

Benchmarking dark matter search using a parity-check protocol with machine-learning optimized pulses

Yu-Han Chang¹, Ilya Moskalenko¹, Marko Kuzmanović¹, Ognjen Stanisavljević¹, Isak Björkman¹, David Díez-Ibáñez², Yikun Gu², Akash V. Dixit³, Igor G. Irastorza², and Gheorghe Sorin Paraoanu¹

¹*Department of Applied Physics and Institute Q, Aalto University, FI-00076 Aalto, Finland*

²*Centro de Astropartículas y Física de Altas Energías, Universidad de Zaragoza, 50009 Zaragoza, Spain*

³*Department of Physics, University of Chicago, Chicago, Illinois 60637, USA*



(Received 30 June 2026; accepted 14 August 2026; published 17 September 2026)

We report on an improved microwave detection protocol for dark matter candidates such as the axion and the dark photon. We employ a superconducting transmon qubit dispersively coupled to a double-cavity system, enabling quantum nondemolition measurements of the photon occupation in a relatively short-lived storage cavity. To reduce the experimental cycle time and enhance sensitivity for axion and dark photon searches, we operate this detector in a regime of increased qubit–cavity coupling, resulting in Stark shifts of 4.6 MHz. In this regime, conventional control pulses suffer from strong frequency-detuning sensitivity and photon-number-dependent errors. We address this limitation by implementing frequency-detuning-robust $\pi/2$ pulses (obtained by machine-learning optimization) that preserve high-fidelity qubit control over a bandwidth of approximately 20 MHz. We experimentally validate this protocol and demonstrate single-photon detection performance comparable to previous implementations, despite significantly reduced qubit coherence times and storage-cavity lifetimes. Using parity-based measurement sequences combined with a hidden Markov model (HMM) analysis, we achieve background rates on the order of $\mathcal{O}(20)$ Hz. In the absence of a magnetic field, we derive exclusion limits on the dark photon model for dark matter, reaching a sensitivity to the kinetic mixing angle of $\epsilon_{95\%} \sim 1 \times 10^{-14}$ at 5.051 GHz. These results establish machine-learning-robust control as a key enabler for faster, more scalable microwave quantum sensors for dark matter searches.

DOI: 10.1103/m92t-xwgt

I. INTRODUCTION

The nature of dark matter (DM) remains one of the central open questions in modern physics. Although it has not yet been directly observed, its existence is strongly supported by astrophysical and cosmological evidence [1–3]. Among the well-motivated candidates, two prominent solutions have garnered significant attention: the axion and the dark photon. The axion, a pseudoscalar particle, was originally proposed to solve the strong- CP problem in quantum chromodynamics (QCD) [4–6], and its detection in laboratory haloscopes relies on the inverse Primakoff effect, in which axions convert into photons in a strong magnetic field [7,8]. The dark photon, in contrast, is a massive vector boson that interacts with the Standard Model primarily through kinetic mixing, allowing for spontaneous conversion into ordinary photons even without an external magnetic field [9–12]. Both candidates can be

searched for using microwave cavity haloscopes that resonantly enhance the conversion rate near the candidate frequency.

Practical microwave haloscopes have historically read out the cavity field with cryogenic linear amplifiers placed at the front of the receiver chain [13–15]. However, linear amplification is fundamentally bounded by the standard quantum limit (SQL), which imposes a minimum added noise equivalent to an occupation $\bar{n}_{\text{SQL}} \approx 1$ [16]. The dark-matter-induced photon flux sits far below this floor: representative models predict an average cavity photon occupation $\bar{n} \lesssim 10^{-5}$ [17–20]. The only remaining handle on sensitivity requires prohibitively long integration times to achieve a sufficient signal-to-noise ratio (SNR), $t \propto \text{SNR}^2$ [21].

The actual system noise temperature T_{sys} realized in such experiments depends sensitively on where the first-stage amplifier sits in the cryogenic chain. With a high-electron-mobility transistor (HEMT) at the 4 K stage serving as the front end, microwave circuits between the cavity and the 4 K stage add appreciable thermal noise, so that the T_{sys} reaches ~ 7 K, well above the nominal HEMT specification [22]. Modern experiments such as ADMX, HAYSTAC, and CAPP therefore deploy a quantum-limited preamplifier

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

[e.g., a Josephson parametric amplifier (JPA) or traveling-wave parametric amplifier (TWPA)] directly at the mixing-chamber stage; the high gain of this cold first stage suppresses the contribution of the 4 K HEMT, bringing T_{sys} down to 0.59 ± 0.31 K and 0.385 ± 0.025 K [23,24]. Even in this case, however, the SQL floor of $\bar{n}_{\text{SQL}} \sim 1$ still bounds how fast such a receiver can scan the parameter space.

Single-photon counting with a superconducting qubit dispersively coupled to the storage cavity offers a route around the SQL bottleneck. By measuring only the cavity field amplitude—the photon number—and forgoing its conjugate phase, photon counting is no longer subject to the quantum back-action that sets the SQL [21]. Moreover, a quantum nondemolition (QND) measurement of the photon population leaves the cavity field intact, so the same photon can be interrogated many times before it decays. The noise floor is set by the residual cavity occupation $\bar{n}_c$ of the cavity [25].

A state-of-the-art realization of this protocol has reported a residual cavity occupation $\bar{n}_c \sim 10^{-3}$ [25], corresponding to an effective storage-mode temperature of a few tens of millikelvin. Reaching this floor relies on the QND character of the readout: each additional parity interrogation on the same cavity state suppresses detector-induced false counts exponentially with the number of repetitions, until the count rate is no longer limited by qubit readout errors but by the residual occupation $\bar{n}_c$ itself. Compared with SQL-limited receivers, the resulting noise floor lies several orders of magnitude below $\bar{n}_{\text{SQL}} \sim 1$, providing the scan-rate advantage that motivates microwave single-photon counting in haloscope searches. The leverage of the protocol, however, is contingent on extracting many parity measurements within a single storage-cavity photon lifetime; the usable number of repetitions is bounded by $N \lesssim T_{1,s}/T_t$, where $T_{1,s}$ is the photon lifetime of the storage cavity and T_t is the duration of a single parity-measurement cycle. Maximizing N therefore requires either a longer-lived cavity or a shorter measurement cycle [26–30].

While record photon lifetimes have been demonstrated in superconducting aluminum and niobium cavities operated in zero magnetic field [28,30–32], such conditions are not representative of axion haloscopes, which require static magnetic fields on the order of several Tesla. In those environments, superconducting cavities suffer substantial degradation of their quality factors due to vortex penetration and magnetic losses [33–35], so practical implementations rely on copper cavities or resonators with thin superconducting or dielectric coatings, with internal quality factors $Q_i \sim 10^4$ – 10^5 in experiments such as ADMX and HAYSTAC [13,14]. Even after dedicated efforts to coat cavities with type-II superconductors [36–40], the in-field quality factor remains restricted: Q_i has been reported to decrease from $\sim 6.6 \times 10^5$ at 4 T to $\sim 5.3 \times 10^5$ at 6 T for a

3.9 GHz cavity, corresponding to photon lifetimes of ~ 27 μs and ~ 22 μs , respectively [35]. Dark photon searches relax the strong-field requirement and can in principle reach longer $T_{1,s}$, but the practical operating regime is still set by the need to scan large frequency ranges with moderately Q -tunable cavities, leaving the same incentive to reduce T_t rather than to chase maximal $T_{1,s}$.

Consequently, rather than pursuing maximally long photon lifetimes under idealized zero-field conditions, our approach targets a realistic operating regime relevant to future axion searches planned, for example, within the RADES Collaboration. In this regime, improving detector performance hinges on minimizing the duration of each parity-measurement cycle, thereby maximizing the number of QND interrogations achievable within a finite photon lifetime.

Motivated by this constraint, we address the problem by optimizing quantum gate operations using neural ordinary differential equations (neural ODEs), a framework previously developed in our group [41]. We design and experimentally validate a robust $\pi/2$ pulse that maintains high fidelity ($\mathcal{F} > 99.9\%$) over a wide range of qubit–cavity detunings, enabling operation deep in the strongly Stark-shifted regime ($2\chi_{\text{qs}} \approx -4.6$ MHz) where conventional rectangular pulses fail. Our main improvement lies in the gate layer: robust neural-ODE pulses enable the parity-measurement cycle to be shortened to $T_t \sim 1$ μs while maintaining $\mathcal{F} > 99.9\%$ across a ± 20 MHz detuning window, lifting the ceiling on N under realistic, short- $T_{1,s}$ operating conditions representative of in-field and multi-photon occupation haloscope cavities [42].

II. METHODS

Our single-photon counter comprises three key elements: a storage cavity (ω_s), a readout cavity (ω_r), and a superconducting transmon qubit (ω_q). The roles of the two cavities are distinct: the readout cavity is coupled to the external environment for efficient signal extraction, while the storage cavity remains isolated, interacting only with the qubit to preserve photon lifetime. Further details of the experimental setup are provided in Appendix A.

The interaction between the transmon qubit and the cavity fields is described by the Jaynes-Cummings Hamiltonian in the dispersive regime ($|\omega_{r,s} - \omega_q| \gg g$). In this limit, the effective Hamiltonian takes the form

$$H = \frac{1}{2}\omega_q\hat{\sigma}_z + (\omega_r + \chi_{\text{qr}}\hat{\sigma}_z)\hat{a}^\dagger\hat{a} + (\omega_s + \chi_{\text{qs}}\hat{\sigma}_z)\hat{b}^\dagger\hat{b}, \quad (1)$$

where $\hat{a}$ ($\hat{b}$) are the annihilation operators of the readout (storage) cavity, and χ_{qr} (χ_{qs}) denote the corresponding dispersive shifts. This dispersive shift defines the photon-number-dependent qubit transition $\omega_{ge}(n) = \omega_{ge} + 2\chi_{\text{qs}}n$;

here $2\chi_{qs}$ is defined as the storage–transmon Stark shift that underlies the parity-mapping interval T_p used.

A. Parity-measurement protocol

The key element of the detection scheme is the QND measurement of the photon occupation of the storage cavity mediated by the transmon qubit. This enables repeated measurements of the cavity state, thereby reducing the impact of quantum errors in single measurements. The protocol is based on a Ramsey-type sequence that is sensitive to the storage–transmon Stark shift $2\chi_{qs}$, inspired by techniques originally developed in atomic physics. The pulse sequence and timing definitions are illustrated in Fig. 1.

Prior to each parity-measurement run, the system is initialized by a two-stage sequence shown in Fig. 1(a): an active reset prepares the qubit in $|g\rangle$, after which a single storage-cavity photon is prepared through the qubit by sequential π pulses on the ω_{ge} and ω_{ef} transitions that drive the qubit to $|f\rangle$, followed by an ω_{f0-g1} sideband drive that implements $|f, 0\rangle \rightarrow |g, 1\rangle$, the excitation to the storage cavity. The detailed calibration of these driving fields is given in Appendix B. A single measurement cycle, with total duration T_t , consists of five distinct stages:

- (1) *Superposition preparation*: The sequence begins with a $\pi/2$ pulse applied to the qubit with duration $T_{\pi/2}$, preparing the superposition state $(|g\rangle + |e\rangle)/\sqrt{2}$.




FIG. 1. Photon counter protocol. (a) Initialization sequence applied once before each parity-measurement run: it consists of an active reset that prepares the qubit in $|g\rangle$, followed by photon injection into the storage cavity via the qubit. (b) Energy-level diagram of the transmon–storage system. The red arrow shows the sideband transfer $|f, 0\rangle \rightarrow |g, 1\rangle$ used for photon injection. (c) Parity-measurement cycle, repeated N times after initialization. Each cycle of total duration T_t comprises two $X_{\pi/2}$ pulses of duration $T_{\pi/2}$ separated by a free-evolution interval T_p , followed by a qubit readout of duration T_r and a delay T_d allowing residual readout photons to decay.

- (2) *Parity mapping (free evolution)*: The qubit undergoes free evolution for a duration T_p , during which it accumulates a photon-number-dependent phase determined by the photon occupation n of the storage cavity. The interaction time T_p is calibrated such that the relative phase accumulation corresponds to a π rotation, mapping the photon-number parity onto the qubit phase. The calibration is determined in Sec. B 4 of Appendix B.
- (3) *State projection*: A second $\pi/2$ pulse of duration $T_{\pi/2}$ maps the accumulated phase information into a population difference in the qubit basis ($|g\rangle$ or $|e\rangle$).
- (4) *Readout*: The qubit state is interrogated via the readout cavity for a duration T_r .
- (5) *Post-measurement delay*: A delay T_d is introduced after each measurement to allow residual photons in the readout cavity to decay before the next cycle.

The total cycle time is given by $T_t = 2T_{\pi/2} + T_p + T_r + T_d$. Maximizing the repetition rate therefore requires minimizing T_t . For the storage lifetime $T_{1,s} = 20.3 \mu\text{s}$ measured in Sec. B 3 of Appendix B (Fig. 10), this cycle time sets the maximum number of parity cycles. We specifically target a reduction of the gate duration $T_{\pi/2}$ while preserving high gate fidelity under a large storage–transmon Stark shift $2\chi_{qs}$, enabled by optimal control techniques.

B. Robustness of $X_{\pi/2}$ gate

In the following characterization, the qubit drive frequency is scanned over a ± 30 MHz window centered on the bare qubit transition ω_{ge} , so that the frequency detuning $\Delta \equiv (\omega_{\text{drive}} - \omega_{ge})/2\pi$ plays the role of the photon-induced Stark shift $2n\chi_{qs}$ that the gate must tolerate during a multiphoton parity readout. Since the photon counting protocol relies exclusively on the $X_{\pi/2}$ rotation of the qubit, the fidelity of this gate directly determines the operational bandwidth of the detector. Specifically, as the photon number n in the storage cavity increases, the qubit transition frequency is shifted by $2n\chi_{qs}$. Conventional rectangular pulses fail once this photon-number-dependent shift exceeds the spectral bandwidth of the control pulse, thereby limiting the maximum detectable photon number.

To quantify the improvement enabled by pulse shaping, we first perform a controlled comparison using a fixed gate duration of $T_{\pi/2} = 100$ ns for all pulse types. Figure 2(a) shows the results of Ramsey interferometry measurements. For the rectangular pulse, rapid loss of fringe contrast is observed, with oscillations washing out beyond detunings of approximately ± 10 MHz. Compared to the rectangular pulse, the DRAG pulse has a significantly wider bandwidth; however, population loss can still be observed beyond ± 20 MHz. In contrast, the robust pulse maintains high-contrast Ramsey fringes over a substantially wider detuning range of up to ± 30 MHz, demonstrating




FIG. 2. Comparison of $X_{\pi/2}$ gate robustness. All pulses share a fixed duration of 100 ns. The frequency detuning on the horizontal axis is defined as $\Delta = (\omega_{\text{drive}} - \omega_{ge})/2\pi$ and is scanned over a detuning range of up to ± 30 MHz. (a)–(c) Ramsey fringe maps as a function of Δ and the interpulse delay, obtained with (a) a rectangular pulse, (b) a DRAG pulse, and (c) the neural-ODE robust pulse. The color scale encodes the measured qubit excited-state population. (d) Gate infidelity $1 - \mathcal{F}$ extracted from quantum process tomography (QPT) as a function of Δ , for the rectangular (blue), DRAG (red), and robust (green) pulses. The robust pulse exhibits a nearly flat profile below 10^{-3} across $-15 < \Delta < 10$ MHz, which is 2 orders of magnitude lower than the rectangular pulse.

enhanced resilience to frequency shifts arising purely from waveform optimization.

This robustness is further corroborated by quantum process tomography (QPT), shown in Fig. 2(d). The rectangular pulse (blue) is unable to implement a high-fidelity $X_{\pi/2}$ rotation across this band: its infidelity remains at the $\mathcal{O}(10^{-1})$ level throughout the scanned ± 30 MHz window, consistent with the loss of Ramsey contrast in Fig. 2(a). The DRAG pulse (red) exhibits the characteristic V-shaped profile expected from a finite-bandwidth correction, reaching $\sim 10^{-2}$ at zero detuning and degrading by an order of magnitude near the edges of the window. By contrast, the neural-ODE robust pulse (green) maintains a nearly flat infidelity below 10^{-3} across $-15 < \Delta < 10$ MHz, suppressing the gate error by roughly 2 orders of magnitude relative to the rectangular pulse over the operationally relevant Stark-shift range.

While the characterization in Fig. 2 uses $T_{\pi/2} = 100$ ns for a fair comparison against the rectangular and DRAG references, the actual photon-counting protocol run in Sec. III is implemented using the same dimensionless neural-ODE waveform at a shorter $T_{\pi/2} = 52$ ns. Both instances share the same on-resonance floor $1 - \mathcal{F} \sim 4 \times 10^{-4}$, and we directly verify the $1/T_{\pi/2}$ scaling of the operational detuning window (Sec. B 6 of Appendix B). This reduction in gate time provides two key advantages: it minimizes decoherence accumulated during the control operation and increases the effective control bandwidth to approximately ± 20 MHz, enabling faster measurement cycles and improved multiphoton compatibility [42].

C. Data processing

The parity-measurement protocol of Sec. II A yields, on each readout, a binary outcome $o \in \{g, e\}$ corresponding to the qubit state $|g\rangle$ or $|e\rangle$. A single execution of the protocol produces a parity sequence of $N = 21$ such outcomes, $\mathbf{o}^{(i)} = (o_1^{(i)}, o_2^{(i)}, \dots, o_N^{(i)})$, indexed by the trial number i . All results in this section are obtained from two independent datasets of $N_{\text{inj}} = N_{\text{bkg}} = 10,000$ trials each: an *injection run*, which loads a single photon into the storage cavity before the parity sequence begins, and a *no-injection run*, identical in every respect except that the injection pulse is omitted.

In the ideal case, a stored photon ($n = 1$) induces a π phase shift (“phase flip”) on the qubit in each parity cycle, producing an alternating outcome sequence ($e, g, e, g, \dots$), whose first measurement outcome is e ; the vacuum state ($n = 0$) preserves the qubit phase and yields a constant sequence ($g, g, g, \dots$). Representative traces of both cases are shown in Figs. 3(a) and 3(b). In practice, these signatures are distorted by qubit readout errors, finite gate fidelity, and—most importantly—stochastic decay of the storage photon over the lifetime $T_{1,s}$, so that measured sequences are neither perfectly alternating nor strictly constant.

To classify a noisy sequence $\mathbf{o}^{(i)}$ against the two photon occupation hypotheses, we employ a hidden Markov model (HMM) [43–45], whose transition and emission probabilities are fixed by the measured storage lifetime $T_{1,s}$, the readout assignment fidelity, and the per-cycle gate infidelity (Appendix C). The forward algorithm evaluates the evidence $P(\mathbf{o}^{(i)}|n_0)$ under each initial-state hypothesis $n_0 \in \{0, 1\}$; with an equal prior, Bayes’s rule gives the




FIG. 3. The HMM analysis of parity measurements. (a),(b) Two representative raw parity-measurement traces of length $N = 21$ after each cycle of the protocol. (a) An injection-run trace exhibiting the alternating signature of a single storage photon ($n = 1$). (b) A no-injection-run trace showing the constant pattern characteristic of the vacuum state ($n = 0$). (c),(d) Evolution of the HMM posterior probabilities $P(n_0 | o_{1:k})$ as the number of parity measurements k included in the analysis increases from 1 to $N = 21$, evaluated for the traces in panels (a) and (b), respectively.




FIG. 4. The likelihood-ratio distributions. Histograms of the per-sequence likelihood ratio λ_i [Eq. (2)] evaluated by the HMM forward algorithm on $N_{\text{inj}} = N_{\text{bkg}} = 10,000$ parity sequences each, for two independent datasets: the injection run (red) and the no-injection run (blue). The separation between the two distributions sets the single-shot discrimination performance of the detector. The vertical dashed line marks the operational threshold $\lambda_{\text{th}} = 4 \times 10^4$ adopted in Sec. III A and discussed in Appendix C.

per-sequence posterior $P(n_0|\mathbf{o}^{(i)})$, from which we define the per-sequence likelihood ratio:

$$\lambda_i \equiv \frac{P(n_0 = 1|\mathbf{o}^{(i)})}{P(n_0 = 0|\mathbf{o}^{(i)})} = \frac{P(\mathbf{o}^{(i)}|n_0 = 1)}{P(\mathbf{o}^{(i)}|n_0 = 0)}. \quad (2)$$

Trial i is classified as a photon detection event if $\lambda_i \geq \lambda_{\text{th}}$, and as a background count otherwise. Equation (2) is evaluated on the full N -outcome sequence; the partial forward posteriors $P(n_0|o_{1:k}^{(i)})$ for $k \leq N$ are intermediate quantities of the HMM, shown in Figs. 3(c) and 3(d).

Figure 4 shows the histogram of λ_i for the two datasets introduced above. The injection run (red) is strongly weighted toward large λ_i , and the no-injection run (blue) is weighted toward small λ_i ; the separation between the two distributions sets the single-shot discrimination capability of the photon counter. Thresholding λ_i at λ_{th} converts these per-trial decisions into the time-averaged detector metrics—the detection efficiency $\eta(\lambda_{\text{th}})$ and the background rate

$R_{\text{bkg}}(\lambda_{\text{th}})$ —reported in Sec. III A; the explicit mapping and the threshold selection criterion are derived in Appendix C.

III. BENCHMARK RESULTS

A. Detector characterization

To evaluate the detector performance, we conducted two independent measurement runs. Each run consists of $N_{\text{inj}} = 10,000$ injection-run sequences and $N_{\text{bkg}} = 10,000$ no-injection-run sequences. The injection run provides the detection efficiency $\eta(\lambda_{\text{th}})$, and the no-injection run provides the residual photon number $\bar{n}_c(\lambda_{\text{th}})$ and the equivalent background rate $R_{\text{bkg}}(\lambda_{\text{th}})$. The present cavity has no additional port directly addressing the storage mode (Appendix A); $\bar{n}_c$ is therefore obtained self-consistently from the no-injection run rather than from a displacement-based reference, and its value is cross-checked against the Bose-Einstein occupation expected from the transmon effective temperature (Appendix C).

The operational threshold λ_{th} is selected from the scan of Fig. 14: it is placed at the onset of the plateau in $R_{\text{bkg}}(\lambda_{\text{th}})$, where further increases of λ_{th} no longer suppress the background rate. Beyond this point, the residual triggers populate the high- λ_i tail of the no-injection distribution (Fig. 4, blue) and cannot be removed by postselection; they instead set the floor at which the detector is genuinely background-limited rather than signal-tail-limited. At $\lambda_{\text{th}} = 4 \times 10^4$, this tradeoff yields a detection efficiency $\eta \approx 25.7\%$ (Run 1) and 28.8% (Run 2), with bare background rates $R_{\text{bkg}}^{\text{bare}} \approx 20.0$ and 25.1 sec^{-1} corresponding to per-run residual cavity occupations $\bar{n}_c \approx 1.6 \times 10^{-3}$ and 1.8×10^{-3} [Eq. (C7)], respectively, integrated over the 0.2 s live time per run. The two runs agree on $\bar{n}_c$ to within statistical uncertainty at the per-run sample size.

The averaged residual occupation $\bar{n}_c \approx 1.7 \times 10^{-3}$ corresponds, via the Bose-Einstein relation [Eq. (C8)], to an effective storage-mode temperature $T_{\text{eff},s} \approx 38 \text{ mK}$, which lies between the mixing-chamber stage temperature $T_{\text{MC}} \approx 12 \text{ mK}$ and the transmon effective temperature $T_{\text{eff},q} \approx 72 \text{ mK}$. The decomposition of $\bar{n}_c$ into a transmon-mediated contribution $\bar{n}_c^q \sim 6.7 \times 10^{-4}$ [Eq. (C12), $\sim 39\%$ of the measured $\bar{n}_c$] and an extraneous bath contribution accounting for the remaining $\sim 61\%$ is derived in Sec. C3 of Appendix C.

TABLE I. Comparison of key performance metrics for microwave single-photon detectors based on repeated QND measurements. Bold text in Table I is used solely to highlight the row corresponding to our own experiment.

Detector	$T_{1,s}$ cavity (μs)	T_t cycle time (μs)	N (maximum number of parity- measurements)	Integration time (s)	Efficiency η (%)	BKG rate (sec^{-1})
[25]	546	10	30	8.33	40.3	1
[46] ^a	65	7.3	25	10.65	15	50
Run 1	20.3	1.092	21	0.2	25.7	20.0
Run 2	...	...	...	...	28.8	25.1

^aFlux-tunable storage cavity across a 22 MHz range.

Table I summarizes the performance of our prototype compared to state-of-the-art QND-based microwave photon detectors. The principal advantage of our system lies in the combination of a comparable detection efficiency ($\sim 27\%$) and high gate fidelity under large-frequency detuning, which enables operation with a higher photon number. Furthermore, the $T_{\pi/2} = 52$ ns robust gate allows a comparable number of measurement cycles ($N = 21$) within the relatively short cavity lifetime ($T_{1,s} = 20.3$ μ s), validating our methodology of minimizing gate duration to maximize N .

B. Exclusion limits on massive dark photons

Since the measurements were performed without an applied magnetic field, the dataset constrains massive dark photons (DP), whose kinetic mixing with the cavity mode is independent of any external field [11,12]. Resonant conversion occurs when the dark photon frequency matches the storage mode, $m_{\text{DP}}c^2 = \hbar\omega_s$, which for our cavity corresponds to $m_{\text{DP}} \approx 20.89$ μ eV ($\omega_s/2\pi = 5.051$ GHz).

Using the statistical framework of Appendix C [Eqs. (C16)–(C18)], we set the 95%-confidence-level (CL) upper limit on the kinetic mixing angle ϵ by combining the two independent datasets through Eq. (C19). The resulting exclusion curve is shown in Fig. 5; the per-run inputs ($N_{\text{obs}}, \eta, \bar{n}_{\text{sig}}^{95\%}$) and outputs ($\epsilon^{95\%}$) are listed in Table IV. The combined limit reaches $\epsilon^{95\%} \approx 1.04 \times 10^{-14}$ at $m_{\text{DP}} = 20.89$ μ eV, extending the search into parameter space previously unexplored by

comparable cavity-based microwave setups operating in this frequency band.

IV. DISCUSSION

The detector performance demonstrated in this work should be interpreted in the context of realistic operating conditions for axion haloscope experiments. While significantly longer photon lifetimes have been reported in superconducting cavities operated in zero magnetic field, such conditions are not compatible with axion detection, where multi-tesla static magnetic fields are required. Under these conditions, cavity quality factors are known to degrade substantially [33–35], resulting in photon lifetimes that are typically limited to the tens-of-microseconds regime.

From this perspective, the storage lifetime realized in our experiment ($T_{1,s} = 20.3$ μ s) is not an incidental limitation but is representative of the performance expected in practical haloscope implementations, including copper cavities and coated resonators operated in strong magnetic fields. In this regime, further increases in sensitivity cannot rely solely on extending cavity lifetimes, but instead require maximizing the information extracted per unit time.

For in-field axion operation, the robustness demonstrated is against quasistatic frequency detuning. Therefore, low-frequency fluctuations—e.g., those originating from trapped vortices—of the qubit transition frequency are absorbed within the calibrated window. Even if changing to a different field strength moves the operating point out of the calibrated range, reoptimizing the waveform is an




FIG. 5. Exclusion limit for dark photons. 95% CL limit derived from the quantum single-photon counter (blue), using Runs 1 and 2. Previous experimental results are shown for comparison. Plot adapted with current constraints (last update: January 2026). The collected limits, references, and plots are available in the GitHub repository: [48] (accessed on June 22, 2026) [12,47,48].

offline step whose cost is negligible against the axion search scanning time and leaves the hardware unchanged.

Our results demonstrate that this can be achieved by reducing the duration of the parity-measurement cycle rather than by pursuing marginal gains in $T_{1,s}$. By operating in a strong dispersive coupling regime with a Stark shift of $2\chi_{qs}/2\pi = -4.6$ MHz and employing neural-ODE-optimized $X_{\pi/2}$ pulses of duration $T_{\pi/2} = 52$ ns that retain $1 - \mathcal{F} < 10^{-3}$ across a ± 20 MHz detuning window (Sec. II B), we complete each parity-measurement cycle in $T_t = 1.092$ μ s and fit $N = 21$ QND interrogations within a single storage lifetime. This shifts the operating bottleneck away from cavity engineering and onto the control layer, where the present platform has substantial remaining headroom. The efficiency reported here ($\eta \approx 27\%$), though competitive, is bounded by the qubit readout T_r , which dominates T_t ; improving the readout thus raises efficiency and shortens the cycle at once. Established approaches, including Purcell-filter-assisted fast readout [49,50] or optimized short-readout pulses for suppressing the measurement-induced state transition (MIST) [51,52], can address this limitation.

The principal advantage of QND photon counting over linear amplifier readout is the elimination of the standard quantum limit (SQL) as a noise floor. Phase-preserving linear amplification carries an unavoidable added noise of $\bar{n}_{\text{SQL}} \sim 1$ photon per cavity mode [16], against which any dark matter signal at $\bar{n} \lesssim 10^{-5}$ must be discriminated by long integration. The residual cavity occupation of our detector at the operating threshold, $\bar{n}_c \approx 1.7 \times 10^{-3}$ [Eq. (C7) and Sec. III A], sits roughly 3 orders of magnitude below this floor:

$$\frac{\bar{n}_{\text{SQL}}}{\bar{n}_c} \approx \frac{1}{1.7 \times 10^{-3}} \sim 6 \times 10^2. \quad (3)$$

The decomposition of $\bar{n}_c$ carried out in Sec. C 3 of Appendix C clarifies the engineering response required to push the detector deeper into the sub-SQL regime. The transmon-mediated dispersive dressing contribution $\bar{n}_c^q \approx 6.7 \times 10^{-4}$ [Eq. (C12)] accounts for $\sim 39\%$ of the measured $\bar{n}_c$ and is set by the geometric ratio $(g_{qs}/\Delta)^2 \simeq 1.4 \times 10^{-2}$ together with the residual transmon population $\bar{n}_q \approx 4.8 \times 10^{-2}$. Reducing this component therefore acts on the qubit side: improved infrared filtering at the qubit chip to suppress quasiparticle generation, and—in field-tolerant variants—gap engineering of the superconducting electrode that hosts the Josephson junction. Increasing $|\Delta|$ provides a complementary handle, but only at the cost of reducing χ_{qs} and thereby the parity-measurement contrast, so the present operating point already trades a modest dressing penalty for a strong $2\chi_{qs}/2\pi = -4.6$ MHz Stark shift.

The remaining $\sim 61\%$ of $\bar{n}_c$ is attributed to extraneous thermal channels not mediated by the qubit (Sec. C 3 of

Appendix C). The dominant suppression levers for this component are environmental rather than qubit-level: additional infrared filters and cryogenic attenuators on the input lines, improved magnetic and electromagnetic shielding of the sample stage. The fact that the two contributions are of comparable magnitude implies that neither lever, applied in isolation, can reduce $\bar{n}_c$ by an order of magnitude; a coordinated suppression on both fronts is required, and this is the explicit subject of follow-up work.

Importantly, the gate-layer strategy underlying this performance remains effective in the presence of large photon-number-dependent frequency shifts, where conventional rectangular or DRAG-based pulses suffer severe fidelity degradation (Fig. 2). The use of neural-ODE-optimized pulses therefore provides a practical route to high-bandwidth, high-fidelity operation that is compatible with the constraints imposed by strong magnetic fields. Taken together, these results indicate that optimizing quantum control and measurement speed—rather than maximizing cavity lifetime under idealized conditions—is a viable and scalable pathway toward quantum-enhanced axion and dark photon searches. The approach demonstrated here is directly applicable to future haloscope experiments operating in realistic-magnetic-field environments and provides a complementary direction to ongoing efforts focused on cavity materials and surface treatments.

V. SUMMARY

In this work, we have demonstrated a prototype single-photon counter for dark matter searches in the microwave domain, based on a superconducting transmon qubit coupled to a 3D aluminum cavity system with storage lifetime $T_{1,s} = 20.3$ μ s. By integrating neural-ODE-optimized control pulses of duration $T_{\pi/2} = 52$ ns, we implemented a robust QND parity-measurement protocol that retains gate infidelity $1 - \mathcal{F} < 10^{-3}$ across a ± 20 MHz detuning window. This broad operational bandwidth provides two benefits: it accommodates the large Stark shifts associated with multiphoton storage states ($2\chi_{qs}/2\pi = 4.6$ MHz), thereby extending the detector's dynamic range beyond the single-photon regime, and it shortens the parity-measurement cycle to $T_t = 1.092$ μ s, fitting $N = 21$ QND interrogations within a single storage lifetime at a detection efficiency of $\eta \approx 27\%$.

Operating in the absence of an applied magnetic field, the prototype exhibits an averaged residual cavity occupation $\bar{n}_c \approx 1.7 \times 10^{-3}$ per cycle, corresponding to the bare background rate $R_{\text{bkg}}^{\text{bare}} \approx 22$ sec^{-1} and an effective storage-mode temperature $T_{\text{eff},s} \approx 38$ mK—roughly 3 orders of magnitude below the standard-quantum-limit noise floor of $\bar{n}_{\text{SQL}} \sim 1$ photon per mode. Combining the two independent datasets, we set a 95% CL upper limit on the kinetic-mixing parameter of $\epsilon^{95\%} = 1.04 \times 10^{-14}$ at the storage-cavity frequency $m_{\text{DPC}} c^2 = \hbar\omega_s$, corresponding to a dark photon

mass of $20.89 \mu\text{eV}$ ($\omega_s/2\pi = 5.051 \text{ GHz}$). Although obtained at zero field, the same protocol applies to axion haloscope conditions; in particular, the storage Compton frequency falls within the axion mass window currently targeted by tunable haloscope programs, establishing this platform as a building block for QCD-axion searches at masses near $20 \mu\text{eV}$.

A microscopic decomposition of $\bar{n}_c$ attributes $\sim 39\%$ of the residual occupation to a transmon-mediated dispersive dressing contribution and the remaining $\sim 61\%$ to extraneous thermal channels, placing the detector in an intermediate regime in which the two routes contribute comparably. The remaining sensitivity headroom is therefore addressable simultaneously at the qubit and the environmental layer of the apparatus rather than through colder operating temperatures, and future work will pursue both: improved infrared filtering at the qubit chip to suppress $\bar{n}_c^q$, together with additional cryogenic attenuation and electromagnetic shielding to suppress the extraneous contribution. Shortening the qubit readout time $T_r = 800 \text{ ns}$ —currently the dominant component of T_t —will further increase the number of QND interrogations achievable per storage lifetime. Together with the scalability and field tolerance of the optimized control scheme, this platform provides a promising pathway toward scanning haloscopes with sensitivity approaching the QCD-axion model predictions.

ACKNOWLEDGMENTS

Our research has been performed as part of the RADES Collaboration, and in particular we are grateful to Laura Seguí, Saiyd Ahyoune, and Jordi Miralda-Escudé for stimulating interactions and comments at different stages of this project. This work is supported by ERC Synergy Grant DarkQuantum (Grant Agreement ID: 101118911 [53]), under the program Horizon.1.1—European Research Council (ERC). This publication is also part of QuantERA project QRADES funded by Agencia Estatal de Investigación Project No. PID2019–108122 GB-C31 (financed by MCIU/AEI/10.13039/501100011033), by Research Council of Finland Project No. 36133, and cofunded by the European Union. The Zaragoza team acknowledges additional support from the Spanish Agencia Estatal de Investigación under Grant No. PID2019–108122 GB-C31 (funded by MCIN/AEI/10.13039/501100011033), as well as support from the “European Union NextGenerationEU/PRTR” (Planes complementarios, Programa de Astrofísica y Física de Altas Energías). The Aalto team thanks the Jane and Aatos Erkkö Foundation for funding under the doctoral program managed by InstituteQ. We would like to acknowledge the use of Servicio General de Apoyo a la Investigación-SAI, Universidad de Zaragoza. We acknowledge the provision of facilities by Aalto University at OtaNano—Micronova Nanofabrication Centre and Low Temperature Laboratory.

We thank Ciaran O’Hare for his contribution to integrating experimental data from research communities.

DATA AVAILABILITY

There are no publicly available research data or software supporting this manuscript. Requests for further information or data should be sent to the authors.

APPENDIX A: EXPERIMENTAL SETUP

1. The cavity

The 3D cavity was machined from a solid block of aluminum (alloy AW-6082; nominal purity 94%–96%) and consists of two physically distinct chambers: a readout cavity and a storage cavity. The dimensions of both cavities are summarized in Table II. A narrow trench milled between the two chambers accommodates a superconducting transmon qubit that bridges them; no direct electromagnetic coupling between the cavities was implemented. Figure 6 shows one half of the assembly with the transmon installed. The sapphire substrate hosting the qubit was mechanically secured within the trench using a small amount of indium. The storage cavity is designed to maximize photon lifetime, with a fundamental mode at $\omega_s = 2\pi \times 5.051 \text{ GHz}$ and a measured quality factor of $Q_s = 6.35 \times 10^5$, corresponding to a photon decay time of $T_{1,s} \approx 20 \mu\text{s}$. The readout cavity operates at $\omega_r = 2\pi \times 6.867 \text{ GHz}$ and is strongly coupled to the external environment. All control pulses, readout tones, and sideband drives are delivered exclusively through the readout cavity via a strongly coupled port.

2. The transmon qubit

The transmon qubit was designed to couple dispersively to both the readout and storage cavities. Its geometry consists of two large antenna pads connected by a Josephson junction, as shown in Fig. 7. To allow post-fabrication selection and optimization, eight transmon designs with varying pad dimensions were fabricated on a single wafer. Electromagnetic simulations predicted an




FIG. 6. Photograph of one half of the aluminum 3D cavity assembly. The superconducting transmon qubit (black structure) is fabricated on a sapphire substrate and mounted in the central trench, bridging the readout cavity (top) and the storage cavity (bottom).

TABLE II. Geometrical dimensions of the readout and storage cavities.

	X (mm)	Y (mm)	Z (mm)
Readout cavity (6.8 GHz)	26.6	5	40
Storage cavity (5 GHz)	45	5	40

anharmonicity of approximately 200 MHz and a storage-transmon Stark shift $2\chi_{qs}$ on the order of 3 MHz. Experimentally, we measured a larger anharmonicity of -272 MHz and a Stark shift of $2\chi_{qs}/2\pi = -4.6$ MHz, consistent with strong coupling to the storage cavity.

Fabrication proceeded in two steps. First, aluminum (Al) antenna pads were patterned using optical lithography and etched via reactive ion etching. Subsequently, the Josephson junction was fabricated using electron-beam lithography and the Dolan bridge technique, yielding an Al/AIO_x/Al trilayer junction. The resulting junctions exhibited critical currents corresponding to Josephson inductances of $L_J \sim 12$ nH and junction capacitances of $C_J \sim 3.5$ pF. After dicing, individual transmon chips were mounted into the cavity assembly. Further fabrication details are provided in Ref. [54].

3. The wiring diagram

The cavity assembly was thermally anchored to a copper plate mounted on the mixing chamber stage of a Bluefors dilution refrigerator. The device was enclosed in a shielded can to protect it from spurious magnetic noise. All measurements were performed at a base temperature of approximately 10 mK. A schematic of the cryogenic wiring and room-temperature control electronics is shown in Fig. 8 (the copper plate is not shown).

Microwave control and readout pulses were generated at room temperature using a Quantum Machines OPX+/Octave system, providing waveform synthesis over a 350 MHz bandwidth. The combined signals were attenuated at each stage of the dilution refrigerator before reaching the cavity, minimizing thermal noise.

The readout signal was first amplified at the mixing chamber stage using a traveling-wave parametric amplifier (TWPA).




FIG. 7. Design of the superconducting transmon qubit. The final dimensions used in the prototype were $L1 = 1$ mm, $H1 = 0.3$ mm, $L2 = 0.15$ mm, $H2 = 0.02$ mm. The Josephson junction is located at the center of the structure, connecting the two antenna pads.




FIG. 8. Schematic of the cryogenic wiring and control setup. The sample is mounted at the mixing chamber stage (≈ 10 mK). Input signals are attenuated at multiple temperature stages to suppress thermal noise. Output signals are amplified by a TWPA at the base temperature, a HEMT amplifier at the 4 K stage and a Narda-Miteq amplifier at room temperature before being digitized by the OPX system.

(TWPA) from VTT Technical Research Centre of Finland, followed by further amplification at the 4 K stage using a high-electron-mobility transistor (HEMT) amplifier from Low Noise Factory. After exiting the refrigerator, the signal was amplified once more by a Narda-Miteq LNA-40-04000800-07-10P amplifier before being demodulated and digitized by the OPX+/Octave system.

APPENDIX B: CALIBRATION OF DRIVING PULSES

1. Basic qubit characterization

We experimentally characterized the transition frequencies of the transmon qubit up to the fourth excited state, from ω_{ge} to ω_{fh} . The detailed parameters are summarized in Table III. Each transition serves a distinct experimental objective.

The fundamental transition ω_{ge} was used to calibrate X_π and $X_{\pi/2}$ control pulses, which constitute the basic building blocks of the parity-measurement protocol. The amplitude and waveform of these pulses are calibrated in Sec. B 4 of

TABLE III. Summary of experimental device parameters.

Device parameter	Symbol	Value
Qubit frequency $ g\rangle \rightarrow e\rangle$	ω_{ge}	$2\pi \times 4.629$ GHz
Qubit frequency $ e\rangle \rightarrow f\rangle$	ω_{ef}	$2\pi \times 4.357$ GHz
Qubit frequency $ f\rangle \rightarrow h\rangle$	ω_{fh}	$2\pi \times 4.055$ GHz
Qubit anharmonicity	α	$-2\pi \times 272$ MHz
Qubit decay time	$T_{1,q}$	16.7 μ s
Qubit decoherence time	$T_{2,q}$	6.2 μ s
Qubit temperature	$T_{\text{eff},q}$	72 mK
Storage frequency	ω_s	$2\pi \times 5.051$ GHz
Storage decay time	$T_{1,s}$	20.3 μ s
Storage-qubit Stark shift	$2\chi_{qs}$	$-2\pi \times 4.6$ MHz
Sideband transition	ω_{f0-g1}	$2\pi \times 3.935$ GHz
Readout frequency	ω_r	$2\pi \times 6.867$ GHz
Readout fidelity $ g\rangle$ and $ e\rangle$	$\mathcal{F}$	0.95739
Readout pulse length	T_r	800 ns
$X_{\pi/2}$ gate pulse length	T_g	52 ns
Delay time between $X_{\pi/2}$	$T_{\pi/2}$	88 ns
Waiting time between cycle	T_d	100 ns
Parity sequence total length	T_t	1.092 μ s

this appendix. At this frequency, we measured a qubit energy relaxation time of ($T_{1,q} = 16.7$ μ s) and a coherence time of ($T_{2,q} = 6.2$ μ s).

The second transition ω_{ef} was characterized to extract the qubit anharmonicity $\alpha = \omega_{ef} - \omega_{ge}$, a vital parameter for optimizing the waveform of the robust $X_{\pi/2}$ pulse. In addition, access to the $|f\rangle$ level is essential for photon injection into the storage cavity via the ω_{f0-g1} sideband transition; the calibration of this transition is described in Sec. B 2 of this appendix.

Finally, the higher-order transition ω_{fh} was measured to verify sufficient spectral separation from all driven tones, ensuring that no spurious excitations occur during device operation. We also measure an effective qubit temperature of 72 mK [55].

2. Photon injection

As discussed in Appendix A, the storage cavity is not directly addressable by an external drive. Photon injection is therefore mediated by the transmon qubit, which couples dispersively to the storage mode. The injection protocol proceeds as follows. Starting from the ground state $|g, n\rangle$, the qubit is sequentially excited to the $|f\rangle$ state by applying resonant π pulses on the ω_{ge} and ω_{ef} transitions. Subsequently, a driven sideband transition at frequency ω_{f0-g1} coherently transfers the excitation from the qubit to the storage cavity, implementing the state transfer $|f, n\rangle \rightarrow |g, n+1\rangle$, thereby increasing the photon number in the storage cavity by 1.

The sideband transition ω_{f0-g1} is not a first-order process and therefore requires relatively high drive power to

achieve efficient population transfer. To this end, a +33 dB power amplifier (Mini-Circuits ZVE-3W-83+) was used to boost the drive amplitude. The measurement in Fig. 9(a) spans a detuning range of ± 5 MHz around $\omega_{f0-g1} = 2\pi \times 3.935$ GHz. At high drive amplitudes, a redshift of the effective transition frequency is observed, which we attribute to an ac Stark shift induced by the strong drive. Owing to the spectral proximity of ω_{f0-g1} to the qubit transition frequencies, a band-reject filter (Wainwright WTRCJV6-4000-5000-3-22-30SS) was installed after the amplifier to suppress spectral components near ω_{ge} and ω_{ef} . This filtering prevents unintended excitation of higher qubit transitions during photon injection. The sideband frequency satisfies the relation $\omega_{f0-g1} = \omega_{ge} + \omega_{ef} - \omega_s$, from which we independently confirm the storage cavity frequency $\omega_s = 5.051$ GHz.

Ramsey interferometry was utilized to infer the photon number in the storage cavity via the photon-number-dependent frequency shift of the transmon qubit. In the dispersive regime, the qubit transition frequency depends on the cavity photon occupation n according to $\omega_{ge}(n) = \omega_{ge} + 2\chi_{qs}n$, where $2\chi_{qs}$ is the storage-qubit Stark shift. Figures 9(b) and 9(c) show representative Ramsey interference fringes measured with the storage cavity prepared in the $n = 0$ and $n = 1$ photon states, respectively. The clear frequency offset between the two fringe patterns directly reflects the dispersive coupling between the qubit and the storage mode. Fitting the center frequency of each Ramsey fringe pattern yields a separation of $2\chi_{qs}/2\pi = -4.6$ MHz listed in Table III.




FIG. 9. Calibration of the $f0-g1$ sideband and the storage-transmon Stark shift. (a) Two-dimensional map of the transmon $|f, 0\rangle \rightarrow |g, 1\rangle$ sideband response: the horizontal axis is the drive detuning from $\omega_{f0-g1} = 2\pi \times 3.935$ GHz, and the vertical axis is the driving amplitude at the AWG output. The chevron pattern at small amplitude collapses into a single sideband resonance whose center redshifts at large amplitude due to an ac Stark effect. (b),(c) Ramsey interferograms of the transmon ω_{ge} transition with the storage cavity prepared in (b) the vacuum ($n = 0$) and (c) the single-photon Fock state ($n = 1$). Horizontal axis: probe detuning from $\omega_{ge} = 2\pi \times 4.629$ GHz. Vertical axis: interpulse delay between the two $\pi/2$ pulses. Color encodes the measured excited-state population P_e (scale at right). The horizontal shift between (b) and (c) yields $2\chi_{qs}/2\pi = -4.6$ MHz.

3. Storage measurement

Following photon injection, we characterized the single-photon lifetime of the storage cavity. As shown in Fig. 10, the photon population was monitored indirectly by tracking the photon-number-resolved qubit transition $\omega_{ge}(n=1)$ using spectroscopy. To spectrally resolve the $\omega_{ge}(n=0)$ and $\omega_{ge}(n=1)$ transitions, we employed a long-duration π pulse (720 ns), thereby reducing the spectral bandwidth of the excitation to well below the -4.6 MHz Stark splitting calibrated in Sec. B 2 above. The population associated with $n=1$ was observed to decay monotonically with increasing delay time, accompanied by a corresponding increase in the $n=0$ population. An exponential fit to the decay curve yields a storage photon lifetime of $T_{1,s} = 20.3 \mu\text{s}$. This lifetime sets the upper bound on the number of parity cycles that can be performed within the photon coherence window, $N \sim T_{1,s}/T_t$ (Sec. I).

4. Drive pulse calibration

The parity-measurement protocol, comprising two $X_{\pi/2}$ gates separated by a free-evolution interval T_p , forms the core of the photon detection scheme. Accurate calibration of both the gate amplitude and the timing parameters is therefore essential. The amplitude of the robust $X_{\pi/2}$ pulse was calibrated using a pseudo-identity sequence, $(X_{\pi/2})^{4n} = (-1)^n I$, where the pulse amplitude was fine-tuned to maximize the return probability to the qubit ground state. This procedure is particularly sensitive to systematic over- or under-rotations and provides a reliable calibration of the pulse area.

Following amplitude calibration, we determine the free-evolution interval T_p used in stage 2 of the parity-




FIG. 10. Storage-cavity photon lifetime $T_{1,s}$. (a) Photon-number-resolved spectroscopy of the transmon ω_{ge} transition as a function of delay after photon injection. Horizontal axis: probe detuning from $\omega_{ge} = 2\pi \times 4.629$ GHz. Vertical axis: delay. Color: excited-state population P_e measured with a long (720 ns) π -probe (scale at right). The two bands at detunings 0 and $2\chi_{qs}/2\pi = -4.6$ MHz correspond to the storage Fock states $n=0$ and $n=1$, respectively. (b)–(d) Horizontal slices through (a) at delays 0, 14, and 100 μs. The $n=1$ peak decays monotonically with delay while the $n=0$ peak grows correspondingly; an exponential fit to the $n=1$ population yields $T_{1,s} = 20.3 \mu\text{s}$.




FIG. 11. Calibration of the parity-mapping delay T_p . A Ramsey-type scan is performed using the parity-measurement sequence [two $X_{\pi/2}$ pulses separated by t , see Fig. 1(c)] with the storage cavity prepared in the single-photon Fock state ($n=1$) using the protocol of Sec. B 2 in Appendix B. The horizontal axis is the interpulse delay t , and the vertical axis is the excited-state population P_e measured after the second $X_{\pi/2}$ pulse. Black markers are data; the red curve is a fit of the form $P_e(\tau) = A + B \cos(2\pi f t + \phi)$ with $f = 4.5$ MHz, consistent with $|2\chi_{qs}|/2\pi = 4.6$ MHz from Sec. B 2 in Appendix B. The first maximum, $t = T_p = 88$ ns, is selected as the parity-mapping interval.

measurement cycle [Sec. II A and Fig. 1(c)]. With the storage cavity prepared in $n=1$ via the protocol of Sec. B 2 of this appendix, we perform a Ramsey-type scan of the parity sequence as a function of delay (Fig. 11). The qubit population exhibits periodic oscillations, with the first maximum occurring at $T_p = 88$ ns. This value is slightly shorter than the theoretical estimate $T_p = 1/(2\chi_{qs}) \approx 109$ ns. The discrepancy arises from additional phase accumulation during the finite duration ($T_{\pi/2} = 52$ ns) of the robust pulses, which is not captured in the instantaneous pulse approximation.

5. Quantum process tomography

Quantum process tomography (QPT) was employed as a diagnostic tool to characterize the experimentally implemented control gates, allowing for a quantitative evaluation of their fidelity with respect to the ideal target operation and a direct comparison between the standard rectangular $X_{\pi/2}$ and neural-ODE-optimized robust $X_{\pi/2}$ pulses. The QPT protocol begins by initializing the transmon qubit in the ground state $|g\rangle$. Four linearly independent input states are then prepared using the gate set $\{I, X_{\pi/2}, Y_{\pi/2}, X_{\pi}\}$, after which the gate under test (GUT) is applied. Finally, the state populations are measured in different bases: $\{|0\rangle\langle 0|, |X\rangle\langle X|, |Y\rangle\langle Y|, |1\rangle\langle 1|\}$.

To ensure a fair comparison, both pulse types were characterized using the same gate duration of




FIG. 12. Quantum process tomography (QPT) of the robust $X_{\pi/2}$ gate under frequency detuning. The panels display the measured populations of the ground state (P_g , blue), excited state (P_e , red), and second excited state (P_f , green) as a function of frequency detuning from -30 to 30 MHz. The rows represent four distinct input states $\{|0\rangle\langle 0|, |X\rangle\langle X|, |Y\rangle\langle Y|, |1\rangle\langle 1|\}$, while the columns correspond to the measurement bases Z, X, and Y. The characteristic oscillations observed in the populations arise from the phase accumulation induced by the frequency detuning during the robust $X_{\pi/2}$ gate operation.

$T_{\pi/2} = 100$ ns, with measurements performed across a detuning range of $\Delta = (\omega_{\text{drive}} - \omega_{\text{ge}})/2\pi = \pm 30$ MHz. At finite detuning, phase accumulation during the gate results in rapid oscillations of the reconstructed process matrix, obscuring the intrinsic gate fidelity, as shown in Fig. 12. To remove this trivial effect, we apply a rotating-frame correction scheme.

First, to separate coherent unitary errors from nonunitary decoherence processes, we fit the experimentally reconstructed process matrix χ_{exp} to the closest unitary process matrix, χ_{fit} . Any single-qubit unitary operation U can be expressed in terms of Pauli operators with coefficients c_k as

$$U = c_0 \sigma_0 - i \sum_{k \in \{x,y,z\}} c_k \sigma_k, \quad (\text{B1})$$

where the diagonal elements of the corresponding process matrix are related to these coefficients by $\chi_{kk} = |c_k|^2$.

Detuning-induced phase accumulation manifests as a rotation in the X/Y subspace of these coefficients. The accumulated phase angle is given by

$$\beta = 2\pi\delta(T_{\pi/2} + \Delta T), \quad (\text{B2})$$

where ΔT accounts for the latency in the control electronics associated with oscillator frequency switching. To recover the intrinsic gate behavior, we apply an inverse rotation to the transverse coefficients:

$$c_x^{\text{corr}} + ic_y^{\text{corr}} = (c_x + ic_y)e^{-i\beta}. \quad (\text{B3})$$

This operation effectively transforms the data into a frame rotating at the detuned frequency, canceling the trivial z -rotation induced by detuning. The gate performance is finally quantified by the infidelity ($\mathcal{I}$), calculated using the corrected process matrix:

$$\mathcal{I} = 1 - \text{Tr}(\chi_{\text{fit}}^{\text{corr}} \cdot \chi_{\text{target}}), \quad (\text{B4})$$

where χ_{target} is the ideal process matrix for the $X_{\pi/2}$ gate. More details about the mathematical procedure can be found in Ref. [41].

6. Pulse-duration scaling of the neural-ODE waveform

The neural-ODE optimization is performed in a dimensionless time coordinate, so that the optimized in-phase and quadrature envelopes $\Omega_{x,y}$ [Fig. 13(a)] define a family of waveforms parametrized by a single physical duration $T_{\pi/2}$. A new physical pulse is obtained by replaying the same




FIG. 13. Pulse-duration scaling of the neural-ODE robust $X_{\pi/2}$ waveform. (a) In-phase $I = \Omega_x$ (red) and quadrature $Q = \Omega_y$ (blue) components of the optimized envelope plotted as a function of the dimensionless time coordinate t/T used during neural-ODE optimization; the physical pulse duration $T_{\pi/2}$ is obtained by mapping this coordinate onto an interval, so the same envelope can be replayed at any $T_{\pi/2}$ without reoptimization. (b) Gate infidelity $1 - \mathcal{F}$ extracted from quantum process tomography (see Sec. B 5 of Appendix B) as a function of the frequency detuning $\Delta = (\omega_{\text{drive}} - \omega_{\text{ge}})/2\pi$, for two instantiations of the envelope shown in (a): $T_{\pi/2} = 100$ ns (red solid) and $T_{\pi/2} = 60$ ns (blue solid). Both pulses reach an on-resonance floor $1 - \mathcal{F} \sim 4 \times 10^{-4}$. The red dashed curve shows the $T_{\pi/2} = 100$ ns trace with its detuning axis rescaled by $100/60 \approx 1.67$; its near-perfect overlap with the blue trace is a direct experimental verification that the operational detuning window scales as $1/T_{\pi/2}$.

dimensionless envelope at different sample points; no reoptimization is required [41].

To verify this scaling experimentally, we performed QPT on two instantiations of the same dimensionless envelope, $T_{\pi/2} = 100$ ns and $T_{\pi/2} = 60$ ns, processed with the rotating-frame correction of Sec. V above. The extracted infidelity is shown in Fig. 13(b). Both curves reach the same on-resonance floor $1 - \mathcal{F} \sim 4 \times 10^{-4}$, confirming that the rescaling preserves the on-resonance gate fidelity. Rescaling the detuning axis of the $T_{\pi/2} = 100$ ns trace by $T_{\pi/2}^{(100)} / T_{\pi/2}^{(60)} = 100/60 \approx 1.67$ [red dashed curve in Fig. 13(b)] brings it into near-perfect agreement with the directly measured $T_{\pi/2} = 60$ ns trace across the entire ± 30 MHz measurement window. This collapse provides a verification that the operational detuning window in physical units obeys the predicted $1/T_{\pi/2}$ scaling, and that the dimensionless waveform inherited from the original neural-ODE optimization remains robust when rescaled to the duration actually used in the parity-measurement protocol.

The lower bound on $T_{\pi/2}$ in our setup is set by the time resolution of Quantum Machines OPX+: with a sample rate of 1 GS/s, the $T_{\pi/2} = 52$ ns instance is rendered with 52 samples across the pulse, which is sufficient to resolve the waveform of the optimized envelope. Below 52 samples, the rescaled waveform is no longer faithfully reproduced; this hardware limit, not the optimization itself, fixes the shortest $T_{\pi/2}$ accessible in the present work.

APPENDIX C: DATA ANALYSIS AND STATISTICAL FRAMEWORK

1. Notation

Two distinct counts both denoted by a capital N appear in the analysis of this appendix. We adopt the convention used throughout the main text: $N = 21$ is the number of consecutive parity measurements composing one sequence [a single execution of the protocol of Fig. 1(c) after one initialization], fixed by the cycle time $T_t = 1.092$ μ s. $N_{\text{meas}} = N_{\text{inj}} = N_{\text{bkg}} = 10,000$ is the number of independent sequences in each of the injection and no-injection runs of Sec. III A. Likelihood ratios λ_i [Eq. (2)] are indexed by the sequence label i within either run; the threshold λ_{th} acts at the run level on the collection $\{\lambda_i\}$.

2. Statistical evaluation of detector performance

To quantify the detector response, we adopt two operational figures of merit, defined entirely from the sequence-level likelihood ratios $\{\lambda_i\}$ of Eq. (2). For a threshold λ_{th} , the number of triggered sequences in the injection and no-injection runs are

$$N_{\text{det}}(\lambda_{\text{th}}) = \sum_{i=1}^{N_{\text{inj}}} \Theta(\lambda_i - \lambda_{\text{th}}), \quad (\text{C1})$$

$$N_{\text{false}}(\lambda_{\text{th}}) = \sum_{j=1}^{N_{\text{bkg}}} \Theta(\lambda_j - \lambda_{\text{th}}), \quad (\text{C2})$$

where Θ is the Heaviside step function. The detection efficiency under controlled injection is

$$\eta(\lambda_{\text{th}}) = \frac{N_{\text{det}}(\lambda_{\text{th}})}{N_{\text{inj}}}, \quad (\text{C3})$$

i.e., the conditional probability that a sequence containing the deterministically prepared single photon yields $\lambda \geq \lambda_{\text{th}}$.

The corresponding figure of merit in the no-injection run is the bare per-sequence trigger probability $N_{\text{false}}(\lambda_{\text{th}})/N_{\text{bkg}}$. Since each measurement sequence duration is limited by the storage lifetime $T_{1,s}$, we define the bare background trigger rate as

$$R_{\text{bkg}}^{\text{bare}}(\lambda_{\text{th}}) = \frac{N_{\text{false}}(\lambda_{\text{th}})}{N_{\text{bkg}} T_{1,s}}. \quad (\text{C4})$$

The bare trigger probability is not, however, equal to the residual photon occupation of the storage cavity, because it is filtered by the same detector response $\eta(\lambda_{\text{th}})$ that also governs the injection run. Removing this common response yields the calibrated residual occupation

$$\bar{n}_c(\lambda_{\text{th}}) = \frac{N_{\text{false}}(\lambda_{\text{th}})}{N_{\text{det}}(\lambda_{\text{th}})} = \frac{1}{\eta(\lambda_{\text{th}})} \frac{N_{\text{false}}(\lambda_{\text{th}})}{N_{\text{bkg}}}, \quad (\text{C5})$$

where the last equality uses $N_{\text{inj}} = N_{\text{bkg}} = 10,000$. We interpret $\bar{n}_c$ as the time-averaged equivalent photon occupation of the storage cavity over one parity-measurement cycle: under the same response function applied to the injected signal, any residual trigger in the no-injection run is converted to a photon count. The corresponding efficiency-corrected background photon rate is

$$R_{\text{bkg}}^{\text{corr}}(\lambda_{\text{th}}) = \frac{\bar{n}_c(\lambda_{\text{th}})}{T_{1,s}}, \quad (\text{C6})$$

the steady-state rate at which incoherent background photons populate and leak out of the storage cavity.

The plateau structure of $\bar{n}_c(\lambda_{\text{th}})$ at intermediate λ_{th} acquires a transparent interpretation under this definition. Let $p_1(\lambda_{\text{th}})$ denote the conditional probability that a sequence with a photon present at $t = 0$ yields $\lambda \geq \lambda_{\text{th}}$. Once λ_{th} is large enough to suppress signal-tail misclassifications in both runs,

$$\frac{N_{\text{det}}}{N_{\text{inj}}} \rightarrow p_1(\lambda_{\text{th}}), \quad \frac{N_{\text{false}}}{N_{\text{bkg}}} \rightarrow \bar{n}_{\text{phys}} p_1(\lambda_{\text{th}}),$$

so the ratio in Eq. (C5) converges to the intrinsic occupation $\bar{n}_{\text{phys}}$ with the response factor explicitly canceled. The




FIG. 14. Detector performance characterization. (a) Detection efficiency η as a function of the likelihood ratio λ_{th} . (b) Bare background rate $R_{\text{bkg}}^{\text{bare}}$ (left axis) and the corresponding average background photon number $\bar{n}$ (right axis) versus λ_{th} . The markers (blue/red) represent independent experimental runs. The dashed vertical line indicates the chosen operational threshold $\lambda_{\text{th}} = 4 \times 10^4$ used to define the background-dominated regime, which maintains a detection efficiency $\eta \approx 0.27$ while suppressing the bare background rate $R_{\text{bkg}}^{\text{bare}} \approx 22 \text{ sec}^{-1}$ and the average background photon number $\bar{n}_c \approx 1.7 \times 10^{-3}$.

operational threshold $\lambda_{\text{th}} = 4 \times 10^4$ (Sec. III A and Fig. 14) is placed at the onset of this plateau, where further increases of λ_{th} no longer change $\bar{n}_c$ within statistical uncertainty.

At the operational threshold $\lambda_{\text{th}} = 4 \times 10^4$, the no-injection runs yield $R_{\text{bkg}}^{\text{bare}} \approx 22 \text{ sec}^{-1}$, which corresponds to the per-run averaged residual cavity occupation

$$\bar{n}_c \equiv \frac{1}{\eta(\lambda_{\text{th}})} \frac{N_{\text{false}}(\lambda_{\text{th}})}{N_{\text{bkg}}} \approx 1.7 \times 10^{-3}, \quad (\text{C7})$$

and equivalently $R_{\text{bkg}}^{\text{corr}} = \bar{n}_c / T_{1,s} \approx 80 \text{ sec}^{-1}$ with $T_{1,s} = 20.3 \mu\text{s}$. The physical decomposition of $\bar{n}_c$ into a transmon-mediated dispersive-dressing contribution and the residual extraneous-bath contribution is treated in Sec. C 3 of this appendix.

3. Source of residual photons

The residual cavity occupation reported in Eq. (C7) corresponds, under the Bose-Einstein relation $\bar{n}_c = [\exp(\hbar\omega_s/k_B T_{\text{eff},s}) - 1]^{-1}$ with $\omega_s/2\pi = 5.051 \text{ GHz}$ ($\hbar\omega_s/k_B \approx 242 \text{ mK}$), to an effective storage-mode temperature

$$T_{\text{eff},s} = \frac{\hbar\omega_s/k_B}{\ln(1 + 1/\bar{n}_c)} \approx 38 \text{ mK}. \quad (\text{C8})$$

This value sits between the mixing-chamber stage of the dilution refrigerator ($T_{\text{MC}} \approx 12 \text{ mK}$, measured at the cold plate) and the transmon effective temperature $T_{\text{eff},q} \approx 72 \text{ mK}$ inferred from the residual qubit population in Table III. The hierarchy $T_{\text{MC}} < T_{\text{eff},s} < T_{\text{eff},q}$ is not a thermodynamic contradiction: in the dispersive limit, the cavity-qubit interaction $H_{\text{int}}/\hbar = 2\chi_{\text{qs}} a^\dagger a \sigma_z/2$ contains no Jaynes-Cummings exchange term $g_{\text{qs}}(\sigma_+ a + \sigma_- a^\dagger)$, so the two modes do not thermalize to each other and equilibrate independently with their respective dissipation channels [25]. The storage cavity, a passive 3D aluminum resonator without a Josephson element, is in particular not directly heated by the quasiparticle bath that dominates the transmon thermal population, allowing $T_{\text{eff},s} < T_{\text{eff},q}$. The remaining offset $T_{\text{eff},s} - T_{\text{MC}} \approx 26 \text{ mK}$ is attributed to extraneous thermal channels, the dominant of which we estimate in what follows.

The transmon population leaks into a residual cavity excitation through the small dispersive dressing between the bare states $|e, 0\rangle$ and $|g, 1\rangle$. Within the single-excitation manifold of the Jaynes-Cummings Hamiltonian, an exact two-by-two diagonalization gives the dressed eigenstates

$$|\tilde{e}\rangle = \cos\theta|e, 0\rangle + \sin\theta|g, 1\rangle, \quad (\text{C9})$$

$$|\tilde{1}\rangle = -\sin\theta|e, 0\rangle + \cos\theta|g, 1\rangle, \quad (\text{C10})$$

with mixing angle determined by $\tan 2\theta = 2g_{\text{qs}}/\Delta$, where $\Delta \equiv \omega_q - \omega_s$ is the qubit-storage detuning. In the dispersive limit $g_{\text{qs}} \ll |\Delta|$, this reduces to

$$\tan\theta \simeq \frac{g_{\text{qs}}}{|\Delta|}. \quad (\text{C11})$$

A quasiparticle-induced tunneling event prepares the transmon in the bare state $|e, 0\rangle$; projected onto the dressed basis, this event is registered by the parity-measurement protocol as a cavity excitation with probability $|\langle\tilde{1}|e, 0\rangle|^2 = \sin^2\theta \simeq (g_{\text{qs}}/\Delta)^2$. The transmon-mediated contribution to the cavity occupation is therefore

$$\bar{n}_c^q = \bar{n}_q (g_{\text{qs}}/\Delta)^2, \quad (\text{C12})$$

with $\bar{n}_q$ being the transmon residual occupation.

For our parameters $\omega_q/2\pi = 4.629 \text{ GHz}$, $\omega_s/2\pi = 5.051 \text{ GHz}$, $\alpha_q/2\pi = -272 \text{ MHz}$, and $\chi_{\text{qs}}/2\pi = -2.3 \text{ MHz}$ (Table III), the transmon-storage coupling inferred from the dispersive shift $\chi_{\text{qs}} = g_{\text{qs}}^2 \alpha_q / [\Delta(\Delta + \alpha_q)]$ is $g_{\text{qs}}/2\pi \simeq 49.8 \text{ MHz}$, giving $(g_{\text{qs}}/\Delta)^2 \simeq 1.4 \times 10^{-2}$ [56]. With $T_{\text{eff},q} \approx 72 \text{ mK}$ at $\omega_q/2\pi = 4.629 \text{ GHz}$ ($\hbar\omega_q/k_B \approx 222 \text{ mK}$), the residual transmon occupation is $\bar{n}_q \approx 4.8 \times 10^{-2}$, and Eq. (C12) yields

$$\bar{n}_c^q \simeq 4.8 \times 10^{-2} \times 1.4 \times 10^{-2} \simeq 6.7 \times 10^{-4}, \quad (\text{C13})$$

accounting for $\sim 39\%$ of the measured residual occupation $\bar{n}_c \approx 1.7 \times 10^{-3}$ [Eq. (C7)]. The remaining $\sim 61\%$ must be attributed to extraneous thermal channels not mediated by the qubit, including blackbody radiation from higher-temperature stages of the dilution refrigerator, insufficient attenuation or thermalization of the input lines, and back-action noise from the readout amplification chain [25].

The detector therefore sits in an intermediate regime in which the two contributions are of comparable magnitude. Three limiting scenarios are worth distinguishing: $\bar{n}_c^q/\bar{n}_c \gtrsim 50\%$ would indicate that qubit-side improvements (active reset, quasiparticle trapping, gap engineering) provide the dominant lever for reducing $\bar{n}_c$; $\bar{n}_c^q/\bar{n}_c \ll 10\%$ would indicate that the environment-side improvements alone (infrared filtering, cryogenic attenuation, magnetic and electromagnetic shielding) would suffice; the present value $\sim 39\%$ places the two routes on equal footing, with the engineering response discussed in Sec. IV.

4. Dark photon model

The expected number of detected events induced by massive dark photons is given by

$$N_{\text{DM}} = \frac{\epsilon^2 \rho_{\text{DM}} Q_{\text{DM}} Q_s G V}{\omega_s} N_{\text{meas}}, \quad (\text{C14})$$

where ϵ is the kinetic mixing angle and $\rho_{\text{DM}} = 0.45 \text{ GeV}/\text{cm}^3$ is the local dark matter density assumed from the standard halo model. The factor $Q_{\text{DM}} \approx 10^6$ corresponds to the expected dark matter signal line width, while Q_s and ω_s are the quality factor and angular frequency of the storage cavity mode, respectively, with $T_{1,s} = Q_s/\omega_s$. The geometric form factor G accounts for the overlap between the cavity electromagnetic mode and the effective electric field induced by the dark photon. Assuming the TE_{101} mode is polarized along the z -axis, G reduces to a one-dimensional overlap of the mode profile with a uniform external field:

$$G = \frac{1}{3} \frac{|\int_V E_z dV|^2}{\int_V |E|^2 dV} = \frac{1}{3} \frac{64}{\pi^4}, \quad (\text{C15})$$

where the factor $1/3$ arises from the isotropic time average over the random polarization direction of the dark photon, and the analytic value $64/\pi^4$ follows from the TE_{101} mode profile $E_z \propto \sin(\pi x/a) \sin(\pi z/c)$ of a rectangular cavity of dimensions $a \times b \times c$ in Table II. The effective cavity volume is $V = 9 \text{ cm}^3$.

5. Confidence-level calculation

The statistical inference on the kinetic mixing angle ϵ uses only the trigger counts above the operational threshold λ_{th} (Appendix C), without further spectral or temporal cuts.

We model the outcome of a single measurement run as a binomial process with N_{meas} independent trials and per-trial success probability $p(\epsilon)$. The probability mass function is

$$B(k|N_{\text{meas}}, p) = \binom{N_{\text{meas}}}{k} p^k (1-p)^{N_{\text{meas}}-k}, \quad (\text{C16})$$

where, for the dark photon hypothesis,

$$p(\epsilon) = \frac{\eta(\lambda_{\text{th}}) N_{\text{DM}}(\epsilon)}{N_{\text{meas}}} = \eta(\lambda_{\text{th}}) \frac{\epsilon^2 \rho_{\text{DM}} Q_{\text{DM}} Q_s G V}{\omega_s}, \quad (\text{C17})$$

with $\eta(\lambda_{\text{th}})$ being the detection efficiency defined in Eq. (C3) at the operating threshold λ_{th} . The cumulative probability of observing no more than N_{obs} trigger events out of N_{meas} trials is then

$$P(\leq N_{\text{obs}}|\epsilon) = \sum_{k=0}^{N_{\text{obs}}} \binom{N_{\text{meas}}}{k} p(\epsilon)^k (1-p(\epsilon))^{N_{\text{meas}}-k}, \quad (\text{C18})$$

where $N_{\text{obs}} \equiv N_{\text{false}}(\lambda_{\text{th}})$ [Eq. (C2)] is the trigger count of the no-injection run. Because no background subtraction is applied, the resulting bound on ϵ is conservative.

To combine the two independent data takings (Run 1, Run 2), we treat them as statistically independent experiments with respective inputs $\{N_{\text{meas},i}, \eta_i, N_{\text{obs},i}\}_{i=1,2}$. Denoting the per-run CDFs by $P_i(\epsilon)$, the joint upper-tail

TABLE IV. Exclusion limits on the dark photon mixing angle at 95% CL for $\rho_{\text{DM}} = 0.45 \text{ GeV}/\text{cm}^3$, obtained from Runs 1, 2, and their combined analysis using the quantum sensor operated as a single-photon counter.

Parameter	Run 1	Run 2	Runs 1 + 2
f_s (GHz)	5.051	5.051	5.051
Q_s	634,727 (0.6M)	634,727 (0.6M)	634,727 (0.6 M)
N_{DM}	4	5	4–5
N_{meas}	10,000	10,000	20,000
Efficiency η	0.257	0.286	0.257–0.286
n^{CL}	0.00358	0.00368	0.00261
ϵ^{CL}	1.2217×10^{-14}	1.2387×10^{-14}	1.0446×10^{-14}

probability factorizes as

$$P_{\text{combined}}(\epsilon) = P_1(\epsilon)P_2(\epsilon), \quad (\text{C19})$$

and the 95% CL exclusion limit is the smallest ϵ for which

$$P_{\text{combined}}(\epsilon^{95\%}) = 0.05. \quad (\text{C20})$$

Equivalently, in terms of the signal-induced mean occupation per cycle $\bar{n}_{\text{sig}}(\epsilon) \equiv N_{\text{DM}}(\epsilon)/N_{\text{meas}}$, the bound can be reported as $\bar{n}_{\text{sig}}^{95\%}$ (Table IV). The numerical results in Table IV are obtained by numerically inverting Eq. (C20); the combined dataset yields $\epsilon^{95\%} = 1.04 \times 10^{-14}$ at $\omega_s/2\pi = 5.051$ GHz.

-
- [1] V. Trimble, Existence and nature of dark matter in the universe, *Annu. Rev. Astron. Astrophys.* **25**, 425 (1987).
 - [2] P. A. Zyla, R. M. Barnett, J. Beringer, O. Dahl, D. A. Dwyer, D. E. Groom, K. S. Lugovsky, E. Pianori, D. J. Robinson, C. G. Wohl *et al.*, Review of particle physics (RPP 2020), *Prog. Theor. Exp. Phys.* **2020**, 1 (2020).
 - [3] R. J. Gaitskell, Direct detection of dark matter, *Annu. Rev. Nucl. Part. Sci.* **54**, 315 (2004).
 - [4] R. D. Peccei and H. R. Quinn, CP conservation in the presence of pseudoparticles, *Phys. Rev. Lett.* **38**, 1440 (1977).
 - [5] S. Weinberg, A new light boson?, *Phys. Rev. Lett.* **40**, 223 (1978).
 - [6] F. Wilczek, Problem of strong P and T invariance in the presence of instantons, *Phys. Rev. Lett.* **40**, 279 (1978).
 - [7] P. Sikivie, Experimental tests of the “invisible” axion, *Phys. Rev. Lett.* **51**, 1415 (1983).
 - [8] P. W. Graham, J. Mardon, and S. Rajendran, Vector dark matter from inflationary fluctuations, *Phys. Rev. D* **93**, 103520 (2016).
 - [9] B. Holdom, Two $U(1)$ ’s and ϵ charge shifts, *Phys. Lett.* **166B**, 196 (1986).
 - [10] P. Arias, D. Cadamuro, M. Goodsell, J. Jaeckel, J. Redondo, and A. Ringwald, Wispy cold dark matter, *J. Cosmol. Astropart. Phys.* **06** (2012) 013.
 - [11] 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).
 - [12] A. Caputo, A. J. Millar, C. A. J. O’Hare, and E. Vitagliano, Dark photon limits: A handbook, *Phys. Rev. D* **104**, 095029 (2021).
 - [13] L. Zhong, S. Al Kenany, K. Backes, B. Brubaker, S. Cahn, G. Carosi, Y. Gurevich, W. Kindel, S. Lamoreaux, K. Lehnert *et al.*, Results from phase 1 of the HAYSTAC microwave cavity axion experiment, *Phys. Rev. D* **97**, 092001 (2018).
 - [14] T. Braine, R. Cervantes, N. Crisosto, N. Du, S. Kimes, L. Rosenberg, G. Rybka, J. Yang, D. Bowring, A. Chou *et al.* (ADMX Collaboration), Extended search for the invisible axion with the axion dark matter experiment, *Phys. Rev. Lett.* **124**, 101303 (2020).
 - [15] K. M. Backes, D. A. Palken, S. A. Kenany, B. M. Brubaker, S. Cahn, A. Droster, G. C. Hilton, S. Ghosh, H. Jackson, S. K. Lamoreaux *et al.*, A quantum enhanced search for dark matter axions, *Nature (London)* **590**, 238 (2021).
 - [16] C. M. Caves, Quantum limits on noise in linear amplifiers, *Phys. Rev. D* **26**, 1817 (1982).
 - [17] J. E. Kim, Weak-interaction singlet and strong CP invariance, *Phys. Rev. Lett.* **43**, 103 (1979).
 - [18] M. A. Shifman, A. Vainshtein, and V. I. Zakharov, Can confinement ensure natural CP invariance of strong interactions?, *Nucl. Phys.* **B166**, 493 (1980).
 - [19] A. P. Zhitnitskii, Possible suppression of axion-hadron interactions, *Sov. J. Nucl. Phys.* **31**, 260 (1980), <https://www.osti.gov/biblio/7063072>.
 - [20] M. Dine, W. Fischler, and M. Srednicki, A simple solution to the strong CP problem with a harmless axion, *Phys. Lett.* **104B**, 199 (1981).
 - [21] S. K. Lamoreaux, K. A. van Bibber, K. W. Lehnert, and G. Carosi, Analysis of single-photon and linear amplifier detectors for microwave cavity dark matter axion searches, *Phys. Rev. D* **88**, 035020 (2013).
 - [22] R. Cervantes, J. Aumentado, C. Braggio, B. Giaccone, D. Frolov, A. Grassellino, R. Harnik, F. Lecocq, O. Melnychuk, R. Pilipenko *et al.*, Deepest sensitivity to wavelike dark photon dark matter with superconducting radio frequency cavities, *Phys. Rev. D* **110**, 043022 (2024).
 - [23] G. Carosi, C. Cisneros, N. Du, S. Durham, N. Robertson, C. Goodman, M. Guzzetti, C. Hanretty, K. Enzian, L. Rosenberg *et al.* (ADMX Collaboration), Search for axion dark matter from 1.1 to 1.3 GHz with ADMX, *Phys. Rev. Lett.* **135**, 191001 (2025).
 - [24] Y. Kim, J. Jeong, S. Youn, S. Bae, K. Lee, A. F. Van Loo, Y. Nakamura, S. Oh, T. Seong, S. Uchaikin *et al.*, Experimental search for invisible dark matter axions around 22 μeV , *Phys. Rev. Lett.* **133**, 051802 (2024).
 - [25] 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).
 - [26] H. Paik, D. I. Schuster, L. S. Bishop, G. Kirchmair, G. Catelani, A. P. Sears, B. Johnson, M. Reagor, L. Frunzio, L. I. Glazman *et al.*, Observation of high coherence in Josephson junction qubits measured in a three-dimensional circuit QED architecture, *Phys. Rev. Lett.* **107**, 240501 (2011).

- [27] M. Reagor, W. Pfaff, C. Axline, R. W. Heeres, N. Ofek, K. Sliwa, E. Holland, C. Wang, J. Blumoff, K. Chou *et al.*, Quantum memory with millisecond coherence in circuit QED, *Phys. Rev. B* **94**, 014506 (2016).
- [28] A. Romanenko, R. Pilipenko, S. Zorretti, D. Frolov, M. Awida, S. Belomestnykh, S. Posen, and A. Grassellino, Three-dimensional superconducting resonators at $T < 20$ mK with photon lifetimes up to $\tau = 2$ s, *Phys. Rev. Appl.* **13**, 034032 (2020).
- [29] S. Chakram, A. E. Oriani, R. K. Naik, A. V. Dixit, K. He, A. Agrawal, H. Kwon, and D. I. Schuster, Seamless high- q microwave cavities for multimode circuit quantum electrodynamics, *Phys. Rev. Lett.* **127**, 107701 (2021).
- [30] O. Milul, B. Guttel, U. Goldblatt, S. Hazanov, L. M. Joshi, D. Chausovsky, N. Kahn, E. Çiftçiyök, F. Lafont, and S. Rosenblum, Superconducting cavity qubit with tens of milliseconds single-photon coherence time, *PRX Quantum* **4**, 030336 (2023).
- [31] T. Takenaka, T. Kubo, I. Mahoob, K. Mizuno, H. Inoue, T. Saeki, and S. Saito, Three-dimensional niobium coaxial cavity with ~ 0.1 second lifetime, *Phys. Rev. Appl.* **25**, 034076 (2026).
- [32] A. E. Oriani, F. Zhao, T. Roy, A. Anferov, K. He, A. Agrawal, R. Banerjee, S. Chakram, and D. I. Schuster, Niobium coaxial cavities with internal quality factors exceeding 1.4×10^9 for circuit quantum electrodynamics, *Phys. Rev. Appl.* **24**, 044080 (2025).
- [33] D. Bafia, B. Abdisatarov, R. Pilipenko, Y. Lu, G. Eremeev, A. Romanenko, and A. Grassellino, Quantifying trapped magnetic vortex losses in niobium resonators at mK temperatures, *Appl. Phys. Lett.* **127**, 152601 (2025).
- [34] T. Braine, G. Rybka, A. A. Baker, J. Brodsky, G. Carosi, N. Du, N. Woollett, S. Knirck, and M. Jones (ADMX Collaboration), Multi-mode analysis of surface losses in a superconducting microwave resonator in high magnetic fields, *Rev. Sci. Instrum.* **94** (2023).
- [35] S. Posen, M. Checchin, O. Melnychuk, T. Ring, I. Gonin, and T. Khabiboulline, High-quality-factor superconducting cavities in tesla-scale magnetic fields for dark-matter searches, *Phys. Rev. Appl.* **20**, 034004 (2023).
- [36] A. Gurevich, Enhancement of rf breakdown field of superconductors by multilayer coating, *Appl. Phys. Lett.* **88** (2006).
- [37] S. Posen, J. Lee, D. N. Seidman, A. Romanenko, B. Tennis, O. Melnychuk, and D. Sergatskov, Advances in Nb₃Sn superconducting radiofrequency cavities towards first practical accelerator applications, *Supercond. Sci. Technol.* **34**, 025007 (2021).
- [38] D. Ahn, O. Kwon, W. Chung, W. Jang, D. Lee, J. Lee, S. W. Youn, H. Byun, D. Youm, and Y. K. Semertzidis, Biaxially textured YBa₂Cu₃O_{7-x} microwave cavity in a high magnetic field for a dark-matter axion search, *Phys. Rev. Appl.* **17**, L061005 (2022).
- [39] G. Marconato, C. Pira, D. Alesini, D. D'Agostino, O. Azzolini, C. Braggio, R. Caforio, E. Chyhyrnyets, A. D'Elia, M. Fracasso *et al.*, NbTi thin-film SRF cavities for dark matter search, *IEEE Trans. Appl. Supercond.* **34**, 1 (2024).
- [40] S. Ahyoune, A. Á. Melcón, S. A. Cuendis, S. Calatroni, C. Cogollos, A. Díaz-Morcillo, B. Döbrich, J. Gallego, J. García-Barceló, B. Gimeno *et al.*, Rades axion search results with a high-temperature superconducting cavity in an 11.7 T magnet, *J. High Energy Phys.* **04** (2025) 113.
- [41] M. Kuzmanović, I. Moskalenko, Y.-H. Chang, O. Stanisavljević, C. Warren, E. Hogedal, A. Aggarwal, I. Ahmad, J. Biznárová, M. Dahiya *et al.*, Neural-network-based design and implementation of fast and robust quantum gates, [arXiv:2505.02054](https://arxiv.org/abs/2505.02054).
- [42] A. Agrawal, A. V. Dixit, T. Roy, S. Chakram, K. He, R. K. Naik, D. I. Schuster, and A. Chou, Stimulated emission of signal photons from dark matter waves, *Phys. Rev. Lett.* **132**, 140801 (2024).
- [43] H. Zheng, M. Silveri, R. Brierley, S. Girvin, and K. Lehnert, Accelerating dark-matter axion searches with quantum measurement technology, [arXiv:1607.02529](https://arxiv.org/abs/1607.02529).
- [44] C. T. Hann, S. S. Elder, C. S. Wang, K. Chou, R. J. Schoelkopf, and L. Jiang, Robust readout of bosonic qubits in the dispersive coupling regime, *Phys. Rev. A* **98**, 022305 (2018).
- [45] S. S. Elder, C. S. Wang, P. Reinhold, C. T. Hann, K. S. Chou, B. J. Lester, S. Rosenblum, L. Frunzio, L. Jiang, and R. J. Schoelkopf, High-fidelity measurement of qubits encoded in multilevel superconducting circuits, *Phys. Rev. X* **10**, 011001 (2020).
- [46] F. Zhao, Z. Li, A. V. Dixit, T. Roy, A. Vrajitoarea, R. Banerjee, A. Anferov, K.-H. Lee, D. I. Schuster, and A. Chou, Flux-tunable cavity for dark matter detection, *Phys. Rev. Lett.* **135**, 201002 (2025).
- [47] I. Batković, A. De Angelis, M. Doro, and M. Manganaro, Axion-like particle searches with IACTs, *Universe* **7**, 185 (2021).
- [48] C. O'Hare, cajohare/axionlimits: Axionlimits, <https://cajohare.github.io/AxionLimits/> (2020).
- [49] 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, *Phys. Rev. Appl.* **17**, 044016 (2022).
- [50] L. Chen, H.-X. Li, Y. Lu, C. W. Warren, C. J. Križan, S. Kosen, M. Rommel, S. Ahmed, A. Osman, J. Biznárová *et al.*, Transmon qubit readout fidelity at the threshold for quantum error correction without a quantum-limited amplifier, *npj Quantum Inf.* **9**, 26 (2023).
- [51] M. F. Dumas, B. Groleau-Paré, A. McDonald, M. H. Muñoz-Arias, C. Lledó, B. D'Anjou, and A. Blais, Measurement-induced transmon ionization, *Phys. Rev. X* **14**, 041023 (2024).
- [52] W.-E. Lin, L.-C. Hsiao, C.-H. Ma, E.-H. Yeh, W.-L. Peng, H.-S. Goan, C.-S. Wu, Y.-N. Chen, Y.-F. Chen, C.-T. Ke, and C.-D. Chen, Mitigation of measurement-induced state transitions via a fast-load and fast-clear readout, [arXiv:2607.23681](https://arxiv.org/abs/2607.23681).
- [53] [10.3030/101118911](https://arxiv.org/abs/10.3030/101118911).

- [54] A. V. Dixit, Searching for dark matter with superconducting qubits, Ph.D. thesis, University of Chicago, 2021.
- [55] A. Sultanov, M. Kuzmanović, A. V. Lebedev, and G. S. Paraoanu, Protocol for temperature sensing using a three-level transmon circuit, *Appl. Phys. Lett.* **119**, 144002 (2021).
- [56] 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).