Rydberg-atom-based single-photon detection for haloscope axion searches

Eleanor Graham[1,2] Sumita Ghosh[3,2,*] Yuqi Zhu,[1,2,†] Xiran Bai[1,2] Sidney B. Cahn,[1,2]
Elsa Durcan,[1,2] Michael J. Jewell,[1,2] Danielle H. Speller,[4] Sabrina M. Zacarias,[1,2]
Laura T. Zhou,[1,2,†] and Reina H. Maruyama[1,2]

[1]*Department of Physics, Yale University, New Haven, Connecticut 06520, USA*
[2]*Wright Laboratory, Department of Physics, Yale University, New Haven, Connecticut 06520, USA*
[3]*Department of Applied Physics, Yale University, New Haven, Connecticut 06520, USA*
[4]*Department of Physics and Astronomy, The Johns Hopkins University,
3400 North Charles Street Baltimore, Maryland 21211, USA*

We propose a Rydberg-atom-based single-photon detector for signal readout in dark matter haloscope experiments between 40 μeV and 200 μeV (10 GHz and 50 GHz). At these frequencies, standard haloscope readout using linear amplifiers is limited by quantum measurement noise, which can be avoided by using a single-photon detector. Our single-photon detection scheme can offer scan rate enhancements up to a factor of 10^4 over traditional linear amplifier readout, and is compatible with many different haloscope cavities. We identify multiple haloscope designs that could use our Rydberg-atom-based single-photon detector to search for QCD axions with masses above 40 μeV (10 GHz), currently a minimally explored parameter space.

I. INTRODUCTION

The axion is a compelling solution to the strong charge-parity (*CP*) problem in quantum chromodynamics (QCD) and a well-motivated dark matter candidate in astrophysics and cosmology [1–7]. Though the axion mass m_a is linked to the energy scale of the spontaneous symmetry breaking f_a, QCD itself does not constrain m_a or f_a [8]. Observations from astrophysics and cosmology constrain $m_a \sim 10^{-6} - 10^3$ μeV [9–13]. The coupling strength between the axion and the standard model depends on the axion mass; for a given value of m_a, there are a range of axion-photon couplings g_γ which are compatible with QCD. This region is typically spanned by two benchmark QCD axion models, the Kim-Shifman-Vainshtein-Zakharov (KSVZ) model [14,15] and the Dine-Fischler-Srednicki-Zhitnitsky (DFSZ) model [16,17]. Axionlike particles (ALPs) with photon couplings outside the range predicted from QCD can also act as dark matter, although they cannot solve the strong *CP* problem [5].

Among laboratory axion searches to date, only microwave cavity haloscope experiments are sensitive to QCD axions [18–22]. Haloscope experiments make use of the axion-photon conversion in a strong magnetic field, using a resonant cavity with a high quality factor to collect the signal [23,24]. The resulting signal is typically amplified with linear amplifiers such as HEMTs or JPAs [18,19]. Even with amplification, the resulting signal rates are very low, necessitating extremely low noise levels in haloscope experiments approaching or surpassing the standard quantum limit (SQL). For instance, in the case of axions with 10 μeV mass and KSVZ coupling strength, the expected signal power for an experiment such as HAYSTAC would be on the order of 10^{-24} W or 0.3 ct/s in terms of event rate [25]. The technology used in the current generation of haloscope experiments can access axions with masses $m_a \sim 0.25$–25 μeV, but higher masses in the range $m_a \sim 40$–200 μeV, corresponding to photons at ~10–50 GHz, are favored for postinflationary QCD axions [13,26–28]. One challenge that must be addressed to perform haloscope searches at higher frequencies is the loss of signal. In order for signal to accumulate in the cavity, the resonant frequency ν_c must match the axion's mass. This leads to a reduction in conversion power at higher frequencies, as the cavity's radius and therefore volume decrease to maintain the resonant condition. In order to maintain sensitivity to QCD axions, haloscopes must either increase the signal by improving the axion-photon conversion rate or lower the noise to achieve

---

sensitivity to smaller photon signals. Several such ideas have been proposed for next-generation experiments [29].

The use of single-photon detectors for readout in next-generation haloscope experiments provides a promising way of addressing this challenge. Unlike linear amplifiers, single-photon detectors are not sensitive to the phase of a microwave field, allowing for measurements of the field which are not subject to the SQL and are instead limited by shot noise from thermal photons [30]. This thermal photon noise is frequency-dependent and at high frequencies $\gtrsim 10$ GHz decreases enough that single-photon detectors will have lower noise power than linear amplifiers under reasonable assumptions [31]. This leads to an overall higher signal-to-noise ratio, enhancing the haloscope's ability to detect small signals. With the advancements in the field of microwave photon detection brought on by quantum computing [32–37], there are a number of potential schemes in development for single-photon detection in axion experiments [38,39].

In this work, we propose a design for a Rydberg atom-based single-photon detector which can be coupled to an axion haloscope to search for high-frequency axions. Rydberg atoms, with extremely high principal quantum number n, have many transitions in the 10–1000 GHz range with large dipole transition moments making them well-suited for single-photon detection at microwave frequencies. A straightforward detection scheme utilizing a transition between two neighboring Rydberg levels has been demonstrated to be sensitive to very small microwave fields around 14 GHz [40].

We begin in Sec. II with an overview of our single-photon detector concept and the advantages of single-photon detection for haloscope experiments in general. In Sec. III, we discuss which Rydberg states are appropriate for single-photon detection and show that photons in the frequency range 10–50 GHz can readily be detected with Rydberg atoms. In Sec. IV, we propose a design for incorporating these Rydberg atoms into a haloscope experiment and estimate its efficiency and noise rate. Finally, in Sec. V, we estimate the axion discovery sensitivity of our single-photon detector when connected to several different haloscope designs and show several designs that feature sensitivity to QCD axions with masses above 40 μeV.

II. OVERVIEW

A. Experimental concept overview

In this paper, we propose a single-photon detection scheme based on Rydberg atoms for haloscope-type axion dark matter experiments. Rydberg atoms have previously been considered for axion searches in a proposed experiment with ^{85}Rb Rydberg atoms in a haloscope tuned to 2.5 GHz [41,42]. Single-photon detection with Rydberg atoms, although applicable to a broad range of frequencies, is best suited to detecting photons at frequencies above

10 GHz [31]. At these higher frequencies, traditional sensors used in haloscope experiments are limited by quantum measurement noise and are less readily available.

A schematic of the single-photon detection scheme incorporated into a haloscope experiment is shown in Fig. 1. As in standard cavity haloscope experiments, axion-photon conversion takes place inside a "conversion cavity"; a microwave cavity placed inside a large magnetic field. Axions interact with the magnetic field and are converted into photons with frequencies $\nu_a \approx m_a c^2/h$. Photons resonant with the conversion cavity are collected, amplified, and coupled out of the cavity. A nonreciprocal transmission line sends the photons to a second cavity, the "detection cavity", which houses the single-photon detector away from the conversion cavity's magnetic field. Both cavities, along with the transmission line, are cooled to sub-Kelvin temperatures with a dilution refrigerator to reduce backgrounds from thermal photons. Single-photon detection takes place in the detection cavity, where photons couple to a beam of Rydberg atoms sent through the cavity. This coupling can take the form of absorption or stimulated emission, corresponding to an increase or decrease in the energy of the Rydberg atom's state. This change in the Rydberg state is detectable by selective field ionization (SFI), where the Rydberg atoms are placed in an electric field strong enough to ionize only one of the two states and an electron detector is used to see whether ionization has occurred. This combination of atom-photon coupling and state readout forms the single-photon detector.

The amount of signal power coupled out of the conversion cavity, P_{sig}, is given by the Sikivie equation:
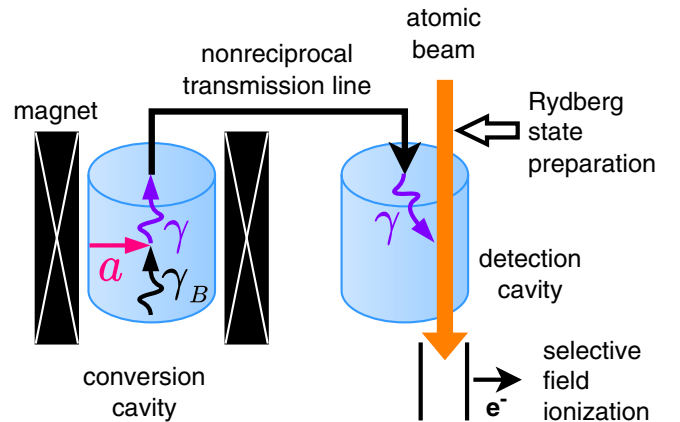
FIG. 1. Single-photon detection scheme for a haloscope experiment. Axions (a) convert into photons (γ) inside a conversion cavity by absorbing virtual photons from a strong magnetic field (γ_B). The axion-converted photon is coupled out of the conversion cavity and into the detection cavity via a nonreciprocal transmission line. Both of the cavities and the transmission line are cooled to sub-Kelvin temperatures by a dilution refrigerator (not shown). In the detection cavity, the photon couples to a beam of Rydberg atoms. The final state of the Rydberg atoms is read out using selective field ionization.

$$P_{\text{sig}} = \left(\frac{2\alpha^2 \hbar^3 c^3}{\pi \mu_0} \frac{g_\gamma^2 \rho_a}{\Lambda^4} \right) \left(B_0^2 V C_{mnl} \nu_c Q_{\text{conv}} \frac{\beta}{1 + \beta} \right) \quad (2.1)$$

in the case where the cavity's resonant frequency exactly matches the axion frequency [23,24]. Physical constants and properties of dark matter appear in the left set of terms: g_γ is the axion-photon coupling, $\rho_a = 0.45$ GeV/cm^3 is the local dark matter density [43], and $\Lambda = 77.6$ MeV describes the relation between the axion mass and the axion decay constant [44]. Experimental parameters appear in the right set of terms; B_0 is the magnetic field strength, V is the conversion cavity volume, C_{mnl} is the cavity form factor for the mode specified by mnl, ν_c is the resonant frequency of the conversion cavity, Q_{conv} is the conversion cavity's loaded quality factor, and β parametrizes the coupling of the cavity and relates the unloaded quality factor Q_{conv}^0 to the loaded quality factor as $Q_{\text{conv}}^0 = Q_{\text{conv}}(1 + \beta)$.

The signal for the single-photon detector is the number of axion-converted photons present in the detection cavity with frequencies inside the cavity's bandwidth. The number of signal counts seen by the single-photon detector in a time τ is related to the signal power P_{sig} and signal count rate R_{sig} by

$$N_{\text{sig,sp}} = \eta_{\text{sp}} \frac{P_{\text{sig}} \tau}{h \nu_a} = \eta_{\text{sp}} R_{\text{sig}} \tau, \quad (2.2)$$

where η_{sp} accounts for all efficiency losses after the photon signal is coupled out of the conversion cavity.

Noise in the single-photon detector takes the form of additional counts in the detector $N_{\text{sys,sp}}$. Assuming that the statistical fluctuations in the number of signal and noise counts are equal to $\sqrt{N}$, the signal-to-noise ratio (SNR) for the single-photon detector is

$$\text{SNR}_{\text{sp}} = \frac{N_{\text{sig,sp}}}{\sqrt{N_{\text{sig,sp}} + N_{\text{sys,sp}}}} = \frac{R_{\text{sig,sp}}}{\sqrt{R_{\text{sig,sp}} + R_{\text{sys,sp}}}} \sqrt{\tau}. \quad (2.3)$$

Alternately, Eq. (2.3) can be solved for τ to give the integration time needed to reach a desired target SNR, which we denote as $\tau_{\text{SNR,sp}}$.

For haloscope axion searches, the standard figure of merit is the scan rate. The scan rate is defined as the time needed to perform an axion search across a given frequency range, reaching a target SNR for every frequency in that range. This incorporates information about the haloscope's frequency tuning that is not captured by the ratio of SNRs at a single frequency.

A haloscope search using a single-photon detector collects counts within the detector bandwidth at each tuning step. The counts collected from this tuning step give sensitivity to axion signals appearing in the entire detection cavity bandwidth $\Delta \nu_{\text{det}}$, but cannot resolve individual frequencies within the bandwidth. Ideally, then,

$\Delta \nu_{\text{det}}$ should be matched to the axion signal's bandwidth $\Delta \nu_a$ to avoid excess noise pickup outside the signal bandwidth while also not attenuating the signal. If $\Delta \nu_{\text{det}} \approx \Delta \nu_a$, the scan rate is given by

$$\left(\frac{d\nu}{dt} \right)_{\text{sp}} = \frac{\Delta \nu_{\text{det}}}{\tau_{\text{SNR,sp}}} = \frac{\Delta \nu_{\text{det}}}{\text{SNR}^2} \frac{R_{\text{sig,sp}}^2}{R_{\text{sig,sp}} + R_{\text{sys,sp}}}. \quad (2.4)$$

The noise in a single-photon detector, $R_{\text{sys,sp}}$, is fundamentally driven by thermal photons produced from black-body radiation (BBR), which produce identical signals to axion-converted photons when inside the detector's bandwidth. However, the number of thermal photons drops exponentially at higher frequencies and lower operating temperatures. For frequencies on the order of 10 GHz and temperatures on the order of 10 mK, the thermal noise background is negligible, leading to high scan rates unachievable with a standard linear amplifier readout [31].

B. Advantages of single-photon detection

Single-photon detection is best suited for high frequencies above 10 GHz, where haloscope experiments must contend with decreased signal power due to shrinking cavity size. Resonance between the conversion cavity and the axion-converted photons requires that the cavity's radius scale with the photon wavelength, while the cavity's length should scale with the cavity's radius to avoid introducing mode crossings into the cavity's modes [45]. Then, $V \propto \nu_a^{-3}$, and the overall signal rate goes as

$$R_{\text{sig}} \propto m_a^{-3}, \quad (2.5)$$

slowing the scan rate of any haloscope searching for axions at higher masses. However, with single-photon detection available at higher frequencies, dramatic enhancements in scan rate become possible, counteracting this drop in signal.

The scan rate of a haloscope using a linear amplifier for readout can be written in the same notation as Eq. (2.4) [46]:

$$\left(\frac{d\nu}{dt} \right)_{\text{la}} = \frac{4}{5} \frac{\Delta \nu_{\text{conv}}}{\text{SNR}^2} \frac{R_{\text{sig,la}}^2}{N_{\text{sys,la}}^2 \Delta \nu_a}. \quad (2.6)$$

The overall form of the single-photon and linear-amplifier scan rates is similar. As we will see, the single-photon enhancement is primarily driven by the difference in the noise rates between the two types of readout.

The noise rate of a linear-amplifier-based haloscope is fundamentally limited by quantum measurement noise. The lower bound on the quantum measurement noise, known as the SQL, is $N_{\text{sys,la}} = 1$, independent of frequency [30,47]. Haloscope experiments have reached SQL-limited noise levels and even used techniques such as squeezed states to achieve measurement noise 4 dB below the SQL at the cost of sacrificing some of their signal [25]. Single-photon

detection avoids measurement noise entirely due to its phase insensitivity, which allows measurements in a basis where the signal amplitude is free of noise without needing to discard any signal power.

Because single-photon detectors can eliminate measurement noise, they are able to more quickly scan for axions across broad frequency ranges, making them ideal for experiments aiming to discover an axion signal at an unknown frequency. While the axion mass is unknown, scan rate is the most important consideration. However, if an axion signal were discovered, linear amplifiers would have an important role to play due to their frequency resolution; although vulnerable to additional sources of noise, linear amplifiers can provide insight into the axion's properties through resolution of the axion lineshape [48]. For this reason, linear amplifier technology should remain a key part of haloscope development.

The remarks in this section are general to all forms of single-photon detection in haloscope experiments. However, in the rest of this paper, we focus on the Rydberg-atom-based single-photon detection scheme outlined in Sec. (II A), showing how it is capable of achieving the performance described here.

III. RYDBERG ATOMS FOR SINGLE-PHOTON DETECTION

The single-photon detection scheme outlined in the previous section uses Rydberg atoms as detectors for microwave photons created from axion-photon conversion. Many transitions between Rydberg states are at microwave frequencies and take place quickly due to the high transition dipole moments of the states, making Rydberg atoms a compelling single-photon detector for this application. Here, we describe the design considerations that inform the choice of which Rydberg transitions to use for single-photon detection and show that axion-converted photons with frequencies from 10 GHz to 50 GHz can be detected with our Rydberg atom design.

A. Choice of atomic species

The first choice that must be made in the design of the Rydberg-atom-based single-photon detector is which atom to use. Alkali atoms are relatively easy to work with, benefiting from well-established techniques for state preparation and spectroscopy. From an ease-of-use standpoint, rubidium atoms are particularly compelling. However, previous attempts to use ^{85}Rb Rydberg atoms for photon detection by the CARRACK Collaboration found that the efficiency of atom-photon coupling was very sensitive to stray electric fields inside the detection region [49,50]. Based on this, CARRACK estimated the size of the stray electric field effect for different species of Rydberg atoms and proposed that using ^{39}K could increase the viability of single-photon detection [50].

We further investigate this effect, comparing the relative susceptibility to stray electric fields for several Rydberg states in ^{85}Rb and ^{39}K with different principal quantum numbers n. First, we estimate each Rydberg state's energy dependence on the electric field $\mathcal{E}$ at low fields via its scalar polarizability α_0, where $\Delta E = \frac{1}{2}\alpha_0\mathcal{E}^2$. Then, we use the alkali Rydberg calculator [51] (ARC) to calculate the scalar polarizability for nS and $nP_{1/2}$ states in ^{85}Rb and ^{39}K. The scalar polarizability has an approximate n dependence of $\alpha_0 \propto n^7$, which is consistent with our ARC calculations and allows us to interpolate across arbitrary values of n [52]. We define $\Delta\alpha_0$ as the difference between the scalar polarizabilities of the nS and nP states. When $\Delta\alpha_0$ is small, the transition energy is robust to stray electric fields regardless of the value of α_0; the energy of each state might change, but the transition energy stays the same.

In Fig. 2, we plot the values of α_0 and $\Delta\alpha_0$ for ^{85}Rb and ^{39}K as a function of n. While the overall polarizability increases with n for both types of atom, the polarizability of the nS and nP states stays close together for ^{39}K and diverges rapidly for ^{85}Rb, leading to a large difference in $\Delta\alpha_0$. This indicates that ^{39}K is well-suited for avoiding the effects of stray electric fields, especially for higher-n states. For this reason, we propose to use ^{39}K Rydberg atoms for our Rydberg-atom-based single-photon detector.

B. Choice of Rydberg States

Rydberg states are prepared by exciting ground-state atoms into a high-n state. We use a two-photon transition for state preparation, passing from the $4S_{1/2}$ ground state to the $5P_{1/2}$ intermediate state, and then from this intermediate state to an $nS_{1/2}$ or $nD_{3/2}$ Rydberg state. This is illustrated in Fig. 3(a) for ^{39}K. Using this state preparation scheme, we have previously created and characterized over fifty different Rydberg states in ^{39}K [53].
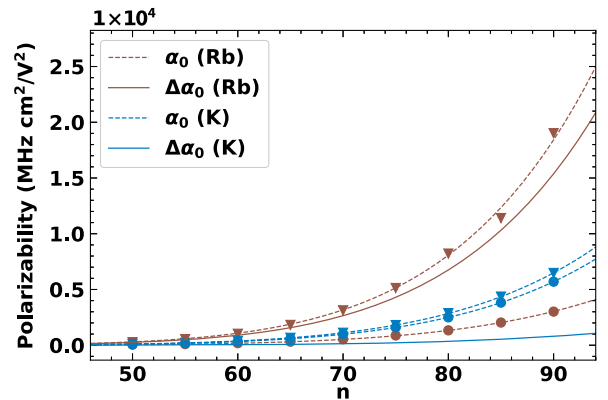

FIG. 2. Scalar polarizability α_0 (calculated with ARC) vs n for nS and $nP_{1/2}$ Rydberg states in ^{85}Rb and ^{39}K. Dashed lines indicate fits to the approximate formula $\alpha_0 \propto n^7$ for nS and $nP_{1/2}$ states, marked with circles and triangles respectively. Solid lines indicate $\Delta\alpha_0$.
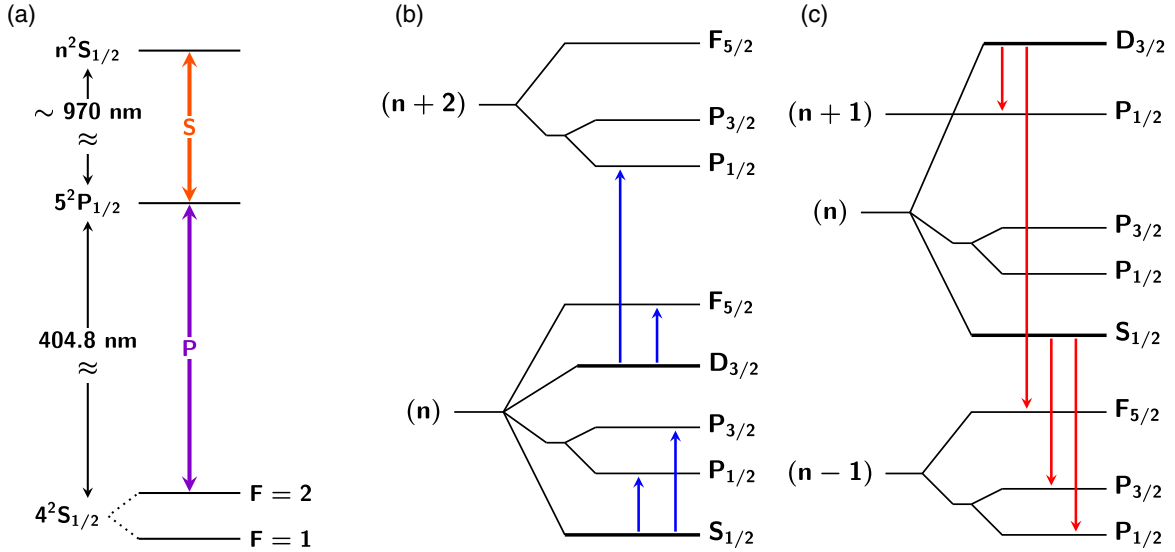
FIG. 3. Partial energy level diagrams of ^{39}K. (a) Energy level diagram showing two-photon excitation used to prepare the Rydberg state $n^2S_{1/2}$. This technique can be used to prepare either $nS_{1/2}$ or $nD_{3/2}$ states. (b) A subset of Rydberg transitions accessible via photon absorption from the prepared states. (c) A subset of Rydberg transitions accessible via photon emission from the prepared states.

From each starting state accessible with state preparation, many different Rydberg-Rydberg transitions are possible, driven by absorption or stimulated emission. Due to this variety of starting and ending states, there are multiple Rydberg transitions resonant with each desired photon frequency. Choosing which transitions to use is mainly motivated by two properties; transition dipole moment and state lifetime. The transition dipole moment d determines how quickly the atom-photon interaction takes place. The detection cavity can only store photons for a short time determined by its quality factor, with typical storage times on the order of a few μs. For a high probability of atom-photon interaction, then, a large transition dipole moment is necessary. A long state lifetime is also necessary to ensure that the atom will not decay out of its Rydberg state during the single-photon detector's operation. Rydberg states are long-lived with typical lifetimes on the order of 100 μs, much longer than the cavity's photon storage time, so the lifetime constraints are primarily set by the time needed for the atoms to travel between the initial state preparation region and the final readout region. The exact requirements for d and τ depend on the experimental design, with the dipole moment more heavily constrained due to the short time that photons spend in the detection cavity.

Achieving suitable d and τ is easier at higher operating frequencies. Transition frequency tends to increase with n, and both dipole moment and lifetime approximately scale with n as $d \propto n^2$ and $\tau \propto n^3$ [52]. For a fixed transition frequency, we can optimize d by looking at several different transitions around that frequency. To find the transitions with the strongest dipole moments, we calculate d using

ARC for all of the allowed transitions from the states that can be created with our state preparation scheme. We then select the transitions with the highest d for further investigation.

We show the properties of the allowed transitions with the highest dipole moments in Table I, choosing transitions with frequencies of 16 ± 0.5 GHz, near the bottom of our proposed frequency range. For each transition, the starting state is listed followed by the ending state. The dipole moment d depends on the difference between the total angular momentum quantum numbers m_j of the starting and ending states. For this table, we focus on $\Delta m_j = \pm 1$ which has the largest dipole moment and therefore the largest transition probability. The Rydberg state lifetimes are listed for the starting states (τ_n) and ending states ($\tau_{n'}$). The absorption transitions in the table are illustrated in Fig. 3(b) and the emission transitions in the table are illustrated in Fig. 3(c).

TABLE I. Rydberg transition frequency, dipole moment, and lifetimes of the starting and ending states near 16 GHz.

Transition	ν (GHz)	$d(ea_0)$	τ_n(μs)	$\tau_{n'}$(μs)
$60S_{1/2} \rightarrow 60P_{3/2}$	15.78	2251	227	781
$63S_{1/2} \rightarrow 62P_{3/2}$	15.73	2230	265	841
$57D_{3/2} \rightarrow 58P_{1/2}$	15.94	2225	453	704
$61D_{3/2} \rightarrow 63P_{1/2}$	16.31	2141	555	909
$68D_{3/2} \rightarrow 67F_{5/2}$	15.79	2062	768	228
$48D_{3/2} \rightarrow 48F_{5/2}$	16.00	2024	271	84
$60S_{1/2} \rightarrow 60P_{1/2}$	15.68	1845	227	781
$63S_{1/2} \rightarrow 62P_{1/2}$	15.81	1813	265	865

C. Rydberg frequency tuning

Performing an axion search using Rydberg-atom-based single-photon detectors will require tuning the single-photon detector to search for axions at many different frequencies. We can broadly "tune" the detection frequency of the Rydberg atoms by choosing different transitions, making large changes in frequency by increasing n and smaller changes in frequency by using a different transition type at a similar frequency. However, an additional tuning mechanism is necessary to fill out the intermediate frequencies between these resonances.

Rydberg transition frequencies can be shifted by either external electric fields (the Stark effect) or external magnetic fields (the Zeeman effect). We propose to use Zeeman shifts as our fine-tuning mechanism, since ^{39}K transition frequencies are relatively robust to external electric fields. For magnetic fields up to a few hundred Gauss, Zeeman tuning produces linear energy shifts that are easy to control. The frequency range of Zeeman tuning is limited by the magnetic field level where higher-order effects begin to dominate, which scales with n as $1/n^4$ [52].

Figure 4 shows the axion frequencies ν_a that can be reached in the 10–50 GHz range (or ~40–200 μeV in mass units), calculated based on the transitions discussed in Sec. (III B) and modified with linear Zeeman tuning. Through Zeeman tuning of absorption transitions, 98% of this mass range is readily accessible, and through Zeeman tuning of emission transitions, 92% of this mass range is readily accessible. Using Zeeman tuning of both absorption and emission transitions, the entire mass range is accessible.

IV. SINGLE-PHOTON DETECTOR DESIGN

We now examine the overall single-photon detection scheme incorporating the Rydberg states discussed in Sec. III. In Sec. IV A and Sec. IV B, we discuss the design considerations informing the choice of detection cavity, atomic beam source, and state readout scheme. Based on this design, we estimate the efficiency and noise rate of the single-photon detector in Sec. IV C and Sec. IV D.

A. Detection cavity quality factor

The quality factor of the detection cavity, Q_{det}, determines how long photon signals will persist in the cavity before decaying. For efficient single-photon detection, Q_{det} should be chosen so that the detector operates in the strong coupling regime, where the overall coupling rate g_{dipole} between Rydberg atoms and the microwave field is greater than any decay rates [54].

The strong coupling conditions require $g_{\mathrm{dipole}} > 1/\tau_{\mathrm{Ry}}$, $1/\tau_\gamma$, where τ_{Ry} and τ_γ are the lifetimes of the atom and the photon respectively. The photon lifetime depends on the detection cavity's loaded quality factor Q_{det} and its resonant frequency ν:

$$\tau_\gamma = Q_{\mathrm{det}}/(2\pi\nu). \tag{4.1}$$

For single-photon detection in an axion experiment, the ideal Q is matched to the axion's quality factor $Q_a \sim \nu_a/\Delta\nu_a \sim 10^6$, where $\Delta\nu_a$ is the width of the axion's spectral distribution [31]. As discussed in Sec. III B,
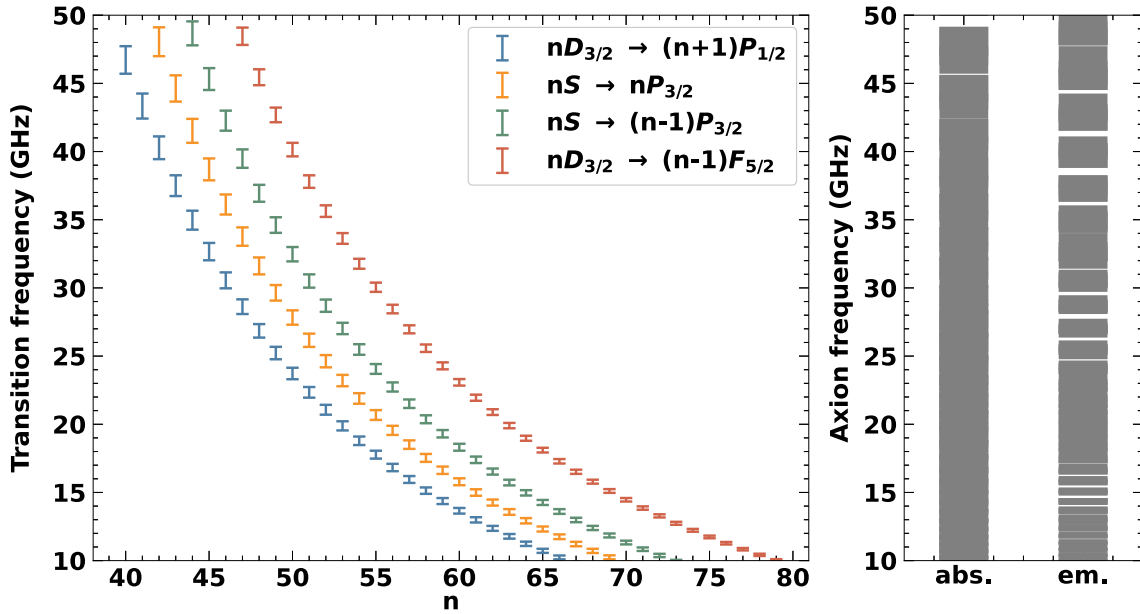

FIG. 4. Left: Zeeman-tuned transition frequency $\nu_{\mathrm{Ry}} + \Delta\nu_{\mathrm{Zeeman}}$ vs. n for the transitions $nD_{3/2} \to (n+1)P_{1/2}$, $nS \to nP_{3/2}$, $nS \to (n-1)P_{3/2}$, and $nD_{3/2} \to (n-1)F_{5/2}$. The vertical bars indicate the maximum Zeeman-tuning range of the individual transitions. Right: Range of accessible ν_a for all absorption transitions (abs.) in Fig. 3(b) and emission transitions (em.) in Fig. 3(c).

$\tau_\gamma \ll \tau_{\mathrm{Ry}}$, so the limiting requirement of the strong coupling condition is

$$g_{\mathrm{dipole}} > \frac{2\pi\nu_a}{Q_{\mathrm{det}}}. \qquad (4.2)$$

It follows that the minimum g_{dipole} required for strong atom-microwave coupling ranges from $g_{\mathrm{dipole}}^{\min} \sim \pi \cdot (20\ \mathrm{kHz})$ for 10 GHz photons to $g_{\mathrm{dipole}}^{\min} \sim \pi \cdot (100\ \mathrm{kHz})$ for 50 GHz photons. We can determine the coupling between a single atom and a single photon using the relation [55]

$$g_{\mathrm{dipole}} = \frac{2d\mathcal{E}_{\mathrm{rms}}}{\pi\hbar}. \qquad (4.3)$$

Here, d is the dipole moment of the Rydberg transition induced in the atom by the photon and $\mathcal{E}_{\mathrm{rms}}$ is the root-mean-square (rms) electric field generated by the photon.

We can estimate $\mathcal{E}_{\mathrm{rms}}$ by assuming a specific cavity geometry. Suppose our cavity size is comparable to the microwave wavelength $\lambda_a = c/\nu_a$ and its volume is $V \sim (\lambda_a/2)^3$. The rms field per photon is then

$$\mathcal{E}_{\mathrm{rms}} \sim \sqrt{\frac{h\nu_a}{\epsilon_0 V}} \propto \nu_a^2, \qquad (4.4)$$

where ϵ_0 denotes vacuum permittivity. Typical values of $\mathcal{E}_{\mathrm{rms}}$ range from $\sim 5\ \mu\mathrm{V/cm}$ at 10 GHz to $\sim 117\ \mu\mathrm{V/cm}$ at 50 GHz.

Combining Eqs. (4.2) and (4.3), we can rewrite the strong coupling condition in terms of d and $\mathcal{E}_{\mathrm{rms}}$ as

$$\frac{d\mathcal{E}_{\mathrm{rms}}}{\hbar} > \frac{\pi^2\nu_a}{Q_{\mathrm{det}}} \qquad (4.5)$$

and use this and Eq. (4.4) to derive a lower bound on d as a function of ν_a and Q_{det}:

$$d > \frac{\hbar\pi^2\nu_a}{\mathcal{E}_{\mathrm{rms}}Q_{\mathrm{det}}} \propto \frac{1}{Q_{\mathrm{det}}\nu_a}. \qquad (4.6)$$

Thus, the minimum d required for strong coupling is inversely proportional to ν_a and Q_{det}. With $\nu_a \sim 16\ \mathrm{GHz}$ and $Q_{\mathrm{det}} \sim 10^6$, this bound is $d > 1600 e a_0$, which is satisfied by all of the transitions given in Table I. At lower frequencies, the resonant states have larger n and therefore larger d. At higher frequencies, the lower bound on d is less strict. Across our proposed frequency range, $Q_{\mathrm{det}} \sim 10^6$ is sufficient for strong coupling.

However, single-photon detection can still take place with lower values of Q_{det}. Increasing the number of Rydberg atoms in the detection cavity, N, can increase the atom-photon coupling rate by a factor of $\sqrt{N}$ due to

collective behavior of the Rydberg ensemble [56–58]. Alternately, the strong coupling condition can be relaxed, corresponding to photon detection with a lower efficiency [59]. For the remainder of this paper, we assume that strong coupling is achievable.

B. Atom-cavity interaction time

Once the Rydberg atom has coupled to a photon, the state of the atom-photon system undergoes Rabi oscillations. For effective single-photon detection, we must control the time the atom spends in the cavity relative to the frequency of these oscillations to maximize the probability that the atom undergoes a change in state. This is primarily accomplished by controlling the speed of the atomic beam passing through the cavity.

The longitudinal velocity of the Rydberg atoms v_L is broadly controlled by the choice of beam source and can be further fine-tuned during the Rydberg state preparation step. Velocity-selective state preparation can be achieved by pumping only specific velocity classes of atoms into the ground state of the Rydberg excitation scheme [60]. In our experiment, $^{39}\mathrm{K}$ atoms at speed v_L would be pumped to the $F = 2$ level of $4^2S_{1/2}$ and, as shown in Fig. 3(a), undergo a two-photon transition to the desired Rydberg state.

In the strong coupling regime, the ideal transient time through the microwave region ($\tau_t = L/v_L$) is

$$\tau_t = \pi/g_{\mathrm{dipole}}, \qquad (4.7)$$

which corresponds to a single π pulse applied to the atom and photon's initial state.

The strong coupling condition [Eq. (4.2)] therefore gives an upper bound for τ_t and a lower bound for v_L. For the cubic cavity geometry assumed in the previous section, $L \sim \lambda_a/2$, giving

$$v_L > \frac{2\nu_a L}{Q_{\mathrm{det}}} \sim \frac{c}{Q_{\mathrm{det}}} \sim 300\ \mathrm{m/s}. \qquad (4.8)$$

This constraint on the transverse velocity informs the choice of an appropriate beam source for this experiment. While the Rydberg state preparation would take place in the cryostat as close to the detection cavity as possible to minimize unwanted interactions and decays, the creation of the atomic beam itself would take place at room temperature. In Table II, we list typical values of v_L and beam flux

TABLE II. Types of beam sources, their typical velocities, and fluxes for alkali atoms.

Type	v_L	Flux
Thermal beam [61]	100–500 m/s	10^{11} atoms/s
2D MOT beam [62]	28–35 m/s	10^{11} atoms/s
Supersonic beam [63]	1–5 km/s	10^{16} atoms/s/sr

for three common types of alkali atom beam source. The v_L estimated above most closely corresponds to the typical v_L of a thermal beam. All three beam sources are capable of producing a high flux of alkali atoms.

Different designs can make other beam sources more favorable than a thermal beam. The above estimate assumes that the strong coupling condition is barely achieved with $g_{dipole} \sim g_{min}$. If the coupling rate is faster, then the ideal interaction time will be lower. Similarly, for a lower value of Q_{det}, the ideal interaction time will also be lower to match the lowered photon lifetime. For modest changes in g_{dipole} or Q_{det} from the above values, a supersonic beam would be more appropriate. However, for $Q_{det} \lesssim 10^5$ a beam with the v_L required for the desired atom-photon interaction would be challenging to produce and control.

C. Photon detector efficiency

We now turn our attention to the overall efficiency of the Rydberg-atom-based single-photon detector. We introduce an overall photon detection efficiency η to account for all losses taking place after the axion-converted photon is coupled out of the conversion cavity.

The conversion cavity is connected to the detection cavity via a shared transmission line with a circulator in between. We add this circulator to ensure that signals only flow from the conversion cavity to the microwave detection cavity. This also prevents the modes from the two cavities from hybridizing, which would make optimizing and tuning the individual modes challenging. This coupling scheme will result in losses from coupling the transmission line to the detection cavity, as well as ohmic losses from the transmission line. We define the total efficiency from this step as

$$\eta_0 = \tilde{\eta}_0 \kappa^{dc}, \qquad (4.9)$$

where $\tilde{\eta}_0$ is the efficiency of the transmission line and κ^{dc} is the coupling efficiency of the detection cavity. With low-loss superconducting cables, we expect that the efficiency of the transmission line will be dominated by the loss in the circulator. Estimating 1 dB loss in the circulator, we expect $\tilde{\eta}_0 \sim 0.8$. The coupling efficiency is given by

$$\kappa^{dc} = \frac{1}{1 + \beta} \qquad (4.10)$$

and has a maximum value of $1/2$, achieved for the critical coupling condition $\beta = 1$.

Once the axion-converted photon has reached the detection cavity, it must couple to the atomic beam and induce a Rydberg transition. In Sec. IV A and Sec. IV B, we described the requirements for Q_{det} and v_L that enable minimal efficiency loss in the interactions between the Rydberg atoms and the detection cavity photons. When

these conditions are met, the efficiency for this first step is $\eta_1 \sim 1$.

After the Rydberg atoms have exited the detection cavity after the appropriate interaction time, their final states are probed using selective field ionization. SFI is a standard technique used for identifying Rydberg states and relies on the fact that the field needed to ionize a Rydberg state depends highly on the state's quantum number [52]. In practice, this is achieved by using a time-varying electric field pulse that ramps up to the ionization field, such that Rydberg states with different ionization energies are ionized at different times [64]. These pulsed sequences are commonly used to resolve multiple different states in an atomic ensemble at different times. However, even if we do not need to resolve different states, using a pulsed electric field provides more control over how our starting states will evolve in the Stark basis, preventing mixing between Rydberg states with similar energy levels. After ionization occurs, the ionized electrons are detected by a charged particle detector such as a channel electron multiplier (CEM) or microchannel plate (MCP).

Efficiency losses in the readout stage can primarily arise from two sources: SFI deadtime (η_2) and electron detection (η_3). SFI deadtime is caused by atoms entering or exiting the ionization region after the pulse has already started. These atoms will not experience the intended pulse sequence and may not ionize properly. Short pulse sequences are less vulnerable to this deadtime effect. Optimizing the shape of the electric field ramp is an effective way to decrease the time needed for SFI: shaped pulse techniques have been used to reliably distinguish S and P Rydberg states in alkali atoms with pulse lengths typically $< 1\ \mu s$ [64–66]. This level of deadtime is negligible compared to the overall repetition rate of the single-photon detector ($\sim 10\ \mu s$), so the ionization efficiency is $\eta_2 \sim 1$. Efficiency after ionization is set by the efficiency of the device used to collect the ionized electrons. For instance, electron collection with a channel electron multiplier can reasonably achieve $\eta_3 > 0.8$ [67,68].

A summary of efficiencies at various steps and associated technical assumptions can be found in Table III. The resulting total efficiency is $\eta = \prod_i \eta_i \sim 0.4$.

D. Total noise rate

Noise in single-photon detection in the 10–100 GHz range is expected to be dominated by BBR [31]. At temperature T and frequency ν_a, the photon occupation number ($\bar{n}_T$) of thermal photons is given by

$$\bar{n}_T = \frac{1}{e^{h\nu_a/k_B T} - 1}. \qquad (4.11)$$

BBR noise in the detector arises from thermal photons occupying the cavities and transmission line. In Fig. 5, we show a schematic of the two-cavity system connected by a

TABLE III. Best-case efficiencies for various loss mechanisms in signal transmission and single-photon detection.

Loss mechanism	Efficiency	Technical assumptions
Transmission line dissipation	$\tilde{\eta}_0 \sim 0.8$	Low-loss cables and circulator
Coupling to transmission line	$\kappa^{dc} = 1/2$	Critical coupling
Atom-microwave coupling	$\eta_1 \sim 1$	Strong coupling
Selective ionization	$\eta_2 \sim 1$	SFI pulse length < 1 μs
Electron detection	$\eta_3 > 0.8$	Channel electron multiplier

transmission line. The transmission line is made nonreciprocal by inserting a circulator with one port terminated by a 50 Ω resistor.

In order to determine the system noise, we track the flow of thermal photons in the system. We start by considering coupling of photons into the transmission line, which is parametrized by κ^{cc}. The photon occupation in the transmission line at the output of the conversion cavity is given by

$$\bar{n}_{T,\text{output}}^{cc} = \bar{n}_T(T_{\text{conv}})\kappa^{cc} + \bar{n}_T(T_{\text{trans}})(1-\kappa^{cc}), \quad (4.12)$$

where the first term represents thermal photons coupled out of the cavity and the second term accounts for thermal photons in the transmission line that reflect off the cavity.

The circulator introduces efficiency losses that replace a fraction of the input with thermal noise from the circulator, assumed to be at the same temperature as the transmission line such that $\bar{n}_{\text{loss}} = \bar{n}_T(T_{\text{trans}})$. The coupling between the circulator loss and the input signal can be modeled as a beam splitter [37,69], in which case the number of thermal photons incident on the detection cavity is given by
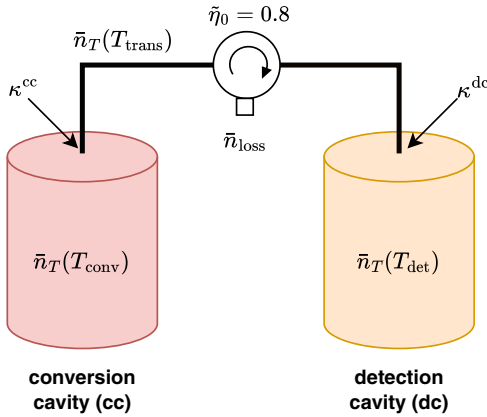

FIG. 5. Signal and noise flow between the conversion cavity and detection cavity. The conversion cavity is at temperature T_{conv} and the detection cavity is at temperature T_{det}. Photons from the conversion cavity flow to the detection cavity via a low-loss transmission line at temperature T_{trans}. A circulator with one port terminated by a 50 Ω resistor (shown as a block) is placed between the two cavities to prevent photon flow in the reverse direction.

$$\bar{n}_{T,\text{input}}^{dc} = \bar{n}_{T,\text{output}}^{cc}\tilde{\eta}_0 + \bar{n}_{\text{loss}}(1-\tilde{\eta}_0). \quad (4.13)$$

Finally, these noise photons couple into the detection cavity, which also contains its own population of thermal photons. The input noise photons are coupled into the detection cavity with efficiency κ^{dc}, while the thermal photons originating in the detection cavity are coupled out of the cavity with efficiency $1-\kappa^{dc}$, leading to a total thermal photon population of

$$\bar{n}_T^{dc} = \bar{n}_{T,\text{input}}^{dc}\kappa^{dc} + \bar{n}_T(T_{\text{det}})(1-\kappa^{dc}). \quad (4.14)$$

We have performed this derivation in the general case where the detection cavity and conversion cavity can have different temperatures. When estimating the noise rate of the experiment, we will assume that the two cavities, the transmission line, and the circulator are approximately thermalized to the same temperature T. In this case, the occupation number of the detection cavity is simply

$$\bar{n}_T^{dc} \approx \bar{n}_T(T). \quad (4.15)$$

So far, we have considered the occupation number of thermal photons with frequency precisely equal to the detection cavity's resonant frequency. However, thermal photons with frequencies across the detection cavity's bandwidth can all couple into the same cavity mode, increasing the number of noise counts ultimately coupled out of the cavity by the single-photon detector. Then, the rate of thermal photons read out of the detection cavity, R_T, is

$$R_T \approx \frac{2\pi\nu_a}{Q_{\text{det}}}\bar{n}_T^{dc} = 2\pi\Delta\nu_{\text{det}}\bar{n}_T^{dc}. \quad (4.16)$$

Although thermal photons are the primary source of noise for the single-photon detector, noise can also enter the single-photon detector during the final state readout. This readout noise would be dominated by dark counts in the electron detector used for SFI. If, for example, a CEM from Sjuts is used for electron detection, the dark count rate is expected to be < 0.02 ct/s [70].

Some of these electron dark counts can be discarded using timing information. With pulsed SFI, the time that the electric field reaches the Rydberg ionization threshold is
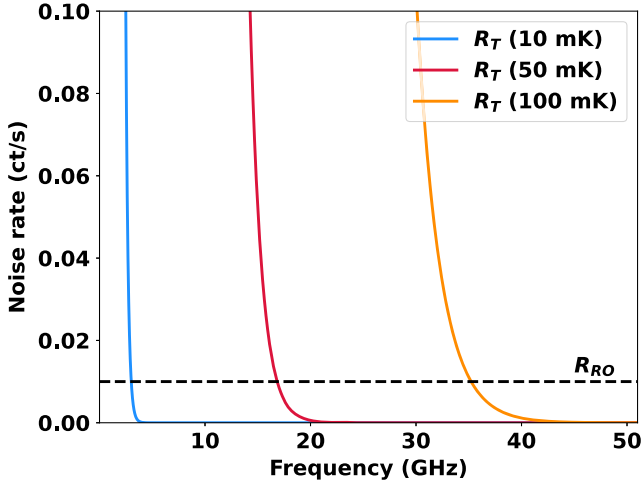
FIG. 6. Rate of thermal photons read out of detection cavity R_T vs frequency, plotted for operating temperatures 10 mK, 50 mK, and 100 mK. The dashed line indicates a conservative estimate for the readout noise from electron detection R_{RO}.

known. If a count shows up when ionization is impossible, it can be identified as a dark count. Conservatively, we estimate that these timing cuts can remove 50% of electron dark counts without affecting the detection efficiency, corresponding to $R_{\mathrm{RO}} \approx 0.01$ ct/s. With careful characterization of the SFI pulses and corresponding electron arrival times, much better dark count rejection may be possible.

The noise rates for different noise sources in the experiment are estimated in Fig. 6. R_T is calculated from Eq. (4.16) for several operating temperatures T, assuming $Q_{\mathrm{det}} = 10^6$ and the efficiencies from Table III. R_{RO} is estimated as 0.01 ct/s. The total noise rate in the single-photon detector is given by

$$R_{\mathrm{sys}} = R_T + R_{\mathrm{RO}}. \tag{4.17}$$

Single-photon detectors perform best above the frequency where the noise contribution from thermal photons is negligible compared to the readout noise. From Fig. 6, we can see the frequency at which thermal photons cease to be the dominant background: 35 GHz at 100 mK, 17 GHz at 50 mK, and 3 GHz at 10 mK.

In many cases, the quantum measurement noise from linear amplifiers is much larger than 0.01 ct/s, and single-photon detectors will outperform linear amplifiers even outside of this ideal operating range. Detailed comparisons of single-photon detectors and linear amplifiers are discussed in the Appendix and indicate single-photon advantage above 15 GHz even when operating at 100 mK.

V. SENSITIVITY

We now estimate the axion sensitivity of a haloscope experiment that uses the single-photon detector described

in the prior sections. We must make several assumptions regarding the overall haloscope design. Since single-photon detectors read out from a haloscope's conversion cavity, they can be incorporated into many different haloscope designs. Rather than proposing a specific haloscope for use with the single-photon detector, we explore the performance of a variety of haloscopes.

The sensitivity of an arbitrary haloscope is improved by a factor of up to 10^4 by switching from linear amplifier readout to single-photon detector readout, as long as the haloscope operates at a frequency and temperature where thermal photon counts are negligible. We describe the scan rate enhancement calculation in more detail in the Appendix. Here, we focus on calculating the axion sensitivity for several specific haloscope designs incorporating single-photon detectors. We begin with a current-generation haloscope as a benchmark and then incorporate a variety of proposed design improvements that could define the next generation of haloscopes.

In Sec. II A, we derived the scan rate of an axion search using a single-photon detector. Now that we have an estimate of our single-photon detector's efficiency and an expression for its noise rate [Eq. (4.17)], we can use Eq. (2.4) to calculate the scan rate of a haloscope that uses this single-photon detector.

Several experimental parameters of the haloscope and single-photon detector that enter into the scan rate are kept constant. The haloscope design parameters V_0 and $C_{mn\ell}$ are set approximately equal to the operational parameters of HAYSTAC, a currently running haloscope experiment that searches for axions around 5 GHz [19]. Q_{det} is set to 10^6 based on the design constraints discussed in Sec. IV A and Sec. IV B. The total efficiency is estimated in Sec. IV C as 0.4 and β is set to 1 for critical coupling out of the conversion cavity. As in Sec. IV D, R_{RO} is conservatively set to 0.01 ct/s. We summarize the design parameters kept constant across all sensitivity projections in Table IV.

We consider four different haloscope design scenarios. These scenarios differ from each other in four ways; the conversion cavity quality factor Q_{conv}, the frequency-volume scaling relation $V(\nu_a)$, the magnetic field strength B_0, and the operating temperature T. Scenario A represents the most conservative case, in which the single-photon detector is attached to a current-generation haloscope and

TABLE IV. Parameters shared between all haloscope designs examined in the sensitivity study.

Parameter	Value	Description
V_0	1.5 L	Volume at 5 GHz
$C_{mn\ell}$	0.5	Form factor
Q_{det}	10^6	Detection cavity Q
β	1	Coupling factor
η	0.4	Total efficiency
R_{RO}	0.01 ct/s	SFI noise rate

the apparatus operates at 100 mK. Scenario B improves over Scenario A broadly across several dimensions: moderately increasing the magnetic field, improving Q_{conv} by a factor of ten, and reducing the operating temperature to 50 mK. Scenario C considers the case of a small-bore, high B-field magnet operating in the upper end of our proposed frequency range, and also features modest improvements to Q_{conv}. Finally, in Scenario D we allow the conversion cavity volume to be independent of frequency, but use only conservative parameters for the magnetic field and quality factor.

We begin the exploration of these scenarios by discussing the quality factor of the conversion cavity. Improving Q_{conv} increases the magnitude of the haloscope signal but poses a challenge due to the strong magnetic fields surrounding the conversion cavity that degrade the quality factor. Typical unloaded quality factors of current haloscope experiments are on the order of 10^4, so with critical coupling ($\beta = 1$) their loaded quality factors are also on the order of 10^4 [18–20]. Cavities coated with superconducting materials have shown higher quality factors on the order of 10^5 even at magnetic fields up to 8 T [71–73]. High values of Q_{conv} are particularly important when using single-photon detectors for readout, because single-photon detectors do not have frequency resolution within the conversion cavity bandwidth. While a linear-amplifier-based haloscope can take advantage of lower values of Q_{conv} to perform a broadband measurement across several frequencies at once, we must develop new approaches when measuring with single-photon detectors.

Next, we consider the conversion cavity geometry. The current generation of haloscope experiments searches for axions at frequencies below 6 GHz. Operating an axion search in the 10–50 GHz frequency range will require

adapting existing conversion cavity designs or developing new approaches altogether. The frequency dependence of the conversion cavity volume is related to the resonant frequency of the available cavity modes at tuning step. We approximate the typical frequency variation by choosing a benchmark frequency and volume and then scaling by the resonance as discussed in Sec. II B:

$$V = V_0 \left(\frac{\nu_0}{\nu_a}\right)^3. \tag{5.1}$$

Some proposed experimental designs seek to eliminate the frequency dependence of the conversion cavity volume altogether, such as dielectric haloscopes and plasma haloscopes [74,75]. The broad bandwidth of dielectric haloscopes (~10 MHz) may be challenging to integrate with narrow-band single-photon detection. However, the bandwidth of plasma haloscopes is similar to the bandwidth of a standard haloscope cavity; plasma haloscope cavities are projected to be capable of achieving unloaded quality factors on the order of 10^4 [76]. Since this is a promising area of axion search development, we consider volume independence in one of our scenarios.

The magnetic field used for a haloscope also has a large impact on the ultimate sensitivity, as the scan rate scales with B_0^4. Current haloscope experiments already use a variety of magnetic fields. For instance, HAYSTAC uses a magnetic field of 8 T [19]. However, magnetic fields up to 18 T have been used by the CAPP axion experiment [77]. Development of magnets for haloscope experiments requires not only producing a large magnetic field, but also maintaining that magnetic field over the volume occupied by the conversion cavity. Typically, higher-field magnets require smaller bores; one design using superconducting
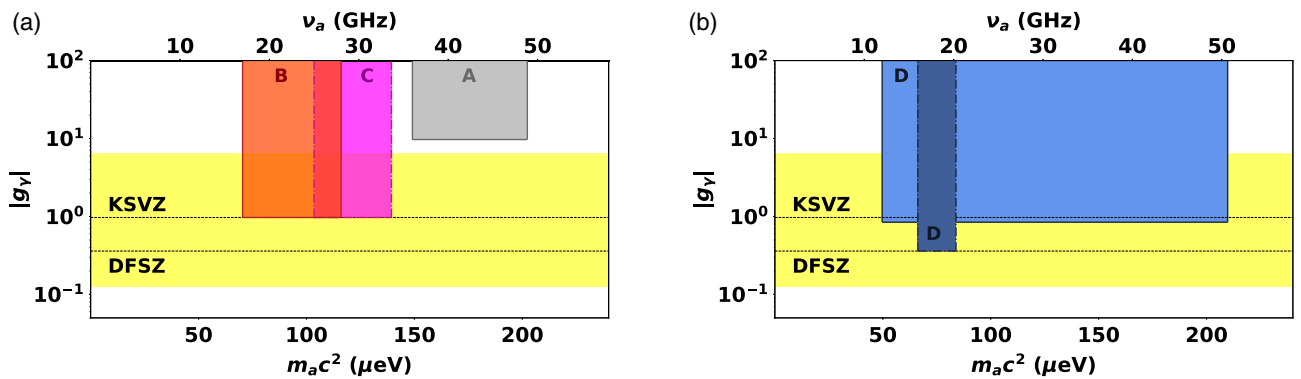

FIG. 7. Axion sensitivity for a haloscope experiment using single-photon detection. Shaded areas show the parameter space that can be covered in five years of operation. Shown in (a) are three scenarios using standard microwave haloscope cavities: A is a conservative design representing the current generation of haloscopes, B is a more ambitious design representing the next generation of haloscopes, C is a design using a small-bore magnet with a high B-field. Shown in (b) are two different searches using a nonstandard cavity whose frequency is independent of the cavity size, where the narrower exclusion is taken with a longer integration time at each tuning in order to achieve DFSZ sensitivity. Experimental parameters shared between the scenarios are given in Table IV and experimental parameters differing between the scenarios are given in Table V. KSVZ [14,15] and DFSZ [16,17] benchmark models are shown with dashed lines while the entire QCD band is highlighted in yellow [79].

TABLE V. Parameters used to specify four different scenarios of a single-photon detector connected to a haloscope.

Scenario	Q_{conv}	$V(\nu_a)$	B_0 (T)	T (mK)
Conservative (A)	1×10^4	$\propto \nu_a^{-3}$	8	100
High Q_{conv}(B)	1×10^5	$\propto \nu_a^{-3}$	12	50
High B_0 (C)	2×10^4	$\propto \nu_a^{-3}$	32	50
Volume-independent (D)	1×10^4	$\propto 1$	8	50

inserts has achieved 32 T fields with a 3.4 cm bore [78]. While this bore size could not hold an experiment like HAYSTAC, which has a diameter of 12.7 cm at 5 GHz, at higher frequencies the cavity diameter would decrease as ν_a^{-1}, allowing such a magnet to be used at frequencies above 20 GHz.

The axion sensitivity of these four hypothetical experiments for five years of data taking is presented in Fig. 7 with the main parameters for each scenario summarized in Table V. As shown in scenario A, searches for axionlike particles can be performed across over ∼40 μeV of parameter space using the single-photon detector with existing haloscope technology. When the single-photon detector is combined with improvements in the magnetic field strength and conversion cavity design, as in scenarios B and C, QCD axion searches across ∼40 μeV of parameter space become feasible. If the frequency-volume scaling can be broken, single-photon detectors allow for thorough QCD axion searches even with modest magnetic field and conversion cavity Q, enabling either a KSVZ axion search across the entire 40 μeV–200 μeV mass range or a DFSZ search across ∼20 μeV of parameter space. Across all of these scenarios, single-photon detectors enhance the discovery potential of their companion haloscopes, opening up regions of parameter space otherwise inaccessible to those haloscopes.

VI. SUMMARY AND CONCLUSION

Searching for axions with a cavity haloscope experiment requires sensitivity to an extremely low photon flux, especially for higher masses above 40 μeV (10 GHz). One way of approaching this challenge is to replace standard linear amplifiers with single-photon detectors, which are not subject to the quantum measurement noise that dominates linear amplifiers at high frequencies and can achieve much higher signal-to-noise ratios as a result.

In this paper, we presented a scheme for a Rydberg-atom-based single-photon detector that can be used for axion searches in the mass range 40–200 μeV. This single-photon detector can be connected to a wide variety of haloscope designs, replacing standard linear-amplifier-based readout schemes and delivering scan-rate enhancement factors of 10^4. Reaching the full potential of single-photon detection with Rydberg atoms will require producing high-Q cavities for atom-photon coupling and developing techniques for

working with Rydberg atoms in cryostats with operating temperatures below 100 mK. When Rydberg-atom-based single-photon detector readout is combined with projected improvements for the next generation of magnetic fields and axion-photon conversion cavities, haloscope searches across 40 μeV (∼10 GHz) of QCD axion parameter space become feasible.

APPENDIX: COMPARISON TO LINEAR AMPLIFIERS

We compare the performance of our single-photon detector to the performance of standard linear amplifier readout by calculating the scan rate enhancement. We define the scan rate enhancement E as the ratio of scan rates for axion searches conducted with a single-photon detector and a linear amplifier attached to the same haloscope. Dividing Eq. (2.4) by Eq. (2.6) under the approximation $\Delta\nu_{\text{det}} \approx \Delta\nu_a$,

$$E = \frac{5}{4} \frac{\Delta\nu_a}{\Delta\nu_{\text{conv}}} \frac{R_{\text{sig,sp}}^2}{R_{\text{sig,la}}^2} \frac{N_{\text{sys,la}}^2 \Delta\nu_a}{R_{\text{sig,sp}} + R_{\text{sys,sp}}}. \qquad (A1)$$

We simplify this expression in a few ways. First, we drop the factor of 5/4. Then, we rewrite $\Delta\nu_a/\Delta\nu_{\text{conv}}$ in its equivalent form Q_{conv}/Q_a. Next, we assume that the two signal rates are approximately equal and drop the subscript. This is equivalent to assuming that the single-photon detector and linear amplifier readout have the same efficiency, which is reasonable according to Sec. IV C. Finally, we set $N_{\text{sys,la}} \approx 1$ for a linear amplifier operating at the SQL, ignoring the relatively small contribution of thermal noise.

Using these simplifications and writing the single-photon noise rate as in Eq. (4.17),

$$E \approx \frac{Q_{\text{conv}}}{Q_a} \frac{\Delta\nu_a}{R_{\text{sig}} + R_T + R_{\text{RO}}}. \qquad (A2)$$

At lower frequencies, the noise affecting the scan rate enhancement primarily comes from BBR. The exact frequency where the single-photon detector starts to outperform the linear amplifier depends on the operating

temperature and the ratio of the conversion cavity Q to the axion Q, consistent with the estimates made in [31]. For low enough temperature and high enough frequency, the BBR noise becomes negligible and the other sources of noise in the single-photon detector take precedence; shot noise in the signal counts and dark counts from the readout mechanism.

At higher frequencies, the dominant source of noise comes from shot noise in the single-photon detector. Although the readout dark count rate can be improved through use of low-noise electron detectors, shot noise comes directly from the statistics of searching for axions by counting photons. Since linear amplifier readout does not have shot noise, this reduces the maximum scan rate enhancement; increasing R_{sig} increases the SNR of the single-photon detector by a factor of $\sqrt{R_{\text{sig}}}$ but increases the SNR of the linear amplifier by a factor of R_{sig}. If all other sources of noise are negligible and $Q_{\text{conv}}/Q_a \approx 0.01$, single-photon detectors are favorable for $R_{\text{sig}} < 100$ ct/s and can achieve a scan rate enhancement of 10^3 for $R_{\text{sig}} < 1$ ct/s.

Signal rates that could significantly reduce the scan rate enhancement achievable with a single-photon detector are well in excess of the signal rates achievable with haloscopes operating above 10 GHz. For instance, the HAYSTAC experiment would see 0.3 ct/s from KSVZ axions when operating around 5 GHz [25]. As indicated in Eq. (2.5), the signal power is proportional to the inverse cube of the frequency, so reaching 1 ct/s at 10 GHz and 50 GHz would require improving the baseline signal power of that haloscope by a factor of 25 and 25000, respectively.

The scan rate enhancement as a function of frequency is shown in Fig. 8 for several different values of T and $R_{\text{sig}}(10 \text{ GHz})$. We estimate $Q_{\text{conv}}/Q_a \approx 0.01$ and
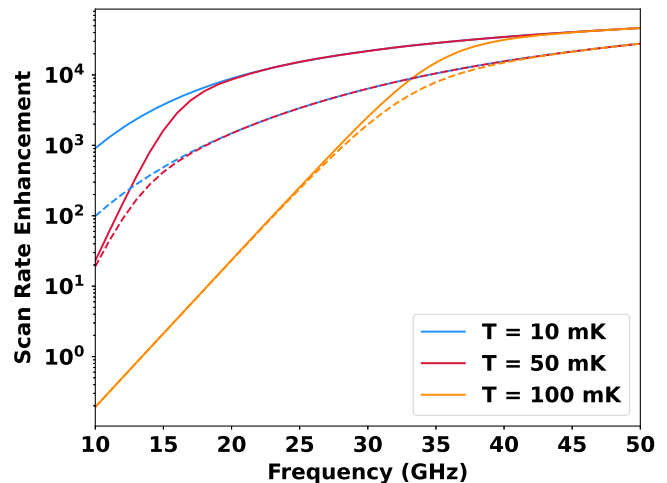

FIG. 8. Scan rate enhancement for a haloscope using the proposed single-photon detector for readout instead of a SQL-limited linear amplifier, shown for a readout dark count rate of 0.01 ct/s and operating temperatures of $T = 10$ mK, 50 mK, and 100 mK. Solid lines show the enhancement for $R_{\text{sig}}(10 \text{ GHz}) = 0.1$ ct/s, while dashed lines show the enhancement for $R_{\text{sig}}(10 \text{ GHz}) = 1$ ct/s.

$R_{\text{RO}} \approx 0.01$ ct/s. At 10 mK, BBR noise is negligible across the entire frequency range. At 50 mK, BBR noise can be seen until around 20 GHz, but the single-photon detector still outperforms the linear amplifier across the entire frequency range. At 100 mK, BBR noise is present until 40 GHz, and the single-photon detector becomes favorable to the linear amplifier around 15 GHz. In the absence of BBR, the scan rate enhancement becomes near-constant with frequency, saturating around 5×10^4 for signal rates of 0.1 ct/s and around 10^4 for signal rates of 1 ct/s.

[1] R. D. Peccei and H. R. Quinn, *CP* conservation in the presence of pseudoparticles, Phys. Rev. Lett. **38**, 1440 (1977).

[2] R. D. Peccei and H. R. Quinn, Constraints imposed by *CP* conservation in the presence of pseudoparticles, Phys. Rev. D **16**, 1791 (1977).

[3] S. Weinberg, A new light boson?, Phys. Rev. Lett. **40**, 223 (1978).

[4] F. Wilczek, Problem of strong *P* and *T* invariance in the presence of instantons, Phys. Rev. Lett. **40**, 279 (1978).

[5] J. Preskill, M. B. Wise, and F. Wilczek, Cosmology of the invisible axion, Phys. Lett. **120B**, 127 (1983).

[6] L. F. Abbott and P. Sikivie, A cosmological bound on the invisible axion, Phys. Lett. **120B**, 133 (1983).

[7] M. Dine and W. Fischler, The not so harmless axion, Phys. Lett. **120B**, 137 (1983).

[8] G. G. di Cortona, E. Hardy, J. P. Vega, and G. Villadoro, The QCD axion, precisely, J. High Energy Phys. 01 (2016) 034.

[9] G. G. Raffelt, Astrophysical axion bounds, in *Lecture Notes in Physics* (Springer, Berlin, Heidelberg, 2008), pp. 51–71, https://doi.org/10.1007/978-3-540-73518-2_3.

[10] A. Arvanitaki, M. Baryakhtar, and X. Huang, Discovering the QCD axion with black holes and gravitational waves, Phys. Rev. D **91**, 084011 (2015).

[11] D. J. Marsh, Axion cosmology, Phys. Rep. **643**, 1 (2016).

[12] V. Cardoso, Ó. J. Dias, G. S. Hartnett, M. Middleton, P. Pani, and J. E. Santos, Constraining the mass of dark photons and axion-like particles through black-hole superradiance, J. Cosmol. Astropart. Phys. 03 (2018) 043.

[13] 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).

[14] J. E. Kim, Weak-interaction singlet and strong *CP* invariance, Phys. Rev. Lett. **43**, 103 (1979).

[15] M. Shifman, A. Vainshtein, and V. Zakharov, Can confinement ensure natural CP invariance of strong interactions?, Nucl. Phys. **B166**, 493 (1980).

[16] M. Dine, W. Fischler, and M. Srednicki, A simple solution to the strong *CP* problem with a harmless axion, Phys. Lett. **104B**, 199 (1981).

[17] A. R. Zhitnitsky, On possible suppression of the axion hadron interactions, Sov. J. Nucl. Phys. **31**, 260 (1980).

[18] C. Bartram *et al.* (ADMX Collaboration), Search for "invisible" axion dark matter in the 3.3–4.2 μev mass range, Phys. Rev. Lett. **127**, 261803 (2021).

[19] 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).

[20] O. Kwon *et al.* (CAPP Collaboration), First results from an axion haloscope at CAPP around 10.7 μeV, Phys. Rev. Lett. **126**, 191802 (2021).

[21] H. Chang *et al.* (TASEH Collaboration), First results from the taiwan axion search experiment with a haloscope at 19.6 μeV, Phys. Rev. Lett. **129**, 111802 (2022).

[22] D. Alesini *et al.*, Search for galactic axions with a high-*q* dielectric cavity, Phys. Rev. D **106**, 052007 (2022).

[23] P. Sikivie, Experimental tests of the "invisible" axion, Phys. Rev. Lett. **51**, 1415 (1983).

[24] P. Sikivie, Detection rates for "invisible"-axion searches, Phys. Rev. D **32**, 2988 (1985).

[25] K. M. Backes *et al.* (HAYSTAC Collaboration), A quantum-enhanced search for dark matter axions, Nature (London) **590**, 238 (2021).

[26] S. Borsanyi *et al.*, Calculation of the axion mass based on high-temperature lattice quantum chromodynamics, Nature (London) **539**, 69 (2016).

[27] G. Ballesteros, J. Redondo, A. Ringwald, and C. Tamarit, Unifying inflation with the axion, dark matter, baryogenesis, and the seesaw mechanism, Phys. Rev. Lett. **118**, 071802 (2017).

[28] R. T. Co, L. J. Hall, and K. Harigaya, Axion kinetic misalignment mechanism, Phys. Rev. Lett. **124**, 251802 (2020).

[29] C. B. Adams *et al.*, Axion dark matter, arXiv:2203.14923.

[30] C. M. Caves, Quantum limits on noise in linear amplifiers, Phys. Rev. D **26**, 1817 (1982).

[31] 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).

[32] 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).

[33] G. Romero, J. J. García-Ripoll, and E. Solano, Microwave photon detector in circuit QED, Phys. Rev. Lett. **102**, 173602 (2009).

[34] Y.-F. Chen, D. Hover, S. Sendelbach, L. Maurer, S. T. Merkel, E. J. Pritchett, F. K. Wilhelm, and R. McDermott, Microwave photon counter based on Josephson junctions, Phys. Rev. Lett. **107**, 217401 (2011).

[35] X. Gu, A. F. Kockum, A. Miranowicz, Y. xi Liu, and F. Nori, Microwave photonics with superconducting quantum circuits, Phys. Rep. **718–719**, 1 (2017).

[36] 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).

[37] Z. Wang, M. Xu, X. Han, W. Fu, S. Puri, S. M. Girvin, H. X. Tang, S. Shankar, and M. H. Devoret, Quantum microwave radiometry with a superconducting qubit, Phys. Rev. Lett. **126**, 180501 (2021).

[38] 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).

[39] A. L. Pankratov, L. S. Revin, A. V. Gordeeva, A. A. Yablokov, L. S. Kuzmin, and E. Il'ichev, Towards a microwave single-photon counter for searching axions, npj Quantum Inf. **8**, 61 (2022).

[40] J. Sedlacek, A. Schwettmann, H. Kübler, R. Löw, T. Pfau, and J. Shaffer, Microwave electrometry with Rydberg atoms in a vapour cell using bright atomic resonances, Nat. Phys. **8**, 819 (2012).

[41] M. Tada, Y. Kishimoto, K. Kominato, M. Shibata, H. Funahashi, K. Yamamoto, A. Masaike, and S. Matsuki, CARRACK II—a new large-scale experiment to search for axions with Rydberg-atom cavity detector, Nucl. Phys. B, Proc. Suppl. **72**, 164 (1999).

[42] 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 Berlin Heidelberg, Berlin, Heidelberg, 2001), pp. 638–645.

[43] J. I. Read, The local dark matter density, J. Phys. G **41**, 063101 (2014).

[44] C. Patrignani *et al.* (Particle Data Group), Review of particle physics, Chin. Phys. C **40**, 100001 (2016).

[45] X. Bai, M. J. Jewell, S. K. Lamoreaux, R. H. Maruyama, and K. van Bibber, On the use of dielectric elements in axion searches with microwave resonant cavities, J. Instrum. **18**, P07017 (2023).

[46] S. Al Kenany *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).

[47] H. A. Haus and J. A. Mullen, Quantum noise in linear amplifiers, Phys. Rev. **128**, 2407 (1962).

[48] H. Zheng, M. Silveri, R. T. Brierley, S. M. Girvin, and K. W. Lehnert, Accelerating dark-matter axion searches with quantum measurement technology, arXiv:1607.02529.

[49] M. Tada, Y. Kishimoto, K. Kominato, M. Shibata, S. Yamada, T. Haseyama, I. Ogawa, H. Funahashi, K. Yamamoto, and S. Matsuki, Single-photon detection of microwave blackbody radiations in a low-temperature resonant-cavity with high Rydberg atoms, Phys. Lett. A **349**, 488 (2006).

[50] T. Haseyama *et al.*, A high-sensitivity microwave-single-photon detector with Rydberg atoms at low temperature, J. Low Temp. Phys. **150**, 549 (2008).

[51] N. Šibalić, J. Pritchard, C. Adams, and K. Weatherill, Arc: An open-source library for calculating properties of alkali Rydberg atoms, Comput. Phys. Commun. **220,** 319 (2017).

[52] T. F. Gallagher, Rydberg Atoms, *Cambridge Monographs on Atomic, Molecular and Chemical Physics* (Cambridge University Press, Cambridge, England, 1994).

[53] Y. Zhu, S. Ghosh, S. B. Cahn, M. J. Jewell, D. H. Speller, and R. H. Maruyama, Electromagnetically-induced-transparency spectroscopy of high-lying Rydberg states in ^{39}K, Phys. Rev. A **105,** 042808 (2022).

[54] R. J. Thompson, G. Rempe, and H. J. Kimble, Observation of normal-mode splitting for an atom in an optical cavity, Phys. Rev. Lett. **68,** 1132 (1992).

[55] J. M. Raimond, M. Brune, and S. Haroche, Manipulating quantum entanglement with atoms and photons in a cavity, Rev. Mod. Phys. **73,** 565 (2001).

[56] S. Haroche and J. Raimond, Radiative properties of Rydberg states in resonant cavities, in *Advances in Atomic and Molecular Physics* (Elsevier, New York, 1985), Vol. 20, pp. 347–411.

[57] M. Gross, P. Goy, C. Fabre, S. Haroche, and J. M. Raimond, Maser oscillation and microwave superradiance in small systems of Rydberg atoms, Phys. Rev. Lett. **43,** 343 (1979).

[58] M. Gross and S. Haroche, Superradiance: An essay on the theory of collective spontaneous emission, Phys. Rep. **93,** 301 (1982).

[59] K. Yamamoto and S. Matsuki, Quantum analysis of the Rydberg atom cavity detector of dark matter axions, Nucl. Phys. B, Proc. Suppl. **72,** 132 (1999).

[60] D. Main, T. M. Hird, S. Gao, E. Oguz, D. J. Saunders, I. A. Walmsley, and P. M. Ledingham, Preparing narrow velocity distributions for quantum memories in room-temperature alkali vapours, Phys. Rev. A **103,** 043105 (2021).

[61] B. Wei, A. Crawford, Y. Andeweg, L. Zhuo, C. Li, and C. Raman, Collimated versatile atomic beam source with alkali dispensers, Appl. Phys. Lett. **120,** 144001 (2022).

[62] J. Catani, P. Maioli, L. D. Sarlo, F. Minardi, and M. Inguscio, Intense slow beams of bosonic potassium isotopes, Phys. Rev. A **73,** 033415 (2006).

[63] R. A. Larsen, S. K. Neoh, and D. R. Herschbach, Seeded supersonic alkali atom beams, Rev. Sci. Instrum. **45,** 1511 (1974).

[64] M. Tada, Y Kishimoto, M Shibata, K Kominato, S Yamada, T Haseyama, I Ogawa, H Funahashi, K Yamamoto, and S Matsuki, Manipulating ionization path in a stark map: Stringent schemes for the selective field ionization in highly excited Rb Rydberg, Phys. Lett. A **303,** 285 (2002).

[65] A. Gürtler and W. J. van der Zande, l-state selective field ionization of rubidium Rydberg states, Phys. Lett. A **324,** 315 (2004).

[66] V. C. Gregoric *et al.*, Perturbed field ionization for improved state selectivity, J. Phys. B **53,** 084003 (2020).

[67] P. Maioli, T. Meunier, S. Gleyzes, A. Auffeves, G. Nogues, M. Brune, J. M. Raimond, and S. Haroche, Nondestructive Rydberg atom counting with mesoscopic fields in a cavity, Phys. Rev. Lett. **94,** 113601 (2005).

[68] M. Shibata, M. Tada, Y. Kishimoto, K. Kominato, T. Haseyama, I. Ogawa, S. Matsuki, S. Yamada, H. Funahashi, and K. Yamamoto, Field-ionization electron detector at low temperature of 10 mK range, Rev. Sci. Instrum. **74,** 3317 (2003).

[69] D. M. Pozar, *Microwave Engineering*, 3rd ed. (Wiley, Hoboken, NJ, 2005).

[70] Dr. Sjuts Optotechnik GmbH, CEM Models / Extended Dynamic Range CEMs, http://www.sjuts.com/CEMModels_EDR.html.

[71] S. Posen, M. Checchin, O. S. Melnychuk, T. Ring, and I. Gonin, Measurement of high quality factor superconducting cavities in tesla-scale magnetic fields for dark matter searches, Phys. Rev. Appl. **20,** 034004 (2023).

[72] D. Alesini *et al.*, Galactic axions search with a superconducting resonant cavity, Phys. Rev. D **99,** 101101 (2019).

[73] 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 yba$_2$cu$_3$o$_{7-x}$ microwave cavity in a high magnetic field for a dark-matter axion search, Phys. Rev. Appl. **17,** L061005 (2022).

[74] A. Caldwell, G. Dvali, B. Majorovits, A. Millar, G. Raffelt, J. Redondo, O. Reimann, F. Simon, and F. Steffen (MADMAX Working Group), Dielectric haloscopes: A new way to detect axion dark matter, Phys. Rev. Lett. **118,** 091801 (2017).

[75] M. Lawson, A. J. Millar, M. Pancaldi, E. Vitagliano, and F. Wilczek, Tunable axion plasma haloscopes, Phys. Rev. Lett. **123,** 141802 (2019).

[76] A. J. Millar *et al.* (ALPHA Collaboration), Searching for dark matter with plasma haloscopes, Phys. Rev. D **107,** 055013 (2023).

[77] Y. Lee, B. Yang, H. Yoon, M. Ahn, H. Park, B. Min, D. Kim, and J. Yoo, Searching for invisible axion dark matter with an 18 T magnet haloscope, Phys. Rev. Lett. **128,** 241805 (2022).

[78] National High Magnetic Field Laboratory, 32 Tesla Superconducting Magnet (SCM-32 T), https://nationalmaglab.org/user-facilities/dc-field/magnets-instruments/superconducting-magnets/scm-32-t/.

[79] L. Di Luzio, M. Giannotti, E. Nardi, and L. Visinelli, The landscape of QCD axion models, Phys. Rep. **870,** 1 (2020).