

Maximizing quantum enhancement in axion dark matter experiments

Chao-Lin Kuo^{1,2}, Chelsea L. Bartram², Aaron S. Chou³, Taj A. Dyson¹, Noah A. Kurinsky², Gray Rybka⁴,
Sephora Ruppert¹, Osmond Wen^{1,5}, Matthew O. Withers¹, Andrew K. Yi², and Cheng Zhang¹

¹*Stanford University, Stanford, California 94305, USA*

²*SLAC National Accelerator Laboratory, 2575 Sand Hill Road, Menlo Park, California 94025, USA*

³*Fermi National Accelerator Laboratory, Batavia, Illinois 60510, USA*

⁴*University of Washington, 1410 NE Campus Parkway, Seattle, Washington 98195, USA*

⁵*California Institute of Technology, 1200 East California Boulevard, Pasadena, California 91125, USA*



(Received 9 February 2025; accepted 22 May 2025; published 10 June 2025)

We provide a comprehensive comparison of linear amplifiers and microwave photon counters in axion dark matter experiments. The study is done assuming a range of realistic operating conditions and detector parameters, over the frequency range between 1 and 30 GHz. As expected, photon counters are found to be advantageous under low background, at high frequencies ($\nu > 5$ GHz), if they can be implemented with robust wide-frequency tuning or a very low dark count rate. Additional noteworthy observations emerging from this study include: (1) an expanded applicability of off-resonance photon background reduction, including the single-quadrature state squeezing, for scan rate enhancements; (2) a much broader appeal for operating the haloscope resonators in the overcoupling regime, up to $\beta \sim 10$; (3) the need for a detailed investigation into the cryogenic and electromagnetic conditions inside haloscope cavities to lower the photon temperature for future experiments; (4) the necessity to develop a distributed network of coupling ports in high-volume axion haloscopes to utilize these potential gains in the scan rate.

DOI: [10.1103/pgtg-3lhd](https://doi.org/10.1103/pgtg-3lhd)

I. INTRODUCTION

QCD axions are excellent candidates for the cosmological dark matter [1–5]. Their presence in a terrestrial laboratory can be detected using a well-established method, the cavity-based haloscope [6–9]. In this scheme, a cavity resonator is immersed in an external dc magnetic field, resonantly converting halo axions into a microwave signal. The resonant frequency of the cavity can be tuned to search for a narrow line, which would indicate the presence of axion dark matter. The cavity is coupled to a low-noise amplifier for readout and data processing.

In recent years, advances in quantum sensing techniques enable new capabilities in axion dark matter searches [10,11]. It is now possible to apply vacuum squeezing techniques to evade the standard quantum limit (SQL) by using a pair of parametric amplifiers operating in a single-quadrature squeezed state setup [12–14]. It has also been demonstrated that a microwave photon counter can be used

in centimeter-wave axion dark matter searches [15–18], which enhances the scan rate by avoiding SQL altogether.

In this paper, we provide an in-depth comparison of the linear amplifiers and the photon-counting microwave detectors in axion dark matter searches. Revisiting this subject is timely given these advances in quantum information technologies. Obviously, such promises can only be realized when auxiliary systems in the haloscope experiments can support operation of these new sensors under appropriate conditions. The purpose of this review is to clarify what these requirements are and to develop a road map for future axion dark matter searches in the frequency range of 1–30 GHz, corresponding to the postinflationary scenario of axion production [5,19–23].

The paper is organized as follows. In Sec. II we derive the scan rate of an axion haloscope read out by either a linear amplifier (with the option of adding a squeezing deamplifier) or a single microwave photon detector. We obtain closed-form expressions of which we take various limits—photon noise background versus dark count rate—as functions of the temperature and search frequencies. In Sec. III, we proceed to use the results to comprehensively compare these two readout systems and their ability to reach theoretical benchmarks such as the Kim-Shifman-Vainshtein-Zakharov (KSVZ) [24,25] and Dine-Fischler-Srednicki-Zhitnitsky (DFSZ) [26,27].

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

We conclude in Sec. IV and discuss the implications of these results on future searches of the postinflationary dark matter axions.

II. THE SCAN RATE CALCULATIONS

To calculate the scan rate R of the haloscope for reaching a certain coupling strength $g_{a\gamma\gamma}$, we closely follow the method and notations adopted in a well-cited paper, Ref. [12]. We consider the experimental setup depicted in Fig. 1. The large square block represents the cavity haloscope, which is coupled to a measuring port m , an internal lossy termination (surface conductive loss) at the haloscope temperature T , and a generator a from axion-to-photon conversion. The coupling strengths are, respectively, κ_m , κ_l , and κ_a .

The measurement port m is connected to a circulator whose output goes to a low-noise amplifier to the right and the input from an impedance matching termination to the left, with the efficiency $\lambda \leq 1$ in either direction. Optionally, a quadrature deamplifier with a deamplification gain G_s can be inserted between the circulator and the termination. We include loss in efficiency due to dissipation, with λ being the efficiency to either the amplifier side or the termination side.

While we closely follow [12], we consider the general case where the termination temperature T_b can be different, usually lower, than the haloscope temperature T . This situation arises in axion experiments because the large surface area of the haloscope cavity and the poorly thermal anchored tuning elements often make it difficult to operate the haloscope at the base temperature T_b of the mixing chamber [13,14]. On the other hand, the small coaxial termination can be much more easily heat sunk to the base temperature T_b of the refrigerator.

Starting with the Heisenberg-Langevin equation for the cavity ladder operator $\hat{A}$ and following the steps based on known solutions to the input-output theory as presented

in Appendix A of [12], we obtain the single-quadrature spectral density at the amplifier,

$$\Sigma_{\text{out},X,m} = \left(n_T + \frac{1}{2}\right)(1 - \lambda) + \frac{\lambda}{B(\omega)} \times \left[\left(n_A + \frac{1}{2}\right)\kappa_a\kappa_m + \left(n_T + \frac{1}{2}\right)\kappa_l\kappa_m + \left(n_b + \frac{1}{2}\right)(1 - \lambda + \lambda/G_s)A(\omega) \right], \quad (1)$$

where $A(\omega) \equiv \omega^2 + (\kappa_m - \kappa_l)^2/4$, $B(\omega) \equiv \omega^2 + (\kappa_m + \kappa_l)^2/4$, and the photon occupation number n for thermal sources is given by $(e^{\hbar\omega/kT} - 1)^{-1}$. The coefficients $A(\omega)$ are simply the reflectivity of the resonator $|\Gamma|^2$ in classical microwave theory and $B(\omega)$ is a normalization factor. The coupling constants κ_m , κ_l can be connected with their microwave equivalents by the following identifications: $\kappa_l/\omega_c = \kappa_l/(2\pi\nu_c) = 1/Q_0$ and $\kappa_m = \beta\kappa_l$, where Q_0 is the intrinsic (unloaded) quality factor of the resonator and β is the coupling parameter of the measurement port. This setup can be analyzed as a cascaded microwave network. At the resonant frequency of the cavity, the amplifier “sees” thermal emission from the cavity (especially if $\beta = 1$). When the cavity becomes reflective coming off resonance, the amplifier sees the radiation from the termination.

The photon occupation number n_A ($\gg 1$) and coupling κ_a can be expressed in terms of physical parameters normally found in the haloscope literature, such as $g_{a\gamma\gamma}$ and the form factor [12]. We reproduce their Eqs. (B10) and (B11) here,

$$n_A = \frac{|g_{a\gamma\gamma}|\rho_a B_0 V}{4\omega_a \Delta_a} \sqrt{\frac{C_{mnl}c^3}{\hbar\mu_0}}, \quad (2)$$

$$\kappa_a = |g_{a\gamma\gamma}|B_0 \sqrt{\frac{C_{mnl}\hbar c^3}{\mu_0}}. \quad (3)$$

The parameters $g_{a\gamma\gamma}$, ρ_a , B_0 , and C_{mnl} are the axion-photon coupling strength, local axion density, external magnetic field, and a geometrical form factor for the eigenmode, respectively. The visibility function $\alpha(\omega)$, or the signal-to-noise ratio, is given by

$$\alpha(\omega) = \lambda n_A \kappa_a \kappa_m \cdot \left\{ \left(n_T + \frac{1}{2}\right) [(1 - \lambda)B(\omega) + \lambda\kappa_l\kappa_m] + \left(n_b + \frac{1}{2}\right) \left(1 - \lambda + \frac{\lambda}{G_s}\right) \lambda \cdot A(\omega) \right\}^{-1}, \quad (4)$$

which can be used in $R = \int \alpha^2(\omega) d\omega$ to find the scan rate figure of merit,




FIG. 1. The experimental setup considered in Secs. II A and II B. The schematic and notation closely follows those of [12]. G_a and G_s are the amplifier and squeezer gains. The coupling coefficients κ_m , κ_l , and κ_a are for the measurement port m , intrinsic (cavity) loss port l , and the axion port a . In this paper, we consider the general case where T_b , the termination temperature, is different from the haloscope temperature T .

$$R_{\text{amp}}(\beta) = \frac{\pi n_A^2 \kappa_a^2 \kappa_l^{-1}}{2(n_T + \frac{1}{2})^2} \cdot \frac{\gamma^2 \beta^2 \sqrt{G_s}}{\sqrt{\lambda + G_s(\gamma + \lambda)(1 - \lambda)}/\lambda \cdot \left\{ \frac{(\beta-1)^2}{4} \left[\frac{\lambda}{G_s} + 1 - \lambda \right] + \gamma\beta + \gamma \frac{(1+\beta)^2}{4} \frac{1-\lambda}{\lambda} \right\}^{3/2}}, \quad (5)$$

where we further define $\gamma \equiv (n_T + \frac{1}{2})/(n_b + \frac{1}{2})$. This scan rate figure of merit will be normalized in later sections relative to the theoretical benchmarks. Here we note that R_{amp} scales as $|g_{a\gamma\gamma}|^4$.

A. Linear amplifier in a standard configuration

We first consider the case in which the squeezer gain G_s is unity and $\gamma = 1$, setting the termination T_b to be the same as the haloscope T . R_{amp} then reduces to $\propto \beta^2/(\beta + 1)^3$, which peaks at $\beta = 2$ [12,28]. At large β , the scan rate scales as $\propto \kappa_l^{-1}/\beta \sim Q_0/(1 + \beta)$. To follow up a candidate signal, one seeks to maximize $\alpha(0) \propto \beta/(1 + \beta)^2$ which peaks at $\beta = 1$. All these are well-known results in the axion literature.

With $G_s = 1$, one can either amplify-digitize both the X and Y quadrature or perform a single-quadrature (X) amplification. The two-quadrature (phase-insensitive) amplification must add an additional half photon to the occupation number [29], and the single-quadrature (phase-sensitive) amplification adds none. However, as pointed out in [12], this apparent advantage for single-quadrature noiseless amplification is canceled by $2\times$ fewer independent measurements and $2\times$ larger noise bandwidth created by the idler. Therefore, phase-sensitive amplification alone does not evade the SQL at the system level.

B. Linear amplifier in a squeezed state readout

When the squeezer gain $G_s > 1$, SQL can be evaded as pointed out in [12] and realized in [13,14]. An enhanced scan rate is achieved by squeezing the photon noise from the termination ($n_b + 1/2$) and increasing the coupling β . As the cavity becomes overcoupled (larger β), it becomes more reflective. As a result, the amplifier sees more photon noise from the termination and less from the cavity. When $\gamma = 1$, our Eqs. (1) and (4) recover Eqs. (A23) and (A24) in [12]. When $\lambda = 1$, Eq. (5) reproduces their Eq. (2). Setting $\lambda = \sqrt{0.69}$, the scan rate enhancement E_t given by the 2D function $E_t = R(\beta, G_s)/R(\beta = 2, G_s = 1)$ is plotted in Fig. 2, which accurately reproduces Fig. 2(b) in [12].

While the above calculations are carried out using quantum mechanics, the final results can be understood classically aside from the $1/2$ photon corresponding to the vacuum state and the appearance of the Planck emission law. We point out the obvious fact that squeezing ($G_s > 1$) improves the scan rate even when $n_T \gg 1$. This is quite clear since n_T drops out completely in E_t when the ratio is

taken. The implication is that a squeezer improves the scan rate regardless of whether the experiment is cooled to the vacuum state. In other words, ADMX at 1 GHz, or even DMRadioGUT [30] at 1 MHz can still benefit from single-quadrature squeezing. While this broader applicability was already obvious from Eq. (2) of Malnou *et al.* [12], it has not been widely appreciated by the axion community because of the emphasis on “quantum” state squeezing.

The reduction of photon background cannot be done by attenuation, because any dissipation would induce emission (Kirchhoff’s law), as is apparent from the effects of λ in Eq. (4). Squeezing achieves this “cooling” by deamplifying the thermal (and vacuum) fluctuations in the quadrature being measured. Another way to achieve a similar effect is to simply cool the termination resistor to a lower temperature than the haloscope temperature ($T_b < T, \gamma > 1$). As discussed earlier, this situation happens because of cryogenic limitations, such as the hot rod problem [13,14]. In such scenario, γ would be playing a similar role to G_s .

The HAYSTAC cavity photons are at around 250 mK, contrasting the termination resistor’s 61 mK [13,14]. At 4.5 GHz, this corresponds to $\gamma = 2.33$, already providing some scan rate enhancement even when there is no squeezing ($E_t \sim 1.55$ over $\gamma = 1$). The ADMX [31] cavity photons typically sit at ~ 300 mK, and the mixing chamber can be as cold as 80 mK. If the termination resistor can be cooled to 80 mK, it would correspond to $\gamma = 3.66$ at 1 GHz. Assuming once again $\lambda = \sqrt{0.69}$, we obtain




FIG. 2. With $\gamma = 1$ and $\lambda = \sqrt{0.69}$, the scan rate enhancement E_t derived from Eq. (5) reproduces results in [12] for all (β, G_s) . The enhancement E_t is independent of n_T since both thermal and quantum noise can be squeezed.




FIG. 3. Squeezing can be used in conjunction of having a lower T_b . In this contour plot, we assume $\gamma = 3.66$, baselining typical performance of ADMX. The efficiency λ remains at $\sqrt{0.69}$. The scan rate peaks at $(\beta, G_s) \sim (8, 20)$, giving an enhancement factor of $3.3\times$. ADMX typically operates at $\beta = 1.8$ with no squeezing, $G_s = 1$.

$E_t = R_{\text{amp}}(\beta = 8, G_s = 20)/R_{\text{amp}}(\beta = 2, G_s = 1) = 3.3$, an enhancement in the scan rate by $3.3\times$ (Fig. 3).

The enhancement can be quite significant for the MADMAX Experiment [32], which plans to run the haloscope at 4 K. A compact dilution refrigerator could be used to cool only the termination connected through a circulator, providing a γ factor significantly larger than 1.

C. Microwave photon-counting detector

One can avoid the SQL by forgoing the phase information and by only “counting” the microwave photons [10]. Notable examples of this approach are Rydberg atoms [33,34], Penning traps [35], and superconducting detectors [36,37]. In recent years, transmon qubits have emerged to be a leading technology to measure single photons in the microwave regime [15–17,36,37]. In [18], such devices are coupled to a haloscope through a coaxial transmission line. This device architecture, which the authors called a single microwave photon detector (SMPD), is particularly appealing because it creates physical separation between the superconductive sensors and the strong magnetic fields associated with an axion experiment. The irreversible photon detection scheme also makes high-duty cycle integration possible. A schematic of the SMPD is shown in Fig. 4, along with a brief explanation of its operating principles.

For the purpose of this optimization/comparative study, we focus on microwave photon counters whose physical parameters are similar to the SMPDs reported in [16–18], although many of the results should be applicable to other types of microwave photon counters. Reference [34] considered several similar issues in the context of




FIG. 4. The SMPD [15–18] consists of an intake resonator (“buffer”) that defines the band of the incoming photons, a resonator that defines the outgoing mode (“waste”), and a transmon qubit coupled to both resonators. An incoming photon combines with a pump photon to induce the qubit transition from $|g\rangle$ to $|e\rangle$ and the generation of a waste photon that subsequently dissipates. The schematic figure is adapted from [16].

Rydberg-atom-based axion searches. We consider the system depicted in Fig. 5, where the linear amplifier in Fig. 1 is replaced by an SMPD. The photon field detected is given by

$$\eta P(\nu) = \left\{ n_T(1 - \eta) + \frac{\eta}{B(\omega)} [n_T \kappa_l \kappa_m + n_b A(\omega)] \right\} \cdot h\nu.$$

For single-mode optics, the noise equivalent power σ associated with $P(\nu)$ consists of contributions from the Poisson noise and the Bose noise,

$$\sigma^2 = \frac{1}{\tau} \left[\int d\nu h\nu \eta P(\nu) + \int d\nu \eta^2 P^2(\nu) \right].$$

Note that Eqs. (9) and (13) in [10], Eqs. (12) in [38], and (41) in [39] all contain the same expression. The readers shall refer to these papers for the derivation. The total noise background is given by




FIG. 5. The experimental setup used in the analysis of Sec. III C. In place of the amplifier, an SMPD is used for readout at the measurement port.




FIG. 6. The contributions to photon noise near the resonance of the haloscope cavity. For a photon-counting detector, the noise within the bandwidth of the detector $\Delta\nu_d$ contributes to measurement uncertainties against a possible axion signal within the bandwidth of the haloscope, $\sim\kappa_l(2\pi)^{-1}(1+\beta)$. Based on the noise integral, Eqs. (6) and (7), the on-resonance cavity emission is independent of $\Delta\nu_d$, while the off-resonance contributions (emission from the termination and the lossy part of the transmission line) are proportional to $\Delta\nu_d$. In many circumstances, $\Delta\nu_d$ needs to be just large enough to cover the resonance and minimize the noise contribution from n_b and $(1-\eta)n_T$. See text for a detailed discussion.

$$\begin{aligned} \sigma_{\text{tot}}^2 = & \frac{1}{\tau} \int_{\nu_c - \Delta\nu_d/2}^{\nu_c + \Delta\nu_d/2} d\nu h^2 \nu^2 \left\{ n_T(1-\eta) \right. \\ & + \frac{\eta}{B(\omega)} [n_T \kappa_l \kappa_m + n_b A(\omega)] \\ & + n_T^2(1-\eta)^2 + \frac{\eta^2}{B^2(\omega)} [n_T^2 \kappa_l^2 \kappa_m^2 + n_b^2 A^2(\omega)] \\ & + 2n_T(1-\eta) \cdot \frac{\eta}{B(\omega)} [n_T \kappa_l \kappa_m + n_b A(\omega)] \\ & \left. + \frac{2\eta^2}{B^2(\omega)} [n_T n_b \kappa_l \kappa_m A(\omega)] \right\}. \end{aligned} \quad (6)$$

The origin of various terms in the integral is clear when one analyzes first the Poisson ($\propto n$) portion of the σ^2 integral. With a photon counter that has a detection bandwidth $\Delta\nu_d$ around ν_c , the integral takes place within $\nu_c \pm \Delta\nu_d/2$. It has a component that comes from the cavity emission and an off-resonance component that comes from the “background,” including the coaxial termination and the lossy transmission line; see Fig. 6.

Note that, when the first stage readout is a linear amplifier with $G_a \gg 1$, the number of photons at detection is increased vastly and the Poisson noise becomes negligible (see Sec. 3C in [39]). Despite the advantage of eliminating the Poisson contribution, the increasing dominance of the $1/2$ photon term in the noise makes linear amplifiers highly undesirable toward high frequencies.

We next include effects of a non-negligible dark count rate (DCR). The dark counts in state-of-the-art SMPDs come from several sources [15,17]. The first contribution to the DCR comes from the photon/microwave environment. Usually this is characterized as a residual effective temperature in the photon field, denoted by T_γ (with photon occupation number n_γ), in the 30–50 mK range. This photon noise would add to the existing background n_b and $(1-\eta)n_T$. The noise integral of all these background sources is proportional to $\Delta\nu_d$.

There are other sources of dark counts associated with the transmon qubit in the SMPD: (a) thermal excitation of the qubit and (b) false positives in the determination of the quantum state, which is ultimately related to the qubit coherence time T_1 . This qubit-related component, which we call $\delta\nu_{\text{DCR}}$, is independent of $\Delta\nu_d$. The Poisson noise from $\delta\nu_{\text{DCR}}$ adds in quadrature to Eq. (6) as an additional variance outside the integral.

Similar to the noise, the average (axion) signal is given by $n_A \kappa_a \kappa_m \eta \int_{-\infty}^{\infty} B^{-1}(\omega) (d\omega/2\pi)$. If $\Delta\nu_d$ is large enough to cover the Lorentzian line, we can use $\int_{-\infty}^{\infty} (\omega^2 + a)^{-1} d\omega = \pi/\sqrt{a}$ and $\int_{-\infty}^{\infty} (\omega^2 + a)^{-2} d\omega = \pi/(2a^{3/2})$ to obtain closed-form results. In this limit, we have

$$\begin{aligned} \int d\omega A(\omega)/B(\omega) &= 2\pi\Delta\nu_d - \kappa_m \kappa_l \int d\omega/B(\omega), \\ \int d\omega A(\omega)/B^2(\omega) &= \int d\omega/B(\omega) - \kappa_m \kappa_l \int d\omega/B^2(\omega), \\ \int d\omega A^2(\omega)/B^2(\omega) &= 2\pi\Delta\nu_d - 2\kappa_m \kappa_l \int d\omega/B(\omega) \\ &\quad + \kappa_m^2 \kappa_l^2 \int d\omega/B^2(\omega). \end{aligned}$$

Using these integrals recursively, we arrive at an expression for the frequency-integrated squared SNR R_{SMPD} , which is directly comparable to the scan rate figure of merit,

$$\begin{aligned} R_{\text{SMPD}} &= 2n_A^2 \kappa_a^2 \kappa_l^{-1} \eta^2 \beta^2 (1+\beta)^{-2} \\ &\quad \{\delta + D\Delta + E\beta(1+\beta)^{-1} + F\beta^2(1+\beta)^{-3}\}^{-1}, \end{aligned} \quad (7)$$

where the DCR appears as $\delta \equiv \delta\nu_{\text{DCR}}(\pi\kappa_l)^{-1}$, a “dark count rate” normalized by the unloaded cavity line width. $\Delta \equiv \Delta\nu_d(\pi\kappa_l)^{-1}$, is a similarly dimensionless measure of the detector bandwidth. The terms in the expression are

organized based on their dependence on Δ and β . The D , E , F functions are

$$\begin{aligned} D(\eta, n_T, n_b) &= n_T(1 - \eta) + n_b\eta \\ &\quad + [n_T(1 - \eta) + \eta n_b]^2 + n_\gamma, \\ E(\eta, n_T, n_b) &= 2\eta[n_T - n_b - 2\eta n_b^2 + 2n_T^2(1 - \eta) \\ &\quad + 2n_T n_b(2\eta - 1)], \\ F(\eta, n_T, n_b) &= 4(n_T - n_b)^2 \eta^2. \end{aligned}$$

In the remainder of this section, various scaling relations are discussed. In Sec. III B, the numerical values supported by the results achieved by ongoing experiments have been inserted into the equations to assess the experimental parameters required to reach various main benchmarks.

Although the results seem verbose, the benefit of obtaining these analytic expressions becomes clear when various limits are taken. For example, whether $\delta\nu_{\text{DCR}}$ significantly affects the scan rate R_{SMPD} depends on how it compares with $n_T \Delta\nu_d$. Completeness also helps clarify the dependence of the scan rate to the haloscope quality factor Q , a subject of great importance in the design of axion search experiments.

Under truly background-free conditions, $n_\gamma, n_b, n_T \rightarrow 0$, or more precisely, $D \ll \delta/\Delta$, the axion experiment becomes limited only by the DCR,

$$R_{\text{SMPD}} \rightarrow 2\pi n_A^2 \kappa_a^2 \eta^2 \beta^2 (1 + \beta)^{-2} / \delta\nu_{\text{DCR}}. \quad (8)$$

Intriguingly, the scan rate does not depend on κ_l and increases weakly with β in this limit. We understand this

observation as follows. Generally, a large Q helps the scan rate by (1) resonantly enhancing the axion-to-photon conversion power and (2) reducing the bandwidth to background photons. On the other hand, high Q hurts the scan rate by (3) reducing the search bandwidth and (4) restricting the output from the haloscope. With no background, (2) is gone, the effects of (1), (3), and (4) cancel, and the benefit of high Q disappears.

When the noise is dominated by photons from the terminating load and other sources of off-resonance background [negligible $\delta\nu_{\text{DCR}}$ and $D\Delta \gg E\beta(1 + \beta)^{-1}$], we have

$$R_{\text{SMPD}} \rightarrow \frac{2\pi n_A^2 \kappa_a^2 \beta^2 (1 + \beta)^{-2}}{D(\eta, n_T, n_b) \Delta\nu_d}. \quad (9)$$

Although the scan rate also *appears* independent of κ_l or Q_0 , it is inversely proportional to the detector bandwidth. If the bandwidth can be engineered, it should be the narrowest possible to still cover the haloscope linewidth ν/Q_L , i.e., $\Delta\nu_d \sim \nu\kappa_l(\beta + 1)$. Equation (9) becomes proportional to $\beta^2(1 + \beta)^{-3}$ in such global optimization, which peaks at $\beta = 2$. In other words, the conventional wisdom for axion searches applies. Figure 7 illustrates how this optimization takes place on the $\Delta - \beta$ plane under various experimental conditions. In the SMPD scheme, $\Delta\nu_d$ is determined jointly by the linewidth of the buffer and the waste resonators [Eq. (D5) in [17]].

In this limit, a competitive axion search must involve synchronized tuning of the SMPD band and the haloscope resonance over a wide-frequency range. So far,




FIG. 7. These contour plots show the scan rate R_{SMPD} given in Eq. (7) versus the coupling parameter β and the photon counter bandwidth $\Delta = \Delta\nu_d(\kappa_l\pi)^{-1}$. As discussed in the text, when Δ is fixed, the scan rate increases monotonically with β as depicted by the arrow at $\Delta = 8$ (or any other horizontal line). When Δ is engineered to be just large enough to cover the resonance, e.g., along the red line [$\Delta = (1 + \beta)$], R_{SMPD} maximizes at the locations of the asterisk: $\beta \sim 8$ (left) when the emission is dominated by the cavity [Eq. (10)], or $\beta \sim 2$ (right) when the photon noise is dominated by the off-resonance background. In between either limit, the scan rate maximizes at $\beta \sim 5$ (middle). The contours are normalized to the maximum in each plot. The efficiency due to dissipation η effectively changes the contribution from the off-resonance background photons. The SMPD “dark count” due to the residual photons T_γ has similar effects. The white-out regions in the lower right are where the SMPD bandwidth $\Delta\nu_d$ starts to be comparable to the linewidth ν_c/Q .

the state-of-the-art SMPD can only be tuned by $< 2\%$. Significant development is still needed on this front.

Next we consider the cavity photon-dominated regime. When $n_T \gg n_b \Delta$ and $\eta \sim 1$, the photon noise is dominated by the haloscope emission. Therefore,

$$R_{\text{SMPD}} \rightarrow \frac{n_A^2 \kappa_a^2 \kappa_l^{-1} \beta^2 (1 + \beta)^{-2}}{n_T \beta (1 + \beta)^{-1} + 2n_T^2 \beta^2 (1 + \beta)^{-3}}. \quad (10)$$

In this limit, the bandwidth of the photon counter $\Delta\nu_d$ is not detrimental to R_{SMPD} because the noise contribution is dominated by the Lorentzian integral around the line (Fig. 6).

A noteworthy result stemming from Eqs. (7)–(10) is their dependence on β . For photon counters with a given bandwidth $\Delta\nu_d$, the scan rate R_{SMPD} increases monotonically in β and asymptotically approaches a maximum, until the resonance width approaches $\Delta\nu_d$. Although, in practice, the rate increases minimally beyond $\beta > 10$.

Consider the numerical values from [18]: $\nu_c = 7.37$ GHz, $Q_L = 2.25 \times 10^5$, $\beta = 3$. The linewidth of the cavity is then 130 kHz, still substantially narrower than the SMPD bandwidth $\Delta\nu_d = 700$ kHz [clear from Fig. 2(a) of [18]]. According to Eq. (7), by going to $\beta = 10$ the scan rate can increase by 47%.

As mentioned earlier, if the SMPD bandwidth can be optimized together with the rest of the experiment, the answer might be different. In Fig. 7, we illustrate the optimization of β in both the fixed—and engineered— $\Delta\nu_d$ scenarios under different loading conditions.

III. OPTIMIZING THE EXPERIMENTAL CONDITIONS

A. Scan rate comparison between AMP and SMPD

We proceed to use Eqs. (5) and (7) to compare the scan rates of two identical haloscopes read out by (1) a linear

amplifier in a single-quadrature squeezed setup and (2) a photon counter such as the SMPD.

If the photon background vanishes, option (2) is vastly superior and the choice between the two is obvious. However, the main axion haloscope projects that had produced meaningful constraints on $g_{a\gamma\gamma}$ still find the effective photon temperature in the cavities to be between 90 and 300 mK, elevated compared to the base temperature of the dilution refrigerator. At this temperature range, the photon noise is not negligible.

To apply the results in the previous sections to a broader range of experiments, T is set between 0.01 and 1 K. The termination temperature T_b is assumed to be $T/3$ unless $T < 0.03$, in which case T_b bottoms out at 0.01 K. Using the haloscope temperature T and the search frequency ν as parameters, we plot the contours of ratio $r \equiv R_{\text{SMPD}}/R_{\text{AMP}}$ (Fig. 8). In the left panel, the SMPD scan rate is compared against the baseline, no squeezing $G_s = 1, \beta = 2$. In the center panel, squeezing is turned on for the amplifier option. With the squeezing on, the contours shift to the right as expected.

In the right panel, the SMPD has a dark count rate $\delta\nu_{\text{DCR}} = 10^3 \text{ s}^{-1}$. At the SQL, the equivalent DCR from an amplifier within the axion signal bandwidth is $\sim \nu/Q_a$, or about 10^4 s^{-1} at 10 GHz. In comparison, the DCR of the state-of-the-art SMPDs (few tens of s^{-1}) is already significantly better. That is why a DCR of 10^3 s^{-1} does not significantly change the overall conclusion except at the low-frequency end.

The $r = 1$ contours represent the same scan rate for the amplifier and the SMPD. From these figures, it is clear that when $T < 150$ mK and at high frequencies, SMPDs become overwhelmingly advantageous as r becomes orders of magnitude greater than 1. However, two things need to happen before the axion experiments broadly adopt SMPDs or similar microwave photon counters for the readout: (a) understanding the cause for the elevated photon




FIG. 8. The contour plots for r , the scan rate ratio $R_{\text{SMPD}}/R_{\text{AMP}}$. Left: for a SQL amplifier operating in the standard configuration ($G_s = 1, \beta = 2$), $\delta\nu_{\text{DCR}} = 100 \text{ s}^{-1}$ for the SMPD. Center: with an amplifier pair in a squeezed state configuration ($G_s = 20, \beta = 7$), $\delta\nu_{\text{DCR}} = 100 \text{ s}^{-1}$. Right: no squeezing on the amplifier but with a DCR of $\delta\nu_{\text{DCR}} = 10^3 \text{ s}^{-1}$.

temperature in the cavity and resolving the issue and (b) upgrading the SMPDs to allow for robust, synchronized tuning with the haloscope. Typically, an axion haloscope has a tuning range $\gtrsim 20\%$ and the microwave photon counter must match that. Before these developments, one can continue to make good progress using SQL amplifiers. Indeed, in Fig. 8, the $T = 0.2$ K line crosses the equality curve ($r = 1$) at 11 (9) GHz and the $r = 2$ line at 14 (12) GHz, with (without) squeezing.

B. Reaching the KSVZ/DFSZ limits for the postinflationary frequencies

The original strong CP -resolving axion model was quickly ruled out by data. The minimal extension that is still consistent with the data is known as the KSVZ model after its proposers Kim [24] and Shifman *et al.* [25]. Another benchmark model that had received broad recognition is the DFSZ model proposed by Dine *et al.* [27] and Zhitnitsky [26]. The DFSZ model predicts a coupling strength ($g_{a\gamma\gamma}$) $2.6\times$ smaller than the KSVZ value [5].

So far the ADMX and CAPP experiments have reached the DFSZ theoretical benchmark at around 1 GHz. The only other project nearing the theoretical limits is HAYSTAC [13,14], coming within a factor of 2 of the KSVZ value over $\sim 2\%$ – 3% of bandwidth between 4 and 5 GHz.

In this study we are most interested in the frequency range ν_a between 1 and 30 GHz (centimeter wave), corresponding to the postinflationary axion production scenario [5,19–23]. One main difficulty in scaling the ~ 1 GHz experiments to high frequencies is due to loss of volume. From Eqs. (2) and (5), the scan rate is $\propto V^2$. Naive dimensional scaling of the cavity geometry in all three dimensions leads to a steep scan rate reduction at high frequencies ($\propto \nu^{-6}$). In addition, Q_0 (or, equivalently, $\nu\kappa_l^{-1}$) decreases at high frequency due to the anomalous skin effect as $\sim \nu^{-2/3}$.

As discussed in Sec. II B, operating the receiver in a squeezed state can improve the scan rate beyond the SQL. However, the main limitation of this approach comes from the dissipative loss between the squeezer and the haloscope, in particular, from the circulator. In practical systems, the gain in scan rate is often limited to around a factor of 2. While this is significant enough to implement on working projects, the sensitivity gap to the DFSZ value at > 5 GHz is too vast for this method to overcome. In principle, the gain could be unlimited if the circulator is lossless or can be avoided altogether. Research and development (R/D) efforts are being made to pursue this but should be characterized as in early stages.

It is possible to improve the scan rates by significantly increasing the quality factor Q using superconductive coatings, such as niobium alloys [40] or tapes of high- T_c superconductors. The most impressive results so far are those by the CAPP group using high- T_c superconducting

tapes that have been shown to survive in strong B fields. It is unclear whether superconducting high- Q cavities can effectively work with photon-counting detectors. As discussed in Sec. II C, the frequency and the bandwidth of the SMPD should match those of the haloscope in the off-resonance background-dominated regime. Synchronized tuning over a wide-frequency range will be extremely challenging for these ultrahigh- Q systems. In the DCR-dominated regime, the scan rate becomes independent of Q [Eq. (8)], rendering null the benefits of having high-quality factors. Even restricting to < 5 GHz and operating with an amplifier face significant challenges in microphonics and cryogenics when it comes to tuning a system with $Q_L \gtrsim 10^6$ over an appreciable search range.

Other groups have proposed to use extremely high-field magnets for axion searches, motivated by the strong dependence of the scan rate to B_0 , i.e., $R \propto B_0^4$. Advanced magnet engineering is used to produce a field strength 20 T or higher, leading to a seemingly vast advantage over standard NbTi magnets at around 8–9 T. One challenge with this approach is that the very-high-field magnets use resistive coils that incur staggering power consumption. Reference [41] reports the 43 T field requires power consumption exceeding 20 MW, which seems prohibitive for typical scales of axion experiments. The CAPP group reports continuous operation of a lower power 18T high T_c magnet [42]. Nevertheless, these high-field magnets often come with relatively small bore diameters, resulting in a figure of merit $B_0^2 V < 0.3 \text{ T}^2 \text{ m}^3$, substantially lower than that of a 8–9 T magnet with an 800 mm bore.

Based on these considerations, we focus the subsequent discussions on two R/D thrusts aiming to reach DFSZ for much of the postinflationary mass range: (1) breaking the volume scaling of ν^{-3} and vacuum state squeezing and (2) developing the SMPD for axion dark matter searches. In the next two sections we discuss these approaches in turn.

1. Volume enhancement and squeezed amplification

The loss in volume at short wavelengths (λ) can be circumvented by adopting novel resonators that fill a volume $\gg \lambda^3$. In [43–45], we propose two such designs: the thin-shell “wedge” cavity and multicell “beehive” cavity (Fig. 9). These geometries can scale up in volume at a fixed resonant frequency and have Q ’s similar to conventional cylindrical cavities. The wedge design has been experimentally demonstrated at 7.5 GHz [46]. This approach is being vigorously pursued by the ADMX-VERA (Volume Enhanced Resonant Axion) working group. Haloscopes based on noncavity resonators can also achieve similar volume enhancements. The most well-developed concepts are the plasma haloscopes [47,48] and the dielectric disk haloscopes [32,49].




FIG. 9. The VERA haloscopes. From left to right: a triple-wedge cavity (5–7 GHz), the cross section view of its field distribution [43,44,46]; a beehive cavity consisting of overlapping cylinders (5–7 GHz) [45], its field distribution; a notional 20-wedge cavity at ~20 GHz, and a beehive cavity at ~20 GHz.

For the purpose of readout comparison, we use the VERA concept as the baseline. Much of the results and conclusions are applicable to other volume-enhanced haloscopes. Scaling from [13], we assume VERA can increase the volume of the resonant cavity at 4.14 GHz by a factor of 7–8 and change the frequency scaling from $V \propto \nu^{-3}$ to $V \propto \nu^{-\alpha}$, with an $\alpha < 3$. In principle, an ideal volume-filling multiwedge cavity or a beehive cavity can completely eliminate the frequency dependence, corresponding to $\alpha = 0$. However, there are still uncertainties regarding how this technology scales in frequency, despite great promises shown at < 8 GHz. Practical limitations at shorter wavelength might induce phase errors that lead to lower form factors and, effectively, loss in volume. Recognizing that, we consider two stages in the VERA program: VERA-1 with $\alpha = 1$ and VERA-2 with $\alpha = 0.5$. Both are normalized at 7 GHz using a three-wedge thin-shell cavity as the baseline model [46]. These two scenarios correspond to retaining ~20% and ~50% of the effective volume at 30 GHz compared to 7 GHz.

To make reliable forecasts, we normalize the scan rate expressions obtained in earlier sections by achieved exclusion limits reported by HAYSTAC in [13]. In 105 days of operation, they ruled out $g_{a\gamma\gamma}$ above 1.38 KSVZ at 90% confidence level over 1.76% of bandwidth around 4.14 GHz. The receiver is operated in a squeezed state setup with $G_s = 21$ and $\beta = 7.1$. The cavity photon temperature T is measured to be 250 mK and the termination T_b is at 61 mK. These parameters are used in normalizing the amplifier scan rate R_{AMP} using Eq. (5) and, by extension, Eq. (7).

For convenience, we normalize the “KSVZ,” “DFSZ,” and other benchmark contours to correspond to 100 days of operation over a 5% fractional bandwidth. The HAYSTAC results have been obtained after several years of optimization. What they have achieved should be considered results of a realistic experimental duty cycle and actual sensitivity and efficiencies of a well-tuned axion experiment using modern cryogenic and readout systems. These

reality-calibrated rates should be a reliable forecast for reaching particular benchmarks over a bandwidth of 30% after a few years of operation.

Using the assumptions in Table I, it can be shown that either VERA-1 or VERA-2 can reach time the KSVZ value benchmark to 8 GHz with SQL amplifiers. Additionally, 0.5 KSVZ is reachable by VERA-1 up to 15 GHz and VERA-2 up to 18 GHz (Fig. 10, left panel). The presence of 1/2 photons in the noise creates two limiting effects that are noticeable: the inability to reach the DFSZ benchmark above ~4 GHz even with VERA-2 and the ineffectiveness in improving the experimental reach by cooling the cavity below 100 mK.

In the middle panel of Fig. 10, squeezing is enabled with VERA-2 with the standard parameters $G_s = 20, \beta = 7$. It produces the expected results of shifting the contours toward high frequencies. The DFSZ limit is reachable in this case up to 4 GHz, while the KSVZ benchmark can be reached < 10 GHz.

TABLE I. The parameters used in Figs. 8–11. The scan rates have been normalized by achieved $g_{a\gamma\gamma}$ limits in [13]. Therefore, these contour plots have built in realistic integration and scanning inefficiencies—including calibration, tuning, thermal settling, etc.

Parameter	Description	Value	Unit
Q_0	Quality factor [13]	$4.7 \times 10^4 (\nu/4.14 \text{ GHz})^{-2/3}$	...
V_h	Cavity volume [13]	$3.94 (\nu/4.14 \text{ GHz})^{-3}$	l
V_{vera1}	Vol. VERA-1	$36 (\nu/4.14 \text{ GHz})^{-1}$	l
V_{vera2}	Vol. VERA-2	$27 (\nu/4.14 \text{ GHz})^{-0.5}$	l
B_0	Magnetic field	8	T
T_b	If $T > 0.03$ K	$T/3$	
T_b	If $T < 0.03$ K	0.01	K
η	Efficiency	$\sqrt{0.7}$	...
$\Delta\nu_d$	Detector bandwidth	7×10^5	Hz
$\Delta\nu_d$	<i>Ibid.</i> , for Fig. 11	$\nu/5$	...




FIG. 10. Combining the VERA approach and quantum techniques, it is possible to cover most of the postinflationary axion mass range (>30 GHz) to DFSZ. Left: VERA-2 cavity with SQL readout. The contours correspond to scan rates to reach DFSZ at 3σ (blue dashed), DFSZ at 90% confidence level (blue), KSVZ at 3σ (red dashed), KSVZ at 90% (red), $2 \times$ KSVZ at 90% (red dotted), and completing a 5% search bandwidth in 100 days of operation around that given frequency. Center: VERA-2 cavity with the squeezed state readout ($G_s = 20, \beta = 7$). Right: VERA-2 cavity with the SMPD, $\text{DCR} = 100 \text{ s}^{-1}$. Other experimental parameters are listed in Table I.

2. Using SMPD in axion searches

Above 4 GHz, the DFSZ limit is stubbornly out of reach even for VERA-2 *with* vacuum state squeezing. The real game changer at these high frequencies must be the introduction of the SMPD in axion searches. In the right panel of Fig. 10 we show the combination of VERA and the SMPD, which shows the DFSZ limit becoming accessible for the entire < 30 GHz range. Impressively, this is achieved even with a 100 s^{-1} DCR in the detector and the cavity running at $\gtrsim 100$ mK.

In Sec. II C, we observe three noise regimes for the SMPD. In the off-resonance background-dominated regime (such as in Fig. 7, right panel), the scan rate is optimized with high Q , $\beta = 2$, and the SMPD bandwidth $\Delta\nu_d$ just large enough to cover ν_c/Q_L . In the on-resonance cavity emission-dominated regime (Fig. 7, left panel), the scan rate increases with β but reaches a diminishing return when $\beta \sim 10$. In both cases, the SMPD band must be tuned together with the haloscope resonator, which usually has a frequency range $\gtrsim 20\%$. In a science integration, the operating efficiency would severely suffer if the tuning range of the detectors is too small compared to the haloscope. We consider a 20% frequency tunability a necessary requirement for adopting SMPDs in axion searching programs aiming at reaching the theoretical benchmark over significant mass ranges. As Fig. 4 depicts, the input band is determined by the tunable resonator on the left, dubbed the buffer resonator in [15–17]. In the current implementation, the tuning range is still limited to $< 3\%$.

The third regime corresponds to the background-free limit, where Eq. (8) applies. In this case, the scan rate of the haloscope is independent of Q in the search mode. Therefore, the SMPD band just needs to cover the cavity haloscope's frequency range, e.g., 20%, and does not need to be tunable. This is also the limit in which the BREAD experiment aims to operate, using ultralow-background

Cooper pair-breaking detectors of single infrared photons [50]. Note that a higher Q offers a measurement of the axion mass (to a precision of $1/Q$) in the event that a positive detection is made.

In Fig. 11, we see the SMPD-based axion experiment enters the background-free regime when T is less than about 30 mK, depending on the frequency. The SMPD band is fixed, with a bandwidth of 20%. The haloscope cavity is tunable within this band, with $\beta = 10$. After completing each 20% range, a new set of cavity and SMPD are swapped in. The three panels represent DCR at 100, 10, and 1 s^{-1} . The DFSZ limit is reachable up to 12 GHz at $\delta\nu_{\text{DCR}} = 1 \text{ s}^{-1}$. This set of forecasts assumes the haloscope volume scales as ν^{-3} from [13], i.e., without using any of the volume enhancement resonators discussed in the previous section. This device-heavy approach carries its own set of risks and benefits compared to VERA and should be viewed as highly complementary.

C. Achieving strong coupling in high-volume axion haloscopes

In the first experimental demonstration of VERA, the coupling parameter β with a single antenna was measured to be ~ 0.1 [46], substantially lower than the optimal value of 2. In earlier sections of this paper, we discuss many scenarios in which overcoupling beyond $\beta = 2$ is desired. It is worth investigating how one can achieve sufficient coupling in a large-volume cavity.

In this section, we argue that achieving the same β is inherently harder in a cavity with $V \gg \lambda^3$ than in a cavity with $V \sim \lambda^3$. With the same loaded Q (determined by β and Q_0), the steady-state electromagnetic energy density U_{EM} on resonance is the same for the typical-volume cavity and the high-volume cavity. For the large-volume cavity, the Ohmic power dissipation rate on the cavity walls κ_l is larger




FIG. 11. It is possible to significantly improve the scan rate to reach DFSZ up to 12.5 GHz even if $V \propto \nu^{-3}$. This can be done by (1) reducing the haloscope photon temperature to $\lesssim 30$ mK to be in the background-free regime and (2) progressively reducing the dark count rate of the SMPD. From left to right shows the experimental reach using a series of cavities directly scaled up in frequency from [13], with DCR from 10^2 , 10, to 1 s^{-1} . In these plots, broadband operation of the SMPD is assumed ($\Delta\nu_d/\nu = 0.2$). In each line, the turning point from horizontal to vertical branches marks the $g_{a\gamma\gamma}$ reach in the background-free regime. These points collectively trace out the $T = h\nu/k$ curve.

due to the larger surface area [Eq. (19.150) in [51]]. The strength of the port κ_m must then be proportionally larger to maintain the same coupling/Ohmic-loss ratio prescribed by $\beta = \kappa_m/\kappa_l$.

Another way to understand this is to consider the ringdown time Q/ω_c , which is the same for both $V \gg \lambda^3$ and $V \sim \lambda^3$ cavities [Eq. (19.153) in [51]]. To drain the much higher total energy $U_{EM} \times V$ stored in the $V \gg \lambda^3$ cavity within the same ringdown time, the physical size of the port must be proportionally larger than the port on the typically sized cavity. This also explains the proportionally larger signal strength from axion-converted photons in a high-volume haloscope at a fixed wavelength.

In Fig. 12, we illustrate these two cavities ($\sim \lambda^3$ vs $\gg \lambda^3$) using the simplest rectangular geometry. The top panel represents a $\sim \lambda^3$ cavity with a coupling at the desired β parameter. The middle and bottom panels represent the $\gg \lambda^3$ cavity with proportionally larger coupling ports to maintain the same target β (and the ringdown time). In the middle panel, it is apparent that the larger port implemented locally would severely distort the field distribution.

In Fig. 13 (left) we demonstrate this behavior using finite element analysis performed in COMSOL-RF. We start with a simple rectangular cavity dimensioned to resonate at 10 GHz. Its volume, $6L = 221\lambda^3$, was chosen to emulate a system with $V \gg \lambda^3$. When coupling to the cavity with a




FIG. 12. Top: a rectangular cavity coupled to a readout port (or a slot antenna) with $V \sim \lambda^3$. We are interested in reading out the fundamental mode that is polarized in the x direction whose frequency is set mostly by the smallest dimension ($\lambda/2$). We assume the port here provides the optimal β parameter. Middle: the volume of the cavity is grown ($V \gg \lambda^3$) by increasing its dimension in x [43]. The fundamental mode with uniform polarization in x still exists; however, to maintain the same β , one must increase the coupling strength of the port proportionally. Doing so with a single large port would drive the mode away from the port instead of achieving the goal (Fig. 13). Bottom: instead, one must implement an array of ports distributed throughout the volume of the cavity and coherently add up the signal in a phase-matching summing network. See text for more discussions.




FIG. 13. Left: electric field contour plots simulated in COMSOL-RF for a $V \gg \lambda^3$ rectangular cavity with two different readout port sizes. When a small readout port is added (top contour plot), coupling is low because energy losses at the cavity walls are proportionally larger than the energy extracted by the port. The cavity mode also begins to localize away from the port location. When the port is grown in an attempt to increase the extracted energy (bottom contour plot), this mode localization increases, leading to further degradation of the coupling. For both plots, the cavity dimensions were $1000 \times 400 \times 15$ mm, and the resonant frequency of the mode was 10 GHz. The readout ports were simulated using rectangular waveguides. Right: receiver chain coupling β achieved with different sized readout ports for the same $V \gg \lambda^3$ cavity. The scale 1 waveguide had a broad wall length of 30 mm for a cutoff frequency of 5 GHz. As expected, increasing the port size leads to a decrease in β as the degree of mode localization increases.

small rectangular waveguide (upper contour plot) the energy lost at the cavity walls is much higher than the energy extracted via the waveguide, leading to a low β . The mode also begins to localize on the opposite end of the cavity. When the size of the port increases in an attempt to extract more energy (lower contour plot), the degree of mode localization increases, further decreasing β .

In Fig. 13 (right) we show simulated β values for a variety of waveguide sizes. In each case, we computed β by taking

$$\beta = \frac{P_{\text{readout}}}{P_{\text{cavity walls}}} = \frac{\omega_c \frac{U}{Q_r}}{\omega_c \frac{U}{Q_0}} = \frac{Q_0}{Q_r} = \frac{\int_{\text{port output}} d\vec{A} \cdot \vec{\hat{n}}}{\int_{\text{cavity walls}} d\vec{A} \cdot \vec{\hat{n}}}, \quad (11)$$

where P_{readout} is the power lost through the waveguide output, $P_{\text{cavity walls}}$ is the power lost at the cavity walls, ω_c is the resonant frequency of the cavity, U is the stored energy, Q_r quantifies coupling to the receiver, $\vec{S}$ is the Poynting vector, and $\hat{n}$ is a unit vector normal to the surface of integration. Instead of extracting more power from the cavity, β falls off rapidly for larger port sizes.

We also ran a series of similar simulations where several identically sized waveguides (scale = 1) were distributed symmetrically across the cavity volume. In these configurations, the mode remained centralized, and, by summing the port outputs, more energy could be extracted from the cavity relative to the losses at the cavity walls, increasing the coupling. The specific coupling achieved depended upon the number and geometric distribution of the waveguides. Our example cases demonstrated couplings

between 0.4 and 5, representing undercoupling, critical coupling, and overcoupling to the cavity system.

In the bottom panel of Fig. 12, we imagine a readout scheme that could permit such a geometric distribution of readout ports. An array of ports similar to the one in the top panel is distributed throughout the volume of the cavity. The signal from each port is summed coherently in a low-loss, impedance-matched tree of transmission lines. This scheme provides a synthesized “port” that is proportional to the volume of the cavity, in phase for the axion signal as long as the extent of the cavity is less than the spatial coherence length of the cosmological axion (a few 100 m).

The large volume of a beehive cavity is achieved by having a large number of cells. It is easy to conclude that, when the coupling of a single port required to reach a certain β (for the entire beehive) exceeds the coupling between cells, the port will draw power only from the cell at which it is located. After reaching that point, multiple ports + signal summing will be needed to utilize the entire volume of the beehive for the axion searches.

Such a summing tree was proposed for an axion experiment in [43] in the context of mode selection. Here we argue that even to maintain β one must implement such structures. It is becoming clear that summing tree implementation should become a focus of R/D in the VERA program.

IV. CONCLUSION

In this paper, we provide a comprehensive comparison of linear amplifiers and microwave photon counters in axion dark matter experiments. We derive expressions for the scan rate of either sensing techniques, assuming a range of

realistic operating conditions and detector parameters, over the frequency range of 1–30 GHz.

As expected, photon-counting detectors such as the SMPD can be overwhelmingly advantageous under low background, at high frequencies ($\nu > 5$ GHz), *if* they can be implemented with appropriate parameters. We also note an expanded applicability of off-resonance photon background reduction, including the single-quadrature state squeezing, for scan rate enhancements and a much broader appeal for operating the haloscope resonators in the over-coupling regime, up to $\beta \sim 10$.

In Sec. III B we investigate the experimental reach for a wide range of scenarios. With the goal of eventually reaching the DFSZ limit up to 30 GHz, we identify several R/D thrusts:

- (1) Develop high-volume ($\gg \lambda^3$) haloscope resonators such as VERA, the plasma haloscopes, and the dielectric disk haloscopes. A good milestone is the reduction in the steepness of frequency scaling from $V \sim \nu^{-3}$ to $\nu^{-0.5}$.
- (2) Develop SMPDs or similar microwave photon-counting detectors to have either of the following properties:
 - (a) bandwidth $\lesssim 0.1\%$ and tunable range $\gtrsim 20\%$; and
 - (b) bandwidth $\sim 20\%$ and very low DCR ($\lesssim 1 \text{ s}^{-1}$).
- (3) Lower the photon temperature of the haloscope resonators and various backgrounds to $\lesssim 30$ mK.

Item (1) alone should represent a robust approach to reach the KSVZ benchmark up to 10 GHz using existing sensor technology in a large- $B_0^2 V$ magnet. The combination (2b) + (3) will reach the DFSZ benchmark up to 12 GHz

and the KSVZ benchmark to 23 GHz (Fig. 11). The combinations (1) + (2a) or (1) + (2b) + (3) will cover the entire < 30 GHz range to the DFSZ limit (Fig. 10).

ACKNOWLEDGMENTS

The authors thank P. Bertet and E. Flurin for suggesting the cooling of the termination as a means to reduce photon noise, motivating overcoupling the cavity haloscope. We acknowledge C. Braggio, G. Carugno, and B. Majorovitz for useful discussions while visiting INFN and MPP. This material is based upon work supported by the National Science Foundation under Grant No. 2209576. We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC), 521528828. C. L. K., C. L. B., T. A. D., and N. A. K.’s investigation on the readout of high-volume axion haloscopes was supported by the Department of Energy, Laboratory Directed Research and Development program at SLAC National Accelerator Laboratory, under Contract No. DE-AC02-76SF00515. A.C. is supported via the Fermi National Accelerator Laboratory (Fermilab), a U.S. Department of Energy, Office of Science, Office of High Energy Physics HEP User Facility. Fermilab is managed by Fermi Research Alliance, LLC (FRA), acting under Contract No. DE-AC02-07CH11359. C. Z. is supported by a KIPAC Porat Fellowship.

DATA AVAILABILITY

No data were created or analyzed in this study.

-
- [1] R. D. Peccei and H. R. Quinn, *CP conservation in the presence of pseudoparticles*, *Phys. Rev. Lett.* **38**, 1440 (1977).
 - [2] S. Weinberg, *A new light boson?*, *Phys. Rev. Lett.* **40**, 223 (1978).
 - [3] F. Wilczek, *Problem of strong P and T invariance in the presence of instantons*, *Phys. Rev. Lett.* **40**, 279 (1978).
 - [4] L. F. Abbott and P. Sikivie, *A cosmological bound on the invisible axion*, *Phys. Lett.* **120B**, 133 (1983).
 - [5] D. J. Marsh, *Axion cosmology*, *Phys. Rep.* **643**, 1 (2016).
 - [6] P. Sikivie, *Experimental tests of the “invisible” axion*, *Phys. Rev. Lett.* **51**, 1415 (1983).
 - [7] N. Du *et al.* (ADMX Collaboration), *Search for invisible axion dark matter with the axion dark matter experiment*, *Phys. Rev. Lett.* **120**, 151301 (2018).
 - [8] L. Zhong *et al.*, *Results from phase 1 of the HAYSTAC microwave cavity axion experiment*, *Phys. Rev. D* **97**, 092001 (2018).
 - [9] A. K. Yi, S. Ahn, C. Kutlu, J. M. Kim, B. R. Ko, B. I. Ivanov, H. Byun, A. F. van Loo, S. Park, J. Jeong *et al.*, *Axion dark matter search around 4.55 μeV with Dine-Fischler-Srednicki-Zhitnitskii sensitivity*, *Phys. Rev. Lett.* **130**, 071002 (2023).
 - [10] 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).
 - [11] A. O. Sushkov, *Quantum science and the search for axion dark matter*, *PRX Quantum* **4**, 020101 (2023).
 - [12] M. Malnou, D. A. Palken, B. M. Brubaker, L. R. Vale, G. C. Hilton, and K. W. Lehnert, *Squeezed vacuum used to accelerate the search for a weak classical signal*, *Phys. Rev. X* **9**, 021023 (2019).
 - [13] K. M. Backes *et al.*, *A quantum enhanced search for dark matter axions*, *Nature (London)* **590**, 238 (2021).
 - [14] M. J. Jewell *et al.* (HAYSTAC Collaboration), *New results from HAYSTAC’s phase II operation with a squeezed state receiver*, *Phys. Rev. D* **107**, 072007 (2023).
 - [15] R. Lescanne, S. Deléglise, E. Albertinale, U. Réglade, T. Capelle, E. Ivanov, T. Jacqmin, Z. Leghtas, and E. Flurin, *Irreversible qubit-photon coupling for the detection of*

- itinerant microwave photons, *Phys. Rev. X* **10**, 021038 (2020).
- [16] E. Albertinale, L. Balembois, E. Billaud, V. Ranjan, D. Flanigan, T. Schenkel, D. Estève, D. Vion, P. Bertet, and E. Flurin, Detecting spins by their fluorescence with a microwave photon counter, *Nature (London)* **600**, 434 (2021).
- [17] L. Balembois, J. Travesedo, L. Pallegoix, A. May, E. Billaud, M. Villiers, D. Estève, D. Vion, P. Bertet, and E. Flurin, Cyclically operated microwave single-photon counter with sensitivity of $10^{-22} \text{ mW}/\sqrt{\text{Hz}}$, *Phys. Rev. Appl.* **21**, 014043 (2024).
- [18] C. Braggio, L. Balembois, R. Di Vora, Z. Wang, J. Travesedo, L. Pallegoix, G. Carugno, A. Ortolan, G. Ruoso, U. Gambardella, D. D'Agostino, P. Bertet, and E. Flurin, Quantum-enhanced sensing of axion dark matter with a transmon-based single microwave photon counter, *Phys. Rev. X* **15**, 021031 (2025).
- [19] S. Borsanyi, Z. Fodor, J. Guenther, K.-H. Kampert, S. D. Katz, T. Kawanai, T. G. Kovacs, S. W. Mages, A. Pasztor, F. Pittler *et al.*, Calculation of the axion mass based on high-temperature lattice quantum chromodynamics, *Nature (London)* **539**, 69 (2016).
- [20] V. B. Klaer and G. D. Moore, The dark-matter axion mass, *J. Cosmol. Astropart. Phys.* **11** (2017) 049.
- [21] P. W. Graham and A. Scherlis, Stochastic axion scenario, *Phys. Rev. D* **98**, 035017 (2018).
- [22] F. Takahashi, W. Yin, and A. H. Guth, QCD axion window and low-scale inflation, *Phys. Rev. D* **98**, 015042 (2018).
- [23] M. Buschmann, J. W. Foster, A. Hook, A. Peterson, D. E. Willcox, W. Zhang, and B. R. Safdi, Dark matter from axion strings with adaptive mesh refinement, *Nat. Commun.* **13**, 1049 (2022).
- [24] J. E. Kim, Weak-interaction singlet and strong CP invariance, *Phys. Rev. Lett.* **43**, 103 (1979).
- [25] M. Shifman, A. Vainshtein, and V. Zakharov, Can confinement ensure natural CP invariance of strong interactions?, *Nucl. Phys. B* **166**, 493 (1980).
- [26] A. R. Zhitnitsky, On possible suppression of the axion hadron interactions. (In Russian), *Sov. J. Nucl. Phys.* **31**, 260 (1980).
- [27] M. Dine, W. Fischler, and M. Srednicki, A simple solution to the strong CP problem with a harmless axion, *Phys. Lett.* **104B**, 199 (1981).
- [28] S. Al Kenany, M. Anil, K. Backes, B. Brubaker, S. Cahn, G. Carosi, Y. Gurevich, W. Kindel, S. Lamoreaux, K. Lehnert *et al.*, Design and operational experience of a microwave cavity axion detector for the 20 – 100 μeV range, *Nucl. Instrum. Methods Phys. Res., Sect. A* **854**, 11 (2017).
- [29] C. M. Caves, Quantum limits on noise in linear amplifiers, *Phys. Rev. D* **26**, 1817 (1982).
- [30] L. Brouwer *et al.* (DMRadio Collaboration), Proposal for a definitive search for gut-scale QCD axions, *Phys. Rev. D* **106**, 112003 (2022).
- [31] C. Bartram *et al.*, Dark matter axion search using a Josephson traveling wave parametric amplifier, *Rev. Sci. Instrum.* **94**, 044703 (2023).
- [32] S. Beurthey, N. Böhmer, P. Brun, A. Caldwell, L. Chevalier, C. Diaconu, G. Dvali, P. Freire, E. Garutti, C. Gooch *et al.*, Madmax status report, [arXiv:2003.10894](https://arxiv.org/abs/2003.10894).
- [33] K. Yamamoto *et al.*, The Rydberg-atom-cavity axion search, in *Dark Matter in Astro- and Particle Physics*, edited by H. V. Klapdor-Kleingrothaus (Springer, Heidelberg, Germany, 2001), p. 638, [arXiv:hep-ph/0101200](https://arxiv.org/abs/hep-ph/0101200).
- [34] E. Graham, S. Ghosh, Y. Zhu, X. Bai, S. B. Cahn, E. Durcan, M. J. Jewell, D. H. Speller, S. M. Zacarias, L. T. Zhou, and R. H. Maruyama, Rydberg-atom-based single-photon detection for haloscope axion searches, *Phys. Rev. D* **109**, 032009 (2024).
- [35] X. Fan, G. Gabrielse, P. W. Graham, R. Harnik, T. G. Myers, H. Ramani, B. A. D. Sukra, S. S. Y. Wong, and Y. Xiao, One-electron quantum cyclotron as a milli-eV dark-photon detector, *Phys. Rev. Lett.* **129**, 261801 (2022).
- [36] D. I. Schuster, A. A. Houck, J. A. Schreier, A. Wallraff, J. M. Gambetta, A. Blais, L. Frunzio, J. Majer, B. Johnson, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Resolving photon number states in a superconducting circuit, *Nature (London)* **445**, 515 (2007).
- [37] 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).
- [38] P. L. Richards, Bolometers for infrared and millimeter waves, *J. Appl. Phys.* **76**, 1 (1994).
- [39] J. Zmuidzinis, Thermal noise and correlations in photon detection, *Appl. Opt.* **42**, 4989 (2003).
- [40] R. Cervantes, J. Aumentado, C. Braggio, B. Giaccone, D. Frolov, A. Grassellino, R. Harnik, F. Lecocq, O. Melnychuk, R. Pilipenko, S. Posen, and A. Romanenko, Deepest sensitivity to wavelike dark photon dark matter with superconducting radio frequency cavities, *Phys. Rev. D* **110**, 043022 (2024).
- [41] P. Pugnât, P. Camus, O. Kwon, R. Ballou, C. Bruyère, H. Byun, W. Chung, T. Grenet, P. Perrier, Y. K. Semertzidis, A. Talarmin, and J. Vessaire, GrAHal-CAPP for axion dark matter search with unprecedented sensitivity in the 1–3 μeV mass range, *Front. Phys.* **12** (2024).
- [42] J. Kim, Y. Kim, S. Yoon, K. Shin, J. Lee, J. S. Jung, J. T. Lee, J.-G. Kim, D. Kim, J. Yoo, H. Lee, S.-H. Moon, and S. Hahn, Design, construction, and operation of an 18 T 70 mm no-insulation (RE)Ba₂Cu₃O_{7-x} magnet for an axion haloscope experiment, *Rev. Sci. Instrum.* **91**, 023314 (2020).
- [43] C.-L. Kuo, Large-volume centimeter-wave cavities for axion searches, *J. Cosmol. Astropart. Phys.* **06** (2020) 010.
- [44] C.-L. Kuo, Symmetrically tuned large-volume conic shell-cavities for axion searches, *J. Cosmol. Astropart. Phys.* **02** (2021) 018.
- [45] M. O. Withers and C.-L. Kuo, A beehive haloscope for high-mass axion dark matter, *Phys. Rev. D* **111**, 072011 (2025).
- [46] T. A. Dyson, C. L. Bartram, A. Davidson, J. B. Ezekiel, L. M. Futamura, T. Liu, and C.-L. Kuo, High-volume tunable resonator for axion searches above 7 GHz, *Phys. Rev. Appl.* **21**, L041002 (2024).
- [47] M. Lawson, A. J. Millar, M. Pancaldi, E. Vitagliano, and F. Wilczek, Tunable axion plasma haloscopes, *Phys. Rev. Lett.* **123**, 141802 (2019).

- [48] R. G. Mani, L. Ghenim, and J. B. Choi, Quantum coherence effects and field-induced localization in InSb, [Phys. Rev. B](#) **43**, 12630 (1991).
- [49] A. J. Millar, G. G. Raffelt, J. Redondo, and F. D. Steffen, Dielectric haloscopes to search for axion dark matter: Theoretical foundations, [J. Cosmol. Astropart. Phys.](#) 01 (2017) 061.
- [50] J. Liu *et al.* (BREAD Collaboration), Broadband solenoidal haloscope for terahertz axion detection, [Phys. Rev. Lett.](#) **128**, 131801 (2022).
- [51] A. Zangwill, *Modern Electrodynamics* (Cambridge University Press, Cambridge, England, 2012).