

Appendix for “A Flux-Tunable Cavity for Dark Matter Detection”

I. CAVITY AND QUANTUM DEVICES FABRICATION

The cavities used in this work are fabricated from high-purity (99.9999%) Aluminium. The full device consists of two microwave cavities (readout and storage), each coupled to the transmon. The storage cavity is coupled with one flux-tunable SQUID device.

The quantum devices were fabricated on 430 μm thick C-plane edge-fed growth Sapphire wafers with a diameter of 50.8 mm. Wafers were annealed in an oxygen-rich environment at 1200 for 3 hours. Afterwards, wafers were cleaned with three steps: 1. solvents clean (Toluene, Acetone, Methanol, Isopropanol, and DI water) in an ultrasonic bath; 2. Nanostrip clean at 45 C for 15 min and DI water rinse for 10min; 3, sulfuric acid clean at 70 C for 10min followed by a thorough rinse in DI water for 20min. Tantalum was deposited after the annealing and cleaning by DC magnet sputtering while maintaining a substrate temperature of 800 C. The features were defined via optical lithography using AZ MiR 1518 photoresist and exposure with a Heidelberg MLA150 Direct Writer. The resist was developed for 1 min in AZ MIF 300 1:1. The features were etched in a Plasma-Thermal inductive coupled plasma etcher using chlorine-based etch chemistry with a plasma consisting of 30 sccm Cl_2 , 30 sccm BCl_3 and 10 sccm Ar. After etching, the photoresist was removed in N-Methylpyrrolidone (NMP) overnight, followed by acetone, isopropanol, and DI water with sonicating for 5 min.

Josephson junctions were patterned using electron-beam lithography. The wafer was first dehydrated by baking at 180 C for 8 mins under low vacuum. Then, a bilayer of MMA EL13 and 950K PMMA A7 was spun with a 5min bake at 180 C following the spinning of each layer. To eliminate charging effects during electron-beam writing, a 10 nm gold anti-charging layer was deposited by electron-beam evaporation. Electron-beam lithography was then performed using a Raith EBPG 5000+ to define the Manhattan pattern for transmon and the Dolan-bridge for SQUID. The resist stack was developed for 90 seconds at 6 C in a solution of 3 parts IPA and 1 part DI water. The wafer was then loaded into the load lock of a Plassys UMS300 electron-beam evaporator. Before deposition, the overlap regions on the pre-deposited capacitors were milled in situ with an Argon ion mill to remove the native oxide. The junctions were then deposited using a three-step electron-beam evaporation and oxidation process. First, an initial 45 nm layer of aluminum was deposited at 1 nm/s. Next, the junctions were exposed to a high-purity mixture of Ar and O₂ (ratio of 80:20) for the first layer to grow a native oxide. Finally, a second 115 nm layer of aluminum was deposited at 1 nm/s. After the second aluminum de-

position, a second static oxidation step was performed using the same Ar:O₂ mixture at 3 mbar for 5 mins in order to cap the surface of the bare aluminum with pure aluminum oxide. The wafer was then coated with a protective resist before dicing into individual chips. Liftoff was then performed by immersing the chips in NMP at 80 C for 3 hours followed by sonication for 2 mins each in NMP, acetone, and isopropanol. After the liftoff, the chips were exposed to an ion-producing fan for 15min to avoid electrostatic discharge of the junctions. The room temperature DC resistance of the Josephson junction on each transmon was measured to select the transmon that corresponds to the target Josephson energy [1].

The cavities underwent a chemical etching treatment using a mixture of phosphoric and nitric acid (Transene Aluminum Enchant Type A) heated to 50 C for etching aluminum 50 μm film. The transmon and coupler are inserted into the slots and clamped by the clamps on the other part of the cavities.

II. COUPLER AND ON-CHIP FILTER DESIGN

The Hamiltonian for the circuit containing the coupler and storage modes in the bare basis is described by

$$H_{cs} = \tilde{\omega}_s \tilde{a}_s^\dagger \tilde{a}_s + \tilde{\omega}_c (\Phi_{\text{ext}}) \tilde{a}_c^\dagger \tilde{a}_c + g (\tilde{a}_s^\dagger + \tilde{a}_s) (\tilde{a}_c^\dagger + \tilde{a}_c). \quad (1)$$

The set $\{\tilde{\omega}_s, \tilde{\omega}_c(\Phi_{\text{ext}})\}$ represents the bare basis frequencies of the storage mode and the coupler, respectively. The corresponding annihilation operators are denoted as $\{\tilde{a}_s, \tilde{a}_c\}$. The capacitive coupling strength is approximated by the constant g within our detector’s operating regime. By applying the rotating wave approximation (RWA), we derive the Jaynes-Cummings Hamiltonian

$$H_{cs} = \tilde{\omega}_s \tilde{a}_s^\dagger \tilde{a}_s + \tilde{\omega}_c (\Phi_{\text{ext}}) \tilde{a}_c^\dagger \tilde{a}_c + g (\tilde{a}_s^\dagger \tilde{a}_c + \tilde{a}_c^\dagger \tilde{a}_s). \quad (2)$$

Treating the coupler mode as a two-level system, storage mode’s first excited state is

$$\omega_s(\Phi_{\text{ext}}) = \frac{\tilde{\omega}_s + \tilde{\omega}_c}{2} + \frac{1}{2} \sqrt{4g^2 + (\tilde{\omega}_s - \tilde{\omega}_c(\Phi_{\text{ext}}))^2}. \quad (3)$$

The storage mode frequency is tunable via DC flux, allowing the full avoided-crossing frequency response to exhibit a maximum tuning range of $2g$. In experimental setups, achieving a coupling strength greater than 100 MHz is straightforward, thereby theoretically enabling a maximum detector tuning range beyond 200 MHz.

The SQUID coupler design is illustrated in Fig. 1(a). This design features two large metal pads on the chip’s

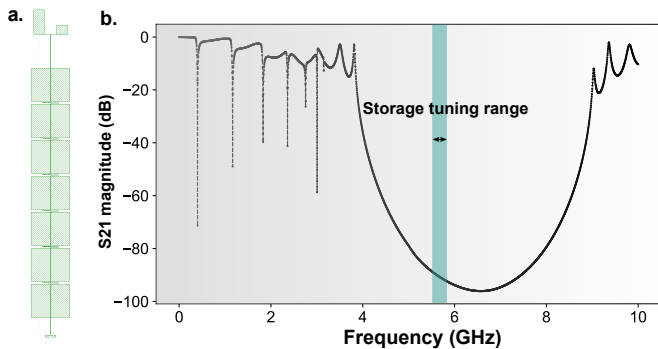


FIG. 1. (a) SQUID coupler and on-chip flux line design. (b) Ansys HFSS simulated filter S_{21} response. The tuning range of the storage mode is highlighted in the same plot.

top side, which provide capacitive coupling to the storage mode. A coupling strength of approximately $g \sim 170$ MHz is extracted from Black-box quantization simulations. The flux loop of the coupler is integrated with a single on-chip low-pass filter. The flux line is configured as a Coplanar Strip Line (CPS) structure. The low-pass filter is comprised of seven sections, alternating between capacitor segments with large pads (offering low impedance Z_{lo}) and inductor segments with meander-shaped lines (providing high impedance Z_{hi}). This configuration establishes an Edge-coupled microstrip CPS filter.

Fig. 1(b) presents the ANSYS HFSS simulation results of our CPS filter. The primary characteristics of our filter design are the stopband center, bandwidth, and depth. These parameters are crucial for protecting the coupler and storage mode coherence. The length of each filter section primarily determines the stopband center frequency. The lengths of the low and high-impedance sections are equal for optimal performance. The stopband bandwidth can be effectively widened by increasing the impedance Z_{hi}/Z_{lo} , which is around 22 for our experiment. However, in our experimental setup, this ratio is constrained by the physical dimensions of the high-impedance sections: the metal width ($2 \mu\text{m}$), the gap between the capacitor pads ($2 \mu\text{m}$), and the overall size of the capacitor parts, which is bound by the chip width. Adding more filter sections increases the stopband depth, although the coupler chip length limits this.

III. CALIBRATION

To conduct the dark matter detection experiment, the following parameters are calibrated: transmon frequency ω_q , single-shot readout fidelity, transmon-storage coupling $\chi(\Phi_{\text{ext}})$, storage frequency $\omega_s(\Phi_{\text{ext}})$, storage displacement drive, storage coherence T_1^s , and transmon and storage background photon populations $\{\bar{n}_q, \bar{n}_s\}$. For the flux-tunable parameters, such as $\omega_s(\Phi_{\text{ext}})$, we calibrate the values at 34 separate flux points Φ_{ext} varying between

Device Parameter	Symbol	Value
Transmon Frequency	ω_q	$2\pi \times 4.811$ GHz
Transmon anharmonicity	α_q	$-2\pi \times 176.7$ MHz
Transmon decay time	T_1^q	$34 \pm 0.5 \mu\text{s}$
Transmon Ramsey time	T_2^q	$31 \pm 0.6 \mu\text{s}$
Transmon background photon	\bar{n}_q	$1.4 \pm 0.2 \times 10^{-2}$
Storage frequency	ω_s	$2\pi \times 5.694$ GHz
Storage decay time	T_1^s	$65 \pm 10 \mu\text{s}$
Storage Ramsey time	T_2^s	$40 \pm 10 \mu\text{s}$
Storage-Transmon Stark shift	2χ	$-2\pi \times 3.31$ MHz
Storage background photon	\bar{n}_s	$8.6 \pm 0.5 \times 10^{-3}$
Readout frequency	ω_r	$2\pi \times 6.733$ GHz
Readout time	t_r	$2.3 \mu\text{s}$
Readout releasing time	t_l	$5 \mu\text{s}$
Reset + Integration time	t_c	$500 \mu\text{s}$
Readout-Transmon Stark shift	$2\chi_{qr}$	$-2\pi \times 0.28$ MHz
Readout fidelity ($ g\rangle$)	F_{gg}	$97.8 \pm 1\%$
Readout fidelity ($ e\rangle$)	F_{ee}	$93.8 \pm 1\%$

TABLE I. Device parameters measured at $\Phi_{\text{ext}} = 0$

-0.4807 and 0 . We perform a linear extrapolation fit for all other values to extract the numbers.

The transmon frequency ω_q is calibrated through the Rabi-Ramsey-Rabi sequence: Using the rough transmon frequency ω_q measured through two-tone spectroscopy, the first transmon Rabi drive signal calibrates the π and $\pi/2$ pulse length. After that, the Ramsey experiments calibrate the transmon frequency, followed by the second transmon Rabi drive signal to fine-adjust the π and $\pi/2$ pulse length. The measured transmon frequency, T_1 , and T_2 at different flux points are shown in Fig. 2 (a) and (b). The measured transmon single shot fidelity (flux-independent) is shown in Fig. 3, where the transmon decay event during readout limits the fidelity.

The transmon-storage coupling strength χ is measured using Ramsey experiments. Initially, we prepare the transmon in the second excited state, $|f\rangle$, by applying two sequential single-transmon rotations. Following this, we drive the transmon port to activate the $|f0\rangle \leftrightarrow |g1\rangle$ sideband transition between the transmon and storage mode (Shown in Fig. 2 (c)). This sideband prepares the state $|1\rangle$ in the storage mode. Subsequently, a Ramsey experiment is conducted on the transmon to determine its frequency. Compared to the storage mode in the $|0\rangle$ state, the observed frequency shift is 2χ . And the parity waiting time $t_p = \pi/|2\chi|$ is simultaneously calibrated.

The storage frequency ω_s is calibrated using Ramsey experiments. For each DC flux point targeted for calibration, the transmon is initially prepared in the state $\frac{1}{\sqrt{2}}(|g\rangle + |f\rangle)$. This is achieved through a $\pi/2$ rotation within the $|g\rangle, |e\rangle$ subspace followed by a π rotation within the $|e\rangle, |f\rangle$ subspace. We then activate the $|f0\rangle \leftrightarrow |g1\rangle$ sideband transition to prepare the state $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ in the storage mode. By varying the waiting time before reapplying the $|f0\rangle \leftrightarrow |g1\rangle$ transition and decoding the storage state into the transmon, we fit the Ramsey fringe data to calibrate ω_s and the storage T_2 .

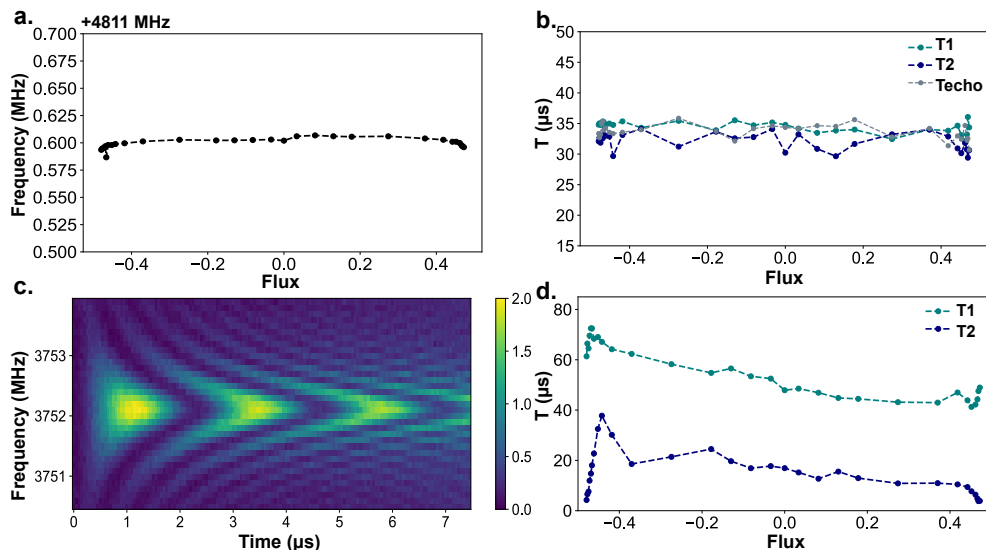


FIG. 2. Transmon and storage mode calibration. Measured transmon frequency ω_q (a) and coherence (b) versus DC flux bias shows that our transmon can be approximated as a flux-independent element in our system. (c) shows the chevron plot of the sideband $|f0\rangle \leftrightarrow |g1\rangle$ between the transmon and the storage. The color bar indicates the photon population in the transmon after different frequencies and duration of the sideband drive. The storage mode coherences at different flux points are shown in (d). Marker sizes are larger than the error bars.

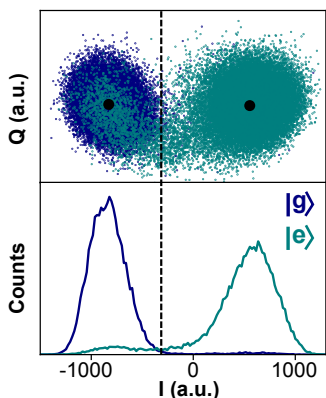


FIG. 3. Single shot readout of transmon states. Transmon is prepared in one of its possible states $\{|g\rangle, |e\rangle\}$ in blue and cyan. 30000 single shot data for each state are measured. The black dash line indicates the demarcation location, where the readout is assigned to $|g\rangle$ when the I channel signal is smaller than the demarcation line and assigned to $|e\rangle$ for the other case. Signals are rotated on the I-Q plane to maximize the contrast in the I channel.

After calibrating ω_s , the storage coherence T_1^s is measured. This involves preparing the storage mode in state $|1\rangle$ using the $|f0\rangle \leftrightarrow |g1\rangle$ sideband transition, waiting for a variable delay, then reapplying the $|f0\rangle \leftrightarrow |g1\rangle$ transition and measure the transmon population. The storage decay and Ramsey times at different flux points are documented in Fig. 2 (d).

To calibrate the correspondence between the drive gain and alpha, we displace the storage cavity using a $0.5\ \mu\text{s}$ square pulse, generating a coherent state $|\alpha\rangle$. The magni-

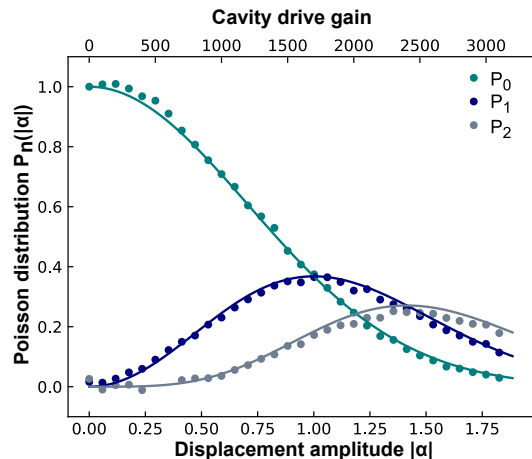


FIG. 4. Mapping drive gain ϵ to storage displacement $\alpha = \epsilon/c$. Storage photon populations in $\{|0\rangle, |1\rangle, |2\rangle\}$ are measured after a constant $0.5\ \mu\text{s}$ storage displacement with a varied drive gain. The data are fitted to extract c . Negative data points are neglected to fit the results better. Marker sizes are larger than the error bars.

tude of α depends on the DAC channel's drive gain ϵ . We establish the ϵ - α mapping by varying ϵ and measuring the resultant displacement $|\alpha\rangle$, as illustrated in Fig. 4. The measurement of α involves applying a photon-number-resolved π pulse to the transmon at frequencies $\omega_q - n\chi$ for various integers n . To improve accuracy, we swept the transmon frequency around $\omega_q - n\chi$ and fitted the signal peak amplitude to extract the cavity Fock state population at each α . For small α , higher cavity Fock states such as $|2\rangle$ have small populations, making the fitted sig-




FIG. 5. Transmon temperature measurement. Transmon $|e\rangle - |f\rangle$ amplitude Rabi signal with (green) and without (purple) initial $|g\rangle - |e\rangle$ rotations are plotted. Dash lines show the fitting results.

nal peak amplitude unstable and occasionally negative. These points were therefore excluded to achieve a more reliable calibration of the ϵ - α mapping. This resolved π pulse maps the population of the Fock state $|n\rangle$ in the coherent state $|\alpha\rangle$ to the excited state $|e\rangle$ of the transmon. The relationship between the population P_n of the Fock state $|n\rangle$ and α is given by:

$$P_n(\alpha = c\epsilon) = \frac{\alpha^{2n} e^{-\alpha^2}}{n!}. \quad (4)$$

Here, c represents the calibration constant that maps the drive gain to the amplitude of the coherent state. By measuring the populations of the Fock states $|0\rangle, |1\rangle, |2\rangle$ for various values of ϵ and fitting these populations, we extract the value of c (See Fig. 4).

To characterize the transmon background population \bar{n}_q , we measure the $|e\rangle - |f\rangle$ Rabi oscillation amplitude ratio with and without the initial transmon $|g\rangle - |e\rangle$ rotation (See Fig. 6). This extracts the transmon population at the ground state $P(g)$ and at the excited state $P(e)$ in its steady state. Ignoring the transmon population above the first excited state, we attain $P(g) = 0.985$ and $P(e) = 0.015$, corresponding to a qubit temperature of 53 mK.

Following the photon counting scheme in Fig. 2(a), two experimental data are shown in Fig.9 by gathering the single-shot data following N parity measurements ($N = 25$).

IV. HIDDEN MARKOV MODEL

We model our storage-transmon-readout system for each storage tuning point as a hidden Markov process. The storage and transmon states are treated as hidden variables emitted as the readout signal [2]. The transition matrix T characterizes how the storage-transmon state evolves, and the confusion matrix E determines the readout signal $\{g, e\}$ given any storage-transmon state:




FIG. 6. Photon counting experiments. Two examples show the measured qubit readout sequences for a cavity initialized in $|N\rangle = 0$ Fock state followed by a small displacement drive. we obtain a series of transmon states for each experiment, which remains constant (left) when the storage mode is in the ground state and displays an oscillating pattern when the storage mode is excited (right).

$$T = \begin{pmatrix} P_{00}P_{gg} & P_{00}P_{ge} & P_{01}P_{ge} & P_{01}P_{gg} \\ P_{00}P_{eg} & P_{00}P_{ee} & P_{01}P_{ee} & P_{01}P_{eg} \\ P_{10}P_{gg} & P_{10}P_{ge} & P_{11}P_{ge} & P_{11}P_{gg} \\ P_{10}P_{eg} & P_{10}P_{ee} & P_{11}P_{ee} & P_{11}P_{eg} \end{pmatrix} \begin{matrix} |0g\rangle \\ |0e\rangle \\ |1g\rangle \\ |1e\rangle \end{matrix} \quad (5)$$

$$E = \begin{pmatrix} g & e \\ F_{gg} & 1 - F_{gg} \\ 1 - F_{ee} & F_{ee} \\ F_{gg} & 1 - F_{gg} \\ 1 - F_{ee} & F_{ee} \end{pmatrix} \begin{matrix} |0g\rangle \\ |0e\rangle \\ |1g\rangle \\ |1e\rangle \end{matrix} \quad (6)$$

The transition matrix T elements are calculated based on the coherence and background photon population in the transmon and storage [2]:

$$\begin{cases} P_{gg} = 1 - P_{ge}, P_{ee} = 1 - P_{eg} \\ P_{ge} = \bar{n}_q[1 - e^{-t_m/T_1^q}] + 1 - e^{-t_p/T_2^q} \\ P_{eg} = 1 - e^{-t_m/T_1^q} + 1 - e^{-t_p/T_2^q} \\ P_{00} = 1 - P_{01}, P_{11} = 1 - P_{10} \\ P_{01} = \bar{n}_s[1 - e^{-t_m/T_1^s}] \\ P_{10} = 1 - e^{-t_m/T_1^s} \end{cases} \quad (7)$$

Here P_{gg}, P_{ee}, P_{00} , and P_{11} correspond to event where no error occurs. Heating and decay in transmon and storage during readout period t_m and transmon dephasing during parity waiting time t_p contribute to the state changes. Here $t_m = t_r + t_l = 7.3 \mu\text{s}$ is the full readout time.

The confusion matrix E is directly measured through the single shot experiments (See Fig. 3). Here F_{gg} and F_{ee} correspond to the case when the measurement result matches the transmon state.

Following Reference [2], the initial storage state probabilities $P(i = 0)$ and $P(i = 1)$ are reconstructed through the backward algorithm after N readout signals (R_1, \dots, R_N) :

$$P(i) = E_{|ig\rangle}^{R_1} T_{|ig\rangle}^{a_1} E_{a_1}^{R_2} T_{a_2}^{a_3} \dots T_{a_{N-1}}^{a_N} E_{a_N}^{R_N} \\ + E_{|ie\rangle}^{R_1} T_{|ie\rangle}^{b_1} E_{b_1}^{R_2} T_{b_2}^{b_3} \dots T_{b_{N-1}}^{b_N} E_{b_N}^{R_N} \quad (8)$$

Here, $\{a_j, b_j\}$ are the indices that are summed over. The larger the ratio $\lambda = P(i=0)/P(i=1)$ is, the less likely the storage mode is excited. The threshold value λ_{thresh} is chosen to determine if the storage is at $|1\rangle$. To characterize the detector's δ and η at different flux tuning points, we inject a list of known weak displacements α into the storage mode and perform the parity measurement with $N = 25$ and $\lambda_{\text{thresh}} = 125$ in the Hidden Markov model analysis. We fit the injected photon population ($\bar{n}_{inj} = |\alpha|^2$) and measured photon population (\bar{n}_{meas}) to the equation:

$$\bar{n}_{inj} = \eta \times \bar{n}_{meas} + \delta \quad (9)$$

The results at different Φ_{ext} (in addition to the data presented in the main text) are shown in Fig. 7. Here $\lambda_{\text{thresh}} = 125$ are chosen based on the efficiency corrected false positive probability δ/η (See Fig. 8): Increasing λ_{thresh} will exponentially suppress the detector's false positive probability δ while reducing the detector efficiency η . The ratio δ/η continues to decrease and become almost insensitive to λ_{thresh} when $\lambda_{\text{thresh}} > 100$. For a large enough threshold $1/(\lambda_{\text{thresh}} + 1) < \delta$, the qubit readout error is no longer the dominant error source and smaller λ_{thresh} helps speed up the detection. In our case, $1/(\lambda_{\text{thresh}} + 1) < \delta$ is satisfied when $\lambda_{\text{thresh}} = 125$. Fig. 9 shows the positive photon counts and mixing angle exclusion data.

The hidden Markov model counts on independent measurements of the probabilities involved in the transition and emission matrices. The elements of these matrices depend on the parameters of the experiment and the device, including the lifetime and dephasing times of the qubit and cavity, qubit population, and readout fidelities. In our dark-photon detection experiments, the detector efficiencies are linearly extrapolated at other Φ_{ext} points.

V. PHOTON DETECTOR EFFICIENCY SIMULATION

To understand the limiting factor of the tunable photon detector efficiency, we perform the master equation simulation using QuTip. We focus on the storage-transmon subsystem that implements the QND measurement of storage excitation and ignore the other components for simplicity. All simulations are carried out in a Hamiltonian of dimension 3×4 (3 levels for the transmon and 4 levels for the storage mode). The Hamiltonian used in the simulation is:

$$H_{qs} = \omega_q a_q^\dagger a_q + \alpha a_q^\dagger a_q^\dagger a_q a_q / 2 \\ + \omega_r a_r^\dagger a_r + 2\chi a_q^\dagger a_q a_r^\dagger a_r \quad (10)$$

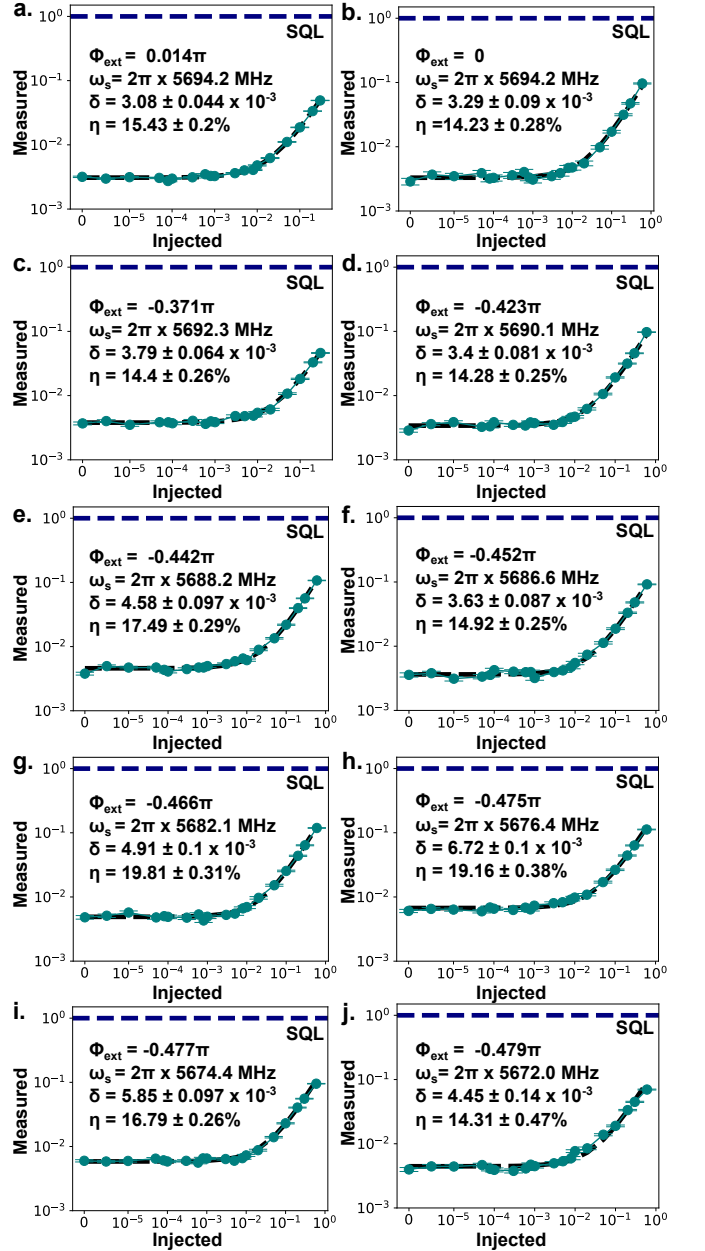


FIG. 7. Detector efficiency characterization at different DC flux biasing points.

Since both transmon and storage modes are only actively used up to the first excited state in the experiment, we approximate the transmon as a Kerr-oscillator. The simulation starts with the first excited state $|1\rangle$ in the storage mode and the ground state $|g\rangle$ for the transmon, followed by a sequence of 25 parity check measurements. Each parity check contains two $\pi/2$ transmon rotations (2σ cutoff Gaussian pulse, $\sigma = 10$ ns) gapped by a waiting time $\frac{\pi}{2\chi}$, and a following waiting time of $7.3 \mu\text{s}$ that represents the total readout time (including readout photon releasing time). To simulate the transmon readout error, we first calculate the ground state probability of the




FIG. 8. Efficiency corrected false-positive probability ($\Phi_{\text{ext}} = 0$).




FIG. 9. (a) A zoom in the our tuning range shows a family of curves describing the calculated exclusion for each cavity tuning. The minimum envelope of these curves describes the dark matter exclusion across our tuning range. (b) Positive photon counts and mixing angle exclusion data. The mixing angle exclusion data has applied the minimum envelope.

transmon P_0 and assign all the other transmon states' probability to the first excited state $P_1 = 1 - P_0$. The transmon state probability vector $\vec{P} = (P_0, P_1)^T$ is further multiplied by the readout confusion matrix \mathbf{F} measured in the experiment: $\vec{\tilde{P}} = \mathbf{F}\vec{P}$. After each parity measurement, the transmon state is assigned based on the $\vec{\tilde{P}}$ probability. A collection of 25 measured transmon states after the parity check sequence is further analyzed using the Markov model analysis discussed before. We choose the same threshold $\lambda_{\text{thresh}} = 125$ as used in the experiment to determine the state of the cavity. The whole simulations are repeated 2000 times, and the average time measuring $|1\rangle$ for the cavity state is the detector efficiency.

We compare the simulated detector efficiency without and with decoherence and plot the results in Fig. 10. The

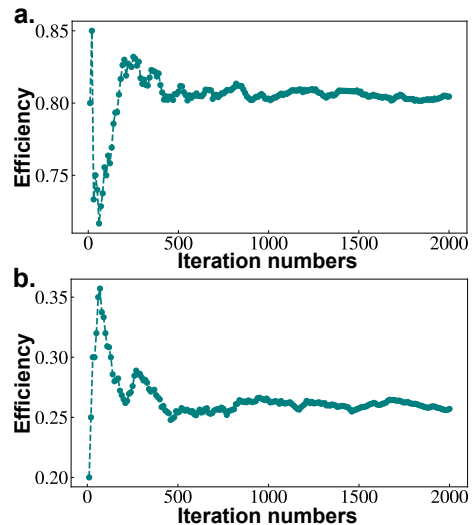


FIG. 10. Simulated detector efficiency without (a) and with (b) decoherence. We use the device parameter experimentally measured at DC flux=0 (Table. 1) in the master equation simulation. The fast transmon $\pi/2$ rotation pulse used in both simulation and experiment is a Gaussian pulse with a 2σ cutoff ($\sigma = 10$ ns).

Parameter	Θ	σ_Θ
Detector efficiency η (%)	[6.81, 19.81]	2.1
Storage quality factor Q_s ($\times 10^6$)	[2.15, 2.98]	0.009
Storage frequency ω_s (MHz)	[5671.4, 5694.2]	0.1
Storage volume V (cm^3)	6.452	0.175
Storage form factor G	0.219	0.003

TABLE II. Mean Θ and standard deviation σ_Θ for experimental parameters. The standard deviation for η is extracted from the injected and measured photon fitting with $\lambda_{\text{thresh}} = 125$. The standard deviation in ω_s is calculated based on the storage T_2 . The storage volume standard deviation is calculated assuming each dimension's machining tolerance of 0.005 inches. The tunable parameters' ranges are listed in the table. We use the largest value for all tunable parameters' standard deviation in the analysis as the lower bound performance of our detector.

relatively slow transmon rabi pulse makes the transmon rotation dependent on the cavity state, which results in an efficiency of 80% even without decoherence. The room temperature control hardware power limits this and can be easily fixed by adding extra amplifiers. The decoherence in the system, especially the T_1 of the cavity, limits the detector efficiency to 25%. Better cavity etching, sealing, and a modified on-chip filter design will help push the efficiency close to the decoherence-free limit.

VI. CONVERTING CAVITY BACKGROUND COUNTS TO HIDDEN-PHOTON EXCLUSION

For a dark matter candidate on resonance with the storage frequency ($m_{DM}c^2 = \hbar\omega_c$), the rate of photons

deposited in the storage cavity by the coherent build-up of the electric field in one storage cavity coherence time is given by [2, 3]:

$$\frac{dN_{HP}}{dt} = \frac{U/\omega_s}{T_1^s} = \frac{2E^2V}{\omega_s} \frac{\omega_s}{Q_s} = \frac{J_{DM}^2 Q_{DM} G V}{2m^2} \quad (11)$$

The storage volume V is $3.556 \times 3.81 \times 0.476 \text{ cm}^3 = 6.452 \text{ cm}^3$. Here, we assume the storage mode quality factor Q_s is larger than the dark matter quality factor Q_{DM} , such that the storage field displacement from dark photons can be approximated as a random walk. The geometric form factor G for our rectangular cavity used in the experiment is [2]:

$$G = \frac{1}{3} \frac{|\int dV E_z|^2}{\int dV |E_z|^2} = \frac{1}{3} \frac{2^6}{\pi^4} \quad (12)$$

The hidden-photon generated current is set by the density of dark matter in the galaxy $\rho_{DM} = 0.45 \text{ GeV}/\text{cm}^3 = 2\pi \times 9.67 \times 10^{19} \text{ GHz}/\text{cm}^3$ [4]:

$$J_{DM}^2 = 2\varepsilon^2 m^4 A'^2 = 2\varepsilon^2 m^2 \rho_{DM} \quad (13)$$

The photon deposition rate now becomes:

$$\frac{dN_{HP}}{dt} = \varepsilon^2 \rho_{DM} Q_{DM} G V \quad (14)$$

The photon rate and the integration time ($T_I(\Phi_{\text{ext}}) = T_1^s N_{\text{meas}} \in [1.29\text{s}, 1.38\text{s}]$) determines the total number of photons expected to be deposited in the storage cavity before each measurement:

$$N_{HP} = \frac{dN_{HP}}{dt} \times T_I(\Phi_{\text{ext}}) = \varepsilon^2 \rho_{DM} Q_{DM} G V T_1^s N_{\text{meas}} \quad (15)$$

By performing single-photon detection at each DC flux point, we perform a hidden-photon search across the storage tuning band. In our experiments, we perform $N_{\text{meas}} = 20,000$ measurements at each tuning point, and the observed counts N_{obs} are shown in supplementary text Figure 11. To convert the counted photons into an exclusion on the hidden-photon mixing angle, we first smooth the measured data to obtain a background of photon counts as a function of the cavity frequency. The smoothing is done by applying a Savitzky–Golay filter with polynomial order 4 and window length of 112, approximately the square root of the number of frequency bins 12695 (See Figure 11). This yields a photon background N_{back}

At each frequency, we calculate the 95% exclusion for the hidden-photon coupling parameter ε , which includes marginalizing over experimental parameter variation (See Table II). Additionally at some frequency points we encounter the situation where the background




FIG. 11. Background photon counts as a function of cavity frequency and the smoothed background for the purposes of background subtraction.

estimated counts are much higher than the observed counts. This leads the unphysical conclusion that there is no hidden-photon candidate that could produce the observed signal. We follow the procedure outlined in [5] and normalize the background + signal hypothesis cumulative distribution by the background only hypothesis distribution. At each frequency, we calculate the confidence level for a given hidden-photon coupling parameter ε , which includes marginalizing over experimental parameter variation as shown in Equation 16.

$$U(\omega, \varepsilon) = \int \prod_i d\Theta'_i F(\Theta_i - \Theta'_i, \sigma_{\Theta_i}) \left[1 - \frac{\sum_0^{N_{\text{obs}}} \binom{N_{\text{back}} + N_{\text{test}}}{N} \frac{e^{-(N_{\text{back}} + N_{\text{test}})}}{N!}}{\sum_0^{N_{\text{obs}}} \binom{N_{\text{back}}}{N} \frac{e^{-(N_{\text{back}})}}{N!}} \right] \quad (16)$$

$$N_{\text{test}} = \frac{\eta \varepsilon^2 \rho_{DM} Q_{DM} Q_c G V_c N_{\text{meas}}}{\omega}$$

$$F(\mu, \sigma) = \frac{e^{-\mu^2/2\sigma^2}}{\sqrt{2\pi}\sigma}$$

The parameters used to calculate the expected photon counts from a given dark matter candidate are measured and their uncertainties must be included in the analysis. Table II reports the uncertainties of all nuisance parameters considered in the analysis. Uncertainties in the transmon frequency are not critical, as they are encompassed within the detector efficiency uncertainty. When the transmon frequency drifts, the corresponding rotation pulses (π and $\pi/2$) will have slightly altered gate fidelities. This leads to minor variations in single-shot fidelity and thus in the detector efficiency. We weight these parameters each with a Gaussian distribution, $F(\mu, \sigma)$, and marginalize over them.

For a given frequency bin, we find the hidden-photon parameter ε below which the observed photon count is produced with more than 5% probability. Hidden-photon candidates which would produce the observed counts with less than 5% probability are excluded at the 95% confidence level. Additionally, we account for the line-shape of the dark matter [6] by convolving the dark matter distribution with the exclusion from each cavity tuning. This yields a family of curves exhibiting the

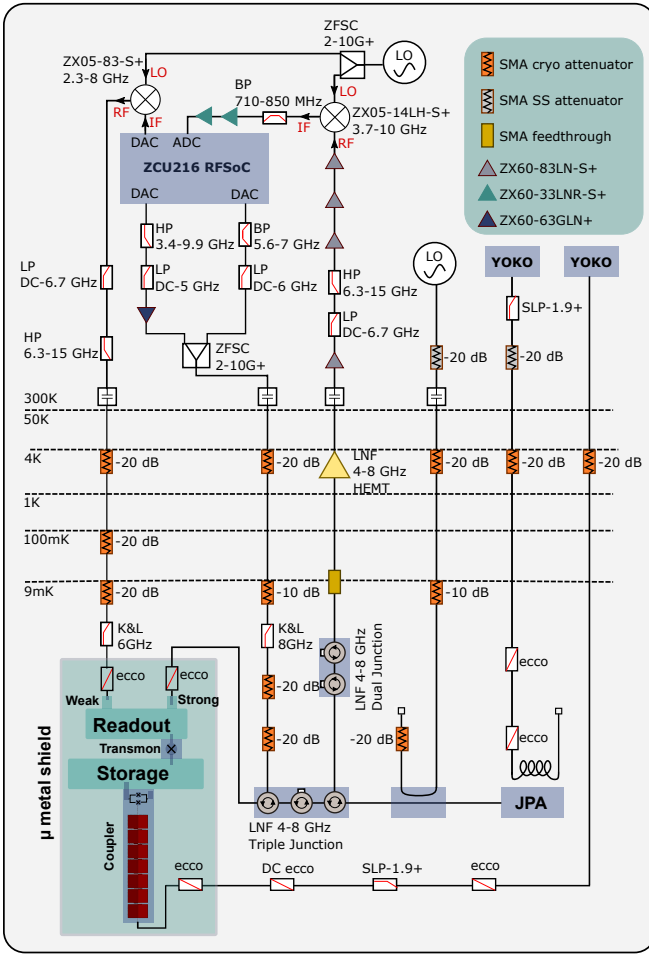


FIG. 12. Wiring diagram inside the dilution refrigerator and the room temperature measurement setup.

dark matter lineshape with maximum sensitivity when the dark matter mass is equal to the cavity frequency, see Figure 9. We now excluded dark matter candidates continuously across our tuning range as the minimum envelope of the family of curves coming from the above described procedure.

VII. MICROWAVE MEASUREMENT SETUP: CRYOGENIC AND ROOM TEMPERATURE WIRING

Fig. 12 shows the room and cryogenic temperature measurement chain. The device is mounted on the mixing chamber plate of the dilution fridge with a 15 mK base temperature. Two layers of μ -metal are used to shield the external magnetic fields. All RF input signals are generated through the ZCU216 RFSoc board and attenuated and thermalized at each temperature stage of the dilution fridge. Two DC sources (Yokogawa GS200) are used to bias the DC flux of the coupler and the Josephson Parametric Amplifiers (JPA). One Signal Core

SC5511A is used as the LO and split into two tones, one used for frequency up-conversion and sent to the device's weakly coupled port as the readout drive pulse; the other tone is used for frequency down-conversion of the readout signal and collected by the RFSoc ADC. The qubit and storage pulses are synthesized directly through different RFSoc DAC channels, combined at room temperature, and sent to the device's strongly coupled port. The readout signal goes through three circulators, then amplified by a JPA with 20dB gain, followed by two circulators and amplified with one LNF High-Electron-Mobility Transistor (HEMT) amplifier. One additional Signal Core SC5511A provides the JPA's 3-wave mixing modulation tones. DC flux tuning line is connected to the coupler chip through an SMA thread and wirebonds. All control lines contain SMA cryo attenuators at the base state, and quantum microwave eccosorbs in the μ -metal can.

-
- [1] 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, [Phys. Rev. A **76**, 042319 \(2007\)](#).
 - [2] A. V. Dixit, S. Chakram, K. He, A. Agrawal, R. K. Naik, D. I. Schuster, and A. Chou, Searching for dark matter with a superconducting qubit, [Phys. Rev. Lett. **126**, 141302 \(2021\)](#).
 - [3] S. Chaudhuri, P. W. Graham, K. Irwin, J. Mardon, S. Rajendran, and Y. Zhao, Radio for hidden-photon dark matter detection, [Phys. Rev. D **92**, 075012 \(2015\)](#).
 - [4] A. Caputo, A. J. Millar, C. A. J. O'Hare, and E. Vitagliano, Dark photon limits: A handbook, [Phys. Rev. D **104**, 095029 \(2021\)](#).
 - [5] A. Read, Modified frequentist analysis of search results (the cl_s method), CERN Report (2000).
 - [6] J. W. Foster, N. L. Rodd, and B. R. Safdi, Revealing the dark matter halo with axion direct detection, [Phys. Rev. D **97**, 123006 \(2018\)](#).