Design and Operational Experience of a Microwave Cavity Axion Detector for the 20 – 100 μeV Range

S. Al Kenany^a, M.A. Anil^b, K.M. Backes^a, B.M. Brubaker^c, S.B. Cahn^c, G. Carosi^d, Y.V. Gurevich^c, W.F. Kindel^b, S.K. Lamoreaux^c, K.W. Lehnert^b, S.M. Lewis^a, M. Malnou^b, D.A. Palken^b, N.M. Rapis^a, J.R. Root^a, M. Simanovskaia^{a,*}, T.M. Shokair^a, I. Urdinaran^a, K.A. van Bibber^a, L. Zhong^c

^a*Department of Nuclear Engineering, University of California Berkeley, Berkeley CA, 94720 USA*

^b*JILA and the Department of Physics, University of Colorado and National Institute of Standards and Technology, Boulder CO, 80309 USA*

^c*Department of Physics, Yale University, New Haven CT, 06511 USA*

^d*Physical and Life Sciences Directorate, Lawrence Livermore National Laboratory, Livermore CA, 94551 USA*

Abstract

We describe a dark matter axion detector designed, constructed, and operated both as an innovation platform for new cavity and amplifier technologies and as a data pathfinder in the 5 – 25 GHz range ($\sim 20 - 100 \mu\text{eV}$). The platform is small but flexible to facilitate the development of new microwave cavity and amplifier concepts in an operational environment. The experiment has recently completed its first data production; it is the first microwave cavity axion search to deploy a Josephson parametric amplifier and a dilution refrigerator to achieve near-quantum limited performance.

Keywords: axion, dark matter, Josephson Parametric Amplifier, microwave cavity, standard quantum limit, superconducting magnet

1. Introduction

The axion is a hypothetical pseudoscalar arising from the Peccei-Quinn mechanism to protect the strong interaction from CP-violating effects. Also, an axion in the 1 – 100 μeV mass range is a compelling dark matter candidate. A comprehensive review of the particle physics of the axion, its cosmological

*Corresponding author

Email address: simanovskaia@berkeley.edu (M. Simanovskaia)

Preprint submitted to Nuc. Instrum. Meth.

February 8, 2017

22 and astrophysical significance, and experimental searches for it can be found in
 23 Ref. [1].

24 The axion, like the π^0 , can couple to two photons, one of which may be vir-
 25 tual. Sikivie thus proposed a practical detection strategy based on the resonant
 26 conversion of the dark matter axion to a single microwave photon carrying its
 27 full energy (mass + kinetic) in a high-Q cavity permeated by a strong magnetic
 28 field [2–4]. The resonant conversion condition is that the axion mass m_a is
 29 close to the resonant frequency ν_c of a cavity mode with an appropriate spatial
 30 profile; more precisely, $|\delta\nu| = |\nu_c - m_a c^2/h| \lesssim \Delta\nu_c$, where $\Delta\nu_c$ is the linewidth
 31 of the mode. The axion kinetic energy distribution is generally assumed to be
 32 Maxwellian, parametrized by the virial velocity $\langle v^2 \rangle^{1/2} \sim 10^{-3} c$. The signal is
 33 thus monochromatic to a part in 10^6 , with linewidth $\Delta\nu_a = m_a \langle v^2 \rangle / h \ll \Delta\nu_c$;
 34 see Fig. 1.

35 The axion conversion power in such a microwave cavity detector is

$$P_{\text{sig}} = \left(g_\gamma^2 \frac{\alpha^2}{\pi^2} \frac{\hbar^3 c^3 \rho_a}{\Lambda^4} \right) \times \left(\frac{\beta}{1 + \beta} \omega_c \frac{1}{\mu_0} B_0^2 V C_{mnl} Q_L \frac{1}{1 + (2\delta\nu/\Delta\nu_c)^2} \right). \quad (1)$$

36 Eq. (1) is valid in any self-consistent set of units; the two sets of parentheses
 37 contain theory and detector parameters, respectively. On the theoretical side,
 38 α is the fine-structure constant, $\Lambda = 77.6$ MeV encodes the dependence of the
 39 axion mass on hadronic physics, and the local axion dark matter density ρ_a and
 40 coupling constant g_γ are the parameters that experiment can constrain.¹ It is
 41 conventional to fix $\rho_a = 0.45$ GeV/cm³ [6] and quote limits on g_γ , a model-
 42 dependent dimensionless coupling that is related to the physical coupling $g_{a\gamma\gamma}$
 43 appearing in the axion-photon Lagrangian by $g_{a\gamma\gamma} = (g_\gamma \alpha / \pi \Lambda^2) m_a$. In the

¹The value of Λ used in our analysis was obtained from chiral perturbation theory [3]. Note also that $\Lambda^4 = \chi(T = 0)$, where χ is the QCD topological susceptibility which may be calculated on the lattice. The most recent lattice calculation, in Ref. [5], obtained $\Lambda = 75.6$ MeV, which would result in an 11% enhancement of the signal power.

44 KSVZ (DFSZ) benchmark axion model, $g_\gamma = -0.97$ (0.36), independent of the
 45 axion mass.

46 The experimental parameters in Eq. (1) include the magnetic field strength
 47 B_0 , the cavity volume V , and several factors characterizing the cavity mode,
 48 with $\omega_c = 2\pi\nu_c$. The coupling between the cavity mode and the receiver used
 49 to detect the signal, parametrized by β , reduces the quality factor from Q_0 to
 50 $Q_L = Q_0/(1 + \beta)$. The form factor $C_{mn\ell}$ parametrizes the overlap between
 51 the cavity mode and the external magnetic field. For the cylindrical geometry
 52 commonly used in cavity axion detectors, $C_{mn\ell}$ may be written

$$C_{mn\ell} = \frac{(\int d^3\mathbf{x} \hat{\mathbf{z}} \cdot \mathbf{e}_{mn\ell}^*(\mathbf{x}))^2}{V \int d^3\mathbf{x} \epsilon(\mathbf{x}) |\mathbf{e}_{mn\ell}(\mathbf{x})|^2}, \quad (2)$$

53 where $\mathbf{e}_{mn\ell}(\mathbf{x})$ is the normalized electric field profile of the mode, $\epsilon(\mathbf{x}) = 1$ is the
 54 dielectric constant inside the cavity, and B_0 is axial and homogeneous. Nodes
 55 in the electric field profile lead to cancellations in the form factor, and thus it
 56 is only appreciable for low-order TM_{0n0} modes.

57 Inserting typical values for the detector described in this work, which are
 58 close to the limits of present technology, the peak signal power is $P_{\text{sig}} \sim 5 \times$
 59 10^{-24} W at KSVZ coupling, a factor of 10^9 below room-temperature thermal
 60 noise power in a typical cavity bandwidth. Thus cryogenic operation and a
 61 low-noise receiver are necessary for any practical realization of a cavity axion
 62 detector. The intrinsic spectral resolution of the linear receivers used in cavity
 63 axion searches to date implies that the relevant noise bandwidth is $\Delta\nu_a$ rather
 64 than $\Delta\nu_c$, and moreover that a measurement at any given cavity frequency si-
 65 multaneously probes $\sim \Delta\nu_c/\Delta\nu_a$ independent values of the axion mass. Further
 66 improvement in the signal-to-noise ratio (SNR) may then be obtained by aver-
 67 aging the cavity noise for a time τ . The preceding discussion is formalized in
 68 the Dicke radiometer equation [7], in which the SNR is defined as

$$\Sigma = \frac{P_{\text{sig}}}{k_B T_{\text{sys}}} \sqrt{\frac{\tau}{\Delta\nu_a}}. \quad (3)$$

69 For any phase-insensitive linear receiver the system noise temperature T_{sys}

70 may be written

$$k_B T_{\text{sys}} = h\nu N_{\text{sys}} = h\nu \left(\frac{1}{e^{h\nu/k_B T} - 1} + \frac{1}{2} + N_A \right), \quad (4)$$

71 where the three additive contributions correspond respectively to a blackbody
 72 gas in equilibrium with the cavity at temperature T , the zero-point fluctuations
 73 of the blackbody gas, and the input-referred added noise of the receiver. The
 74 price we pay for spectral resolution is the quantum limit $N_A \geq 1/2$ on the
 75 added noise of any phase-insensitive linear receiver [8], which together with
 76 the second term in Eq. (4) implies a “standard quantum limit” $N_{\text{sys}} \geq 1$ for
 77 microwave cavity axion detection. These limits imply that units of quanta,
 78 used throughout this work, are more appropriate than temperature units for
 79 sufficiently low-noise receivers.

80 The axion mass is a priori unknown, so the cavity must be tuned in small
 81 discrete steps after each measurement interval τ ; the SNR for any given value
 82 of the axion mass is then given by a quadrature sum of terms with the form of
 83 Eq. (3). An average scan rate is obtained by solving for $1/\tau$ and multiplying
 84 by the frequency step size. With the simplifying assumption that the detector
 85 parameters remain approximately constant over its tuning range, the scan rate
 86 is

$$\frac{d\nu}{dt} \approx \frac{4}{5} \frac{Q_L Q_a}{\Sigma^2} \left(g_\gamma^2 \frac{\alpha^2}{\pi^2} \frac{\hbar^3 c^3 \rho_a}{\Lambda^4} \right)^2 \times \left(\frac{1}{\hbar \mu_0} \frac{\beta}{1 + \beta} B_0^2 V C_{mn\ell} \frac{1}{N_{\text{sys}}} \right)^2, \quad (5)$$

87 where $Q_a = (\langle v^2 \rangle / c^2)^{-1}$ is the “quality factor” of the axion signal, and the
 88 factor of $4/5$ comes from a sum over the squared Lorentzian factors character-
 89 izing the effect of the axion’s changing detuning from the cavity mode. This
 90 is a good approximation for any frequency step size $\lesssim \Delta\nu_c/2$; qualitatively, a
 91 smaller step size implies that the same sensitivity can be achieved with smaller
 92 τ , and the two effects cancel out in the scan rate. Note also that, while Eq. (1)

is maximized at critical coupling ($\beta = 1$) for a given Q_0 , the scan rate is maximized for an overcoupled cavity with $\beta = 2$. Eq. (5) is the most useful figure of merit for the cavity axion search.

Figure 1: Schematic of the microwave cavity search for dark matter axions. The axion signal is designated by the narrow peak (red) within the bandpass of the cavity (pink).

The detector described in this work builds upon experience from the Axion Dark Matter eXperiment (ADMX), a larger platform designed to operate initially in the sub-GHz range. ADMX entered into operation in 1996 with a physical temperature of $T \approx 1.5$ K [6]. The readout chain, initially led by a high electron mobility transistor amplifier (HEMT) [6], was later upgraded to include a microstrip-coupled SQUID amplifier (MSA) [9]. A dilution refrigerator has recently been incorporated into the setup, and commissioning is underway.

Our detector was conceived as both an innovation platform and a data pathfinder for the microwave cavity axion search in the 5 – 25 GHz range ($\sim 20 - 100 \mu\text{eV}$ axion masses; $1 \text{ GHz} = 4.136 \mu\text{eV}$). The next decade in frequency brings new challenges for both microwave cavities and receivers. On the cavity side, preserving the aspect ratio, the volume of the cavity $V \sim \nu^{-3}$, resulting in an increasing penalty in conversion power with increasing frequency. This requires insightful cavity designs to maintain useful volume at higher frequencies without sacrificing the form factor or incurring an unacceptable density of intruder TE modes in the spectrum. The absolute machining and alignment tolerances will also become tighter for smaller structures, to avoid mode localization. Another challenge is that MSAs do not work well above a gigahertz;

114 to circumvent this problem we have introduced Josephson parametric amplifiers
 115 (JPA) that are ideal for the 5 GHz range. Together with a dilution refrigera-
 116 tor (DR), built into our design from the beginning, the JPA allows us to push
 117 down towards the Standard Quantum Limit and thus remain competitive with
 118 lower-frequency limits despite the loss of sensitive volume.

119 At the same time, this frequency range presents an opportunity to deploy
 120 entirely novel technologies that promise dramatic new capability for the mi-
 121 crowave cavity search. A technology that will be investigated in the near future
 122 is a receiver based on squeezed states of the vacuum to circumvent the Stan-
 123 dard Quantum Limit. Innovations to be tested in the microwave cavity domain
 124 include photonic band gap resonators, designed to eliminate the spectrum of
 125 TE modes which otherwise mix with the TM_{010} mode of interest as it is tuned,
 126 resulting in a loss of sensitivity at mode crossings and making it more difficult
 127 to track the TM_{010} mode. Other schemes to be investigated are the application
 128 of distributed Bragg reflector based resonator schemes, and thin-film Type II
 129 superconducting coatings to improve the quality factor Q of the cavity.

130 2. Description of Experimental Setup

131 Fig. 2 presents an overview of the experiment, sited at the Wright Labo-
 132 ratory of Yale University, and its integration. The microwave cavity and the
 133 magnetically-shielded canister housing the JPA are assembled on a gantry sus-
 134 pended from the dilution refrigerator. The gantry assembly is lowered by crane
 135 into the bore of the magnet, which is located in a room below the floor level of
 136 the lab containing the electronics and computer control.

137 The DR precooling system and the magnet cryogenic system use pulse tube
 138 and Gifford-McMahon cryocoolers, respectively. This allows operation of our
 139 detector without external liquid cryogenics, providing major simplification of
 140 operations, and a substantial reduction in operating costs. However, such a
 141 cryogen-free system relies on an uninterrupted supply of electrical power as the
 142 consequences of a magnet quench can be severe (see section 2.2).

Figure 2: Overview of the experiment and its integration.

The individual systems are described in detail below.

2.1. Dilution Refrigerator

The DR was manufactured in 2007 by VeriCold Technologies (subsequently acquired by Oxford Instruments), and was among the first commercially available external-cryogen-free DRs. The only major modification was the replacement of the VeriCold pulse tube with a Sumitomo RDK-415D Gifford-McMahon cryocooler (in retrospect this was a poor choice because of excessive vibration). In addition, instead of using the DR's Lakeshore temperature bridge to control the mixing chamber plate, we have incorporated a custom temperature PID-feedback controller at the mixing chamber level, resulting in reduced delay and improved performance.

The cooling power of the DR is about $150 \mu\text{W}$ at 127 mK, which is the system operating temperature. This operating temperature was chosen because we observed strong modulation of the JPA gain by vibrational fluctuations at lower temperatures; the temperature-dependence appears to originate in the circulator's magnetic shielding, whose susceptibility and heat capacity both fall

below 4 K. In addition, operating at higher temperature results in reduced temperature excursions from mechanical motions and actuation of the microwave switch, due to the increased heat capacity of the mixing chamber plate. The cavity, JPA, mixing chamber, and still temperatures are monitored with calibrated ruthenium oxide sensors.

The support gantry is a tripod with copper alloy legs and copper rings at both ends, which are clamped to the bottom of the DR's mixing chamber plate and to the upper endcap of the cavity, respectively. The equilibration time for changes in the mixing chamber temperature to propagate to the cavity temperature (measured at the lower endcap) is on the order of a few minutes. This might be surprising as the cavity is constructed of stainless steel, which has very low thermal conductivity. However, the cavity has a thick Cu plating (0.125 mm) which provides a sufficient thermal link for a short equilibration time.

The cavity and gantry are thermally shielded with an extension of the DR still shield. This extension comprises two demountable sections with dimensions 17.8 cm OD $\times$ 45 cm and 13.0 cm OD $\times$ 53.3 cm for the upper and lower sections, respectively. At the bottom of the still shield extension, there is a G-10 fiberglass disk attached via a stainless steel rod, which centers the shield extension and maintains the 0.5 cm gap between the still shield and magnet bore. The still shield extension was originally constructed from 0.16 cm sheet copper with welded seams, but that was deformed when we experienced a magnet quench (more information in section 2.2). Subsequently, we constructed a replacement using heavily plated stainless steel (0.125 – 0.250 mm), prompted by our experience that a heavy copper plating provides a good thermal link.

The DR and magnet share a single common insulating vacuum space. Extensions of the 4 K and 77 K DR stages overlap the corresponding stages at the top of the magnet, and thermal contact is provided by finger stock that presses between the extensions and the corresponding magnet thermal stages and also by blackbody radiation between the overlapping surfaces. This design allows the DR/cavity/still shield assembly to be easily inserted into and removed from

190 the magnet by use of an overhead hoist. The gas flow and electronic control
 191 lines are long enough that only a few cables must be disconnected to insert or
 192 remove the cryostat. The DR itself is supported from above, and the magnet
 193 is on a cart equipped with jack screws. When the DR is lifted, the magnet can
 194 be rolled away and the DR lowered again to its support, allowing easy and safe
 195 access to the assembly below the DR. To insert the DR, the magnet is rolled
 196 to the appropriate location, the DR and extension are lowered into the magnet,
 197 and the jack screws are used to raise the magnet until the corresponding DR
 198 and magnet cryostat vacuum flanges contact, after which the flanges are bolted
 199 together.

200 The DR operating alone has a cooldown time of about 14 hours; with the
 201 gantry, cavity, and magnet, the cooldown time to 127 mK cavity temperature
 202 is 72 hours.

203 *2.2. Magnet*

204 The magnet and its controller were designed and manufactured by Cryomag-
 205 netics, Inc., based on an existing design for a gyrotron system with similar field
 206 requirements. As mentioned already, the magnet is a pulse-tube cooled system,
 207 requiring no external cryogenics or helium gas, and operates in a persistent mode.
 208 The magnetic field is ramped up to a maximum of 9 T over 8 hours; the rate
 209 is limited by heating of the magnet itself and by eddy current heating of the
 210 cavity/gantry/still shield assembly. An interesting feature of the heating due to
 211 the ramping is that the heat load inferred from the rise in cavity temperature
 212 is largest when the field is between about 4 and 7 T, which we attribute to
 213 electrons being polarized in the stainless steel cavity body.

214 Our design specification was that the field needs to be homogenous at a level
 215 such that the perpendicular (radial) field at the inner surface of the cylindrical
 216 cavity barrel is less than 50 G. For future versions of this axion detector, we
 217 are interested in coating this inner surface with superconducting thin films to
 218 improve Q_0 . The mentioned design specification is a requirement for a thin film
 219 superconducting layer to remain effective in the presence of a strong magnetic

field. The idea is that the vortices formed by the perpendicular component of the field would be of sufficiently low density and mobility to not contribute significantly to radiofrequency power loss in the film [10]. The modification of the gyrotron design includes a set of superconducting coils in series with the main magnet that cancel the axial field over a 15.2 cm region to less than 50 G. This is the location of the JPA and its shield.

The magnet bore diameter is 14 cm inside the main solenoid and steps up to 18.4 cm further up, where the bucking coils are supported. The effective and homogeneous region in the field maximum extends axially over at least 25 cm.

An unfortunate disadvantage of a cryogen free system is that it relies on an uninterrupted source of power. The time to quench after a power loss is about 4 minutes for our system. In principle, Yale has fast acting emergency backup power, but due to system upgrades and other issues, this emergency power was not available in early March 2016, when an unscheduled power outage resulted in a magnet quench, causing significant mechanical/structural damage to the DR and coaxial microwave lines. The repairs of this damage took two months, and fortunately none of the DR's ^3He was lost; in spite of the mechanical damage, the gas flow lines were in no way compromised. Most of the forces due to this quench were caused by the large copper rings at the lower part of the gantry, and by the copper still shield extension. As noted above, we have replaced the copper still shield extension with another made from copper-plated stainless steel. In future designs, most of the gantry components will be constructed from copper-plated stainless steel, which should lead to more than a 100-fold reduction in forces during a quench.

2.3. Microwave cavity

The initial microwave cavity for the experiment consists of a right circular cylinder of stainless steel, electroplated with oxygen-free high conductivity (OFHC) copper. After annealing, the cavity achieves a near-theoretical maxi-

Figure 3: Layout of the experiment. The red volume is the DR vacuum shield, the blue volume is the magnet body, and the orange shaded region is the volume bounded by the still shield and its extensions. The magnet’s 70 K and 4 K shields are not shown below their interface with the fridge shields.

248 mum value of the quality factor Q , as limited by the anomalous skin depth [11],

$$\delta_{\text{anom}} = \left(\frac{\sqrt{3} c^2 m_e v_F}{8\pi^2 \omega n e^2} \right)^{1/3}, \quad (6)$$

249 where m_e is the mass of the electron, n the electron density in the metal, v_F the
 250 Fermi velocity, and ω the angular frequency. For copper, $n = 8.5 \times 10^{22} \text{ cm}^{-3}$,
 251 and $v_F = 1.57 \times 10^8 \text{ m/s}$.

252 The cavity employed for the initial data run was 25.4 cm in height and
 253 10.2 cm in diameter. A single large diameter rod, $\varnothing 5.1 \text{ cm}$, was used to tune
 254 the cavity (see Fig. 2). The rod pivoted around an off-axis ceramic axle, such
 255 that its radial position could be adjusted from touching the wall to centered
 256 within the cavity. This produced a dynamic range of the TM_{010} -like mode of
 257 3.6 – 5.8 GHz. To make finer frequency steps, we used a tuning vernier, a 3.2

mm diameter alumina rod, of variable insertion depth into the cavity. With the tuning rod centered, the cavity is more properly described as an annular cavity, which has the largest volume for a TM_{010} -like mode at a given frequency. The upper endcap had ports for two antennas, one with weak coupling used to measure the cavity's response in transmission and one whose insertion depth (and thus coupling β) was variable. The penetrations and services of the cavity are shown in Fig. 4. Around 5.8 GHz, typical values were $C_{010} \sim 0.5$ for the form factor and $Q_0 \sim 30,000$ for the unloaded cavity Q.

Our cavity development program has involved precision metrology and modeling in simulation environments such as CST Microwave Studio. Fig. 5 displays a simulation of the TM_{010} -like mode and a mode map of the cavity. Another innovation of this axion search is the use of the bead perturbation method to measure field profiles in a microwave cavity. To date, these have been longitudinal profiles along the full length of the cavity at a single transverse position, derived by translating a small dielectric bead (alumina, $\epsilon \sim 10$) on a Kevlar line controlled by a stepper motor and some pulleys. For a dielectric bead, the shift in mode frequency in perturbation theory is given by:

$$\frac{\Delta\omega}{\omega} = \frac{-(\epsilon - 1)}{2} \frac{V_{\text{bead}}}{V_{\text{cav}}} \frac{E(r)^2}{\langle E(r)^2 \rangle_{\text{cav}}} \quad (7)$$

where V_{bead} and V_{cav} indicate the volume of the dielectric bead and cavity, respectively [12]. For the TM_{010} -like mode, the electric field should be longitudinally invariant, and thus the shift in frequency as a function of longitudinal distance should be constant. In the microwave cavity axion search, the bead perturbation technique is useful in the design and testing of new resonators, to determine the symmetry of a mode, to confirm or correct alignment and machining tolerances to eliminate mode localization, and to identify mode-crossings regions where the modes have become substantially admixed, and thus the form factor C_{mnl} of the TM_{010} -like mode has begun to diminish (Fig. 6). A detailed study of microwave cavities for the dark matter axion experiment bringing together metrology, simulation, and EM characterization will be the subject of a

286 future publication. For the present, a few observations in relation to experiments
 287 an order of magnitude lower in frequency [6, 9] may be useful.

Figure 4: Top and side view of the cavity upper endcap, detailing the positions and design of the main tuning rod, tuning vernier, signal injection port, and the variable-coupling antenna.

288 First, for cavities of characteristic dimension $O(10\text{ cm})$, machining and align-
 289 ment tolerances to avoid significant mode localization [13] are approaching the
 290 limit of good machine shop standards and readily available components. Even
 291 the transverse play of the inner race of the bearing for the axle of the tuning
 292 rod, of order $50 - 75\text{ }\mu\text{m}$, can cause a tilt in the TM_{010} axial electric field profile
 293 of several percent. Second, modeling revealed that for the large tuning rod used
 294 in this experiment ($r/R = 0.5$), the form factor C_{010} only approached the value
 295 associated with the idealized case (i.e. the rod electrically joined to the top
 296 and bottom endcaps) for gaps $G < 250\text{ }\mu\text{m}$. Third, it is becoming clear that
 297 engineered features necessary for practical cavities (e.g. penetrations for the an-
 298 tenna coupling, diagnostic ports, etc.) have a significant effect on the mode, and
 299 are likely to be more important at higher frequencies given the larger relative
 300 size of mechanical and electrical components penetrating the cavity.

301 2.4. Mechanical Controls

302 The experiment has three mechanical systems: the rotary cavity tuning, the
 303 adjustment of the vernier, and the adjustment of the antenna. All mechan-
 304 ical controls are based on Applied Motion Products HT-23-595D NEMA 23,

Figure 5: (a) CST Microwave Studio simulation of the TM_{010} -like mode with the tuning rod in the middle of the cavity. (b) Measured mode map of the cavity, as a function of the pivot angle of the tuning rod within the cavity. The frequencies of the TM_{01n} modes decrease steeply with increasing radial distance of the tuning rod from cavity center. TE modes (indicated by yellow arrows) become apparent approaching mode crossings with TM modes; the TE mode frequencies are largely insensitive to the position of the tuning rod.

305 high-torque, double-shaft stepping motors, driven by model 5000-235 STR2 mi-
 306 crostepping controllers. One side of each stepping motor double shaft is coupled
 307 to a ten turn potentiometer that encodes the net rotation angle. The vernier and
 308 antenna stepping motors are coupled to the experiment using 10:1 worm gear
 309 reductions, while the cavity tuning motor is coupled directly. The rotary mo-
 310 tions are coupled into the vacuum using three K.J. Lesker FMH-25A Dynamic
 311 O-Ring Shaft Seal feedthroughs.

312 Within the cryostat, the antenna and vernier are controlled by 0.36 mm
 313 Kevlar thread lines that are wound directly onto the rotary feedthrough shaft

Figure 6: Cavity spectra for two different rod positions, as the TM_{010} -like mode approaches a crossing with a high-order TE mode. For each rod position, the plots above the main panels show the shifts in frequency of the two modes, as a small dielectric bead is translated axially through the cavity. The frequency shift at each axial position Z is proportional to the square of the local electric field according to perturbation theory (see Eq. 7), thus elucidating the nature of the mode. (The spectra in the main panels are with the bead removed from the cavity.) (a) At a separation of 8 MHz, the purity of the TM_{010} -like mode is confirmed by the flat axial field profile shown on the top left. (b) At a separation of 2 MHz, the lower-frequency mode now clearly reflects a strong admixture of both a high-order TE mode along with the TM_{010} -like mode.

314 that extends into the vacuum. At each stage of the DR, the lines pass through
 315 3.2 mm ID, 2.5 cm long tubes that are thermally linked to the DR stage and
 316 serve as radiation shields. The lines are routed to the cavity by use of nylon
 317 pulleys with ball bearings (McMaster-Carr 3434T13) that have been ultrasoni-
 318 cally cleaned to remove all lubricants that would freeze at low temperatures. It
 319 should be noted that Kevlar expands on cooling, and the spring tensioning of
 320 the lines must absorb the length change lest the lines loosen and possibly fall
 321 out of the pulley grooves.

322 The antenna and vernier are mounted on fixtures that allow motion along
 323 the cavity axis only. A spring is compressed when the antenna or vernier is
 324 pulled out of the cavity by the Kevlar thread, ensuring smooth and reversible

325 motion.

326 Rotary motion for the tuning rod system is delivered to the mixing chamber
327 level by use of a cryogenic G-10 tube (6.4 mm diameter, 0.79 mm wall). The
328 ends of the G-10 tube are fitted with glued-in brass extensions which allow the
329 use of set screws to couple to the feedthrough at the top, and to the mechanics
330 at the mixing chamber end. At each DR stage, a 2.5 cm section of brass tubing
331 (6.4 mm ID, 0.4 mm wall) is glued to the G-10, and a loose-fitting brass tube
332 (1.9 cm long, 0.4 cm ID) is slipped over each 2.5 cm section and coupled to the
333 stage using copper braid. At the 4 K stage, the G-10 tube inside has a brass
334 blackbody radiation block.

335 Just below the mixing chamber, a pulley and torsion spring system is used to
336 transfer rotary motion from the upper G-10 tube to a cryogenic G-10 tube (also
337 6.4 mm) along the DR and magnet axis, supported at the mixing chamber end by
338 a ceramic bearing. A kevlar line connects the brass extension of the upper shaft
339 to the 10.2 cm diameter pulley, and the torsion spring ensures that the line is
340 always pulled taut. The lower G-10 shaft runs through the JPA magnetic shield
341 and down to the cavity, where a 1.4:1 anti-backlash gear reduction provides the
342 final radial displacement required to couple to the tuning rod axle.

343 These mechanical systems have generally provided the levels of control needed
344 for the antenna and vernier; all three required appropriate selection of the mi-
345 crostepping resolution and had negligible heat loads at the operating tempera-
346 ture. The cavity bearings had some stiction due to alignment issues, and this
347 resulted in sometimes erratic positioning, and a slow drift to final equilibrium
348 after stepping. In operation of the experiment, the fine frequency control was
349 done using the vernier, with less frequent larger rotations of the tuning rod
350 which required up to 15 minutes of waiting for the rod to come to its equi-
351 librium position. We plan to eliminate this problem by improving the cavity
352 mechanics and replacing the room-temperature tuning rod drive system with
353 an AttoCube piezoelectric rotator at base temperature.

354 2.5. Josephson Parametric Amplifier

355 A JPA is a nonlinear LC resonator capacitively coupled to a transmission
 356 line with an array of SQUIDs playing the role of the nonlinear inductance (see
 357 Fig. 7). By applying a DC flux through the SQUID loops, the resonant frequency
 358 can be tuned over several gigahertz. Driving the system near its resonance with
 359 a strong pump tone enables amplification of a weak signal nearby in frequency.
 360 The gain of this amplification can be traded off with the bandwidth over which
 361 it occurs by adjusting the pump power and frequency.

Figure 7: Microphotograph of JPA circuit. The SQUID array (approximately 150 μm long) is highlighted in red on the left; the circuit's resonance is determined by the SQUID inductance and the geometric capacitance (blue). The circuit is coupled to a 50 Ω transmission line through a smaller capacitance (green). The surrounding superconducting ground plane is waffled in order to pin magnetic flux vortices in place and keep them from the SQUID array.

362 When the signal of interest is centered about the pump tone, the amplifi-
 363 cation process is noiseless (insofar as the JPA's internal loss is negligible) but
 364 phase-sensitive. In this mode of operation, the signal quadrature in phase with
 365 the pump is amplified by $\sqrt{G}$, while the one 90 degrees out of phase is squeezed
 366 by $1/\sqrt{G}$, where G is the single-quadrature power gain. When detuned com-

367 pletely to one side of the pump, the signal is amplified independent of its phase,
 368 but the JPA's intermodulation gain implies that an extra noise term enters
 369 from the image frequency, symmetric with the signal about the pump. In this
 370 configuration the added noise of a lossless JPA is equal to the thermal noise
 371 at the image frequency, giving rise to the quantum limit $N_A = 1/2$ at zero
 372 temperature.

373 It is important to note that operating in the phase-sensitive mode by itself
 374 does not eliminate the half-photon of noise associated with the zero-point motion
 375 of the blackbody gas in the cavity (i.e., the second additive term in Eq. (4)).
 376 Moreover, the factor of two improvement in noise temperature for a lossless
 377 JPA in the phase-sensitive mode is exactly canceled by the loss of information
 378 in one quadrature. Thus, phase-sensitive operation offers no improvement in
 379 axion search sensitivity unless we can also eliminate the zero-point motion of
 380 the input noise, e.g., by initializing the cavity in a squeezed state [14]. R&D
 381 with this aim is ongoing within the collaboration; for the present, we operate
 382 the JPA in the phase-insensitive mode.

383 The JPA currently installed in the experiment has a maximum resonant
 384 frequency of 6.5 GHz and a gain bandwidth product $\sqrt{GB} = 26$ MHz. Half a
 385 flux quantum threaded through the area of a single SQUID loop in the array
 386 tunes the circuit through its full 2 GHz range; this extreme flux sensitivity
 387 necessitates specialized magnetic shielding, detailed in section 2.6. The circuit
 388 was fabricated using a Nb/Al-Ox/Nb trilayer process; the nonlinear inductance
 389 comes from 20 SQUIDs in series, where each SQUID comprises two Josephson
 390 junctions of critical current $6 \mu\text{A}$ in parallel. The flux bias is delivered through
 391 a coil comprising about 20 turns of superconducting wire wound around the
 392 copper box housing the JPA chip.

393 The JPA can be turned into a passive mirror very simply by tuning off the
 394 pump tone and tuning the resonant frequency ~ 10 linewidths from the frequency
 395 of interest via the flux bias; we use this process to calibrate the absolute JPA
 396 gain. We set a gain target of $\simeq 21$ dB, below the onset of bifurcation inferred
 397 from deviations from a gaussian distribution in the JPA's quadrature noise

398 spectra, but large enough to overwhelm the ~ 20 quanta added noise of the
 399 HEMT amplifier which follows the JPA in the receiver chain. At this operating
 400 point the JPA bandwidth is about 2.3 MHz.

Figure 8: Schematic representation of the parameter space for JPA biasing. All features are intended for illustrative purposes only; they do not represent real measurements or calculations. The pump power is on the vertical axis and the detuning Δ between the pump frequency and the 0-power LC resonance of the JPA circuit is on the horizontal axis. The black curves are contours of constant gain, and the black arrows represent the path taken by the bias procedure outlined in the text, intersecting each gain curve at the minimum-power point. Beyond the critical point (Δ_c, P_c) , the system begins to oscillate and can be bistable or multistable. We operate at lower power (red dot) to avoid this region of parameter space.

401 The procedure used to bias the JPA to this target gain is illustrated schemat-
 402 ically in Fig. 8. The detuning Δ between the LC resonance and the pump fre-
 403 quency is adjusted by varying the flux bias. The gain is optimized with respect
 404 to detuning at constant pump power, and the pump power is increased if the
 405 gain remains too low after this optimization; the pump frequency remains fixed
 406 throughout the biasing process. With this procedure we always obtain within
 407 0.2 dB of the target gain, and always operate at the point of lowest pump power
 408 for a given gain, where the JPA is most stable. The biasing procedure is fully
 409 automated and incorporated into the LabVIEW code that controls the data ac-
 410 quisition (see section 2.8). It typically takes ~ 6 s to adjust the bias parameters

411 after each cavity tuning step.

412 Maintaining high gain throughout the tuning range requires the pump power
 413 and flux bias to be controlled to 0.01 dB and 2 parts in 10^5 , respectively. The
 414 latter condition corresponds to 300 nA current resolution with our present bias
 415 coil, which we easily obtain using a 20-bit ADC with 10 μ V resolution and a
 416 homemade current source with 1 mA/V transconductance.

417 2.6. Magnetic Shielding

418 Shielding of the sensitive receiver components, in particular the JPA, from
 419 the stray field of the 9 T magnet presents a complicated engineering problem.
 420 Also of concern are magnetic gradients at the location of the JPA that lead
 421 to excess amplifier noise due to gain fluctuations in the presence of mechani-
 422 cal vibrations. We have successfully attained the degree of magnetic shielding
 423 required for stable JPA operation, as shown in Fig. 9(b).

Figure 9: (a) Magnetic shielding of the JPA, including three superconducting persistent coils and (b) JPA frequency vs. flux bias current, showing no change in frequency as the main magnetic field is ramped from 0 – 9 T. The residual flux through the JPA is inferred to be $\Delta\Phi < 0.01$ flux quantum.

424 The outer elements of the JPA shield can be seen in Fig. 9(a). The first
 425 level of shielding is from the bucking coils described in section 2.2. The next

level of shielding comprises three persistent superconducting coils that provide a quasi-active shield against flux changes; given that they become superconducting before the 9 T field is ramped up, the flux inside will be maintained at the original ambient level. The axial spacing of the coils is not uniform because the bucked field gradient is asymmetric, with slightly higher flux at the lower end of the shield compared to the upper. The axial center of the shield is located at the field minimum.

The superconducting shield coils each comprise 100 turns of Supercon Inc. SC-T48B-M-0.7mm Cu-clad NbTi SC wire, with 0.43 mm diameter NbTi core (nominal), 0.70 mm diameter Cu cladding, and 0.75 mm diameter HML (Polyimide) insulation. The windings are on 9.3 cm ID brass hoop forms and potted with Stycast 2850 FT. After winding, the coil wire ends are trimmed to about 15 cm, and the last 5 cm of wire are annealed by heating to red heat using a MAPP gas torch; this also burns off the Polyimide insulation. The heated regions are cleaned and smoothed using 400 grit alumina emery paper. The wire ends were bonded [15] to make a persistent superconducting connection by use of a Koldweld KBM-9 wire coldwelding machine; the dies we had available were too large, so the outer copper cladding was built up through acid copper electroplating. After electroplating, the ends were trimmed back 1.6 mm, and the wire ends bonded. A micrograph analysis of an axially mechanically sectioned bond shows that the NbTi material flowed together at the center and pushed the copper uniformly in a radial direction from the center.

A similarly wound coil, with 24 cm OD, is clamped to the bottom of the DR mixing chamber plate. The fringe field of the 9 T magnet is about 300 G in this region; without the coil this is large enough to affect operation of the shielded cryogenic circulators (see section 2.7).

The ferromagnetic shield comprises a nested pair of cylinders made of 1.5 mm thick Amumetal 4 K, annealed per Amunreal's proprietary process. The bottoms of the can have welded-on disks, while the tops are closed with tight-fitting lids that overlap 1.25 cm of the cylinder walls. The outer shield (FS1) has length 15.2 cm and OD 8.9 cm, while the inner shield (FS2) has length 13.3

457 cm and OD 7.6 cm. Both top lids have holes for the tuning rod (1.25 cm) and
458 also for the JPA thermal link and transmission line. There are also three small
459 holes, placed at 120° in both the top and bottom, that allow Kevlar tuning lines
460 to pass through both shields.

461 Between the two ferromagnetic shields is a Pb superconducting shield, fabri-
462 cated from 1.6 mm Pb sheet, which fits closely on the outside of FS2 and has a
463 removable close-fitting lid. The Pb shield can is about 12.0 cm long, and covers
464 FS2 to just below the lid. The Pb lid is about 3.8 cm long, giving an overlap
465 with the cylinder of 2.5 cm. The inside of the Pb can is coated with clear acrylic
466 paint to avoid electrical contact and thermal currents, and is glued to FS2 using
467 Loctite 680.

468 The final passive shield component is a 0.13 mm thick Nb sheet (99.8 %
469 purity, Alfa-Aesar) that lines the inner surface of FS2. This sheet forms an
470 “open cylinder” in that it has no endcaps, and there is an effective axial slit
471 along its length which allows flux that would otherwise be trapped to escape.
472 The sheet is long enough so that the ends overlap by about 1.25 cm, and the
473 sheet is coated with acrylic paint, again to prevent the formation of electrical
474 contacts. The sheet is glued to the inner surface using Loctite. This sheet
475 enforces a boundary condition that the magnetic field must be axially homoge-
476 neous along its surface. The ferromagnetic boundary condition at the bottom
477 and at the lid of FS2 is that the magnetic field is perpendicular to the surface,
478 which matches perfectly the boundary condition imposed by the Nb sheet. The
479 combined effect of the superconducting and ferromagnetic boundary conditions
480 is to substantially reduce field gradients at the JPA.

481 A G-10 disk and nylon standoffs are used to support the JPA at the axial
482 center of the shield, with a radial offset to allow the tuning rod to pass. The
483 standoffs are screwed to 8-32 studs that extend from OFHC Cu spacers between
484 the two ferromagnetic shields; the cooling of the shield assembly is through these
485 studs and spacers which clamp the shields to each other and to a Cu plate affixed
486 to the gantry. The JPA itself is thermally linked to the gantry via a Cu braid
487 that exits through the top of the FS2/Pb shield assembly.

488 A final active layer of magnetic shielding is provided by feedback to the JPA
 489 flux bias coil, discussed in section 2.7.

490 2.7. Receiver and microwave measurements

491 The cryogenic microwave layout and room-temperature microwave/IF lay-
 492 out are depicted schematically in Fig. 10 and Fig. A.14 in the appendix, re-
 493 spectively. The receiver signal path is shown in blue in both diagrams. This
 494 section will describe the microwave portion of the receiver, the RF input lines,
 495 and other detector functionality which is best understood in reference to Fig. 10
 496 and Fig. A.14. The IF part of the receiver chain is discussed in section 2.8 below.

Figure 10: The cryogenic microwave layout. Blue arrows indicate the receiver signal path from the cavity to room temperature; black arrows indicate other paths used for network analysis, noise calibration, and JPA biasing. Component part numbers and manufacturers are listed in Table A.1, in the Appendix.

497 The input of the receiver is the cryogenic microwave switch S1, which may be
 498 toggled to connect the receiver to either the cavity or a termination thermally
 499 anchored to the DR still plate at $T_H = 775$ mK. This switch allows us to
 500 calibrate the receiver’s added noise *in situ* via a procedure described in detail
 501 in section 2.9.

502 The JPA, playing the role of the preamplifier, comprises the magnetically
 503 shielded JPA circuit described in section 2.5, a directional coupler for the pump
 504 tone input, and a commercial ferrite circulator to separate input and output
 505 signals. Two other circulators are required to isolate the JPA from both the
 506 backaction of the second-stage amplifier and its own backaction in reflection from
 507 the cavity. Signals exiting the JPA are amplified further by a HEMT amplifier
 508 at 4 K and another low-noise transistor amplifier at room temperature, and then
 509 routed both to a commercial vector network analyzer (VNA) and to an IQ mixer
 510 (M2) serving as the input of a homemade spectrum analyzer. The former path
 511 is used for cavity and receiver characterization, the latter for periodic noise
 512 calibrations and the cavity noise measurements constituting the axion search
 513 dataset.

514 Swept tones produced by the VNA may be directed via software-controlled
 515 switches S2 to any of three fridge input lines, through which they are transmitted
 516 through the cavity, reflected off the cavity, or sent directly to the JPA, bypassing
 517 the cavity. The first two lines are used to measure the cavity parameters ν_0 ,
 518 Q_L , and β ; the third line is used for JPA biasing and gain measurements, as
 519 discussed in section 2.5. Room-temperature attenuators were chosen to equalize
 520 sweep power incident on the JPA through each path. Attenuators at 4 K and
 521 base temperature reduce room-temperature thermal noise and the phase noise
 522 of the JPA pump generator to $\sim$ mK contributions we can safely ignore.

523 0.085'' NbTi/NbTi coaxial cables are used in the receiver signal path be-
 524 tween base temperature and 4 K; these cables (marked SC in Fig. 10) remain
 525 superconducting in the 9 T field due to flux pinning. Stainless 0.085'' coax is
 526 used between room temperature and 4 K in all four lines and down to base tem-
 527 perature in the input lines. All four coaxial lines are thermalized at each stage

528 of the fridge with gold-plated copper clamps; the output line inner connectors
 529 are thermalized via a bias tee at 4 K with its DC input shorted to ground. NiCr
 530 resistors are used in all cryogenic attenuators and terminations.

531 The room-temperature microwave chain includes three signal generators
 532 from Keysight in addition to the VNA. An 8340B serves as the local oscillator
 533 (LO) for both the spectrum analyzer system and the flux feedback system dis-
 534 cussed below. An E8257D with an ultra-low phase noise option (-120 dBc/Hz
 535 at 100 kHz detuning from the carrier) provides the pump tone for the JPA.
 536 Finally an N5183B is used to inject synthetic axion-like signals into the cavity
 537 transmission line. All three generators, along with the VNA and the digitizer
 538 board used in the spectrum analyzer, lock to a common 10 MHz reference pro-
 539 vided by a Stanford Research Systems (SRS) FS725m rubidium source (not
 540 pictured).

541 To generate synthetic axion signals with the expected linewidth (~ 5 kHz
 542 for a $25 \mu\text{eV}$ axion), we inject band-limited white noise into the FM port of
 543 the N5183B; the linewidth is then controlled by the modulation depth. The
 544 synthetic axion power in the cavity is only known to ± 3 dB due to unknown
 545 cryogenic insertion losses of individual components. This uncertainty prevents
 546 us from using the synthetic axion system to precisely calibrate our sensitivity,
 547 but it is still useful as a way to validate both the data acquisition and analysis
 548 procedures. Our data acquisition software includes functionality for blind in-
 549 jection of axion signals at random frequencies throughout a data run; Fig. 11
 550 illustrates one such injection in our first run, with a nominal intracavity power
 551 level of -190 dBm.

552 The room-temperature signal paths used exclusively by the JPA flux feed-
 553 back system mentioned at the end of section 2.6 are shown in pink in Fig. A.14.
 554 The JPA flux bias is modulated at 26 Hz by adding an oscillating current (con-
 555 trolled by an Agilent 33220A function generator) to the DC bias current set in
 556 software. The modulation frequency is limited by eddy current shielding in the
 557 Cu JPA enclosure; the amplitude is set to yield a modulation depth of ~ 0.1
 558 flux quantum. When the DC flux bias is offset from the value that maximizes

Figure 11: (a) 100 Hz-resolution power spectra from individual 15-minute integrations around a frequency at which a synthetic axion signal was injected. (b) The optimally weighted combined spectrum, still at 100 Hz resolution. (c) The combined spectrum after rebinning to 5 kHz with weighting that takes into account the axion lineshape; the synthetic axion is clearly visible. See [16] for an overview of the analysis procedure, which will also be described more thoroughly in a forthcoming publication.

the JPA gain at fixed pump power and frequency, the JPA gain is modulated at the same frequency as the flux with phase determined by the sign of the offset. When the bias modulation is centered on the optimal value, only higher harmonics are present in the gain variation. Thus the JPA gain variation at the modulation frequency may be used as an error signal in an analog feedback system to stabilize the flux bias.

During noise measurements we use the VNA to inject a weak CW probe tone through the pump line at 30 kHz detuning from the pump. The function of the feedback circuitry in the upper left corner of Fig. A.14 is to provide an LO at the appropriate IF frequency to bring the probe tone to DC on the IF side of the mixers M3. Both quadrature outputs are amplified and squared using AD633 analog multipliers; the sum of squared signals is a measure of the

received power in the probe tone and thus of the JPA gain. A Stanford Research Systems SR510 lock-in amplifier measures the variation of the probe tone power at the modulation frequency. By adding the SR510 output to the bias current we obtain a simple proportional feedback loop; the SR510 gain, phase, and filtering are chosen to provide a stable feedback signal. During JPA biasing and sweep measurements the probe tone is not present because the VNA is otherwise occupied, so both modulation and feedback are interrupted by switching off the signal and reference outputs of the 33220A in software.

2.8. Data acquisition and operations

The spectrum analyzer comprises the discrete components between M2 and ADC in Fig. A.14, along with LabVIEW code that computes FFTs, implements image rejection, and averages power spectra in parallel with timestream data acquisition. During all noise measurements the LO for M2 is set 780 kHz above the TM_{010} mode frequency, the JPA pump is set 820 kHz below the mode, and both IF channels are sampled by the GaGe Oscar CSE4344 PCIe digitizer board at 25 MS/s.

Each IF channel is sensitive not only to the RF frequencies of interest in the lower sideband of the LO but also to unwanted image frequencies in the upper sideband. The 90° relative phase shift between the two otherwise identical IF outputs of an IQ mixer may be exploited for image rejection: adding the Q output to the I output with a $+(-)90^\circ$ phase shift suppresses the upper (lower) LO sideband. This is the operating principle of commercial image reject mixers, which only work at particular IF frequencies because they implement the required 90° phase shift in hardware.

We implement image rejection in the frequency domain in software: this scheme works at any IF frequency, with the only possible drawback being that amplitude and phase mismatches in the discrete IF components in the two channels can limit the degree of image rejection. By taking FFTs of both the I and Q channels and defining $X(\omega) = (\text{Re}[I(\omega)] - \text{Im}[Q(\omega)]) + i(\text{Im}[I(\omega)] + \text{Re}[Q(\omega)])$ we obtain rejection of the upper sideband better than 20 dB at all IF frequencies

601 of interest in the power spectrum $|X(\omega)|^2$.

602 More specifically, 14-bit ADCs on the GaGe board digitize both channels
603 simultaneously in 5 s segments, then transfer each segment to PC RAM. The
604 total data in each segment across both IF channels is 437.5 MB, and the time
605 required to transfer this data to RAM is 1.2 s, which caps the data acquisition
606 efficiency at 80%. The *in situ* processing code divides each segment into 500 non-
607 overlapping 10 ms records in each channel, computes the FFT of each record
608 with no windowing, combines the I and Q FFTs corresponding to the same
609 10 ms time slice to implement image rejection, constructs a power spectrum
610 from each sample of $X(\omega)$, and averages all 500 power spectra corresponding to
611 each segment. All processing for each 5 s segment occurs in RAM in parallel
612 with the acquisition of the next segment. At the end of the data acquisition
613 period (typically 15 minutes), the power spectra from all segments are averaged
614 together to obtain a single spectrum obtained from $100 \text{ Hz} \times 15 \text{ minutes} = 9 \times 10^4$
615 averages, which is written to disk.

616 Thus the output at each tuning step is a heavily averaged power spectrum
617 with 100 Hz resolution extending from DC to the 12.5 MHz Nyquist frequency.
618 The usable IF bandwidth extends to $\sim 2.5 \text{ MHz}$, limited by the bandwidth of
619 the low-pass filters F2. The sampling rate is set so far above this to eliminate
620 the need for a high-order filter with significant passband ripple on the output of
621 the IF amps A2. The function of the F1 filters is to attenuate RF leakage into
622 the IF chain.

623 The 100 Hz spectral resolution is much smaller than the $\sim 5 \text{ kHz}$ expected
624 linewidth of a virialized axion signal – thus data from this detector can be
625 used to set more stringent limits on the abundance of non-virialized axions
626 with $Q_a \lesssim 10^7$ without any additional hardware. A technical advantage of the
627 narrow resolution is that it enables us to better identify and eliminate indi-
628 vidual IF bins contaminated by narrowband interference. Empirically some of
629 these “IF spikes” are associated with ground loops and others are due to room-
630 temperature electronics. The DC blocks and baluns in Fig. A.14 were added
631 in commissioning to eliminate as many of the IF spikes as possible; remaining

spikes are flagged and removed as part of the analysis procedure.

Figure 12: Diagram illustrating the IF setup for the experiment. Note that the IF frequency axis is reversed relative to the RF frequencies of these features. The JPA gain profile and TM_{010} Lorentzian profile are plotted using real data and a fit to real data, respectively. Both plots have logarithmic y axes; the absolute scale of the sensitivity plot is arbitrary. The red dot-dashed arrow indicates the probe tone created on the opposite side of the pump by the JPA’s intermodulation gain. Images of the axion-sensitive Fourier components around the cavity, created on the other side of the pump by the same process, are omitted in the diagram for clarity.

Fig. 12 is a simplified diagram of the IF setup showing the most relevant features discussed above, and illustrating that we further limit our analysis to 1.302 MHz centered on the cavity. The width of this analysis band is roughly twice the maximum cavity linewidth; axion conversion at larger detunings from the cavity mode frequency contributes negligibly to the SNR. The LO frequency must lie outside the analysis band to avoid superimposing different analysis band Fourier components, and the pump frequency must lie outside the analysis band for the JPA to operate in the desired phase-insensitive mode (see section 2.5). The LO and pump frequencies also cannot be equal, or else the I and Q channels will pick out (not necessarily equal) linear combinations of the JPA’s amplified and squeezed quadratures, and image rejection will fail.

These constraints indicate that the analysis band should be centered roughly halfway between the LO and pump in each spectrum. If the analysis band were too close to the pump, the spectrum would be contaminated by both the feedback probe tone and the pump tone’s phase noise. If the analysis band were too close to DC, $1/f$ noise would dominate, and the relative contribution of the

649 4 K HEMT amplifier to the receiver’s added noise would grow with detuning
 650 from the pump frequency. The precise positioning of the analysis band was
 651 tweaked to exclude bins in which IF interference was most persistent.

652 Data acquisition for the experiment is fully automated and controlled by a
 653 LabVIEW program. At the beginning of each iteration, the DAQ program tunes
 654 the TM_{010} mode, measures ν_0 and Q_L with a VNA sweep through the cavity
 655 transmission line, then sets the LO and pump frequencies based on the new value
 656 of ν_0 . It optimizes the JPA gain as described in section 2.5, measures the gain
 657 profile over 5 MHz centered on the pump, turns on the flux feedback system,
 658 and makes the 15-minute power spectrum measurement described above. At the
 659 end of the power spectrum measurement, the JPA gain and cavity transmission
 660 are each measured again, to identify drifts of the cavity mode associated with
 661 mechanical relaxation (see section 2.3) and unusually large bias flux jumps that
 662 the feedback system was unable to correct for; roughly 0.6% of spectra were cut
 663 from our initial analysis due to anomalous flux or frequency drifts. Finally, β is
 664 measured with a reflection sweep. Noise calibrations (see section 2.9) are also
 665 performed intermittently to constrain variation of T_S throughout the run. The
 666 average live time efficiency during the run is 72%.

667 All sweep data is saved to disk along with the 100 Hz averaged power spec-
 668 trum and critical parameter values such as the LO frequency; this amounts to
 669 about 3 MB per iteration. Data is stored locally at Yale and also transferred
 670 to a remote server at Berkeley for back-up and long term storage. The offline
 671 analysis procedure is outside of the scope of this paper, and will be the subject
 672 of a forthcoming publication.

673 *2.9. Noise calibration*

674 A simple way to measure the added noise of any microwave receiver is via a
 675 Y-factor measurement, in which we connect the receiver input to two $50\ \Omega$ loads
 676 at known temperatures T_C and T_H , and measure the hot/cold noise power ratio

$$Y = \frac{P_H}{P_C} = \frac{N_H + N_A}{N_C + N_A}, \quad (8)$$

where $N_C(N_H)$ is the thermal noise from the cold (hot) load including the zero-point contribution, and N_A is the receiver's added noise. We can then solve Eq. (8) for N_A to obtain

$$N_A = \frac{N_H - Y N_C}{Y - 1}. \quad (9)$$

In our experiment, N_H comes from a 775 mK termination, and N_C comes from the cavity (on resonance) or from a terminated port on the directional coupler in the reflection input line, reflected off the cavity (off resonance; see Fig. 10). N_A includes the added noise of the JPA preamplifier and also the added noise of subsequent amplifiers referred to the JPA input; of those, in practice only the 4 K HEMT contribution is non-negligible.

Eq. (3) indicates that the effect of imperfect power transmission efficiency η between the cavity and the JPA is to rescale the total noise by $1/\eta$. It is convenient to define $\eta = \eta_0 \eta_1$, with η_0 between the cavity and S1 and η_1 between S1 and the JPA input. Then N_A measured by the Y-factor method will also include an additional additive term

$$N_\eta = \frac{1 - \eta_1}{\eta_1} (N_C + N'_A), \quad (10)$$

where N'_A indicates the sum of amplifier contributions. This term arises because the lossy elements that contribute to η_1 are at temperature T_C (see Fig. 10), and thus attenuate both the hot load noise and the axion signal but not the cold load noise. To derive Eq. (10) more formally it is important to note that the lossy elements also generate thermal noise given by $(1 - \eta_1)N_C$. Thus our measurement of N_A yields only the net effect of the JPA added noise, HEMT noise, and η_1 , not any of these components individually. The effect of η_0 on the SNR cannot be measured *in situ*; we estimate 0.6 dB from the long superconducting cable and connector losses, and take this loss into account in calculating our exclusion limits.

The general expression in Eq. (9) also must be modified to account for several specific features of our receiver. First, as noted in section 2.5, in our present

703 configuration the JPA's added noise is simply the thermal noise on the opposite
 704 side of the pump from the cavity, coupled into the analysis band via image
 705 gain. Thus, the JPA contribution to the added noise differs by a known amount
 706 in the two switch configurations: if we define N_A to be the added noise in
 707 cold load measurements, the added noise with the switch pointed at the hot
 708 load is $N_A + N_H - N_C$, since the loss and HEMT contributions do not change
 709 with temperature. Second, both sides of Eq. (9) are in principle functions of IF
 710 frequency within any given measurement; in particular the lower JPA gain at the
 711 low-IF-frequency end of the analysis band (see Fig. 12) implies less suppression
 712 of the HEMT noise and thus larger N_A . Third, actuating the switch affects
 713 the standing wave pattern due to slight impedance mismatches on the cryogenic
 714 transmission lines, resulting in a small change in the pump power delivered to
 715 the JPA. We rebias the JPA every time we actuate the switch, but small changes
 716 in the JPA gain and bandwidth between switch configurations are still possible.
 717 We take all three of these effects into account by generalizing Eq. (9) to

$$N_A(\nu) = \frac{2[A(\nu)N_H - Y(\nu)N_C]}{Y(\nu) - A(\nu)} + N_C, \quad (11)$$

718 where $A(\nu) = G_H(\nu)/G_C(\nu)$ is obtained from the measured hot/cold JPA gain
 719 profiles.

720 During commissioning of the experiment, we consistently obtained $N_A \simeq$
 721 1.35 quanta from Y-factor measurements far detuned from the cavity mode
 722 independent of RF frequency. This result is consistent with the added noise of
 723 an ideal nondegenerate JPA at T_C (0.63 quanta), $\simeq 0.2$ quanta input-referred
 724 HEMT noise, and an additional $\simeq 0.5$ quanta which we can attribute to $\simeq 2$ dB
 725 loss before the JPA. This is a plausible value result for η_1 , if a little on the high
 726 side.

727 In cold load noise measurements near the TM_{010} frequency we observe an
 728 additional Lorentzian excess $N_{\text{cav}}(\nu)$ centered on the IF frequency of the mode.
 729 Naively applying Eq. (11) to such measurements would treat the N_{cav} as a
 730 contribution to the added noise (because this expression assumes the input

Figure 13: A representative noise measurement. N_C (black dotted line) is obtained from thermometry, N_A (red dot-dashed line) is obtained from the average of off-resonance Y-factor measurements, and N_{cav} (blue dashed line) is from a single Y-factor measurement during the data run. N_{sys} (pink solid line) is the sum of these contributions.

731 noise in the cold load is completely characterized by the single number T_C), and
 732 therefore overestimate the effect on the total noise N_{sys} , since N_A appears in
 733 both the numerator and denominator of Eq. (8), whereas N_C appears only in
 734 the denominator.

735 Therefore, we make the ansatz that the added noise is the same on-resonance
 736 as off-resonance, which allows us to quantify the peak excess contribution $N_{cav}(\nu_0) \simeq$
 737 1 quantum. During detector commissioning we examined the dependence of
 738 N_{cav} on the mixing chamber plate temperature T_C , and observed that the ex-
 739 cess vanished for $T_C \gtrsim 550$ mK. This behavior indicates that N_{cav} has a thermal
 740 origin as opposed to being an artifact of feedback between the JPA and cavity
 741 mode. Studies of the dependence of N_{cav} on JPA gain and on β further disfavor
 742 any sort of feedback explanation, which validates the above ansatz.

743 Having established that N_{cav} is of thermal origin, the most likely culprit is
 744 a poor thermal link between the tuning rod and the cavity barrel, as the only
 745 thermal connection in the present design is through a thin-walled alumina shaft.

746 An improved thermal link through the axle of the tuning rod has been designed
 747 and is being implemented; it should bring the tuning rod into equilibrium with
 748 the mixing chamber and improve the linear coupling sensitivity by 20% ².

749 In practice, our *in situ* noise calibration procedure consists of the following
 750 steps, repeated for each switch configuration: we bias up the JPA to target gain,
 751 measure the gain profile, and take 5 s of noise data, from which we construct
 752 a power spectrum with 10 kHz resolution, for a total of 130 data points in
 753 the analysis band. We wait 2 minutes for thermal transients to die down after
 754 actuating the switch in either direction; the heat due to each switch actuation
 755 is 10 mJ.

756 From these measurements we obtain a typical total noise of $N_{\text{sys}} \simeq 3$ quanta
 757 on resonance, falling to $\simeq 2.2$ quanta at the edges of the analysis band. The
 758 overall fractional error $\delta N_{\text{sys}} \sim 6\%$ is dominated by $\sim 17\%$ uncertainty in N_{cav} ;
 759 δN_A is only $\sim 4\%$, and the change in N_C is negligible even allowing for a ± 20 mK
 760 calibration error in the mixing chamber thermometry. Individual contributions
 761 to N_{sys} are plotted in Fig. 13 for a representative measurement. This is the
 762 lowest noise demonstrated to date in a microwave cavity axion search.

763 3. Summary and Conclusions

764 We have demonstrated *in situ* near-quantum-limited noise performance in
 765 the first operation of this experiment, and achieved sensitivity to cosmologically
 766 relevant QCD axion models with $g_\gamma \geq 2.3 \times g_\gamma^{\text{KSVZ}}$ [16]. That sensitivity this
 767 close to the KSVZ line can be reached in an experiment of 1.5 L volume (c.f. the
 768 200 L volume of the ADMX detector [6, 9]) underscores the role of technology
 769 in the microwave cavity experiment. This platform is now poised to explore
 770 innovative concepts in both amplifiers and cavities to increase sensitivity and
 771 decrease scan time. A receiver based on injecting a squeezed state of vacuum
 772 into the cavity by one JPA and reading it out with another will be the first

²Subsequent to the review of this paper, this thermal link has been implemented and the thermal effect has indeed been substantially mitigated.

773 such innovation to be explored [14]. Beating the Standard Quantum Limit with
774 such a receiver has been demonstrated on the bench [17], but this will be the
775 opportunity to validate the feasibility of such a receiver in the rigors of an actual
776 operating environment.

777 On the cavity side, a photonic band gap resonator is being designed to
778 eliminate all interfering TE modes. As the TM mode of interest is tuned in
779 frequency, avoided crossings with TE modes block a significant fraction of fre-
780 quency coverage. This loss of coverage has been the bane of rapid and efficient
781 mass coverage to date. Such structures have been thoroughly explored in accel-
782 erator physics [18, 19]; their successful adaptation to the microwave cavity axion
783 experiment will largely pivot on an effective and reliable means of tunability.
784 Similarly, cavities are being explored adapting design principles of distributed
785 Bragg reflector structures, which attain much higher quality factors Q , providing
786 greater sensitivity, and thereby scan rate, for the experiment [20].

787 4. Acknowledgements

788 This work was supported under the auspices of the National Science Founda-
789 tion, under grants PHY-1067242, and PHY-1306729, the Heising-Simons Foun-
790 dation under grants 2014-181, 2014-182, and 2014-183, and the U.S. Department
791 of Energy by Lawrence Livermore National Security, LLC, Lawrence Livermore
792 National Laboratory under Contract DE-AC52-07NA27344. We thankfully ac-
793 knowledge the critical contributions by Matthias Bühler of Low Temperature
794 Solutions UG for design of and upgrades to the cryogenic system. MS is sup-
795 ported by the National Science Foundation Graduate Research Fellowship Pro-
796 gram under Grant No. DGE-1106400.

797 References

- 798 [1] P. Graham, I. Irastorza, S. Lamoreaux, A. Lindner, K. A. van Bib-
799 ber, Experimental Searches for the Axion and Axion-Like Particles, An-

- 800 nual Review of Nuclear and Particle Science 65 (2015) 485–514. doi:
801 10.1146/annurev-nucl-102014-022120.
- 802 [2] P. Sikivie, Experimental Tests of the “Invisible” Axion, Phys. Rev. Lett.
803 51 (1983) 1415. doi:10.1103/PhysRevLett.51.1415.
- 804 [3] P. Sikivie, Detection rates for “invisible”-axion searches, Phys. Rev. D 32
805 (1985) 2988–2991. doi:10.1103/PhysRevD.32.2988.
- 806 [4] L. Krauss, J. Moody, F. Wilczek, D. E. Morris, Calculations for cosmic
807 axion detection, Phys. Rev. Lett. 55 (1985) 1797–1800. doi:10.1103/
808 PhysRevLett.55.1797.
809 URL <http://link.aps.org/doi/10.1103/PhysRevLett.55.1797>
- 810 [5] S. Borsanyi, Z. Fodor, J. Guenther, K. H. Kampert, S. D. Katz, T. Kawanai,
811 T. G. Kovacs, S. W. Mages, A. Pasztor, F. Pittler, J. Redondo, A. Ring-
812 wald, K. K. Szabo, Calculation of the axion mass based on high-
813 temperature lattice quantum chromodynamics, Nature 539 (7627) (2016)
814 69–71.
815 URL <http://dx.doi.org/10.1038/nature20115>
- 816 [6] H. Peng, S. Asztalos, E. Daw, N. Golubev, C. Hagmann, D. Kinion,
817 J. LaVeigne, D. Moltz, F. Nezrick, J. Powell, L. Rosenberg, P. Sikivie,
818 W. Stoeffl, N. Sullivan, D. Tanner, M. Turner, K. van Bibber, Cryogenic
819 cavity detector for a large-scale cold dark-matter axion search, Nucl. In-
820 strum. Meth. A 444 (2000) 569–583.
- 821 [7] R. H. Dicke, The measurement of thermal radiation at microwave frequen-
822 cies, Review of Scientific Instruments 17 (7) (1946) 268–275.
- 823 [8] C. M. Caves, Quantum limits on noise in linear amplifiers, Phys. Rev. D
824 26 (8) (1982) 1817.
- 825 [9] S. Asztalos, G. Carosi, C. Hagmann, D. Kinion, K. van Bibber, M. Hotz,
826 L. Rosenberg, G. Rybka, A. Wagner, J. Hoskins, C. Martin, N. Sullivan,

827 D. Tanner, R. Bradley, J. Clarke, Design and performance of the ADMX
828 SQUID-based microwave receiver, Nucl. Instrum. Meth. A 656 (2011) 39–
829 44. doi:10.1016/j.nima.2011.07.019.

830 [10] D. Tanner, private communication (2011).

831 [11] C. Kittel, Introduction to Solid State Physics, 8th Edition, John Wiley &
832 Sons, Inc., New York, 2005.

833 [12] J. C. Slater, Microwave Electronics, Rev. Mod. Phys. 18 (1946) 441–512.
834 doi:10.1103/RevModPhys.18.441.
835 URL <http://link.aps.org/doi/10.1103/RevModPhys.18.441>

836 [13] C. Hagmann, P. Sikivie, N. Sullivan, D. Tanner, S. Choi, Cavity design for
837 a cosmic axion detector, Rev. Sci. Instrum. 61 (3) (1990) 1076–1085.

838 [14] H. Zheng, M. Silveri, R. T. Brierley, S. M. Girvin, K. W. Lehnert, Acceler-
839 ating dark-matter axion searches with quantum measurement technology
840 (Jul. 2016). arXiv:1607.02529.

841 [15] G. Luderer, P. Dullenkopf, G. Laukien, Superconducting joint between mul-
842 tifilamentary wires, Cryogenics 14 (9) (1974) 518–519.

843 [16] B. M. Brubaker, et al., First results from a microwave cavity axion search at
844 $24\ \mu\text{eV}$ (submitted to Physical Review Letters) (2016). arXiv:1610.02580.

845 [17] F. Mallet, M. Castellanos-Beltran, H. Ku, S. Glancy, E. Knill, K. Irwin,
846 G. Hilton, L. Vale, K. Lehnert, Quantum state tomography of an itinerant
847 squeezed microwave field, Phys. Rev. Lett. 106 (2011) 220–502. doi:10.
848 1103/PhysRevLett.106.220502.

849 [18] E. Smirnova, C. Chen, M. Shapiro, J. Sirigiri, R. Temkin, Simulation of
850 photonic band gaps in metal rod lattices for microwave applications, J.
851 Appl. Phys. 91 (2002) 960–968. doi:10.1063/1.1426247.

- 852 [19] E. Nanni, S. Lewis, M. Shapiro, R. Griffin, R. Temkin, Photonic-Band-Gap
853 Traveling-Wave Gyrotron Amplifier, Phys. Rev. Lett. 111 (2013) 235101.
854 doi:10.1103/PhysRevLett.111.235101.
- 855 [20] C. Flory, R. Taber, High Performance Distributed Bragg Reflector Mi-
856 crowave Resonator, IEEE Transactions on Ultrasonics, Ferroelectrics and
857 Frequency Control 44 (1997) 486–495. doi:10.1109/58.585133.
- 858 [21] G. E. Theriault, A. J. Leidich, T. H. Campbell, Application of the CA3018
859 Integrated-Circuit Transistor Array, Tech. Rep. AN5296.0, Intersil (2012).

Table A.1: Component part numbers.

	Label	Type	Supplier: Part #
Cryogenic			
	BT1	Bias tee	Mini-Circuits: ZX85-12G+ (ferrite removed)
	BT2	Bias tee	Anritsu: K250
	C	Circulator	Pamtech Inc.: CTH1184K18
	D	Directional coupler	Pasternack: PE2211-20
	HEMT	HEMT amplifier	Low Noise Factory: LNF-LNC4.8A
	S1	Switch	Radiall: R577443005
	SC	NbTi/NbTi coax	Keycom: UPJ07
Room-temperature			
	A1	RF amplifier	Miteq: AMF-4F-04001200-15-10P
	A2	IF amplifier	Homemade: based on Fig. 2 in [21]
	A3	IF amplifier	Mini-Circuits: ZFL-500LN-BNC+
	A4	IF amplifier	Stanford Research: SR445A
	A5	IF amplifier	Stanford Research: SR560
	AT1	Step attenuator	Agilent Technologies: 8496H
	AT2	Step attenuator	Agilent Technologies: 8494H
	B	Balun	North Hills: 0017CC
	DC	DC block	Inmet: 8039
	F1	Low-pass filter	Mini-Circuits: VLFX-80
	F2	Low-pass filter	Mini-Circuits: SLP-1.9+
	F3	Low-pass filter	Mini-Circuits: BLP-2.5+
	F4	Low-pass filter	Homemade: 3.39 kHz single pole RC filter
	I1	Isolator	Ditom Microwave: D314080
	I2	Double isolator	Ditom Microwave: D414080
	M1	Mixer	Marki Microwave: M1-0408
	M2	IQ mixer	Marki Microwave: IQ0307LXP
	M3	Mixer	Mini-Circuits: ZAD-8+
	PS1	Power splitter/combiner	Mini-Circuits: ZX10-2-71-s+
	PS2	Power splitter/combiner	Mini-Circuits: ZFRSC-2050+
	S2	Switch	Mini-Circuits: ZFSWA2-63DR+