

Microwave cavity searches for low-frequency axion dark matter

Robert Lasenby ^{*}

Stanford Institute for Theoretical Physics, Stanford University, Stanford, California 94305, USA



(Received 23 March 2020; accepted 12 June 2020; published 14 July 2020)

For low-mass (frequency \ll GHz) axions, dark matter detection experiments searching for an axion-photon-photon coupling generally have suppressed sensitivity, if they use a static background magnetic field. This geometric suppression can be alleviated by using a high-frequency oscillating background field. Here, we present a high-level sketch of such an experiment, using superconducting cavities at \sim GHz frequencies. We discuss the physical limits on signal power arising from cavity properties, and point out cavity geometries that could circumvent some of these limitations. We also consider how backgrounds, including vibrational noise and drive signal leakage, might impact sensitivity. While practical microwave field strengths are significantly below attainable static magnetic fields, the lack of geometric suppression, and higher quality factors, may allow superconducting cavity experiments to be competitive in some regimes.

DOI: [10.1103/PhysRevD.102.015008](https://doi.org/10.1103/PhysRevD.102.015008)

I. INTRODUCTION

While dark matter (DM) has so far only been observed through its gravitational interactions, many theoretical candidates have additional interactions with the Standard Model (SM). This has motivated a wide range of laboratory searches for DM, most famously through the WIMP direct detection program. Another well-motivated dark matter candidate is the QCD axion (and more generally, axionlike particles), for which an extensive experimental program also exists [1].

A major part of this program involves searches for axion DM through the $aF_{\mu\nu}\tilde{F}^{\mu\nu}$ axion-photon-photon coupling. This coupling is a generic prediction of QCD axion models [2,3], and arises in many other theories of axionlike DM [4]. For axion masses corresponding to frequencies \gtrsim GHz, simple experimental designs such as cavity haloscopes (e.g., the ADMX experiment [5,6]) or dielectric haloscopes [7,8] can achieve almost-optimal [9] axion-mass-averaged power absorption from the DM field.

However, the situation is different for much smaller axion masses, corresponding to Compton wavelengths significantly longer than the length scale of the experiment (either the shielding scale or the extent of the magnetic field, whichever is smaller). The EM fields at these frequencies are naturally in the quasistatic regime, and for axion detection experiments using a static background

magnetic field, this means that the absorbed power is parametrically suppressed by $\sim(m_a L)^2$, where L is the experimental length scale. The origin of this suppression is discussed in a number of papers [9–11]; roughly speaking, it arises because the electric field oscillations associated with current fluctuations are suppressed compared to the magnetic field oscillations, reducing the interaction with the axion-sourced effective current. This suppression affects low-frequency axion detection proposals such as ABRACADABRA [12] and DM Radio [13], reducing their sensitivity.¹

One way to avoid these issues is to use an oscillating background magnetic field. If this is oscillating at a frequency ω_0 , then it combines with the axion field oscillation to generate an effective current oscillating at $\omega_0 \pm m_a$. By taking $\omega_0 \gtrsim L^{-1} \gg m_a$, the response excited by this effective current is no longer in the quasistatic regime, and the absorbed power is no longer suppressed. Since the signal excitation is at a much higher frequency than the axion field, this approach is referred to as “up-conversion” [14]. The idea of using an oscillating background magnetic field in axion detection experiments was first proposed in [15], but they were mostly concerned with GHz-scale axion frequencies, and did not consider the parametric scaling at low axion masses.

Up-conversion experiments using optical-frequency background fields have been proposed in [16–18]. However, these encounter a number of issues. The most serious is that achievable magnetic field strengths at

^{*}rlasenby@stanford.edu

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI. Funded by SCOAP³.

¹While the sensitivity of experiments such as ABRACADABRA is not directly related to absorbed power, the g_{arr} sensitivity is still suppressed by $\sim(m_a L)$, compared to its theoretical scaling at higher frequencies [9].

optical frequencies are very small, compared to static magnetic fields. Another is that, even if other noise sources are overcome, the shot noise suppression coming from absorbing fewer, but higher-energy, optical photons degrades the theoretical sensitivity limits still further. Consequently, it would be very difficult to match the sensitivity of static-field experiments, despite the lack of a $(m_a L)^2$ suppression.

We can address both of these issues by using lower-frequency magnetic field oscillations. In particular, reasonably large magnetic fields (~ 0.2 T) are routinely attained at \sim GHz frequencies inside superconducting (SRF) cavities [19,20], and as we discuss below, it may be possible to achieve even higher fields. The lower frequency also alleviates shot noise issues. Thus, while the average magnetic fields will still be lower than those used in static-field experiments, the relative $(m_a L)^2$ enhancement may be enough to make up-conversion experiments interesting.

SRF up-conversion experiments were first proposed in [15], but as stated above, they were mainly concerned with \sim GHz axion frequencies. Up-conversion experiments for low-mass axions were considered in [14], but their sensitivity estimates are affected by some calculational errors [21], and seem to violate theoretical bounds [9].

In this paper, we discuss the basic physics and design considerations involved in microwave cavity up-conversion experiments. One of the most obvious questions is how to choose the cavity geometry. We derive constraints on the signal power attainable from cavities, and show that for simple geometries (in which all of the walls are visible from an interior point) the RMS magnetic field is limited by the magnetic field at the walls. Since, for superconducting cavities, this is limited by the superconductor's material properties, the signal power from such cavities is bounded. We also show how this bound can be circumvented using cavities with more complicated shapes, and illustrate that, to probe significant QCD axion parameter space with small cavity volumes, such geometries may be practically necessary.

In addition to such considerations, which would be important for more advanced experiments, we also outline a nominal first-generation experiment, aiming to be as simple as possible while still having interesting reach. We give a high-level discussion of the noise issues that might arise, and derive representative sensitivity estimates. This discussion is not intended as a design study, and a realistic experiment might be significantly more complicated. Instead, our goal is to illustrate the physical parameters that might be required, and to motivate further study of this experimental direction.

II. AXION DM UP-CONVERSION

We will assume that dark matter consists of a light axionlike particle a , which couples to the SM via

the electromagnetic $F_{\mu\nu}\tilde{F}^{\mu\nu}$ operator. This has Lagrangian²

$$\begin{aligned}\mathcal{L} &\supset \frac{1}{2}(\partial_\mu a)^2 - V(a) - \frac{1}{4}g_{a\gamma\gamma}aF_{\mu\nu}\tilde{F}^{\mu\nu} \\ &= \frac{1}{2}(\partial_\mu a)^2 - V(a) + g_{a\gamma\gamma}aE \cdot B,\end{aligned}\quad (1)$$

where $V(a)$ is the potential for the axion—in general, only the mass term $V(a) = \frac{1}{2}m_a^2 a^2$ will be important for us.

For light DM ($m_a \ll$ eV), the occupation number in the Milky Way is $\gg 1$, and almost all cosmological histories result in its state today being a coherent, classical-like oscillation [2,22–25]. Since $g_{a\gamma\gamma}$ (and other couplings) are constrained to be very small, interactions with a detector will have a negligible effect on the DM's state. Consequently, for the purposes of detection, we can treat the DM oscillation as a fixed classical background field.

Under integration by parts, the interaction term is equivalent to $\mathcal{L} \supset -\frac{1}{2}A_\mu J_{(a)}^\mu$, where the “axion current” $J_{(a)}^\mu$ is given by

$$J_{(a)}^\mu = g \begin{pmatrix} -\nabla a \cdot B \\ \dot{a}B + \nabla a \times E \end{pmatrix} \quad (2)$$

Since axion DM in the galaxy is nonrelativistic, with typical velocity $\sim 10^{-3}$, the dominant term is the spatial current $J_{(a)} \simeq g\dot{a}B$. Hence, the effects of the axion field are equivalent to those of an oscillating current density, with profile set by the background magnetic field.

As described in the introduction, we will consider up-conversion experiments, in which the background magnetic field is oscillating at a frequency $\omega_0 \gg m_a$. Writing the background magnetic field as $B_0(t, x) \simeq B_0 \cos(\omega_0 t)b(x)$, and given a single-frequency axion oscillation, $a(t) = a_0 \cos(m_a t)$, the effective current $J_{(a)}$ will have frequency components at $\omega_0 \pm m_a$. As illustrated in Fig. 1, the basic idea is to arrange things so that there is another EM mode with resonant frequency $\simeq \omega_0 + m_a$ (or $\omega_0 - m_a$), so that it can efficiently absorb power from the effective current oscillation. The larger this absorbed power, the easier it is to detect an axion DM signal (other things being equal).

For an EM mode with electric field profile $E_1(x)$, the instantaneous power input from the axion current is $P_a = -\int dV E_1 \cdot J_{(a)}$. Considering only a single frequency component of $J_{(a)}$, say $\omega_J = m_a + \omega_B$, the cycle-averaged input power at this frequency is

²We take the $(+---)$ signature, and use the convention $\epsilon_{0123} = -1$. Except where indicated, we use natural units with $c = \hbar = 1$. In general, we will abbreviate $g_{a\gamma\gamma} = g$.




FIG. 1. Schematic illustration of an up-conversion experiment for axion DM detection. The left-most panel (“cavity modes”) illustrates a cavity in which we are interested in two modes, one with magnetic field profile B_0 , and another with electric field profile E_1 . The adjacent graph (“mode frequencies”) illustrates the frequencies of these modes as we vary the shape of the cavity (here, the length d), showing that they become degenerate at some d . When the cavity length is tuned to be close to this point, the frequency splitting between the modes will be small. If the ω_0 mode is driven with high amplitude, then in the presence of an axion DM oscillation with $m_a \simeq |\omega_1 - \omega_0|$, power will be transferred to the ω_1 mode (illustrated in the “interaction” panel). This signal can be detected as excess power, above thermal (or quantum) fluctuations, in the ω_1 mode, as indicated in the “signal” panel. While some of the power in the ω_0 mode may leak into the detected signal (indicated as a dashed peak in the signal panel), the frequency difference between ω_0 and ω_1 helps to distinguish this from the axion signal (see Sec. III C). The mode profiles shown here are schematic—Fig. 2 shows actual mode profiles for a particular cavity geometry.

$$\bar{P}_a = \frac{1}{4} \cos \alpha g \omega_J a_0 B_0 \int dV E_1 \cdot b \quad (3)$$

where α is the relative phase of the electric field response and the $-J_{(a)}$ oscillation. The average power dissipated is $P_{\text{diss}} = \omega_J U / Q_l$, where $U = \frac{1}{2} \int dV E_1^2$ is the stored energy in the mode (at this frequency), and Q_l is the (loaded) quality factor of the mode. In a steady state, the absorbed and dissipated powers will be equal. Equating $\bar{P}_a = P_{\text{diss}}$, we find that the cycle-averaged absorbed power, once fully rung up, is

$$P_{\text{sig}} = \frac{1}{8} \cos^2 \alpha (g a_0 B_0)^2 m_a^2 \frac{Q_l (\int dV b \cdot E_1)^2}{\omega_J \int dV E_1^2}. \quad (4)$$

We can define the geometric overlap factor

$$C_{01} \equiv \frac{(\int dV b \cdot E_1)^2}{(\int dV E_1^2)(\int dV b^2)} \quad (5)$$

which measures the degree of overlap between the background magnetic field and the E_1 electric field; this has $C_{01} \leq 1$, with equality iff $E_1(x) \propto b(x)$. Using this, we can write the signal power as

$$P_{\text{sig}} = \frac{1}{4} \cos^2 \alpha (g a_0)^2 m_a^2 \frac{Q_l}{\omega_J} C_{01} U_0 \quad (6)$$

where $U_0 = \frac{1}{2} B_0^2 V_b$. For a high quality factor mode, if the signal frequency ω is close to the resonance frequency ω_1 , then

$$\cos^2 \alpha \simeq \frac{1}{1 + Q_l^2 \frac{(\omega^2 - \omega_1^2)^2}{\omega_1^4}} \simeq \frac{1}{1 + 4Q_l^2 (\frac{\omega}{\omega_1} - 1)^2}. \quad (7)$$

The expression in Eq. (6) is equivalent to Eq. (3) in [15], and gives the cycle-averaged power absorbed by a particular mode (assuming that the mode bandwidth is small compared to m_a). In most of the cases we are interested in, only one mode of the cavity will have a resonance frequency close to the signal frequency, so only that mode will absorb appreciable power from the axion current.

In many circumstances (see Appendix B), we are interested in the signal power averaged over different axion masses. If we integrate over an axion mass range Δm significantly larger than the bandwidth of the target mode, then the average power absorbed is

$$\bar{P} \simeq \frac{\pi}{4} g^2 \frac{\rho}{\Delta m} C_{01} U_0. \quad (8)$$

This formula is valid even for low-frequency B_0 oscillations, as long as the integration time is long enough to resolve the variation of B . However, if $\omega_0 \ll L^{-1}$, where L^{-1} is the linear scale of the shielded experimental volume (or the magnetic field extent, whichever is smaller), then C_{01} is generally suppressed compared to the theoretical limit, with $C_{01} \sim (\omega_0 L)^2$. This is because the EM fields in the volume are in the quasistatic regime, as discussed in [9]. Consequently, to avoid this geometrical suppression, we want to take $\omega_0 \gtrsim L^{-1}$, i.e., $\gtrsim \text{GHz}$ for laboratory-scale experiments. In such cases, if C_{01} is close to 1, then the absorbed power can be close to the theoretically obtainable limit, for a given background magnetic field energy [9].

A common limitation on the rate at which large magnetic fields can be varied is the large amount of

field energy stored—taking some nominal parameters, $\text{Tesla}^2 \times \text{m}^3 \sim \text{MJ}$. Feeding that amount of energy into and out of a system must generally be done rather slowly. For example, the currents through high-field superconducting magnets generally take minutes or more to build up to their full values, with faster changes damaging the system. Faster rates of change are possible with specially-designed superconducting systems [26], but the dissipated power increases with frequency, and is generally prohibitive for rates of change $\gtrsim 1$ T/s. Resistive conductors can tolerate more heating / stresses, but sustained operation at high field strengths dissipates very large powers. Similarly, varying the fields from magnetic materials will either involve mechanical motion, or hysteretic energy losses. In all of these cases, achieving strong $\sim \text{GHz}$ frequency magnetic fields will not be possible.

A. SRF cavities

We can get around these issues by using cavities with high quality factors, in which magnetic field energy is exchanged back and forth with the electric field energy inside the cavity, rather than needing to be transferred in and out each cycle. Filling the cavity to high field amplitude is still a slow process, but once this amplitude has been established, a high-field oscillation can be maintained with only a small energy input, to counteract the small dissipation rate.

The basic setup of a cavity up-conversion experiment is illustrated in Fig. 1. The cavity’s shape is tuned so that there are two modes with frequency difference $|\omega_1 - \omega_0| \simeq m_a$. One of these modes is driven to a high field amplitude, and in the presence of this background oscillation, an axion DM oscillation would lead to signal power in the other mode, according to Eq. (6). To maximise the signal power, we want C_{01} to be close to one.

To achieve a high quality factor, and allow for high drive fields without excessive energy dissipation, the cavity walls should be superconducting. SRF (superconducting radio frequency) cavities have been extensively developed for particle acceleration [19]. They are also starting to be used directly in hidden sector particle searches—the Fermilab DarkSRF project [27] is constructing a dark photon detection experiment using SRF cavities, and there have been other proposals for axion detection [28,29]. For axion DM detection experiments, oscillating background magnetic fields inside SRF cavities were first proposed in [15]. In this section, we will review some of the important properties of SRF cavities, which will affect the design and operation of an up-conversion experiment.

- (i) *Peak surface magnetic field*: if the magnetic field at the walls of the cavity becomes too large, the behavior of the superconducting material will change. Type-I superconductors generally have rather low critical fields (e.g., for Aluminium, $H_c \simeq 0.01$ T [30]), so are unsuitable for high-field cavities. SRF cavities are

fabricated using Type-II superconductors, and almost always use niobium. This has a critical field value H_{c1} above which vortices penetrate; if this happens, then the radio-frequency oscillation of these vortices will lead to increased dissipation, and generally runaway thermal instability [19]. It is actually possible to operate slightly above H_{c1} , in a metastable “Meissner state”—while it energetically favorable to have flux deep within the bulk, establishing this configuration involves penetrating a surface energy barrier [19]. Vortices start penetrating the surface at H_{sh} , which for niobium is $\simeq 0.2$ T. Consequently, the magnetic field at the cavity walls should always be $\lesssim 0.2$ T. As we will see below, this puts limits on the achievable fields inside the cavity, and correspondingly, on the signal power attainable from axion DM.

- (ii) *Surface resistance*: the quality factor of a cavity mode is set by its magnetic fields at the cavity walls (which determine the wall currents), and the surface resistance there. This resistance has a “BCS” component, and a “residual” component,

$$R_s = R_{\text{BCS}} + R_{\text{res}} \quad (9)$$

The BCS component, for oscillations at frequency ω , is given approximately by

$$R_{\text{BCS}} \simeq \omega^2 \lambda^3 \sigma_n \frac{\Delta}{T} \log\left(\frac{2.246T}{\omega}\right) e^{-\Delta/T} \quad (10)$$

where λ is the (effective) penetration depth, σ_n is the normal-state conductivity, and Δ is the superconducting gap energy [31]. The residual resistance is the component that persists as $T \rightarrow 0$, and is measured to be $\sim \text{few n}\Omega$ for good niobium cavities (physically, it is not entirely clear what this residual resistance is dominated by [19,31]). The $R_{\text{BCS}} \propto \omega^2$ dependence means that, for frequencies $\gtrsim 3$ GHz, the BCS resistance starts to dominate. Since it is an increasing function of temperature, this can lead to thermal instability problems [19]. Hence, SRF cavities are generally operated at lower frequencies.

- (iii) *Cooling*: The quality factor of a mode sets the power dissipated, for a given energy stored. As we will see in the next section, taking the H_{sh} limit and R_s values from above, and applying them to a simple laboratory-scale (~ 60 L) cavity, gives $P_{\text{diss}} \sim 30$ W (with $Q \simeq 2 \times 10^{11}$). Since this heat eventually needs to be dissipated to the $T_H \sim 300$ K environment, the maximum efficiency of the cooling system is $\eta_C = \frac{T_0}{T_H - T_0} \simeq T_0/T_H \simeq 3 \times 10^{-3} T_0/\text{K}$, where T_0 is the temperature of the cavity. So, we would need at least 10 kW of electrical power to cool the cavity to $T_0 = 1$ K. Taking a typical thermal efficiency of $\eta_T \simeq 0.2$, this becomes $\gtrsim 50$ kW. These

figures illustrate that cooling the cavity to significantly sub-kelvin temperature would be prohibitively power-hungry. The high cooling powers required generally necessitate the use of liquid helium cooling systems. The pump machinery involved leads to mechanical vibrations of the cavity, which can introduce noise and tuning issues, as we discuss in Sec. III B.

- (iv) *Field emission*: if the electric fields at the cavity walls are high enough, then electrons can escape from the surface via tunneling. The field emission rate can be approximated via a modified Fowler-Nordheim formula [19],

$$I_F \simeq 10^{-7} \frac{\phi}{\text{eV}} \frac{A_e (\beta E_0)^2}{\phi^2} \exp(-6\phi^{3/2} \sqrt{m_e} / (\beta E_0)) \quad (11)$$

where ϕ is the work function of the wall material, A_e is the effective emitting area, E_0 is the electric field at the wall, and β is a phenomenological ‘‘field enhancement factor.’’ For pure niobium, $\phi \simeq 4.3$ eV [32], so $6\phi^{3/2} \sqrt{m_e} \simeq 60$ GV/m $\simeq 200$ T. As we will see in Sec. II B, for the cavities of interest to us, the peak electric field at the walls will be \lesssim the peak magnetic field, which is restricted to $H_{sh} \simeq 0.2$ T. Consequently, if $\beta \sim 1$, we would naively expect field emission to be negligible. However, experimentally, field emission is observed at significantly lower electric fields, down to ~ 10 MV/m $\simeq 0.03$ T [33]. This appears to arise from defects (especially foreign objects, such as metallic fragments) on the walls of the cavity, which can have β up to ~ 700 [33]. Too much field emission can lead to quality factor degradation, as EM energy is lost to electrons. Even if this is not a concern, we may be worried about much smaller levels of field emission, if the electrons deposit energy into the signal mode. We discuss this noise source in Sec. III D.

- (v) *Tuning*: to change the axion mass that we are sensitive to, we need to change the frequency splitting between the drive and signal modes, which entails changing the shape of the cavity. If the axion mass is small compared to the mode frequencies, then only a small fractional change in the frequency of each mode is needed to cover an $\mathcal{O}(1)$ axion mass range, simplifying the tuning problem. The usual method of tuning SRF cavities is simply via elastic deformation of the cavity walls, through an external forcing (such as a piezoelectric transducer, and/or a mechanical screw) [19,34]. To stay within the (low-temperature) elastic limit of niobium, the material strain should be $\lesssim \text{few} \times 10^{-3}$ [35]. For typical \sim GHz cavities, this usually translates to a tunable range of a few hundred kHz [19]—for higher cavity

modes, the absolute range may be greater. As well as static deformations, we also need to worry about vibrations, as discussed in Sec. III B.

These points only represent a simple sketch of the issues that an experiment would encounter, and much more detailed analysis would be needed for a real implementation. However, they give some idea of the properties of SRF cavities that are relevant to our setup.

B. Cavity geometry

In designing an SRF up-conversion experiment, the most obvious question is what cavity geometry to use. To detect low-frequency axions, we need a pair of almost degenerate modes to act as the drive and signal modes, and these should have large C_{01} . If we want to use low-lying cavity modes, this generally requires a tuning of the cavity shape. For example, if we take a cylindrical cavity, and consider e.g., the TE_{011} mode, the mode this is naturally degenerate with is TM_{111} (i.e., they are degenerate at all height-to-radius ratios). However, due to the m mismatch, these modes have $C_{01} = 0$. We can attain good overlap with e.g., the $\text{TE}_{011}/\text{TM}_{020}$ mode pair, but the mode frequencies are only degenerate at $d/R \simeq 0.79$, where d is the height of the cylinder and R is its radius [14]. For more symmetrical cavities, such as a sphere, the overlaps for degenerate modes are always zero.

If we are only concerned with thermal noise, the best possible sensitivity is determined by the signal power, and the bandwidths of the drive and signal modes (see Sec. III A). Since the signal power scales as B_0^2 , where B_0 is the magnetic field strength in the driven mode, the simplest way to increase the signal power is to increase the amplitude of the driven mode. However, the power dissipated increases $\propto B_0^2$ as well, and even if this is not a problem (or other issues such as field emission), at some point the maximum magnetic field at the cavity walls will increase past $H_{sh} \simeq 0.2$ T. This means that, for a given cavity geometry, there is a maximum achievable signal power.

Given some constraints on available volume and cooling power, we can ask how large a signal power can be obtained by optimising the cavity geometry. The volume constraint is necessary; since dissipation is a wall effect, whereas signal scales with volume (for fixed field amplitudes), if we increase the cavity dimensions by a factor α , while reducing the field amplitudes to keep the signal power constant, the dissipated power scales $\propto 1/\alpha$. Hence, by scaling up the cavity, we can always reduce the required cooling power for a given signal level.

Taking a cylindrical cavity geometry as our example, which height-to-radius ratio gives rise to a degenerate mode pair with the best (axion-mass-averaged) signal power per volume? If we hold $\max_{\partial V} B^2$ constant (where V denotes the volume of the cavity, and ∂V its boundary walls), then among the low-lying modes, the best choice is to drive

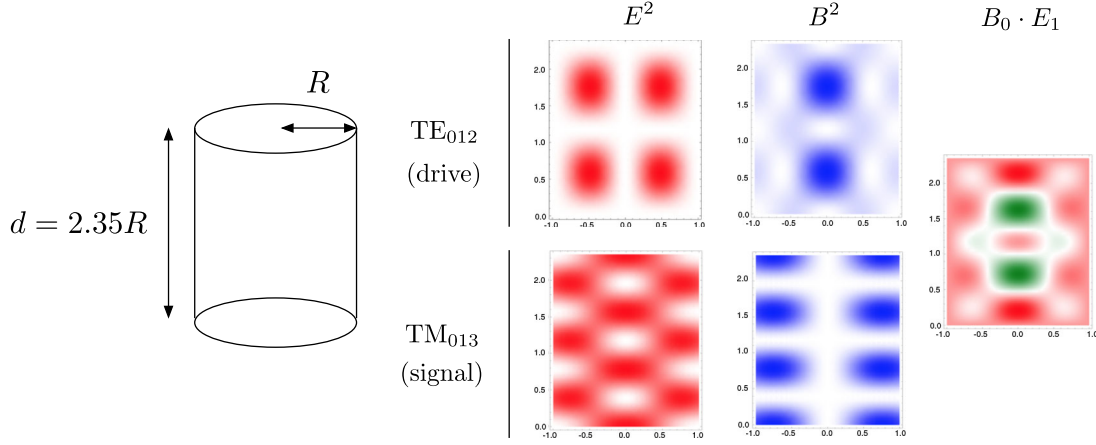


FIG. 2. Illustration of the optimum drive and signal modes for a cylindrical cavity, as discussed in Sec. II B. The $B_0 \cdot E_1$ panel shows the dot product of the drive mode's magnetic field with the signal mode's electric field, with green indicating a positive value, and red a negative value—the integrated value over the volume gives $C_{01} \simeq 0.19$, as defined in Sec. II.

TE_{012} , and pick up in TM_{013} , which are degenerate at $d/R \simeq 2.35$. This mode pair has $C_{01} = 0.19$, and maximizes $U_0 C_{01}/V$ (for example, driving TM_{013} and picking up in TM_{012} gives the same C_{01} , but the stored energy in the TM mode is smaller, if we restrict the magnetic field at the walls to be <0.2 T). The field profiles for these modes are illustrated in Fig. 2. In addition to giving high signal power, this mode pair also has attractive noise-rejection properties, as discussed in Sec. III. Consequently, we will use it as our nominal experimental setup for most of this paper. Some features of this mode are summarized in Table I, where its properties for $R = 20$ cm and $R = 50$ cm are given (at the degenerate d/R ratio). Properties for other sizes can be obtained by scaling.

1. General constraints

We can also consider more general cavity geometries. Both the cooling power and H_{sh} limitations are based on the magnetic fields of the drive mode at the cavity walls. It is not immediately obvious that these cannot be made small by a clever choice of cavity geometry. For example, the wall *electric* fields for the cylindrical TE_{012} mode are everywhere zero (with the consequence that field emission can be highly suppressed; see Sec. III D). However, as

shown in Appendix A, the wall fields can be related to the energy stored in the cavity via

$$U = \frac{1}{2} \left\langle \oint_{\partial V} dA (B^2 - E^2) (x \cdot n) \right\rangle \quad (12)$$

where x is the vector from some origin to the wall location, n is the outward-pointing normal to the wall, and the angle brackets denote time averaging. Since $\oint dAx \cdot n = 3V$, if we can choose an origin for which $x \cdot n \geq 0$ everywhere, then

$$U \leq \frac{3}{2} V \max_{\partial V} \langle B^2 - E^2 \rangle \leq \frac{3}{4} V \max_{\partial V} B^2 \quad (13)$$

for a harmonic oscillation, so the magnetic field energy inside the cavity can be bounded by the maximum magnetic field at the walls. Similarly, the power dissipated is

$$P_{\text{diss}} = R_s \left\langle \oint dAB^2 \right\rangle \quad (14)$$

Since, if $x \cdot n \geq 0$ everywhere,

TABLE I. Parameters for the nominal SRF cavity experiments referred to in Fig. 5. In each case, the maximum magnetic field at the cavity walls is taken to be 0.2 T, and the surface resistance of the cavity walls is taken to be $R_s = 5n\Omega$. The various quantities are defined in Sec. II. Note that the quality factors given are in the sense of dissipation, i.e., $P_{\text{diss}} = \omega U/Q$, rather than frequency stability. The latter will depend on how well vibrations can be controlled, as per Sec. III B 1. The B_{rms} column shows the RMS magnetic field, averaged over time and space.

	V	f	U_{max}	$C_{01} U_{\text{max}}$	P_{max}	Q_0	Q_1	B_{rms}
'60 L cylinder'	60 L	1.1 GHz	690 J	132 J	27 W	1.8×10^{11}	9.5×10^{10}	0.12 T
'900 L cylinder'	920 L	440 MHz	11 kJ	2 kJ	170 W	1.8×10^{11}	9.5×10^{10}	0.12 T
'8m toroid'	~ 2700 L	2.8 GHz	~ 210 kJ	~ 210 kJ	~ 650 W	$\sim 6 \times 10^{12}$	$\sim 6 \times 10^{12}$	~ 0.3 T




FIG. 3. Illustration of a toroidal corrugated waveguide, as discussed in Sec. II B 2. The radius of curvature R is taken to be much larger than the waveguide radius a .

$$U \leq \frac{1}{2} (\max x \cdot n) \left\langle \oint dA B^2 \right\rangle \quad (15)$$

we have

$$P_{\text{diss}} \geq \frac{2R_s U}{\max x \cdot n} \quad (16)$$

From Eq. (8), the signal power is bounded by the energy in the drive mode. Consequently, for “simple” cavities, for which there is an interior point from which all of the walls are visible, a given signal power implies a lower bound on $\max_{\partial V} B^2$, and on P_{diss} (in the latter case, assuming a given linear extent for the experiment, as per the scaling discussion above).

Comparing these limits to the properties of the $\text{TE}_{012}/\text{TM}_{013}$ mode pair, the maximum signal power per volume, for a given $\max_{\partial V} B^2$, is $\lesssim 10$ times larger, using the above limits. Similarly, for a cavity with the same maximum linear extent, the signal power is $\lesssim 12$ times larger, for a given dissipated power. For realistic geometries, the limits are probably significantly lower.

2. Nonconvex geometries

However, Eq. (12) also makes it fairly obvious how to get around these limitations. If $\oint dA |x \cdot n| \gg 3V$, then U can be large even if the wall fields are small. As an explicit example, we can consider a toroidal cavity (illustrated in Fig. 3), formed by bending a corrugated cylindrical waveguide of radius a around to meet itself, resulting in a toroid with overall radius R . If we assume quarter-wavelength-deep corrugations, then at frequencies for which $\omega a \gg 1$, the modes of a corrugated waveguide are dominantly transverse, and the wall fields are suppressed by $\sim (\omega a)^{-1}$ relative to the fields in the interior [36–38]. Taking an explicit example, the linearly polarized HE_{11} modes have transverse fields $|E_{\perp}|, |B_{\perp}| \sim J_0(k_0 r)$, where $k_0 a$ is the first zero of J_0 , and have

$$\left| \frac{B(r=a)}{B(r=0)} \right| \simeq \frac{1.25}{\omega a} \simeq \frac{0.2}{a/\lambda} \quad (17)$$

where λ is the free-space wavelength of the mode (and similarly for the electric fields) [36]. Consequently, by making ωa large, the interior fields can be made parametrically larger than the wall fields.

For a linear waveguide, the problem comes at the end caps—here, the wall fields are $\sim |B(r=0)|$. Bending the waveguide into a toroid eliminates these end caps. Of course, in order to preserve the smallness of the wall fields, the radius of curvature R must be large compared to a . Performing a naive perturbative calculation [39], in which we take both R/a and ω/a to be large, and treating the corrugated wall as a surface with uniform effective reactance [36], the correction to the wall fields from the waveguide’s curvature is approximately

$$\frac{\delta B(r=a)}{B(r=a)|_{R=\infty}} \simeq 0.5 \frac{a}{R} (a\omega)^2 \quad (18)$$

The constant factor in this estimate should be treated as an $\mathcal{O}(1)$ estimate, since we do not calculate the actual behavior in the corrugations. However, the parametric form should hold, and illustrates that as long as $R \gg a$, it is possible to make $a\omega$ large and obtain interior fields significantly larger than the wall fields.

Taking some illustrative numbers, if we assume a toroid with waveguide radius $a = 10$ cm, and take $R = 4$ m, then choosing as high as mode frequency as practical, $\omega \simeq 2\pi \times 3$ GHz, gives $a\omega \simeq 6$. The total energy stored in the HE_{11} mode, for wall fields $\lesssim 0.2$ T, is $\simeq 40$ kJ. For a $\text{TE}_{012}/\text{TM}_{013}$ cylinder of comparable volume, $U \simeq 10$ kJ.

To use the toroidal cavity for axion detection, we need almost degenerate modes with good overlap. Conceptually, this is very simple—we simply use the HE_{11} modes with orthogonal polarizations, offset by a quarter-wavelength along the toroid. This is in exact analogy to the optical up-conversion experiments proposed in [16–18], which convert linearly-polarized optical photons to the orthogonal polarization. It has the advantage of having, in the large- R limit, perfect overlap between the drive and signal modes. Consequently, the advantage in axion signal strength is even larger than the advantage in stored energy—using the same nominal parameters as above, the toroid’s signal strength would be ~ 20 times higher (a more detailed calculation would be required to find the proper finite- R behavior). This compares to the factor 10 limit derived above for simple cavities, showing that the nontrivial geometry is necessary for such improvements. If we instead used a linear corrugated waveguide, then the drive amplitude would be constrained by the end caps, and the signal strength would be \sim the same as a $\text{TE}_{012}/\text{TM}_{013}$ cylinder of the same volume.

The toroidal cavity serves as a particularly symmetrical, and thus easy-to-analyze, example of a cavity with large $\oint dA|x \cdot n|$. There are, of course, many other possibilities. For example, we could instead take a linear corrugated waveguide, and replace its end caps by large-area reflectors. In [40], it is claimed that this can reduce the peak magnetic field at the walls by a factor ~ 2 ; for a long waveguide, this would result in a signal power per volume ~ 4 times larger than for the TE₀₁₂/TM₀₁₃ cylinder pair. It is not immediately clear whether the enhancement can be made parametrically large, as for the toroid example.

These kinds of cavities may be significantly more complicated to fabricate than the simple cylindrical cavities discussed above. In addition, as we will discuss in the next section, they lack some of the noise-rejection properties of more symmetrical cavities, and would be more complicated to tune and control. Consequently, we will not attempt to analyse them in detail. They do, however, serve as an example of how larger signal powers could potentially be realized.

In Table I, we list some estimated parameters for a $R = 8$ m, $a = 13$ cm toroidal cavity used as an up-conversion experiment. This size is chosen make significant QCD axion sensitivity possible, at least theoretically (see section III A and figure 5). Even these large sizes are significantly smaller than the axion coherence length for the masses of interest, $l_a \sim 10^3 \nu_a^{-1} \sim 300 \text{ m} \frac{\text{GHz}}{\nu_a}$, so our approximation of the axion field as spatially constant will be valid.

III. BACKGROUNDS & SENSITIVITY

A. Thermal and amplifier noise

The signal of axion DM in our experiment would be excess power in the signal mode (as illustrated in Figure 1), and any noise present at these frequencies will make this harder to detect. As mentioned in section II A, the power dissipated at high drive fields means that cooling the cavity to sub-kelvin temperatures is impractical. Since $2\pi \text{GHz} \simeq 50$ mK, this means that the physical temperature is always significantly higher than the signal mode frequency, and thermal noise needs to be taken into account. In addition, the system we use to read out the signal—generally, a chain of microwave amplifiers—will introduce its own noise.

In appendix B, we review the theory of signal detection for a high- Q target mode, assuming readout via an amplifier isolated behind a circulator (as for most microwave systems). If the noise associated with the amplifier system corresponds to a smaller temperature than the physical temperature of the cavity, then it is favorable to ‘overcouple’ the signal mode to the output port (i.e., have it lose more power to the output port than to environmental dissipation) [41]. This reduces the loaded quality factor of the target mode, which is naively bad, but also dilutes the

thermal noise reaching the amplifier; the latter effect turns out to dominate.

The high quality factors attainable in SRF cavities mean that, even given this overcoupling, the bandwidth of the target mode will be $\ll \nu_a$, for axion masses of interest. This means that, to cover an $\mathcal{O}(1)$ range in axion masses, we need to operate the experiment in multiple different configurations, with different frequency splittings between the drive and target modes (as discussed briefly in Sec. II A). As demonstrated in Appendix B, in the limit of long integration times, any sufficiently dense, and roughly equal, spacing of these frequency splittings will give approximately the same expected SNR,

$$\text{SNR} \simeq \sqrt{0.5 \frac{(P_1/Q_1)^2 t_{\text{tot}} Q_a Q_1 \omega_1}{T_0 T_n m_a \Delta m_a}} \quad (19)$$

Here, T_0 is the physical temperature of the cavity walls (assumed to be $\gg \omega_1$), T_n is related to the noise of the amplifier system (see Appendix B), Δm_a is the range of axion masses we want to cover (assumed to be $\lesssim \mathcal{O}(1)$), $Q_a \simeq 10^6$ is the inverse fractional bandwidth of the axion signal, t_{tot} is the total integration time for all configurations, Q_1 is the unloaded quality factor of the target mode, ω_1 is the frequency of the target mode (assumed to change by a small fractional amount between configurations), and P_1 is the power absorbed by the (unloaded) target mode when on-resonance with a monochromatic signal [see Eq. (6)]. Since $P_1 \propto g^2$, the smallest coupling we have sensitivity to scales as $g_{\text{sens}} \propto t_{\text{tot}}^{-1/4} Q_1^{-1/4} T_0^{1/4} T_n^{1/4}$, as expected [41].

From [9], we know that for fixed t_{tot} , this improvement with increasing Q_1 must saturate at some point. Fairly obviously, this will happen when the time spent in each configuration is not long enough to resolve the loaded bandwidth of the target mode. If, as is the case in most of our parameter space, $\Delta \omega_a \simeq m_a/Q_a$ is such that $\Delta \omega_a \gg \omega_1/Q_l$, where Q_l is the loaded quality factor of the target mode, then the minimum spacing of frequency splittings is $\sim \Delta \omega_a$ (otherwise some axion masses will fall into the gaps). This means that the time spent in each configuration is $t_1 \simeq t_{\text{tot}} \frac{\Delta \omega_a}{\Delta m_a} \simeq \frac{t_{\text{tot}} m_a}{Q_a \Delta m_a}$. If $t_1 \lesssim Q_l/\omega_1$, then the mode cannot fully ring up in the time available, and it would be more favorable to reduce Q_l by overcoupling further. In this regime, the best attainable SNR is

$$\text{SNR} \lesssim 0.2 \frac{\bar{W}}{T_n} \quad (20)$$

where $\bar{W} \equiv \bar{P} t_{\text{tot}}$ is the expected power absorbed from the axion signal.

In Fig. 4, the dot-dashed line shows the thermal-noise-limited sensitivity for the nominal 60 L cavity discussed above, given a total integration time of one year to cover an e -fold in axion mass range. It assumes that the physical

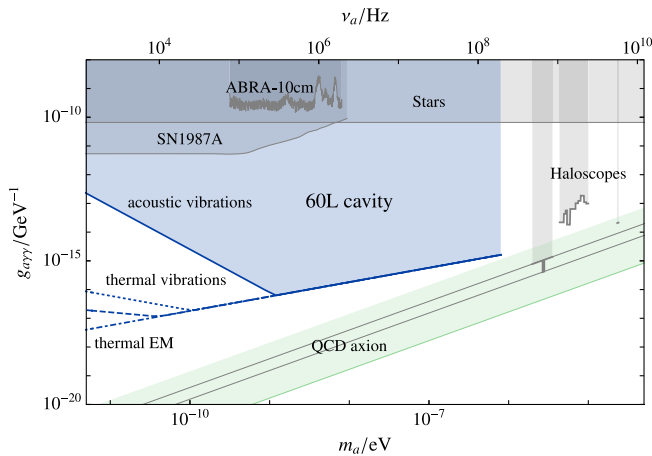


FIG. 4. Sensitivity projection for an up-conversion experiment using a 60 litre (20 cm radius) cylindrical cavity, where the TE_{012} mode is driven, and signals are picked up in the TM_{013} mode. We assume that the maximum magnetic field for the drive mode at the cavity wall is $H_{sh} = 0.2$ T, and an integration time of one year per e-fold in axion mass range. The sensitivity threshold is set at an expected SNR value of 3. The “thermal EM” line shows the sensitivity limit in the presence of thermal EM noise, at the assumed physical temperature of 1.4 K (see Sec. III A). The “thermal vibrations” line is an estimate of the effect of thermal vibrations of the cavity walls, which up-convert power from the signal mode to the drive mode (Sec. III B). The “acoustic vibrations” line also incorporates extra acoustic noise from external sources. The dashed line between them is an estimate of the effects of time-varying drive signal leakage (Sec. III C). Taking all of these noise sources into account, the estimated sensitivity reach of the experiment is given by the blue shaded region. It should be noted that, while the “thermal EM” limit is set by basic physical parameters, the vibrational limits depend on guesses about the properties of the cavity system, and could be very different in a real experiment. Also, while we have extended the projected reach to axion frequencies ~ 100 MHz, scanning an order one axion mass range at these frequencies would be difficult (see Sec. II A). In Fig. 5, this reach is compared to other proposed experiments. The gray shaded regions correspond to the parameter space ruled out by observations of horizontal branch stars [43–45], SN1987A [46], existing cavity haloscope experiments [5,47–51], and the ABRACADABRA-10 cm experiment [52]. The green diagonal band corresponds to the “natural” range of $g_{a\gamma\gamma}$ values at each QCD axion mass—if we write $g_{a\gamma\gamma} = \frac{g_{EM}}{2\pi f_a} (\frac{E}{N} - 1.92)$ [2], then the upper edge of the band is at $E/N = 5$ [3], and the lower edge at $E/N = 2$ [2]. The gray diagonal lines indicate the Kim-Shifman-Vainshtein-Zakharov (KSVZ) (upper, $E/N = 0$) and Dine-Fischler-Srednicki-Zhitnitsky (DFSZ) (lower, $E/N = 8/3$) models.

temperature of the cavity is $T_0 = 1.4$ K, and the amplifier system is quantum-limited, so $T_n = \omega_1$. Near-quantum-limited amplifiers have been demonstrated at microwave frequencies, and incorporated in the ADMX [5,6] and HAYSTAC [42] axion DM experiments. The loaded quality factor, for the optimal overcoupling level, is $Q_l \simeq 2 \times 10^9$, so it takes ~ 2 s to resolve the signal mode bandwidth, and a

total integration time of $\gtrsim 2 \times 10^6$ s $\simeq 24$ days to be in the regime of Eq. (19). Since the coupling sensitivity scales like $g_{sens} \propto m_a^{1/2}$, but $g \propto m_a$ for the QCD axion, lower-mass QCD axions are still harder to detect. As the figure shows, an experiment with these parameters could not attain QCD axion sensitivity over an $\mathcal{O}(1)$ axion mass range, even at higher masses (without violating our assumptions, e.g., by injecting a highly nonclassical state into the target mode). Moreover, though we extend the sensitivity projection up to $\nu_a \sim 100$ MHz, tuning a cavity over an $\mathcal{O}(1)$ range of frequency splittings would be difficult at such high masses, as discussed in Sec. II A.

The simplest way to obtain higher sensitivities would be to use larger cavities. Figure 5 shows the thermal-noise-limited sensitivities for two such examples. The first is a simple scaling-up of the TE_{012}/TM_{013} configuration, using a cylinder of radius 50 cm (giving a volume of ~ 900 L). Its properties are summarized in Table I—compared a smaller cylinder, the signal power scales with the volume. Even an experiment of this size could not reach KSVZ axion sensitivity over an $\mathcal{O}(1)$ axion mass range, at reasonable axion masses. At this size, the mode frequency is ~ 440 MHz, which is at the lower end of the frequency range generally used in SRF cavities.

To illustrate the kind of parameters that would be required for significant DFSZ axion sensitivity, Fig. 5 also shows the approximate thermally-limited sensitivity for a corrugated toroidal cavity, as introduced in Sec. II B 2. This is taken to have waveguide radius $a = 13$ cm, and $R = 8$ m; the resulting properties are summarized in Table I. The naive estimate of the signal mode quality factor is high enough to put us in the regime of Eq. (20), even for a total integration time of a year.

As emphasized in Sec. II B 2, this kind of projection should not be taken as a concrete experimental proposal—for that, one would need to understand the control issues, noise problems etc. associated with these cavity designs. Instead, it illustrates what would be necessary, in principle, to probe very small couplings.

B. Vibrations

While thermal noise must be taken into account in all axion detection experiments, up-conversion experiments have the additional challenge of a very large-amplitude background EM oscillation. If there are any environmental oscillations at the axion frequency, and any nonlinear processes which couple these to the drive and signal modes, then power can be delivered to the signal mode, representing an additional noise source.

The most obvious such process is mechanical vibrations of the cavity walls. If we write $\delta x(t, x)$ as the displacement of the cavity wall from initial position x , and n as the outward pointing normal to the initial wall position, then to linear order in δx , the interaction Hamiltonian of the wall displacement with the EM fields is

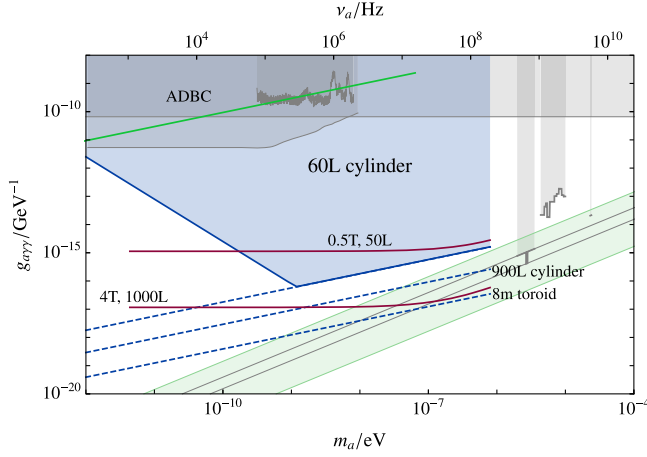


FIG. 5. Comparison of sensitivity projections for different kinds of low-frequency axion DM detection experiments. The “60 L cylinder” region corresponds to the 20 cm cavity up-conversion experiment described in the text, with parameters given in Table I (we assume an integration time of 1 year per e-fold in axion mass range). At high frequencies ($\nu_a \gtrsim 200$ kHz), its sensitivity is limited by thermal EM noise from the cavity walls, while at lower frequencies it is limited by (a rough estimate of) up-converted acoustic noise from external sources (see Sec. III B and Fig. 4). We also display the thermal-noise-limited sensitivities for the larger up-conversion experiments listed in Table I. The “900 L cylinder” line corresponds to a scaled-up $\text{TE}_{012} \rightarrow \text{TM}_{013}$ experiment, with 50 cm radius. The “8 m toroid” line corresponds to a corrugated toroidal cavity, as described in Sec. II B 2; we assume an overall radius of 8 m, a waveguide radius of 13 cm, and a frequency of 2.8 GHz, resulting in the parameters listed in Table I. We leave an analysis of the less fundamental noise sources for these larger up-conversion experiments to future work. The red lines are thermal noise limited sensitivity projections for representative static-magnetic-field experiments, modeled on the Dark Matter Radio proposal [13,53]. The “0.5 T, 50 L” line corresponds to a 50 liter experiment, with a background magnetic field of 0.5T, and a physical temperature of 10 mK—the resonator quality factor is taken to be 10^6 , and the geometric overlap factor to be 0.2 [13] (in the sense that $P_{\text{sig}} \simeq 0.2^2 g^2 B_0^2 V Q_m^2 (mL)^2$, where L is the shielding length scale [9]). The “4 T, 1000 L” line corresponds to the same assumptions, but with a cubic metre volume, and a background magnetic field of 4 T. The quasistatic $\sim(mL)^2$ suppression used is only a dimensional estimate; the actual value would depend on the geometry of the experiment. We also show an example of a proposed optical-frequency up-conversion experiment, the “ADBC” proposal from [18]. This assumes a 2 meter Fabry-Perot cavity, with a circulating optical power of 10 kW. The existing haloscope limits, QCD axion band, and astrophysical constraints are as per Fig. 4.

$$H_{\text{int}} = \frac{1}{2} \int dA (B^2 - E^2) n \cdot \delta x \quad (21)$$

The resulting interaction between the drive and signal modes is

$$H_{\text{int}} = \int dA (B_0 \cdot B_1 - E_0 \cdot E_1) n \cdot \delta x \quad (22)$$

This interaction actually represents the signal mechanism for proposed SRF gravitational-wave detectors, such as MAGO [54–56]—the effect of a gravitational wave can be treated as a deformation of the cavity walls, which up-converts photons from the drive mode to the signal mode.

We can see from Eq. (22) that, if the drive and signal modes of the unperturbed cavity are orthogonal at the walls, then to linear order in δx , wall vibrations do not couple them. This is one reason why the $\text{TE}_{012}/\text{TM}_{013}$ mode pair introduced in Sec. II B is attractive—in an ideal cylinder, these modes have $E_0 \cdot E_1 = 0$, $B_0 \cdot B_1 = 0$ throughout the whole cavity. However, it will not be possible to fabricate the cavity shape perfectly, and some deformations will result in nonorthogonal fields at the walls. For example, if the cavity cross-section is elliptical rather than cylindrical, then the deformed TE_{012} mode has a small electric field at the cylinder’s walls,

$$E_r \simeq 40 \text{ kV m}^{-1} \frac{f}{10^{-3}} \sin(2\phi) \sin\left(\frac{2\pi z}{d}\right) \quad (23)$$

where f is the flattening of the ellipse ($f = (a - b)/a$, where $a \geq b \geq 0$ are the axis lengths), and the normalization is set by taking the maximum magnetic field at the walls to be = 0.2 T. Since the TM_{013} mode has $E_r \propto \sin(3\pi z/d)$ at the walls, the effect of a cavity vibration is controlled by

$$\int dA \sin(2\phi) \sin\left(\frac{2\pi z}{d}\right) \sin\left(\frac{3\pi z}{d}\right) n \cdot \delta x \equiv 0.09 C A_w x(t) \quad (24)$$

where A_w is the area of the cavity walls, and we have normalized so that a vibration with $\delta x = \sin(2\phi) \sin\left(\frac{2\pi z}{d}\right) \sin\left(\frac{3\pi z}{d}\right) x(t)$ has $C = 1$. Then, for monochromatic oscillations, the power transferred to the signal mode is

$$P \simeq 7 \text{ W } C^2 Q_1 \left(\frac{f}{10^{-3}}\right)^2 \left(\frac{x}{\text{mm}}\right)^2 \quad (25)$$

where we have taken the parameters of the 20 cm cylindrical cavity listed in Table I. Consequently, the displacement noise in the relevant bandwidth must be very small, in order not to overwhelm the axion signal power.

The example of an elliptical cavity illustrates the general feature that, for a cavity deformation of fractional size $\sim f$, we expect the wall fields to be perturbed by $\sim f$ [57,58], and in general to be nonorthogonal. It may be possible to deform the cavity post-fabrication to alleviate such issues, but we do not attempt to consider such possibilities here.

Conversely, for other cavity geometries, such as a corrugated waveguide, the wall fields are nonorthogonal even for an unperturbed cavity.

The other information we need to determine the vibration-induced noise spectrum is the vibration spectrum of the cavity's walls. At high enough frequencies, this will probably be dominated by thermal vibrations. If ultrasonic waves from outside the cavity are sufficiently attenuated, then the amplitude will be set by the cavity's physical temperature. The ultrasonic attenuation length scale in liquid helium is \sim cm for frequencies \gtrsim 10 MHz [59,60], though it may be significantly smaller for metals [59,61]. Investigation would be required to determine the actual properties of a cavity setup.

At lower frequencies, external acoustic noise will not be strongly attenuated, and in addition, above-thermal noise sources (such as machinery) will be present. A nominal spectrum for this external displacement noise, in a reasonably "quiet" setting, is

$$\sqrt{S_{xx}(\nu)} \simeq 10^{-7} \frac{\text{cm}}{\sqrt{\text{Hz}}} \left(\frac{10 \text{ Hz}}{\nu} \right)^2 \quad (26)$$

for $\nu \gtrsim 10 \text{ Hz}$ [62]. The measured acoustic noise in the vicinity of the MAGO prototype (Fig. 14 of [56]) also shows a $S_{xx} \sim 1/\nu^4$ spectral density for $\nu \gtrsim 3 \text{ kHz}$, with displacement amplitude a factor ~ 10 higher than Eq. (26). At lower frequencies, especially with the helium cooling system in operation, there is a complicated, spiky spectrum with significantly larger amplitude.

We can compare these acoustic spectra to the spectrum from thermal oscillations by noting that, if we can treat the material as a bulk medium of large extent, with a linear acoustic dispersion relation, then the PSD of surface displacements is given by [63]

$$S_{xx} \simeq \frac{2pT}{9\pi\rho v_s^3} \quad (27)$$

where ρ is the density of the material, T is the temperature, v_s is the sound speed, and p is a dimensionless constant governing the interaction of the surface with bulk phonons (e.g., for aluminium, $p \simeq 2.3$ [63]). Plugging in representative parameters for niobium, we obtain

$$\sqrt{S_{xx}^{\text{therm}}} \simeq 10^{-17} \frac{\text{cm}}{\sqrt{\text{Hz}}} \left(\frac{T}{1.4 \text{ K}} \right)^{1/2} \quad (28)$$

suggesting that, very roughly, thermal noise will dominate for $\nu \gtrsim \text{MHz}$.

Properly calculating the effects of vibrations would require assumptions about the imperfections of the cavity shape, and modeling the mechanical modes of the cavity-cryostat system. Here, we only attempt to make a rough estimate of these effects. To calculate the vibration-induced

noise power as a function of frequency, we need the PSD S_{xx}^v for oscillations with profile set by $B_0 \cdot B_1 - E_0 \cdot E_1$ at the cavity walls. We will assume that this profile is reasonably smooth, corresponding to slowly-varying deviations from a cylindrical shape (e.g., the ellipse example above).

Decomposing the cavity vibrations into weakly-damped modes, there will be contributions to S_{xx}^v from modes with different resonant frequencies. At ω much higher than the resonant frequency of low-order mechanical modes, the vibrational modes with similar frequencies will have small overlap with the (assumed smooth) profile, while the low-order vibrational modes will have much smaller frequencies. For thermal vibrations, the contribution from the latter will be $S_{xx} \sim \frac{1}{Q_m} \frac{T}{M\omega^3}$, where M is the mass of the cavity walls, and Q_m is the mechanical quality factor of the low-order modes. The modes with resonant frequency $\sim \omega$ will contribute $S_{xx} \sim (\lambda/L)^c \frac{T}{M\omega^3}$, where $\lambda \sim v_s/\omega$ is the wavelength of the modes, and $(\lambda/L)^c$ corresponds to the suppressed spatial overlap with the profile (with higher c corresponding to smoother profiles). At frequencies low enough that individual modes stop overlapping in frequency, the contributions to S_{xx} will be spikier. For ω around that of the low-lying mechanical modes, it will have maximum value $S_{xx} \sim Q_m \frac{T}{M\omega^3}$.

For acoustic, rather than thermal, noise, the same general considerations will apply, but the effective T in the above equations will be frequency dependent. In Fig. 4, we display an example of the vibration-noise-limited sensitivity, using the estimates discussed above. Taking acoustic noise with spectrum as per Eq. (26), we find that vibration-induced noise dominates thermal EM noise for $\nu_a \lesssim 300 \text{ kHz}$ (we assume that $Q_m \simeq 10^3$ and $c = 2$, in the notation of the previous paragraph, and assume 3 mm thick cavity walls, giving a cavity mass of $\sim 20 \text{ kg}$). The thermal vibrations, assuming $T = 1.4 \text{ K}$, are subdominant. We emphasize that all of these estimates are best viewed as guesses, and a much more careful treatment would be necessary for a realistic experiment.

1. Frequency variability

So far, we have considered how vibrations can give rise to a coupling between the drive and signal modes. However, vibrations will also change the energy of each of these modes themselves. Typically, this "detuning" will (before any compensation) be $\mathcal{O}(10 \text{ Hz})$, corresponding to $\sim \text{nm}$ -scale displacements of the cavity walls [64–66]. Most of this variation will come from low-frequency ($\lesssim \mathcal{O}(10 \text{ Hz})$) vibrations.

If these vibrations cause the frequency splitting between the drive and signal modes to change over time, by more than their bandwidths, then we are effectively changing the up-conversion tuning over time. This change in "scan

strategy” can affect the thermal noise limited sensitivity, as discussed in Sec. III A.

For SRF accelerator cavities, large frequency detunings can increase the power needed to drive the cavity [67], as well as moving the accelerating voltage out of phase with the electron bunches [66]. Consequently, feedback systems are general employed which measure the detuning of the cavity, and apply a mechanical deformation (via a piezo transducer) to correct this. Using such systems, frequency stabilization to \lesssim Hz has been demonstrated [66]. In our case, measurement of the changing mode frequencies is very important, since any *unknown* variation of the frequencies over time will reduce sensitivity. Further investigation of how well such a measurement system (and feedback control system) could perform in our setup would be required.

In principle, it may also be possible to measure the higher-frequency mechanical vibrations of the cavity discussed above, and use this to compensate for, or subtract out, the induced noise in the signal mode. We do not attempt to model the feasibility of such measurements here.

C. Drive leakage

In an ideal experiment, the cavity input power would couple only to the drive mode, and the detector would be coupled only to the signal mode. For mode pairs such as our TE₀₁₂/TM₀₁₃ example, it is easy to see how this could be accomplished in principle. Input and output is usually accomplished through small slots in the cavity wall, which connect to waveguides. For a slot much smaller than the wavelength of a cavity mode, the coupling is approximately $\propto B_{\text{slot}} \cdot B_{\text{cav}}$, where B_{cav} is the mode’s magnetic field at the cavity wall [68]. Since the signal and drive modes are everywhere orthogonal, a correctly polarized input/output port should, to a good approximation, couple only to one of them. The different spatial profiles of the modes (see Fig. 2) provide an extra discriminator—for example, we could place an input port at a wall location where the signal mode’s magnetic field is very small, and an output port where the drive mode is small.

However, as was the case for the vibrational couplings discussed above, nonperfect fabrication of the cavity will spoil these assumptions. For example, if the output slot was slightly misaligned, it would have a small coupling to the drive mode. Generally, for a Fourier-domain signal $S_I(\omega)$ at the input port, and assuming that the drive and signal modes are the only ones at nearby frequencies, the output signal will be

$$S_O(\omega) \simeq C_{oD}G_D(\omega)C_{Di} + C_{oS}G_S(\omega)C_{Si}S_I(\omega). \quad (29)$$

Here, C_{oD} represents the coupling between the drive mode and the output port, and $G_D(\omega)$ is the response function of the drive mode, etc. For a well-constructed cavity, C_{oD} and C_{Si} will be small, but still nonzero. Since the signal and

drive modes have very narrow-bandwidth response functions, we expect $S_O(\omega)$ to be peaked around ω_0 and ω_1 .

Naively, the output signal near ω_0 is not an issue—we can just Fourier transform the output data, and ignore it. However, the nonideal nature of the output electronics may introduce problems. Most directly, if the amplitude of the $\sim\omega_0$ noise, relative to the $\sim\omega_1$ signal, is too large, then the dynamic range of the amplifier may be exceeded. One way to address this is to place a frequency filter on the output port, to reject frequencies too far from ω_1 . However, if this is not enough (which is especially likely at smaller m_a), it may be necessary to reduce C_{oD} through some feedback control mechanism.

Using the nominal experimental parameters from above, the energy stored in the drive mode of the 60 L cavity is ~ 700 J. If the output port were critically coupled to the drive mode, this would result in ~ 30 W output power. At the thermal-noise-limited sensitivity (Sec. III A), the axion-sourced signal power is $P_{\text{sig}} \simeq 2 \times 10^{-23}$ W for $\nu_a \sim$ MHz (assuming an integration time of 1 year per e-fold in axion mass range). If we assume an amplifier with a dynamic range of 50 dB, then to avoid swamping this signal, the output power needs to be a factor of $\sim 10^{-20}$ lower than the 30 W level, corresponding to a $\sim 10^{-10}$ amplitude suppression. These numbers make it clear why leakage suppression might be challenging. Furthermore, in addition to leakage through the cavity, we would also need to worry about leakage through the laboratory environment. An experimental rule of thumb is apparently [69] that suppressions of more than $\sim 10^{-19}$ in power are difficult to achieve, for electronics in the same laboratory space.

These kinds of issues were encountered by the MAGO SRF gravitational-wave detector mentioned in the previous section. Their setup used two identical cavities with a small coupling between them, giving rise to almost degenerate symmetric and antisymmetric modes. To drive the symmetric mode, a magic-tee was used to split the drive signal into in-phase components, while to sample the antisymmetric mode, another magic-tee was used to take the difference of outputs from each cavity. Since the drive and signal modes had the same profiles in each cavity, this relative phase was the only way of distinguishing between them. With this setup, they achieved a power suppression of ~ 48 dB for the output, corresponding to an amplitude suppression of $\sim 1/250$ [54]. The output power was dominated by the $\sim\omega_0$ peak.

To improve this, they implemented a feedback system, using a variable phase shifter to control the phase of one of the input ports (and another to control the phase of an output port). This phase shifter was controlled by a feedback loop, designed to suppress the output power. Using this system, they achieved an output power suppression of ~ 140 dB [54] (which would seem to indicate that the amplitudes at the two output lines were matched to one part in $\sim 10^7$ by the geometry of the setup, unless they also

applied feedback control the amplitudes). One could imagine other ways of implementing feedback schemes—for example, combining the (attenuated) drive signal with the output directly—but the MAGO scheme provides a concrete example of the kind of control system that may be required. In our case, the extra geometric rejection provided by the profiles of the drive and signal modes, along with a high- Q bandpass filter on the output, make it seem plausible that sufficiently high rejections could be obtained.

If the drive signal has narrow bandwidth compared to the frequency splitting between the drive and signal modes, then leakage at frequencies close to ω_1 will be suppressed. However, it may interfere with the axion signal we hope to detect, so more care is required in dealing with it. Since we can measure the input drive signal, it does not necessarily represent irreducible noise—if we know the transfer function from the input to the output, e.g., via the signal mode of the cavity, then we can simply subtract it out. This can either be done in software, or via analogue mixing (though in the former case, dynamic range issues may still arise).

The more worrying case is if the transfer function varies over time, in an *a priori* unknown way (for example, MAGO observed that temperature fluctuations in their input and output cables affected the phase shift experienced by the signal [56]). As a quantitative example, suppose that the time delay $\delta(t)$ along the signal path is time-varying, $\delta(t) = \delta \cos \omega_\delta t$. For $\omega_0 \delta \ll 1$, we can use the Jacobi-Anger expansion [70],

$$e^{i\omega_0(t+\delta \cos \omega_\delta t)} = e^{i\omega_0 t} (J_0(\omega_0 \delta) + 2iJ_1(\omega_0 \delta) \cos \omega_\delta t + \dots) \quad (30)$$

So, given leakage at frequencies $\sim \omega_0$, the amplitude of the noise component at frequencies close to ω_1 is $\sim (\omega_0 \delta)$ times the constant leakage amplitude. This gives a noise power, for our nominal setup, of

$$P_{\text{noise}} \sim 3 \times 10^{-21} \text{W} \left(\frac{|C_{Si}|}{10^{-3}} \right)^2 \left(\frac{\text{MHz}}{\nu_a} \right)^2 (\omega_0 \delta)^2 \quad (31)$$

where we take $|C_{Si}|$ to be the amplitude suppression, compared to critical coupling, of the output ports to the drive mode. Translating this into an equivalent noise temperature,

$$T_{\text{eff}} \simeq 1 \text{ K} \left(\frac{|C_{Si}|}{10^{-3}} \right)^2 \left(\frac{50 \text{ kHz}}{\nu_a} \right)^3 Q_\delta (\omega_0 \delta)^2 \quad (32)$$

for a δ PSD of fractional bandwidth $1/Q_\delta$, and total power δ^2 . Consequently, unless the temporal variation in transfer characteristics is substantial, this noise contribution is likely to be smaller than e.g., that of the acoustic noise considered in the previous section. In Fig. 4, we show the effect on sensitivity for $Q_\delta \simeq 1$, $\omega_0 \delta \simeq 0.1$, illustrating that,

for these parameters, it is fairly similar to our estimate for thermal vibrations.

Of course, as well as this in-principle issue, there may be technical issues in suppressing the $\sim \omega_1$ leakage, especially for low-dynamic-range amplifiers. For the purposes of our sensitivity estimates, we will assume that these are surmountable at high axion masses, and subdominant to vibrational noise at lower ones.

If there are other modes close-to-degenerate with the drive and signal modes, then these may affect signal leakage. For example, in a perfect cylindrical cavity, the TE_{012} mode is exactly degenerate with the TM_{112} mode. While such degeneracies will be lifted by the inevitable cavity shape imperfections (e.g., the $\sim 10^{-3}$ fractional deformations we were assuming would generically give $\sim \text{MHz}$ splittings), if they are still a problem, then it may be simplest to give the cavity an intentional, slight deformation, as was done for the cylindrical MAGO prototype [54].

D. Free charges

As well as wall effects, such as the vibrational backgrounds considered in Sec. III B, EM noise can also arise through currents inside the volume, i.e., from the movement of free charges inside the cavity. These charges may originate from the cavity walls via field emission, or may come from outside the cavity.

For initially low-energy charged particles, such as those arising from field emission, the oscillating drive field inside the cavity will have a significant effect on their motion, and it is necessary to solve for their trajectories inside the cavity. However, for particles with sufficiently high initial kinetic energy, the effect of the drive field on the trajectory is unimportant, and the particle effectively travels in straight line. This is the case for e.g., cosmic ray muons, which provide a simple test case we can work out.

We are interested in the effect of the charged particle on the signal mode. Working in a gauge where $A_0 = 0$, we can write the interaction Hamiltonian as

$$H_{\text{int}} = \int dV A \cdot J = qv(t) \cdot \hat{A}(x(t)) \quad (33)$$

If we write the vector potential for the signal mode as $\vec{A}(t, x) = A_s(t) \vec{a}(x)$, then the dynamics of A_s are analogous to those of a harmonic oscillator with “mass” $= \int dV a^2 \equiv V_a$ [9]. Writing $H_{\text{int}} = (qv(t) \cdot a(x(t))) A_s(t) \equiv j(t) A_s(t)$, the linear response function governing the response of A_s to the forcing j is $\tilde{\chi}(\omega) \simeq \frac{V_a^{-1}}{\omega_1^2 - \omega^2 + i\omega\omega_1/Q_1}$. The expected energy delivered to the signal mode is

$$\langle W \rangle = \frac{1}{2\pi} \int_{-\infty}^{\infty} d\omega \omega |j(\omega)|^2 \text{Im} \tilde{\chi}(\omega) \quad (34)$$

where the averaging is taken over different phases of the signal mode oscillation. For a high quality factor mode, $\text{Im}\tilde{\chi}$ will be very narrowly peaked compared to \tilde{j} , so

$$\langle W \rangle \simeq \frac{1}{2V_a} |\tilde{j}(\omega_1)|^2 \quad (35)$$

Consequently, for a low-lying mode of the cavity, a single relativistic transit of a charged particle will deposit, on average, $\sim Cq^2$ photons into the signal mode, where q is the charge of the particle, and C is a dimensionless geometric overlap factor. It should be noted that, in the presence of an oscillation of definite phase in the signal mode (for example, due to thermal fluctuations), there will also be an $\mathcal{O}(q)$ contribution to the energy absorbed. However, the energy uncertainty of a coherent state is also larger, and the relative detectability of the perturbation is still set by $\langle W \rangle/T$.

The flux of cosmic ray muons at sea level is $\sim 1/((10 \text{ cm})^2 \text{ s})$ [71]. They have an average energy of $\sim 4 \text{ GeV}$, so are certainly high-energy for the purposes of our calculation. Considering our nominal 20 cm-radius cylindrical cavity, there are $\lesssim 20$ muons passing through the cavity per second, corresponding to a delivered power of $\sim \bar{C} \times 10^{-24} \text{ W}$ to the signal mode, where \bar{C} is the average value of C over different trajectories. A large value of C , e.g., for a trajectory along the central z -axis, is $C \simeq 0.17$. In comparison, the thermal-noise-limited signal power, assuming an integration time per e-fold in axion mass range of one year, is $P_{\text{sig}} \simeq 2 \times 10^{-23} \text{ W} \frac{\nu_a}{\text{MHz}}$. Since we do not expect the flux of other high-energy charged particles to be significantly larger than the muon flux, we can conclude that cosmic rays will not be a significant background at higher ν_a . These calculations would apply to standard cavity haloscopes as well, with similar results (see also [72]).

For cavities with small enough electric fields at the walls, field emission should be negligible. Taking the previous example of a slightly elliptical cavity, if the TE_{012} drive mode has maximum magnetic field of 0.2 T at the walls, then from equation (23), the maximum wall electric field is $E_{\text{max}} \simeq 40 \text{ kV/m} \frac{f}{10^{-3}}$, where f is the flattening of the ellipse. Since electric fields \gtrsim few MV/m are required for field emission, even fairly loose mechanical tolerances should be enough to suppress it. This is contrast to the modes used in SRF cavities for particle acceleration, such as the TM_{010} mode—to obtain an accelerating voltage down the axis of the cavity, these have large electric fields at the walls.

Other cavity geometries, such as a corrugated waveguide, do not necessarily have such suitable drive modes with small electric fields at the walls. In these cases, proper simulations would need to be carried out to determine whether the resulting electron trajectories transfer too much energy to the signal mode. If the rate of field emission is

high enough, then multiple electrons could contribute coherently to such energy transfer, potentially worsening the problem.

IV. SENSITIVITY COMPARISONS

As mentioned in the Introduction, there are a number of different approaches to low-frequency axion DM detection through the $F\tilde{F}$ coupling. Whether SRF up-conversion experiments are worth pursuing depends on their plausible sensitivity, relative to these alternatives.

Static background field experiments, such as the ABRACADABRA [12] and DM Radio [13] proposals, offer the best theoretical sensitivity at higher axion masses. In Fig. 5, we show thermal-noise-limited sensitivity projections for nominal static field experiments, modeled on the DM Radio proposals³ [53], and compare them to the sensitivity projections for the nominal SRF experiments discussed in the previous section, with parameters summarized in Table I. Compared to the nominal static field parameters, the corresponding SRF experiments have significantly smaller RMS magnetic fields, but gain by avoiding the quasistatic geometric suppression, and by having higher target mode quality factor. Correspondingly, they have different scaling of sensitivity with axion mass.

Figure 5 illustrates that, for significant QCD axion sensitivity in a SRF up-conversion experiment, either very large (many cubic meters) conventional cavities, or few cubic meter advanced cavities, would be required (ignoring potential quantum enhancements). This is true even if we only take thermal noise into account—further study would be required to determine whether other noise sources could be mitigated. As discussed in the previous section, such mitigations will likely work best at higher axion masses; more realistic sensitivity projections for the larger cavities would likely have shapes similar to the small-scale cylinder projection of Fig. 4.

The projections of [73], which proposes similar SRF up-conversion experiments, are presented somewhat more optimistically, particularly regarding QCD axion sensitivity. This seems to be mostly due to two factors. The first is that they take their sensitivity threshold as $\text{SNR} = 1$, rather than $\text{SNR} = 3$ as we use. The second is that they present their projections for a nominal cavity with peak spatial-RMS magnetic field of 0.2 T, and an overlap factor $C_{01} = 1$. This gives a signal power per volume ~ 6 times better than our $\text{TE}_{012}/\text{TM}_{013}$ cylindrical mode pair (which has \sim the same signal power as the corrugated cylindrical cavities discussed in [73]). As we discussed in Sec. II B,

³As discussed in [9], the quantum-limited sensitivity of ABRACADABRA-style experiments may theoretically be better, but it is less clear whether it is plausibly reachable. The DM Radio projections are more closely related to the SQL-limited sensitivity for a static-field experiment, so provide a simpler point of comparison here.

such a high signal power likely requires more complicated, nonconvex cavity geometries, such as our corrugated toroids.

We can also compare SRF up-conversion experiments to those at optical frequencies, as proposed in [16–18]. In Fig. 5, we show the projections for the ADBC experiment [18], taking the 2 meter cavity version for commonality with the other meter-scale experiments on the plot. They assume a circulating optical power of 10 kW, giving a stored energy of $\sim 10^{-4}$ J. This is many orders of magnitude below the cavity energies of the SRF experiments, and the shot-noise-limited sensitivity shown in the plot is correspondingly many orders of magnitude worse. The ADBC projection does not consider experimental issues such as polarizer efficiency, vibrational noise, etc, so the sensitivity of a realistic experiment may be worse. Of course, optical experiments may also be easier and/or cheaper to implement than SRF experiments.

V. CONCLUSIONS

The nominal SRF experiments we considered above have signal power limited by the maximum magnetic field at the cavity walls. Apart from increasing the cavity size, one possibility for improving this is to use a different superconducting material for the walls. For example, Nb_3Sn has $H_{sh} \simeq 0.41$ T, and has been extensively investigated as a potential alternative to niobium [19,74]. While existing fabrication methods lead to worse high-field performance than niobium (potentially due to defects), research into improving these is ongoing [75]. Another potential benefit of alternative materials is that they can have higher T_c (e.g., $T_c \simeq 18$ K for Nb_3Sn [19]); this can help by suppressing R_{BCS} , or by allowing operation at higher temperatures, where cooling systems are more efficient.

If the thermal noise limited sensitivity projections can actually be achieved, then further progress (for the same cavity geometry) would depend on “quantum engineering.” Preparing the target mode in a nonclassical state, such as a squeezed state or a Fock state, can improve sensitivity. At microwave frequencies, such technologies are being developed as part of the \gtrsim GHz axion detection program. For example, squeezed state injection is being incorporated into Phase II of the HAYSTAC experiment [42].

Looking beyond axion DM detection experiments, the more complicated cavity geometries we have discussed may be useful in other situations where large, high-frequency magnetic fields are required. For example, microwave frequency light-through-wall experiments to search for hidden sector particles, independent of whether they are DM, have been proposed [29,76], and a dark photon experiment of this form is being constructed at Fermilab [27]. The signal strength in these experiments depends on how large a field amplitude can be sustained in

the driven cavity. Understanding whether or not high-field cavities can be constructed with the correct geometry to source hidden-sector particles would require further work.

Coming back to the prospects for low-frequency axion DM detection, a pathfinder SRF experiment that covers significant ALP parameter space, along the lines of our nominal 60 L cavity, seems plausibly feasible. In principle, larger volumes and/or more advanced cavity geometries could allow for QCD axion sensitivity, at least insofar as fundamental noise limits are concerned. Further work would be required to understand whether other backgrounds, such as drive leakage, could be mitigated. As an alternative to static background field approaches, SRF up-conversion would present very different experimental challenges, and may be worth further investigation. Unlike SRF experiments aimed at \sim GHz frequency axions [15], up-conversion of low-frequency axions does possess parametrically better scaling than static-field approaches, though (as we have reviewed) there are many compensating difficulties. To draw an analogy, resonant bar detectors and laser interferometers represented very different approaches to gravitational wave detection. While static-field experiments may well prove to be more practical, it is important to understand the alternatives, especially as technology (such as superconducting materials) evolves.

Similar experimental concepts to those we have discussed are also proposed in [73], which we became aware of while this work was in progress.

ACKNOWLEDGMENTS

We thank Saptarshi Chaudhuri, William DeRocco, Sebastian Ellis, Peter Graham, Junwu Huang, Kent Irwin, Emilio Nanni, Chris Nantista, Raffaele Tito D’Agnolo, Sami Tantawi, Jesse Thaler, and Paul Wellander for useful discussions. This research was supported by the Munich Institute for Astro- and Particle Physics (MIAPP) which is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy—EXC-2094–390783311.

APPENDIX A: CAVITY WALL FIELDS

The Maxwell SET is traceless, with $T_{\mu}^{\mu} = 0$, so $T^{00} = T^{ii}$. Consequently, the total energy inside a cavity is

$$U = \int dV T^{00} = \int dV T^{ii} \quad (\text{A1})$$

We can use the standard trick [77,78] of rewriting T^{ij} as follows,

$$T^{ij} = \partial_k(x^i T^{kj}) - x^i \partial_k T^{kj} \quad (\text{A2})$$

$$\begin{aligned} T^{ij} &= \partial_k(x^i T^{kj}) - x^i \partial_k T^{kj} \\ &= \partial_k(x^i T^{kj}) + x^i \partial_i T^{0j} \end{aligned} \quad (\text{A3})$$

where we have used conservation of the SET, $\partial_\mu T^{\mu\nu} = 0$. Then, time-averaging over an oscillation period of the fields inside the cavity,

$$U = \int dV T^{ii} = \left\langle \oint_{\partial V} dS_k x^i T^{ki} \right\rangle. \quad (\text{A4})$$

Using the boundary conditions at the conducting walls, $n \times E = 0$, $n \cdot B = 0$, and the form of the Maxwell SET, we have

$$n_k T^{ik} = -n_k \left(E_i E_k + B_i B_k - \frac{1}{2} (E^2 + B^2) \delta_{ik} \right) \quad (\text{A5})$$

$$= \frac{1}{2} (B^2 - E^2) n_i \quad (\text{A6})$$

where n is the outward-pointing normal to the cavity wall. Thus,

$$U = \frac{1}{2} \left\langle \oint_{\partial V} dA (B^2 - E^2) (x \cdot n) \right\rangle. \quad (\text{A7})$$

Physically, this is an integrated version of the Slater formula for the energy change due to cavity wall deformations [79]. If we imagine dilating the cavity dimensions from their initial value to zero, while keeping the wall fields fixed, we obtain Eq. (A7).

APPENDIX B: SNR OF RESONANT SEARCHES

In the text, we emphasized that the physical temperature T_0 of the cavity will generally be significantly greater than the signal frequency ω_1 . However, as discussed in [9,41], it is possible to reduce the thermal noise contaminating the signal, by overcoupling the signal mode to a ‘‘cold’’ detector. Here, we review these SNR calculations, as they apply to our setup.

Figure 6 illustrates the readout system we assume, with an amplifier isolated behind a circulator, and a cold load absorbing the amplifier’s back-action noise. If the output port is coupled to the signal mode ξ times more strongly than environmental dissipation (i.e., a mode fluctuation loses ξ times more of its energy to the output port), then the transmission coefficient for thermal noise from the walls is $A(\omega) = \frac{4\xi}{(1+\xi)^2} \cos^2 \alpha(\omega)$, where $\cos^2 \alpha$ is as per Sec. II. If we assume that the output line is impedance matched to its load, so that no reflections are sent back to the cavity, then the (single-sided) noise PSD at the detector input is

$$S_{\text{in}} = \frac{4\xi}{(1+\xi)^2} \cos^2 \alpha (S_{T_0} - S_{T_c}) + S_{T_c} \quad (\text{B1})$$

where T_c the effective temperature of the back-action noise from the detector. Here, $S_T(\omega) = n_T(\omega)\omega$ [80], where




FIG. 6. Signal diagram of a cavity readout system using an amplifier isolated by a circulator, as discussed in Appendix B. The T_0 load represents the thermal noise from the environment (i.e., the cavity walls etc), while the T_c load is a cold load that absorbs the amplifier’s back-action noise.

$n_T(\omega) \equiv (e^{\omega/T} - 1)^{-1}$ is the thermal occupation number, so $S_T \simeq T$ for $T \gg \omega$. As expected, if $T_0 = T_c$, then the system is in equilibrium, and the noise PSD is the same at all frequencies.

Integrating over frequencies close to the resonance, the energy flux from the environment to the detector is approximately $P_{01} = (T_0 - T_c) \frac{\gamma_0 \gamma_1}{\gamma_0 + \gamma_1} = (T_0 - T_c) \gamma \frac{\xi}{(1+\xi)^2}$, where γ_0 and γ_1 are the damping rates to the environment and the detector, and $\gamma = \gamma_0 + \gamma_1$. This is generically the result obtained when a single harmonic oscillator interacts with baths at different temperatures [81,82]. The energy in the cavity mode corresponds to an effective temperature of $T_{\text{eff}} = \frac{\gamma_0 T_0 + \gamma_1 T_c}{\gamma_0 + \gamma_1}$, and there is an energy flux $(T_{\text{eff}} - T_c) \gamma_1 = P_{01}$ to the detector port, as we would expect from a naive analysis.

Assuming a high-gain amplifier, the noise PSD at the amplifier output also has contributions from the amplifier’s output noise, and from the amplification of vacuum fluctuations at the input. We can refer the amplifier’s output fluctuations to its input by dividing by the power gain G ,

$$S_n \equiv \frac{S_{\text{out}}}{G} \simeq S_{\text{in}} + S_{\text{vac}} + S_{\text{amp}}. \quad (\text{B2})$$

For a phase-insensitive amplifier, if the input state is coherent, then $S_{\text{vac}} = \frac{\omega}{2}$. For a SQL-limited amplifier, $S_{\text{amp}} = \frac{\omega}{2}$ as well, so vacuum plus amplifier noise combine to give ‘‘a single photon’’ of output noise, in the usual phrasing [80]. Below, we will write $S_a \equiv S_{\text{vac}} + S_{\text{amp}}$, which for a phase-insensitive amplifier with coherent input state is $\geq \omega$.

The PSD of absorbed power from the axion field can be written as

$$S_{\text{abs}} \simeq S_{jj} \frac{P_0}{|j_0|^2} \cos^2 \alpha(\omega) \equiv S_0 \cos^2 \alpha(\omega) \quad (\text{B3})$$

where $j(t) = \dot{a}(t)B_0(t)V_b$ (in the notation of Sec. II), and P_0 is the power that would be absorbed on-resonance from a monochromatic $j(t) = j_0 \cos^2(\omega t)$ oscillation. For a top-hat j spectrum of bandwidth $\delta\omega_a$, we would have $S_{jj}/|j_0|^2 = 2\pi/\delta\omega_a$. A fraction $\frac{\xi}{1+\xi}$ of this will enter the detector port. Consequently, the ratio of signal to noise PSDs, referred to the amplifier input, is

$$\frac{S_{\text{sig}}}{S_n} \simeq \frac{\frac{\xi}{1+\xi} \cos^2 \alpha S_0}{\frac{4\xi}{(1+\xi)^2} \cos^2 \alpha (S_{T_0} - S_{T_c}) + S_{T_c} + S_a}. \quad (\text{B4})$$

From the Dicke radiometer formula [83], the SNR contribution from a small frequency bin, of bandwidth $\delta\nu$, is $\text{SNR} \simeq \frac{S_{\text{sig}}}{S_n} \sqrt{t_1 \delta\nu}$, where t_1 is the integration time. The contributions from different frequency bins add in quadrature, so assuming that the integration time is long enough to resolve the spectral features of S_{sig} and S_n , we have

$$\text{SNR}^2 \simeq t_1 \int d\nu \left(\frac{S_{\text{sig}}}{S_n} \right)^2 \quad (\text{B5})$$

At this point, the obvious question is what value of ξ we should select to maximize the total SNR, in different circumstances. To answer this, it is helpful to extract the ξ dependence from P_0 and $\cos^2 \alpha$; this gives

$$\frac{S_{\text{sig}}}{S_n} \simeq \frac{S_{jj}}{|j_0|^2} \frac{\xi P_1}{4\xi(S_{T_0} - S_{T_c}) + (S_{T_c} + S_a)((1+\xi)^2 + 4Q_1^2(\frac{\omega}{\omega_1} - 1)^2)} \quad (\text{B6})$$

where Q_1 is the unloaded quality factor, and P_1 is the power that would be absorbed on-resonance in the unloaded case (cf. Eq. (6) with $Q_l = Q_1$). Plugging this into Eq. (B5) reproduces the form of Eq. (167) in [41].

If S_{jj} is even narrower than the unloaded bandwidth of the signal mode, then to maximize sensitivity at the axion mass we are tuned to, we simply want to maximize S_{sig}/S_n on-resonance, which is achieved at $\xi = 1$ (i.e., the usual critical coupling [80]). If, on the other hand, S_{jj} is wide compared to the unloaded signal mode bandwidth, then we want to maximize

$$\text{SNR}^2 \simeq t_1 \left(\frac{S_{jj} P_1}{|j_0|^2} \Big|_{\omega_1} \right)^2 \int d\nu \left(\frac{\xi}{4\xi(S_{T_0} - S_{T_c}) + \dots} \right)^2 \quad (\text{B7})$$

We write ξ_{opt} for the value of ξ that maximizes this integral. For example, if $T_c = T_0$, then $\xi_{\text{opt}} = 2$. Another case of physical interest is if $T_0 \gg T_c, S_a$. For an SRF cavity, while cooling the cavity walls to below 1 K would be prohibitively difficult, realising a cold load at significantly lower temperatures, and an amplifier noise temperature $\ll 1$ K, is feasible. In this case, $\xi_{\text{opt}} \simeq \frac{2S_{T_0}}{S_{T_c} + S_a}$. The best achievable parameters for a SQL-limited amplifier are $T_c = 0$, $S_a = \omega$, in which case $\xi_{\text{opt}} \simeq 2T_0/\omega_1$. This is the optimum overcoupling found in [41].

In the $T_0 \gg T_c, S_a$ case, with $\xi = \xi_{\text{opt}}$, the loaded quality factor of the signal mode is $Q_l \simeq Q_1 \frac{\omega_1}{2T_0}$. The equivalent quality factor for the expression in Eq. (B6) is $Q_s \simeq Q_l/\sqrt{3}$, labeled the ‘‘sensitivity quality factor’’ in

[41]. Physically, overcoupling by ξ_{opt} reduces the loaded quality factor of the mode, but also dilutes the thermal noise reaching the amplifier down to $S_{\text{in}} \simeq \frac{\omega_1}{2} \cos^2 \alpha$. Increasing ξ further results in S_a dominating over S_{in} , so we reduce the quality factor for little gain.

In [41], the improvement in scan-averaged sensitivity coming from using $\xi = \xi_{\text{opt}} \simeq 2T_0/\omega_1$ versus $\xi \simeq 1$ is phrased as gaining sensitivity ‘‘outside of the resonator bandwidth.’’ For scattering-type setups, this applies to the *unloaded* resonator bandwidth $\sim \omega_1/Q_1$. At $\xi = \xi_{\text{opt}}$, the *loaded* resonator bandwidth (i.e., the physical bandwidth of mode fluctuations in the experiment) is, to within a $\mathcal{O}(1)$ factor, the same as the sensitivity bandwidth, as per the previous paragraph. The improved scan-averaged sensitivity comes from this bandwidth being parametrically larger than the unloaded bandwidth, but the on-resonance SNR being only a $\mathcal{O}(1)$ factor smaller. Overcoupling the signal mode reduces the on-resonance signal power, but reduces the thermal fluctuations in the signal mode by parametrically the same amount, while increasing the signal mode’s loaded bandwidth.

This situation can be contrasted with the case of an amplifier used in ‘‘op-amp’’ mode [80], which is discussed in [13,41] for the case of flux-to-voltage amplifiers. In the limit of very large amplifier power gain, the power lost from the mode to the amplifier necessarily vanishes, and the effect on the signal mode’s quality factor is very small. As discussed in [13,41], one can obtain an analogous increase in scan-averaged sensitivity by overcoupling to the amplifier, by the same $\xi_{\text{opt}} \simeq 2T_0/\omega_1$ factor relative to the ‘‘critical’’ coupling that optimizes on-resonance sensitivity.

In this case, however, the sensitivity bandwidth is $\sim \xi_{\text{opt}}$ times larger than the *physical* resonator bandwidth, and the mode's fluctuations are larger than those expected from the temperature T_0 .

Returning to the scattering case, if we take $\xi = \xi_{\text{opt}}$, then for an axion signal with bandwidth $\delta\omega_a \gg \omega_1/Q_l$, Eq. (B7) tells us that the SNR from a single tuned configuration, with frequency splitting within the axion bandwidth, is

$$\text{SNR}^2 \simeq 0.7 \frac{(P_1/Q_l)^2 t_1 Q_a^2 Q_l \omega_1}{T_0 T_n m_a^2} \quad (\text{B8})$$

where $T_n \equiv S_{T_c} + S_a$. P_1 is evaluated for a monochromatic j oscillation with a_0 set by the dark matter density, and B_0 set by the RMS magnetic field in the drive mode.

More generally, unless the axion mass is small compared to the sensitivity bandwidth, covering an $\mathcal{O}(1)$ range in axion masses requires running the experiment in multiple different configuration—for up-conversion, with multiple different frequency splittings between the drive and signal modes. A scan strategy specifies which frequency splittings to choose, and how long to stay in each of them. In order to cover the axion mass range equally, the frequency splittings should be spaced equally,⁴ at frequencies differing by $\lesssim \nu_b$, where $\nu_b = \max(\delta\nu_a, \nu_1/Q_s)$.

If the SNR formulas derived above apply, then any such scan strategy will give approximately the same SNR.

⁴for $\lesssim \mathcal{O}(1)$ frequency ranges, the choice of axion masses prior, e.g., linear vs log-linear, does not make a significant difference.

The simplest example we can consider is to take S_{jj} to be a top-hat of width $\delta\omega_a$, and take the different frequency splittings to be spaced equally at $\delta\omega_a$, such that only a single configuration responds strongly at each possible axion mass. Then, the time spent in each configuration is $t_1 \simeq t_{\text{tot}} \frac{\delta\omega_a}{\Delta m_a}$, and the SNR from the responding configuration is

$$\text{SNR}^2 \simeq 0.7 \frac{(P_1/Q_l)^2 t_{\text{tot}} Q_a Q_l \omega_1}{T_0 T_n m_a \Delta m_a}. \quad (\text{B9})$$

For a denser set of frequency splittings, multiple configurations will have significant response at each axion mass. Since the absorbed power, averaged over axion masses, is given by Eq. (8), any sufficiently dense set of \sim equal spacings will give \sim the same signal power at each axion mass, so the SNR will still be given by Eq. (B9).

However, once the time spent covering each frequency becomes too small, these SNR formulas become invalid. In the case of a single configuration, once $t_1 \lesssim Q_l/\nu_1$, we cannot resolve the signal mode bandwidth, and the axion signal does not have time to fully ring up the mode.

We can gain some insight into this behavior by rewriting Eq. (B9) in terms of the average energy absorbed over the lifetime of the experiment, $\bar{W} \simeq \bar{P} t_{\text{tot}}$. This gives

$$\text{SNR}^2 \simeq 0.2 \frac{Q_l}{t_1 T_0 T_n} \frac{\bar{W}^2}{\omega_1} \simeq 0.4 \frac{Q_l}{\omega_1 t_1} \frac{\bar{W}^2}{T_n^2}. \quad (\text{B10})$$

As discussed in [9], SNR/\bar{W} stops growing for $Q_l \gtrsim t_1 \nu_1$, attaining a maximum value of $\text{SNR} \simeq 0.2 \bar{W}/T_n$.

-
- [1] P. W. Graham, I. G. Irastorza, S. K. Lamoreaux, A. Lindner, and K. A. van Bibber, *Annu. Rev. Nucl. Part. Sci.* **65**, 485 (2015).
 - [2] G. Grilli di Cortona, E. Hardy, J. Pardo Vega, and G. Villadoro, *J. High Energy Phys.* **01** (2016) 034.
 - [3] L. Di Luzio, F. Mescia, and E. Nardi, *Phys. Rev. Lett.* **118**, 031801 (2017).
 - [4] J. Jaeckel and A. Ringwald, *Annu. Rev. Nucl. Part. Sci.* **60**, 405 (2010).
 - [5] N. Du *et al.* (ADMX Collaboration), *Phys. Rev. Lett.* **120**, 151301 (2018).
 - [6] T. Braine *et al.* (ADMX Collaboration), *Phys. Rev. Lett.* **124**, 101303 (2020).
 - [7] A. Caldwell, G. Dvali, B. Majorovits, A. Millar, G. Raffelt, J. Redondo, O. Reimann, F. Simon, and F. Steffen (MADMAX Working Group Collaboration), *Phys. Rev. Lett.* **118**, 091801 (2017).
 - [8] M. Baryakhtar, J. Huang, and R. Lasenby, *Phys. Rev. D* **98**, 035006 (2018).
 - [9] R. Lasenby, [arXiv:1912.11467](https://arxiv.org/abs/1912.11467).
 - [10] J. Ouellet and Z. Bogorad, *Phys. Rev. D* **99**, 055010 (2019).
 - [11] M. Beutter, A. Pargner, T. Schwetz, and E. Todarello, *J. Cosmol. Astropart. Phys.* **02** (2019) 026.
 - [12] Y. Kahn, B. R. Safdi, and J. Thaler, *Phys. Rev. Lett.* **117**, 141801 (2016).
 - [13] S. Chaudhuri, K. D. Irwin, P. W. Graham, and J. Mardon, [arXiv:1904.05806](https://arxiv.org/abs/1904.05806).
 - [14] M. Goryachev, B. Mcallister, and M. E. Tobar, *Phys. Dark Universe* **26**, 100345 (2019).
 - [15] P. Sikivie, [arXiv:1009.0762](https://arxiv.org/abs/1009.0762).
 - [16] W. DeRocco and A. Hook, *Phys. Rev. D* **98**, 035021 (2018).
 - [17] I. Obata, T. Fujita, and Y. Michimura, *Phys. Rev. Lett.* **121**, 161301 (2018).
 - [18] H. Liu, B. D. Elwood, M. Evans, and J. Thaler, *Phys. Rev. D* **100**, 023548 (2019).
 - [19] H. Padamsee, J. Knobloch, and T. Hays, *RF Superconductivity for Accelerators*, Wiley Series in Beam Physics and Accelerator Technology (Wiley, New York, 1998).

- [20] A. Grassellino *et al.*, *Supercond. Sci. Technol.* **30**, 094004 (2017).
- [21] M. Tobar (private communication).
- [22] M. Kawasaki, K. Saikawa, and T. Sekiguchi, *Phys. Rev. D* **91**, 065014 (2015).
- [23] A. Ringwald and K. Saikawa, *Phys. Rev. D* **93**, 085031 (2016); **94**, 049908(A) (2016).
- [24] R. T. Co, L. J. Hall, and K. Harigaya, *Phys. Rev. Lett.* **120**, 211602 (2018).
- [25] P. W. Graham, J. Mardon, and S. Rajendran, *Phys. Rev. D* **93**, 103520 (2016).
- [26] M. Tinkham, *Introduction to Superconductivity*, International series in pure and applied physics (McGraw Hill, New York, 1996).
- [27] A. Grassellino, <https://indico.fnal.gov/event/19433/session/2/contribution/2/material/slides/0.pdf> (2019).
- [28] Z. Bogorad, A. Hook, Y. Kahn, and Y. Soreq, *Phys. Rev. Lett.* **123**, 021801 (2019).
- [29] R. Janish, V. Narayan, S. Rajendran, and P. Riggins, *Phys. Rev. D* **100**, 015036 (2019).
- [30] J. F. Cochran and D. E. Mapother, *Phys. Rev.* **111**, 132 (1958).
- [31] A. Gurevich, *Rev. Accel. Sci. Technol.* **05**, 119 (2012).
- [32] H. B. Michaelson, *J. Appl. Phys.* **48**, 4729 (1977).
- [33] H. Padamsee and J. Knobloch, *AIP Conf. Proc.* **474**, 212 (1999).
- [34] Y. Pischalnikov, E. Borissov, I. Gonin, J. Holzbauer, T. Khabiboulline, W. Schappert, S. Smith, and J.-C. Yun, in *Proceedings, 6th International Particle Accelerator Conference (IPAC 2015): Richmond, Virginia, USA, May 3–8, 2015 (JACoW, Geneva, Switzerland, 2015)*, p. WEPTY035.
- [35] C. Antoine, M. Foley, and N. Dhanaraj, <https://doi.org/10.2172/1022786> (2011).
- [36] J. L. Doane, in *IN: Infrared and Millimeter Waves. Volume 13 (A86-37151 17–33). Orlando*, Vol. 13 (Academic Press, Orlando, FL, 1985), pp. 123–170.
- [37] P. J. Claricoats and A. D. Olver, *Corrugated Horns for Microwave Antennas* (IET, Peter Peregrinus, London, 1984).
- [38] E. J. Kowalski, D. S. Tax, M. A. Shapiro, J. R. Sirigiri, R. J. Temkin, T. S. Bigelow, and D. A. Rasmussen, *IEEE Trans. Microwave Theory Tech.* **58**, 2772 (2010).
- [39] M. Miyagi, K. Harada, and S. Kawakami, *IEEE Trans. Microwave Theory Tech.* **32**, 513 (1984).
- [40] G. B. Bowden, C. Chang, J. Neilson, M. Shumail, S. G. Tantawi, and C. Pellegrini, *Conf. Proc. C* **110904**, 3328 (2011).
- [41] S. Chaudhuri, K. Irwin, P. W. Graham, and J. Mardon, [arXiv:1803.01627](https://arxiv.org/abs/1803.01627).
- [42] A. Droster and K. van Bibber (HAYSTAC Collaboration), in *13th Conference on the Intersections of Particle and Nuclear Physics (CIPANP 2018) Palm Springs, California, USA, May 29–June 3, 2018* (2019).
- [43] N. Vinyoles, A. Serenelli, F. L. Villante, S. Basu, J. Redondo, and J. Isern, *J. Cosmol. Astropart. Phys.* **10** (2015) 015.
- [44] H. Schlattl, A. Weiss, and G. Raffelt, *Astropart. Phys.* **10**, 353 (1999).
- [45] A. Ayala, I. Domínguez, M. Giannotti, A. Mirizzi, and O. Straniero, *Phys. Rev. Lett.* **113**, 191302 (2014).
- [46] A. Payez, C. Evoli, T. Fischer, M. Giannotti, A. Mirizzi, and A. Ringwald, *J. Cosmol. Astropart. Phys.* **02** (2015) 006.
- [47] B. M. Brubaker *et al.*, *Phys. Rev. Lett.* **118**, 061302 (2017).
- [48] L. Zhong *et al.* (HAYSTAC Collaboration), *Phys. Rev. D* **97**, 092001 (2018).
- [49] W. Wuensch, S. De Panfilis-Wuensch, Y. K. Semertzidis, J. T. Rogers, A. C. Melissinos, H. J. Halama, B. E. Moskowitz, A. G. Prodell, W. B. Fowler, and F. A. Nezrick, *Phys. Rev. D* **40**, 3153 (1989).
- [50] C. Haggmann, P. Sikivie, N. S. Sullivan, and D. B. Tanner, *Phys. Rev. D* **42**, 1297 (1990).
- [51] S. J. Asztalos *et al.* (ADMX Collaboration), *Phys. Rev. Lett.* **104**, 041301 (2010).
- [52] J. L. Ouellet *et al.*, *Phys. Rev. Lett.* **122**, 121802 (2019).
- [53] K. Irwin, <https://irwinlab.sites.stanford.edu/dark-matter-radio-dm-radio> (2019).
- [54] P. Bernard, G. Gemme, R. Parodi, and E. Picasso, *Rev. Sci. Instrum.* **72**, 2428 (2001).
- [55] R. Ballantini, P. Bernard, E. Chiaveri, A. Chincarini, G. Gemme, R. Losito, R. Parodi, and E. Picasso, *Classical Quantum Gravity* **20**, 3505 (2003).
- [56] R. Ballantini *et al.*, [arXiv:gr-qc/0502054](https://arxiv.org/abs/gr-qc/0502054).
- [57] P. Bernard, A. Chincarini, G. Gemme, R. Parodi, and E. Picasso, [arXiv:gr-qc/0203024](https://arxiv.org/abs/gr-qc/0203024).
- [58] D. Meidlinger *et al.*, *Proceedings of SRF2009, Berlin, Germany*, THPPO005 (2009).
- [59] W. M. Whitney, *Phys. Rev.* **105**, 38 (1957).
- [60] W. A. Jeffers and W. M. Whitney, *Phys. Rev.* **139**, A1082 (1965).
- [61] N. Tsuda and T. Suzuki, *J. Phys. Chem. Solids* **28**, 2487 (1967).
- [62] P. R. Saulson, *Fundamentals of Interferometric Gravitational Wave Detectors* (World Scientific, Singapore, 2017).
- [63] R. L. Weaver and O. I. Lobkis, *Phys. Rev. Lett.* **87**, 134301 (2001).
- [64] M. Ge, N. Banerjee, J. Dobbins, R. Eichhorn, F. Furuta, G. Hoffstaetter, M. Liepe, P. Quigley, J. Sears, and V. Veshcherevich, in *Proceedings, 28th International Linear Accelerator Conference (LINAC16): East Lansing, Michigan, September 25–30, 2016 (JACoW, Geneva, Switzerland, 2017)*, p. TUPLR010.
- [65] M. P. Kelly, J. D. Fuerst, M. Kedzie, S. I. Sharamentov, K. W. Shepard, and J. Delaysen, *Conf. Proc. C* **030512**, 1291 (2003).
- [66] A. Neumann, W. Anders, O. Kugeler, and J. Knobloch, *Phys. Rev. ST Accel. Beams* **13**, 082001 (2010).
- [67] H. Padamsee, in *Proceedings, CAS—CERN Accelerator School: Course on Superconductivity for Accelerators (CAS 2013): Erice, Italy, April 24–May 4, 2013* (CERN, Geneva, 2014), pp. 141–169.
- [68] D. Alesini, in *Proceedings, CAS—CERN Accelerator School: RF for Accelerators, Ebeltoft, Denmark, 8–17 Jun 2010* (CERN, Geneva, 2011).
- [69] S. Chaudhuri (private communication).
- [70] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, ninth dover printing, tenth gpo printing (Dover, New York, 1964).

- [71] M. Tanabashi *et al.* (Particle Data Group Collaboration), *Phys. Rev. D* **98**, 030001 (2018).
- [72] F. Caspers and M. Betz, Benchmarking microwave cavity dark matter searches using a radioactive source, <http://cds.cern.ch/record/1979598/files/Poster-2015-463.pdf> (2015).
- [73] A. Berlin, R. T. D’Agnolo, S. A. R. Ellis, C. Nantista, J. Neilson, P. Schuster, S. Tantawi, N. Toro, and K. Zhou, [arXiv:1912.11048](https://arxiv.org/abs/1912.11048).
- [74] S. Posen and D.L. Hall, *Supercond. Sci. Technol.* **30**, 033004 (2017).
- [75] S. Posen, S. K. Chandrasekaran, J. Lee, O. Melnychuk, D. Sergatskov, B. Tennis, and Y. Trenikhina, in Proceedings, 9th International Particle Accelerator Conference (IPAC 2018): Vancouver, BC Canada, April 29-May 4, 2018 (JACoW, Geneva, Switzerland, 2018), p. WEPML016.
- [76] P. W. Graham, J. Mardon, S. Rajendran, and Y. Zhao, *Phys. Rev. D* **90**, 075017 (2014).
- [77] S. Weinberg, *Gravitation and Cosmology* (John Wiley and Sons, New York, 1972).
- [78] L. Hui and A. Nicolis, *Phys. Rev. Lett.* **105**, 231101 (2010).
- [79] J. Slater, *Microwave Electronics*, Bell Telephone Laboratories series (Van Nostrand, Princeton, 1950).
- [80] A. A. Clerk, M. H. Devoret, S. M. Girvin, F. Marquardt, and R. J. Schoelkopf, *Rev. Mod. Phys.* **82**, 1155 (2010).
- [81] Z. Rieder, J. L. Lebowitz, and E. Lieb, *J. Math. Phys. (N.Y.)* **8**, 1073 (1967).
- [82] H. C. Fogedby and A. Imparato, *J. Stat. Mech.* (2011) P05015.
- [83] R. H. Dicke, *Rev. Sci. Instrum.* **17**, 268 (1946).