drive amplitude and detuning, is in very good agreement with the experimental data, see triangle markers in Fig. 3(c) and solid line in Fig. 2 of the main text, respectively. In the latter case, deviations at large detunings arise from the proximity of a single-cavity transition between $|0\rangle|1\rangle|g\rangle$ and $|0\rangle|0\rangle|f\rangle$, which significantly enhances Bob's self-Kerr and its cross-Kerr coupling to Alice [25]. We attribute this discrepancy to uncertainties in some Hamiltonian parameters, particularly g_b, which were calibrated solely from independently measured experimental observables without additional fitting.

VII. CALIBRATION OF COHERENT EXCHANGE RATE

To confirm the location of the coherent exchange interaction ω_e, and calibrate the exchange strength, g_1, we measure the population of Alice, Bob, and the coupler, as a function of the drive duration and frequency (ω_d), as shown in Fig. 4.

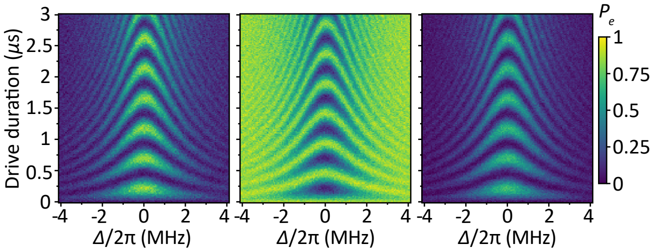

FIG. 4. **Coherent exchange of excitations between the coupler and the oscillators**. Left: probability of finding Alice in vacuum. Middle: probability of finding Bob in vacuum. Right: probability of the coupler being excited. System was initialized in $|1\rangle|0\rangle|g\rangle$. $\Delta = 0$ corresponds to $\omega_d = \omega_e$, where exchange occurs in $\sim 200\,$ns. The coupler configuration corresponds to Fig. 2(a) in the main text.

The experiment starts by preparing the state $|1\rangle|0\rangle|g\rangle$ through a numerical GRAPE pulse on the side transmon q_a and Alice. A microwave drive is applied to the coupler with variable duration t and frequency ω_d. After the drive pulse, the system populations are measured. The coupler population is obtained via direct dispersive readout of its readout resonator, while the cavity populations are measured using selective π-pulses on their auxiliary transmon, conditioned on each cavity being in vacuum. The cavity data is normalized by the selective π-pulse contrast, and the coupler population is corrected by the readout contrast.

The data exhibits the characteristic chevron pattern expected for a coherent exchange between $|1\rangle|0\rangle|g\rangle$ and $|0\rangle|1\rangle|e\rangle$. At zero detuning, the populations oscillate sinusoidally in time with a frequency $2g_1$ while detuning from the exchange condition leads to faster oscillations with reduced contrast. A line cut yields the oscillations shown in Fig. 2(a) of the main text, from which we can extract $g_1/2\pi = 1.024 \pm 0.004\,\mathrm{MHz}$.

VIII. EXTENSION TO MULTI-PHOTON CODES

Due to the bosonic enhancement factor in the coupling strength, we can increase the gate speed by populating higher Fock states in the control cavity. We demonstrate this by measuring the phase acquired by the $|+\rangle$ state during gate operation when the control cavity is in $|2\rangle$ and comparing it to the equivalent case when the control is in $|1\rangle$. The phases acquired as a function of time for both cases are shown in Fig. 5(a). Representative Wigner tomography snapshots at selected times for initial state $|+\rangle|2\rangle$ are shown in Fig. 5(b). We observe that the CZ gate time is reduced from $\sim 5.5\,\mu$s to $\sim 2.5\,\mu$s.

It is important to note that the gate time is reduced by more than a factor of two as the resonant condition is shifted by $n\chi_b$ when the control cavity is prepared in $|n\rangle$ in accordance with Eq. 8. Since the gate detuning is defined with respect to the single-photon case and is kept constant for both experiments, the effective detuning becomes smaller for the two-photon state, leading to a gate rate that exceeds the simple factor-of-two scaling expected from bosonic enhancement alone. This entails that when the engineered dynamics is applied to multi-photon states, we must account for this difference in effective detuning appropriately.

IX. PAULI BARS FOR OTHER STATES

Here, we present the reconstructed two-cavity density matrices in the Pauli decomposition, after applying 0 to 6 CZ gates to the initial states $|+\rangle|0\rangle$, $|1\rangle|+\rangle$ and $|+\rangle|+\rangle$, see Fig. 6. For the initial state $|+\rangle|+\rangle$ we further distinguish between the case of drive on and off (i.e. $g_{CZ} = 0$). These data complement Fig. 3 of the main text, where the state fidelity is shown as a function of applied gates.

The density matrices are reconstructed from the measurement data using ORENS and Bayesian inference, see Section X for further details. The colored bars correspond to the experimentally reconstructed Pauli coefficients. The black outlined bars indicate the ideal values, while the dashed lines show the expected coefficients obtained from numerical simulations. The simulations are performed using the effective Hamiltonian $H_{\mathrm{eff}} = g_{ab}\, a^\dagger a\, b^\dagger b$ with a gate time of $5.5\,\mu$s. We perform the simulations including the experimentally-calibrated cavity decoherence parameters when the drive is on, with $T_{1a} = 210\,\mu$s, $T_{\phi a} = 70\,\mu$s, $T_{1b} = 388\,\mu$s and $T_{\phi b} = 52\,\mu$s. See Section XI for further details.

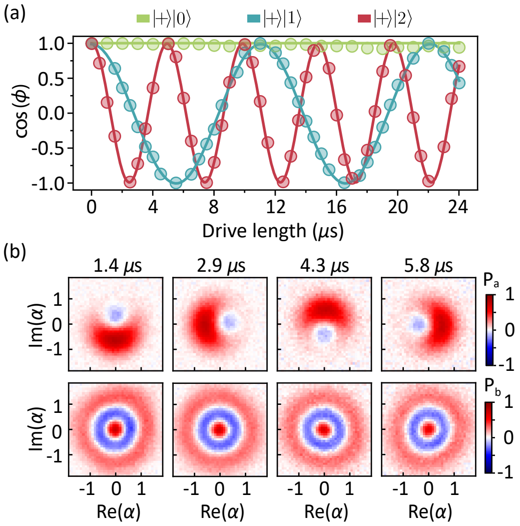

FIG. 5. **Bosonic enhancement of the gate rate.** (a) Phase acquired by the target cavity initialized in $|+\rangle$ during the gate when the control cavity is prepared in $|0\rangle$, $|1\rangle$ and $|2\rangle$. The gate time is reduced from $\sim 5.5\mu$s to $\sim 2.5\,\mu$s when populating the control cavity with two photons. (b) Representative Wigner tomography snapshots of the target cavity at selected times for the initial state $|+\rangle|2\rangle$.

Furthermore, Fig. 7 shows the reconstructed Pauli decompositions after a single CZ gate applied to the initial states $|-\rangle|-\rangle$, $|i\rangle|i\rangle$ and $|-i\rangle|-i\rangle$, which ideally produce maximally entangled states under the engineered interaction. As above, the Pauli coefficients are obtained from measurements using ORENS combined with Bayesian inference. The experimental results are compared to both the ideal states (black outline) and to numerical simulations including decoherence (dashed lines) using the same effective Hamiltonian and parameters described above.

X. QUANTUM STATE RECONSTRUCTION

We perform quantum state reconstruction following the Optimized Reconstruction with Excitation Number Sampling (ORENS) protocol. We extended the scheme in Ref. [12] for reconstructing two-mode cavity states. It involves applying a set of optimized displacement points on both cavities and performing photon number measurements. The measurement results are then processed to estimate the two-mode density matrix. Up to this point, the resulting density matrix might not be physical due to the presence of noise in the experimental data. Thus, we employ Bayesian inference to get a statistically accurate and physical estimated density matrix. The following two sections provide details on the implementation of the ORENS protocol and the Bayesian inference procedure used to obtain the final estimated density matrix.

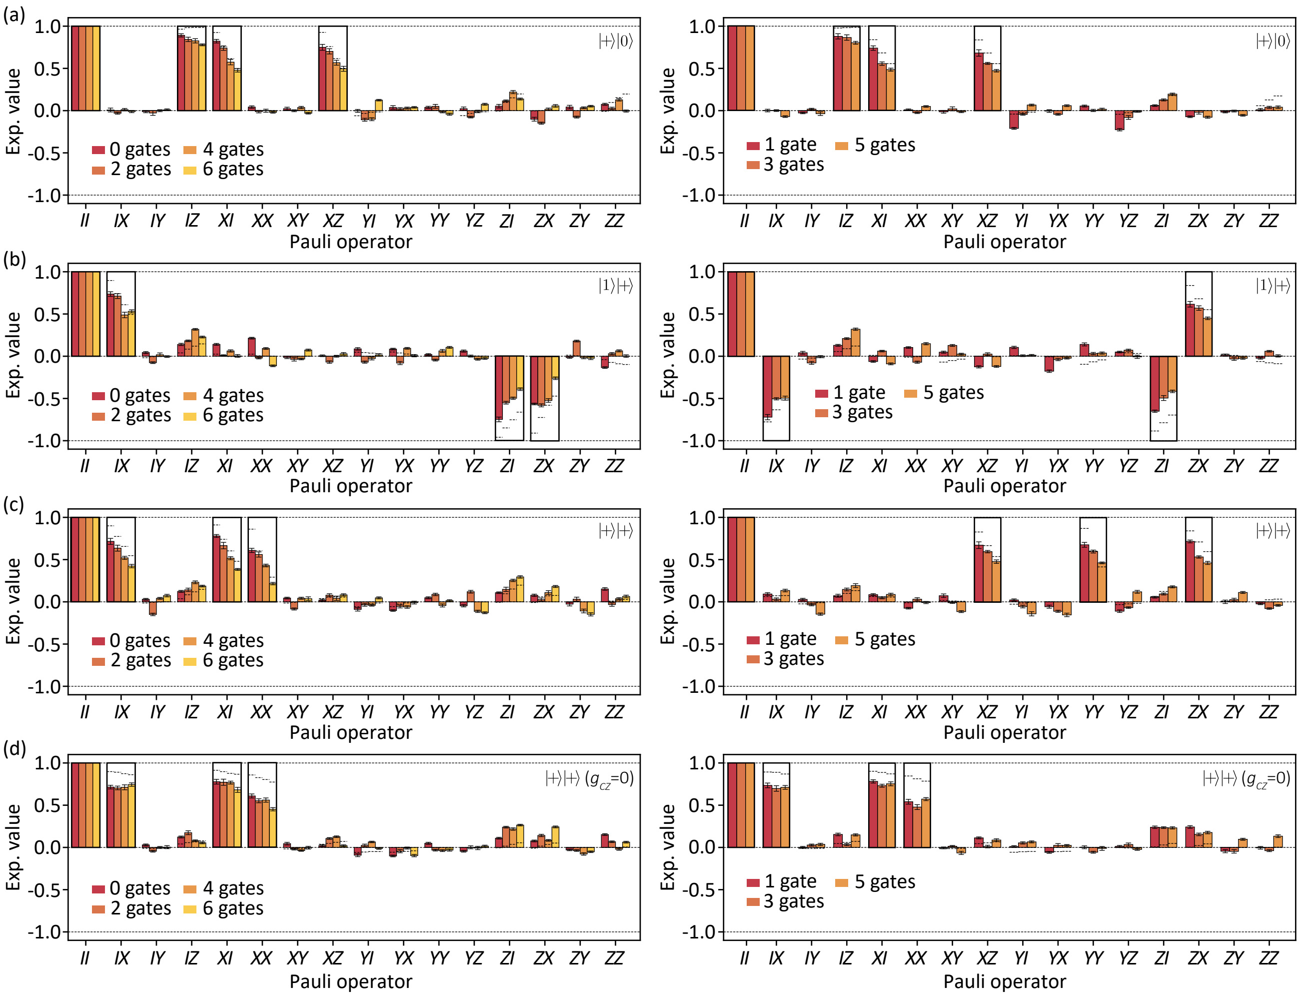

FIG. 6. **Reconstructed density matrices after several CZ gates.** Reconstructed two-cavity states in the Pauli operator representation after 0–6 applied CZ gates, for the initial states (a) $|+\rangle|0\rangle$, (b) $|1\rangle|+\rangle$ (b), and (c) and (d) $|+\rangle|+\rangle$ with drive on and off, respectively. Black outlines indicate the ideal values and dashed lines the numerical simulations including cavity decoherence.

A. State reconstruction via ORENS

ORENS enables efficient quantum state tomography using the theoretical minimum number of measured single-valued observables, which are excitation (photon) number measurements performed after a set of optimized phase-space displacements. For reconstructing a single bosonic mode of dimension D, the least number of observables is $N_{\rm obs} = D^2 - 1$. These are obtained using $N_{\rm obs}$ optimized displacement points, each followed by a measurement of a single photon number n. The set of observables can be written as $X_i = {\rm tr}(|n\rangle\langle n|\hat{D}^\dagger(\alpha_i)\rho\hat{D}(\alpha_i))$, where $i = 1, 2, \cdots, N_{\rm obs}$ and $\{\alpha_i\}$ is the set of optimized displacement points obtained by minimizing the measurement matrix via a gradient descent algorithm [12].

An arbitrary density matrix ρ of dimension D can be parametrized by its D^2 real parameters (without the normalization condition). In our case, we take the diagonal elements of ρ as well as its real and imaginary off-diagonal (upper triangular) elements, which we shall arrange in a vector form $\vec{Y}$. The parameters $(\vec{Y})$ and observables (also arranged in a vector form $\vec{X}$) are linearly related [26]. In our case, we write this as $\vec{Y} = M\vec{X}$. Here, M is a mapping matrix, which can be computed given the set of optimized displacement points $\{\alpha_i\}$ and photon number n to be measured. To get an estimate of the parameters $\vec{Y}_{\rm est}$, given the experimental data $\vec{X}$, we apply the inverse of the matrix equation $\vec{Y}_{\rm est} = M^+\vec{X}$ with $M^+ = (M^\dagger M)^{-1}M^\dagger$ being the left Moore-Penrose pseudoinverse. From the estimated parameters, we then construct the density matrix $\rho_{\rm LS}$. Given this density matrix, which might be unphysical, we employ Bayesian inference to get the final physical density matrix (see next section).

To reconstruct a two-mode cavity state in the present case, we use a total dimension of $D = 4$ (2 for each mode) and $D = 9$ (3 for each mode) for the case of 0/1 encoding and 0/2 encoding, respectively. This truncation is justi-

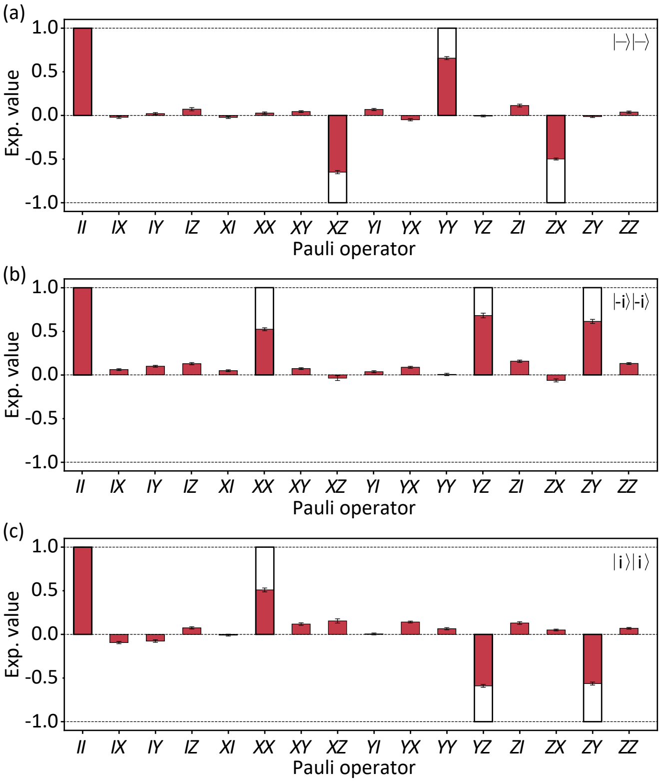

FIG. 7. **Reconstructed density matrices for other cardinal states.** Reconstructed two-cavity states in the Pauli operator representation after a single CZ gate applied to the initial states (a) $|-\rangle|-\rangle$, (b) $|-i\rangle|-i\rangle$, and (c) $|i\rangle|i\rangle$. Black outlines correspond to ideal states, and dashed lines to numerical simulations including cavity decoherence.

fied by the absence of any observed population in $|2\rangle$ and $|3\rangle$ after the gate, within readout error. The displacement operations we apply are of the form $\{\hat{D}_a(\alpha_i) \otimes \hat{D}_b(\alpha_j)\}$, where $\{\alpha_i\}$ refers to the optimized displacements obtained via optimization for the single-mode tomography case. For the 0/1 and 0/2 encoding, we use a total of 25 and 100 different displacements, respectively. This exceeds the theoretical minimum of $D^2 - 1$ in both encoding scenarios, providing an overcomplete set of observables that improves robustness to uncertainties and reduces the averaging required per observable. After each displacement, we measure the probability of observing the joint photon number $|n\rangle_a \otimes |n\rangle_b$ in the two-mode cavity state. This is obtained through individual single-shot measurements of photon number $|n\rangle$ in Alice and Bob, by applying selective π-pulses on their side transmons at frequencies $\omega_{q_a} - n\chi_{a\,q_a}$ and $\omega_{q_b} - n\chi_{b\,q_b}$, respectively, and measuring the resulting excited state probabilities. By correlating the outcomes of these measurements for each experimental run, we obtain the joint probability of observing $|n\rangle_a \otimes |n\rangle_b$ photons.

Note that as we use parametrization with D^2 parameters (without enforcing the normalization condition), the density matrix obtained via linear inversion (least square) ρ_{LS} might not have unit trace and positive eigenvalues. We shall keep ρ_{LS} as a data-driven estimator that faithfully reflects the measurement outcomes, and use it solely as the reference point for a subsequent Bayesian inference procedure that restores physicality in the final estimated density matrix.

Two-mode state reconstruction has previously also been achieved by measuring the joint parity $\langle \hat{P}_a \otimes \hat{P}_b \rangle$ [27]. However, this approach requires closely matched dispersive shifts χ between transmon and both cavities, placing stringent constraints on the hardware. One can in principle perform separated single shot parity measurements for $\hat{P}_a$ and $\hat{P}_b$, and correlate them to get $\langle \hat{P}_a \otimes \hat{P}_b \rangle$, as we did for photon number in this work. However, the traditional parity measurement $\pi/2$-wait-$\pi/2$ is known to suffer from coherent as well as incoherent errors [12]. A standard technique employed to counter this problem is to perform a corrected parity measurement, in which $\langle \hat{P} \rangle_{\mathrm{corr}} = (\langle \hat{P} \rangle - \langle \hat{P} \rangle_{\mathrm{rev}})/2$ is measured. Here $\langle \hat{P} \rangle_{\mathrm{rev}}$ is obtained using a $\pi/2$-wait-$(-\pi/2)$ sequence. Implementing this technique in our device would require correlating 4 different measurements: $\hat{P}_a$ with $\hat{P}_b$; $\hat{P}_a$ with $\hat{P}_{B,\mathrm{rev}}$; $\hat{P}_{A,\mathrm{rev}}$ with $\hat{P}_b$; and $\hat{P}_{A,\mathrm{rev}}$ with $\hat{P}_{B,\mathrm{rev}}$, therefore multiplying the number of required measurements by a factor of 4. This tomography method adds a factor of 2^N to the total number of displacement points for N-mode state reconstruction, which is certainly less favorable compared to the method implemented in this work.

B. Bayesian inference

To obtain a physical (positive semidefinite) estimated density matrix, a standard technique consists of performing linear inversion (as explained in the previous section) followed by maximum likelihood estimation [28]. However, this method suffers from the problem, that the estimated physical density matrix often has zero eigenvalues [29]. This is not justified by finite measurement statistics, as the absence of observed events cannot support the conclusion that the corresponding probabilities are exactly zero. Instead, we employ Bayesian inference, which gives a statistically more accurate estimate. Incorporating Bayes' rule, Bayesian inference provides a posterior probability distribution given the measured experimental data and prior knowledge. The posterior distribution can then be used (sampled) to get an estimate for any function of the density matrix. We follow an efficient Bayesian inference protocol introduced in Ref. [30].

We use uniform prior, i.e., no prior assumption on the density matrix, such that the posterior distribution is proportional to the pseudo likelihood function

$$\mathcal{L}(\rho) \propto \exp\left(- \frac{N}{2} ||\rho - \rho_{LS}||^2 \right) \tag{23}$$

where $|| \cdot ||$ is the Frobenius norm and N is the total number of events (number of repetitions times number of applied displacements). We then draw a total of RT samples according to the Crank-Nicolson Metropolis-Hastings procedure, with a thinning parameter $T = 2^7$ to reduce serial correlation in the subsequent samplings and keep only $R = 2^{10}$ samples. The final reconstructed state is given by the Bayesian mean estimator:

$$\rho_{\mathrm{BME}} = \frac{1}{R}\sum_{r=1}^{R}\rho_r \tag{24}$$

where $\{\rho_r\}$ are the retained physical density matrices sampled from the posterior distribution. We calculate the fidelity between the density matrix ρ_{BME} and a target density matrix ρ_{tar} using:

$$\mathcal{F} = \left(Tr\sqrt{\sqrt{\rho_{\mathrm{tar}}}\rho_{\mathrm{BME}}\sqrt{\rho_{\mathrm{tar}}}}\right)^2, \tag{25}$$

where ρ_{tar} is the target density matrix generated by simulating the numerical GRAPE pulses under realistic decoherence mechanisms, followed by propagation through an ideal gate model with no decoherence, where the effective Hamiltonian used is $\hat{H}_{\mathrm{eff}} = g_{ab}\,\hat{a}^\dagger\hat{a}\,\hat{b}^\dagger\hat{b}$. This procedure isolates the sources of gate imperfections, such as system decoherence and spurious coupler excitations, from errors of the state preparation and measurement. A quantitative error budget is provided in Section XII.

XI. DRIVE-INDUCED DECOHERENCE

A. Strong drive effects on cavity coherences

In this work, we drive the coupler with a strong microwave tone to activate the cross-Kerr interaction. When operating in the Raman regime, the coupler is never physically populated, thereby mitigating errors arising directly from its decoherence mechanisms. Nevertheless, we observe a substantial reduction of the oscillator coherences as a function of both the drive strength and the detuning from the resonance condition described in the main text. To quantify the cavity properties under this strong drive, we characterize the driven single-photon lifetime of the oscillators by loading a photon in Alice (Bob) and replacing the delay time by a strong off-resonant drive applied to the coupler. We observe a degradation of T_1 times from 700–900 μs (800–1000 μs) to $\sim 210\,\mu$s ($\sim 390\,\mu$s). The dephasing time of the oscillators was characterized via a Ramsey experiment where Alice (Bob) is initialized in a superposition $(|0\rangle + |1\rangle)/\sqrt{2}$ and the delay time is replaced with a strong off-resonant drive. We observe a reduction in the oscillator T_2^* times from 500–600 μs (700–900 μs) to $\sim 40\,\mu$s ($\sim 60\,\mu$s).

The observed reduction in coherence is consistent with previous reports in strongly driven cQED systems [8, 31, 32]. We attribute the enhanced oscillator decay to drive-induced dissipation [23, 32], arising from the dressing of the coupler by the off-resonant coherent exchange interaction. A simple model of this effect can be obtained by considering the first-order perturbative correction to the coupler decay operator, $|g\rangle\langle e| \rightarrow |g\rangle\langle e| + \left[\hat{S}, |g\rangle\langle e|\right]$, where $\hat{S}$ is the SWT generator introduced in Section V. The resulting commutator generates an additional term $\beta\, ab^\dagger\left(|g\rangle\langle g| - |e\rangle\langle e|\right)$. Likewise, dressing the coupler dephasing operator yields $\beta\left(a^\dagger b|g\rangle\langle e| + ab^\dagger|e\rangle\langle g|\right)$. Operating the coupler away from its sweet spot, where $T_\phi^c \sim 2\,\mu$s, may therefore lead to an effective oscillator dressed decay time of $T_1^{a,b} \sim \beta^{-2} T_\phi^c \sim 200\,\mu$s, consistent with the observed driven coherence times. In addition, the applied drive induces AC Stark shifts of the coupler transitions, which can approach the cavity frequencies at large drive amplitudes. Furthermore, the off-resonant drive enhances hybridization between the modes, thereby increasing the inherited cavity decay [25, 33]. Similar to [32], there is a tradeoff between the gate time, which scales with Δ, and the sideband-dressed oscillator decay, which scales as Δ^{-2}. While this simple model provides a qualitatively consistent description of our observations, a more rigorous treatment is required to systematically account such drive-induced decoherence effects. This remains an active area of research and lies beyond the scope of the present work.

We further attribute the degradation of the oscillator dephasing time to two distinct mechanisms. First, dressed dephasing [33], arising from an elevated ancilla temperature in the Floquet basis. This contribution depends primarily on the drive-coupler detuning ($\sim 280\,$MHz) and is therefore insensitive to the interaction detuning ($\Delta \sim 6\,$MHz), resulting in a nearly flat photon-shot noise spectrum [32]. Second, the interplay between the enhanced oscillator decay rates and the direct cross-Kerr coupling between the oscillators leads to an additional, cross-Kerr dependent dephasing channel.

B. Driven coupler decoherence

When the system is driven at the resonance condition, we observe a coherent exchange of excitations between the states $\{|1\rangle|0\rangle|g\rangle, |0\rangle|1\rangle|e\rangle\}$, as shown in Fig. 2(a) of the main text. The corresponding coupler population oscillations are presented in Fig. 8. In the presence of decoherence, dephasing reduces the oscillation contrast, while energy relaxation skews the dynamics toward the ground state. Following [8], we model the coupler oscillations as a function of the drive duration using

$$P_0(t) = Ae^{-\kappa_1 t}\Big(1 + e^{-\kappa_\phi t}\cos(2g_1 t + B)\Big) + C, \tag{26}$$

where A, B, and C are amplitude, initial phase and offset fitting parameters to compensate for state preparation and measurement (SPAM) imperfections; κ_ϕ denotes the effective dephasing rate within the exchange subspace; and κ_1 characterizes population decay to the vacuum state. From this fit, we extract effective coherence

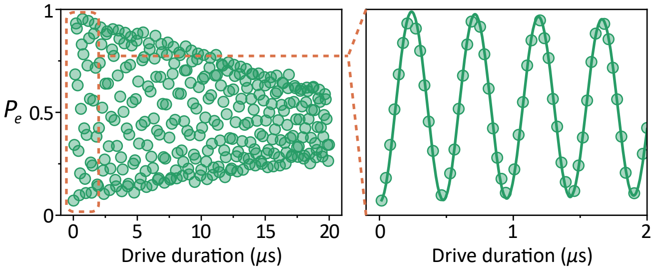

FIG. 8. **Coherent coupler oscillations.** Experimentally measured exchange oscillations of the coupler mode (green markers), used to fit Eq. 26 and extract the corresponding decay and dephasing times. The right panel shows a zoomed-in view of the first 2 μs of the oscillations. The solid line corresponds to master-equation simulations using the cavity and coupler decoherence rates, with the coupler dephasing time measured via a CPMG sequence.

times of $\kappa_1^{-1} = 94 \pm 3\,\mu$s and $\kappa_\phi^{-1} = 28.0 \pm 0.4\,\mu$s. Because the excitations involved in the interaction are delocalized across both cavities and the nonlinear coupler, these rates cannot be directly interpreted as averages of the individual cavity coherence times, in contrast to Ref. [8].

Instead, we perform full master equation simulations using the independently measured cavity and coupler decoherence rates. Importantly, using the bare coupler dephasing time extracted from Ramsey experiments predicts a significantly faster decay of the oscillations than is observed experimentally. This discrepancy indicates that the relevant coherence properties during the exchange process differ from those inferred from static measurements. It is well established that, under continuous driving, the effective coherence times of superconducting transmons are modified relative to their free-evolution values [34, 35]. To account for this, we adopt a coupler T_2 time of 13 μs, corresponding to the experimentally measured decoherence time obtained via a Carr–Purcell–Meiboom–Gill (CPMG) sequence [36]. This choice is justified by the fact that the continuous drive effectively rotates the transmon Bloch vector, producing dynamics analogous to a repeated echo sequence. With this modification, the simulations are in very good agreement with the experimental data, as shown in the right panel of Fig. 8 and in Fig. 2(a) of the main text.

XII. ERROR BUDGETING

In this section, we break down and estimate the different sources of infidelity affecting our experimental data. To this end, we perform full master equation simulations in which individual error mechanisms are introduced one at a time, allowing us to isolate their respective contributions. While this approach provides a useful estimate of each contribution to the infidelity, it does not capture potential correlations between error channels. We therefore also present simulations that incorporate all error sources simultaneously for comparison. Overall, both approaches yield a consistent reproduction of the experimentally measured values.

We begin by considering errors arising from SPAM in an idealized setting without decoherence. We simulate the numerical GRAPE pulses used to prepare the initial states using the Hamiltonian of Eq. 2, followed by an application of a single CZ gate, modeled by the effective Hamiltonian

$$H_{eff} = g_{ab} a^\dagger a b^\dagger b - \chi_a\, a^\dagger a\, |e\rangle\langle e| - \chi_b\, b^\dagger b\, |e\rangle\langle e|. \quad (27)$$

We then simulate the ORENS tomography using a pulse-based description, with realistic selective π-pulses and displacement operations, while neglecting readout errors. The infidelity of the reconstructed simulated final state is quantified as $1 - |\langle \psi_{\mathrm{sim}} | \psi_{\mathrm{target}} \rangle|^2$, where $|\psi_{\mathrm{target}}\rangle$ denotes the ideal state after the gate (e.g. $|1\rangle|-\rangle$ for an initial state $|1\rangle|+\rangle$). The infidelity obtained from this decoherence-free simulation (see turquoise bars in Fig. 9) arises primarily from imperfections in the state preparation, that generate small populations outside the 0/1 subspace, which cannot be captured by the $D = 2$ ORENS tomography. These leakage populations are sufficiently small not to be detected in experiments.

We then repeat the previous analysis, now including decoherence during both the initial state preparation and the ORENS tomography protocol. In addition, measurement errors are incorporated by applying the confusion matrices described in Section II D to the expected excited state populations of the side transmons. This contribution constitutes the dominant source of infidelity in our experiments (see red bars in Fig. 9) and is ultimately limited by the coherence times of the side transmons. We also consider infidelities arising from residual coupler population after the gate (green), corresponding to the residual excitations shown as green crosses in Fig. 2(b) of the main text, as well as from the driven decoherence of Alice (orange) and Bob (purple) during the gate, using the values described in Section XI. Each of these contributions is simulated by including only a single error mechanism at a time, without accounting for SPAM errors. Finally, we compute the total infidelity by including all decoherence channels simultaneously (dotted lines). In some cases, this results in a lower infidelity than that obtained by summing the individual contributions, indicating that certain error mechanisms are not independent. For instance, cavity decoherence can partially suppress spurious populations outside the tomography subspace, thereby mitigating their contribution to the total infidelity. For the states in the 0/2 subspace, we further simulate the effect of post-selection by projecting the simulated states into the even-parity subspace (dotted lines) before passing them into the ORENS tomography simulation. Overall, for a single gate, the SPAM errors dominate the infidelity, whereas for increasing numbers of gates (Fig. 3(e) in the main text), cavity decoherence becomes the leading error mechanism.

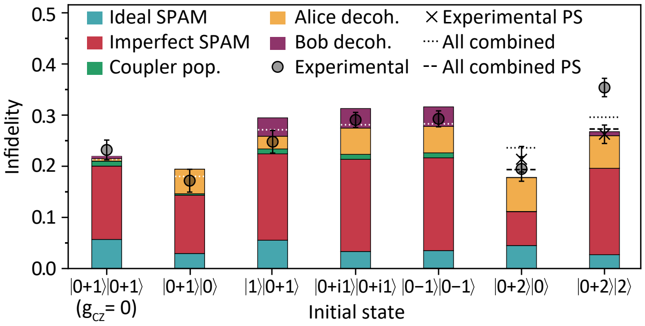

FIG. 9. **Error budgeting for the CZ gate.** Stacked bars show simulated contributions to the infidelity arising from SPAM (with and without decoherence), coupler population, and cavity decoherence for several initial states. Dotted lines indicate simulations that include all error channels simultaneously, while gray markers denote the experimentally measured infidelities. For the last two states, we show the simulated (dashed line) and experimental (crosses) infidelities after post-selecting (PS) the data on the cavities being in an even-parity state.

XIII. CROSS-KERR IN 3-WAVE MIXING

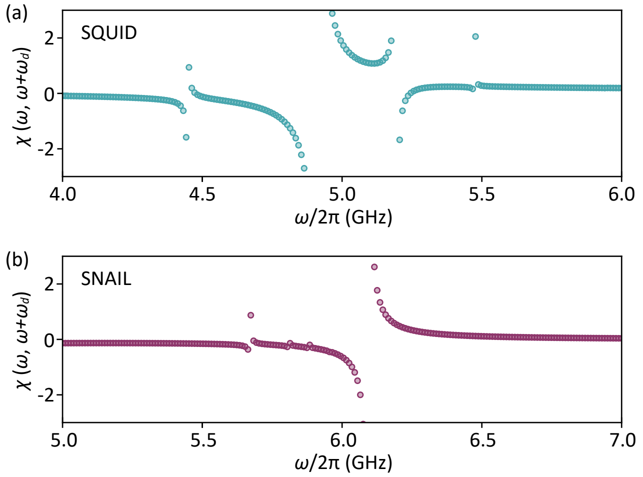

FIG. 10. **SQUID and SNAIL susceptibilities.** Real part of the susceptibility of a (a) SQUID, (b) SNAIL, quantifying the coupler response to a probe at ω when driven at ω_d. The SNAIL, biased at its Kerr-free point, exhibits fewer spurious resonances than the SQUID for the same drive amplitude.

The Raman-assisted protocol presented in this work employs a SQUID coupler and concatenates two 4-wave mixing processes to engineer an effective 6th-order Hamiltonian term (involving 2 drive photons). While this activates the desired cross-Kerr interaction, it also introduces parasitic effects. The always-on dispersive coupling between the cavities and the coupler modifies the effective drive detuning from the coherent exchange condition as a function of the cavities photon number. In addition, AC Stark shifts arising from the 4-wave mixing dynamics can lead to undesired frequency collisions under strong drives. Finally, the nonlinear mixing between all the modes that participate in the coupler junction gives rise to a spectrum of spurious transitions that constrains the choice of drive parameters. Although spurious transitions involving cavity modes are spectrally narrow, sufficient detuning is required to suppress off-resonant second-order processes.

To mitigate these imperfections, we envision replacing the SQUID coupler with a tailored 3-wave mixing element. For example, a linear inductive coupler (LINC) [37] or a Superconducting Nonlinear Asymmetric Inductive eLement (SNAIL) biased to the operating point where it only exhibits 3- and 5-wave mixing [38]. For the remainder of this section we focus on the latter as a concrete example. The SNAIL Hamiltonian is

$$H^{\mathrm{S}}/\hbar = \omega_s \hat{s}^\dagger \hat{s} + g_3\left(\hat{s}+\hat{s}^\dagger\right)^3 + g_5\left(\hat{s}+\hat{s}^\dagger\right)^5 + \cdots, \tag{28}$$

where $\hat{s}$ is the annihilation operator of the SNAIL mode, and g_3 and g_5 are its third- and fifth-order nonlinearities. Coupling the SNAIL to both oscillators and driving at a frequency $\omega_d^{(S)} = \omega_c - \omega_a$ engineers the interaction

$$\hat{H}_d^{(3)}/\hbar = g_3\left(\hat{a}^\dagger\hat{s}+\hat{a}\hat{s}^\dagger\right). \tag{29}$$

At this same resonant condition, the 5-wave mixing term

$$\hat{H}_d^{(5)}/\hbar = g_5\left(\hat{a}^\dagger\hat{b}^\dagger\hat{b}\hat{s}+\hat{a}\hat{b}^\dagger\hat{b}\hat{s}^\dagger\right) \tag{30}$$

is also resonant. As in the SQUID case, driving with a detuning Δ from this resonance activates a Raman-assisted interaction between $\hat{H}_d^{(3)}$ and $\hat{H}_d^{(5)}$, yielding an effective Hamiltonian

$$\hat{H}^{\mathrm{S}}_{\mathrm{eff}}/\hbar = \frac{g_3 g_5}{\Delta}\hat{a}^\dagger\hat{a}\,\hat{b}^\dagger\hat{b}\left(|g\rangle\langle g| - |e\rangle\langle e|\right). \tag{31}$$

Here we assume that, despite operating at its Kerr-free point, the SNAIL retains a small but finite anharmonicity from its 3-wave mixing term. Provided the drive amplitude and bandwidth remain small compared to this anharmonicity, leakage to higher excited states is suppressed, and the coupler can be approximated as an effective two-level system, with $\hat{s} = |g\rangle\langle e|$.

Using a 3-wave mixing element would not only solve the problem of photon-number–dependent cross-Kerr couplings, but also significantly reduce drive-induced frequency shifts and suppress many of the spurious multi-photon processes inherent to 4-wave mixing. To illustrate the latter, we compare the susceptibility $\chi(\omega,\, \omega{+}\omega_d)$ of a SQUID and a SNAIL coupler, following the procedure of [25], see Fig. 10. Here, the susceptibility characterizes the frequency-converting response of the coupler, relating a probe at frequency ω to a response at $\omega{+}\omega_d$, where ω_d is the drive frequency. When the cross-Kerr drive is on, the SNAIL exhibits fewer activated resonances across the probe frequency spectrum, resulting in a significantly cleaner susceptibility profile.

---

[1] Y. Y. Gao, B. J. Lester, Y. Zhang, C. Wang, S. Rosenblum, L. Frunzio, L. Jiang, S. Girvin, and R. J. Schoelkopf, Programmable interference between two microwave quantum memories, Physical Review X **8**, 021073 (2018).

[2] X. Zhao, P. Liao, F. A. Mele, U. Chabaud, and Q. Zhuang, Complexity of quantum tomography from genuine non-gaussian entanglement, Nature Communications (2025).

[3] J. Guo, S. Liu, M. Fadel, and Q. He, Non-gaussian entanglement hierarchy based on the schmidt number, arXiv preprint arXiv:2605.18605 (2026).

[4] S. Rosenblum, Y. Y. Gao, P. Reinhold, C. Wang, C. J. Axline, L. Frunzio, S. M. Girvin, L. Jiang, M. Mirrahimi, M. H. Devoret, *et al.*, A cnot gate between multiphoton qubits encoded in two cavities, Nature communications **9**, 652 (2018).

[5] N. Mehta, J. D. Teoh, T. Noh, A. Agrawal, A. Anderson, B. Birdsall, A. Brahmbhatt, W. Byrd, M. Cacioppo, A. Cabrera, *et al.*, Bias-preserving and error-detectable entangling operations in a superconducting dual-rail system, arXiv preprint arXiv:2503.10935 (2025).

[6] Y. Xu, Y. Ma, W. Cai, X. Mu, W. Dai, W. Wang, L. Hu, X. Li, J. Han, H. Wang, *et al.*, Demonstration of controlled-phase gates between two error-correctable photonic qubits, Physical review letters **124**, 120501 (2020).

[7] K. Reuer, J.-C. Besse, L. Wernli, P. Magnard, P. Kurpiers, G. J. Norris, A. Wallraff, and C. Eichler, Realization of a universal quantum gate set for itinerant microwave photons, Physical Review X **12**, 011008 (2022).

[8] Y. Lu, A. Maiti, J. W. Garmon, S. Ganjam, Y. Zhang, J. Claes, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, High-fidelity parametric beamsplitting with a parity-protected converter, nature communications **14**, 5767 (2023).

[9] B. J. Chapman, S. J. De Graaf, S. H. Xue, Y. Zhang, J. Teoh, J. C. Curtis, T. Tsunoda, A. Eickbusch, A. P. Read, A. Koottandavida, *et al.*, High-on-off-ratio beamsplitter interaction for gates on bosonically encoded qubits, PRX Quantum **4**, 020355 (2023).

[10] Z. K. Minev, Z. Leghtas, S. O. Mundhada, L. Christakis, I. M. Pop, and M. H. Devoret, Energy-participation quantization of josephson circuits, npj Quantum Information **7**, 131 (2021).

[11] R. W. Heeres, P. Reinhold, N. Ofek, L. Frunzio, L. Jiang, M. H. Devoret, and R. J. Schoelkopf, Implementing a universal gate set on a logical qubit encoded in an oscillator, Nature communications **8**, 94 (2017).

[12] T. Krisnanda, C. Y. Fontaine, A. Copetudo, P. Song, K. X. Lee, N.-N. Huang, F. Valadares, T. C. Liew, and Y. Y. Gao, Demonstrating efficient and robust bosonic state reconstruction via optimized excitation counting, PRX Quantum **6**, 010303 (2025).

[13] M. Hutchings, J. B. Hertzberg, Y. Liu, N. T. Bronn, G. A. Keefe, M. Brink, J. M. Chow, and B. Plourde, Tunable superconducting qubits with flux-independent coherence, Physical Review Applied **8**, 044003 (2017).

[14] F. Valadares, N.-N. Huang, K. T. N. Chu, A. Dorogov, W. Chua, L. Kong, P. Song, and Y. Y. Gao, On-demand transposition across light-matter interaction regimes in bosonic cqed, Nature Communications **15**, 5816 (2024).

[15] J. Zimmerman, Sensitivity enhancement of superconducting quantum interference devices through the use of fractional-turn loops, Journal of Applied Physics **42**, 4483 (1971).

[16] Y. Sunada, S. Kono, J. Ilves, S. Tamate, T. Sugiyama, Y. Tabuchi, and Y. Nakamura, Fast readout and reset of a superconducting qubit coupled to a resonator with an intrinsic purcell filter, Physical Review Applied **17**, 044016 (2022).

[17] D. T. McClure, H. Paik, L. S. Bishop, M. Steffen, J. M. Chow, and J. M. Gambetta, Rapid driven reset of a qubit readout resonator, Physical Review Applied **5**, 011001 (2016).

[18] 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, Physical review letters **115**, 137002 (2015).

[19] J. R. Schrieffer and P. A. Wolff, Relation between the anderson and kondo hamiltonians, Physical Review **149**, 491 (1966).

[20] J. Koch, T. M. Yu, J. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Charge-insensitive qubit design derived from the cooper pair box, Physical Review A—Atomic, Molecular, and Optical Physics **76**, 042319 (2007).

[21] S. O. Mundhada, A. Grimm, S. Touzard, U. Vool, S. Shankar, M. H. Devoret, and M. Mirrahimi, Generating higher-order quantum dissipation from lower-order parametric processes, Quantum Science and Technology **2**, 024005 (2017).

[22] M. Grifoni and P. Hänggi, Driven quantum tunneling, Physics Reports **304**, 229 (1998).

[23] X. You, A. C. Li, T. Roy, S. Zhu, A. Romanenko, A. Grassellino, Y. Lu, and S. Chakram, Floquet-engineered fast snap gates in weakly coupled circuit-qed systems, Physical Review Applied **24**, 034072 (2025).

[24] N. Lambert, E. Giguère, P. Menczel, B. Li, P. Hopf, G. Suárez, M. Gali, J. Lishman, R. Gadhvi, R. Agarwal, *et al.*, Qutip 5: The quantum toolbox in python, Physics Reports **1153**, 1 (2026).

[25] Y. Zhang, J. C. Curtis, C. S. Wang, R. J. Schoelkopf, and S. M. Girvin, Drive-induced nonlinearities of cavity modes coupled to a transmon ancilla, Physical Review A **105**, 022423 (2022).

[26] T. Krisnanda, H. Xu, S. Ghosh, and T. C. Liew, Tomographic completeness and robustness of quantum reservoir networks, Physical Review A **107**, 042402 (2023).

[27] C. Wang, Y. Y. Gao, P. Reinhold, R. W. Heeres, N. Ofek, K. Chou, C. Axline, M. Reagor, J. Blumoff, K. Sliwa, *et al.*, A schrödinger cat living in two boxes, Science **352**, 1087 (2016).

[28] J. A. Smolin, J. M. Gambetta, and G. Smith, Efficient method for computing the maximum-likelihood quantum state from measurements with additive gaussian noise, Physical review letters **108**, 070502 (2012).

[29] R. Blume-Kohout, Optimal, reliable estimation of quantum states, New Journal of Physics **12**, 043034 (2010).

[30] J. M. Lukens, K. J. Law, A. Jasra, and P. Lougovski,

A practical and efficient approach for bayesian quantum state estimation, New Journal of Physics **22**, 063038 (2020).

[31] L. Carde, P. Rouchon, J. Cohen, and A. Petrescu, Flux-pump-induced degradation of t 1 for dissipative cat qubits, Physical Review Applied **23**, 024073 (2025).

[32] T. Kim, T. Roy, X. You, A. C. Li, H. Lamm, O. Pronitchev, M. Bal, S. Garattoni, F. Crisa, D. Bafia, *et al.*, Ultracoherent superconducting cavity-based multi-qudit platform with error-resilient control, arXiv preprint arXiv:2506.03286 (2025).

[33] Y. Zhang, B. J. Lester, Y. Y. Gao, L. Jiang, R. Schoelkopf, and S. Girvin, Engineering bilinear mode coupling in circuit qed: Theory and experiment, Physical Review A **99**, 012314 (2019).

[34] F. Yan, S. Gustavsson, J. Bylander, X. Jin, F. Yoshihara, D. G. Cory, Y. Nakamura, T. P. Orlando, and W. D. Oliver, Rotating-frame relaxation as a noise spectrum analyser of a superconducting qubit undergoing driven evolution, Nature communications **4**, 2337 (2013).

[35] G. Ithier, E. Collin, P. Joyez, P. Meeson, D. Vion, D. Esteve, F. Chiarello, A. Shnirman, Y. Makhlin, J. Schriefl, and G. Schön, Decoherence in a superconducting quantum bit circuit, Physical Review B **72**, 134519 (2005).

[36] H. Y. Carr and E. M. Purcell, Effects of diffusion on free precession in nuclear magnetic resonance experiments, Physical review **94**, 630 (1954).

[37] A. Maiti, J. W. Garmon, Y. Lu, A. Miano, L. Frunzio, and R. J. Schoelkopf, Linear quantum coupler for clean bosonic control, PRX Quantum **6**, 040326 (2025).

[38] N. Frattini, U. Vool, S. Shankar, A. Narla, K. Sliwa, and M. Devoret, 3-wave mixing josephson dipole element, Applied Physics Letters **110** (2017).