

Large-scale microwave cavity search for dark-matter axions

S. Asztalos, E. Daw, H. Peng, and L. J. Rosenberg

Department of Physics and Laboratory for Nuclear Science, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, Massachusetts 02139

C. Hagmann, D. Kinion, W. Stoeffl, and K. van Bibber

Lawrence Livermore National Laboratory, Physics and Advanced Technology Directorate, 7000 East Avenue, Livermore, California 94550

P. Sikivie, N. S. Sullivan, and D. B. Tanner

Department of Physics, University of Florida, Gainesville, Florida 32611

F. Neznick

Fermi National Accelerator Laboratory, Batavia, Illinois 60510-0500

M. S. Turner

NASA/Fermilab Center for Astrophysics, Fermi National Accelerator Laboratory, Batavia, Illinois 60510-0500 and Departments of Astronomy & Astrophysics and Physics, Enrico Fermi Institute, The University of Chicago, Chicago, Illinois 60637-1433

D. M. Moltz and J. Powell

Nuclear Science Division, Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, California 94720

M.-O. André, J. Clarke, and M. Mück

Department of Physics, University of California and Lawrence Berkeley National Laboratory, Berkeley, California 94720

Richard F. Bradley

National Radio Astronomy Observatory, Charlottesville, Virginia 22903

(Received 11 April 2001; published 5 October 2001)

We have built and operated a large-scale axion detector, based on a method originally proposed by Sikivie, to search for halo axions. The apparatus consists of a cylindrical tunable high- Q microwave cavity threaded axially by a static high magnetic field. This field stimulates axions that enter the cavity to convert into single microwave photons. The conversion is resonantly enhanced when the cavity resonant frequency is near the axion rest mass energy. The experiment is cooled to 1.5 K and the electromagnetic power spectrum emitted by the cavity is measured by an ultra-low-noise microwave receiver. The axion would be detected as excess power in a narrow line within the cavity resonance. The apparatus has achieved a power sensitivity better than 10^{-23} W in the mass range 2.9–3.3 μeV . For the first time the rf cavity technique has explored plausible axion models, assuming axions make up a significant fraction of the local halo density. The experiment continues to operate and will explore a large part of the mass in the range of 1–10 μeV in the near future. An upgrade of the experiment is planned with dc superconducting quantum interference device microwave amplifiers operating at a lower physical temperature. This next generation detector would be sensitive to even more weakly coupled axions contributing only fractionally to the local halo density.

DOI: 10.1103/PhysRevD.64.092003

PACS number(s): 14.80.Mz, 95.35.+d, 98.35.Gi

I. AXIONS IN QCD, ASTROPHYSICS, AND COSMOLOGY**Introduction**

Many experiments, notably the negative searches for a neutron electric dipole moment [1], imply that CP symmetry is conserved to a remarkable degree in quantum chromodynamics (QCD). This fact is a puzzle in the standard model of elementary particles because CP symmetry is observed to be violated in the weak sector of the theory, e.g., in the mixing of the K_L and K_S mesons. As discussed below, one would expect CP violation in the weak sector to feed into the strong sector through the intermediary of the QCD θ angle.

A solution to this “strong CP problem” was proposed by Peccei and Quinn [2], which involves the spontaneous break-

ing of a $U_{PQ}(1)$ symmetry and a concomitant quasi-Nambu-Goldstone particle [3], called the axion. The axion solution to the strong CP problem is rich in experimental, observational, and cosmological implications. If its mass is of order 10^{-5} eV, the axion is a good candidate for the dark matter of the Universe. We have carried out an experiment that searches for Milky Way halo axions by their conversion into microwave photons in an electromagnetic cavity permeated by a strong magnetic field. The hardware, data analysis, and results of the experiment are described in Sec. II. A future upgrade is briefly discussed in Sec. III. The remainder of this Introduction describes the strong CP problem and its solution through the existence of an axion, the mechanisms by

which cold dark-matter axions are produced in the early universe, the cavity detector method, and the structure we expect to find in an axion signal.

Several detailed reviews of the theory of the axion and its cosmological and astrophysical implications are found in Ref. [4]. The experimental situation is reviewed in Ref. [5], which discusses methods and constraints on axions from laboratory, astrophysical, and cosmological searches.

1. The strong CP problem and the axion

Consider the Lagrangian of QCD [6]:

$$\mathcal{L}_{\text{QCD}} = -\frac{1}{4} G_{\mu\nu}^a G^{a\mu\nu} + \sum_{j=1}^n [\bar{q}_j \gamma^\mu i D_\mu q_j - (m_j q_{Lj}^\dagger q_{Rj} + \text{H.c.})] + \frac{\theta g^2}{32\pi^2} G_{\mu\nu}^a \tilde{G}^{a\mu\nu}. \quad (1)$$

The last term, the product of the gluon field strength tensor G with its dual, is a four-divergence and hence does not contribute in perturbation theory. That term does, however, contribute through nonperturbative effects [7] associated with QCD instantons [8]. Using the Adler-Bell-Jackiw anomaly [9], one can show that QCD must be θ dependent if none of the current quark masses vanish. If the θ dependency were absent, QCD would have a $U_A(1)$ symmetry and would predict the mass of the η' pseudoscalar meson to be less than $\sqrt{3}m_\pi \approx 240$ MeV [10], contrary to observation. One can further show that QCD depends upon θ only through the difference of θ and the argument of the quark mass matrix:

$$\bar{\theta} = \theta - \arg(m_1, m_2, \dots, m_n). \quad (2)$$

If $\bar{\theta} \neq 0$, QCD violates P and CP . The absence of a P and CP violation in strong interactions, therefore places an upper limit upon $\bar{\theta}$. The best constraint follows from the experimental bound [1] on the neutron electric dipole moment, which yields $\bar{\theta} < 10^{-9}$.

The question then is: why is $\bar{\theta}$ so small? In the standard model of particle interactions, the quark masses originate in the electroweak sector of the theory, which violates P and CP . There is no reason why the overall phase of the quark mass matrix should match the value of θ from the QCD sector to yield one part in 10^9 or better. In particular, if CP violation is introduced in the manner of Kobayashi and Maskawa [11], the Yukawa couplings that give masses to the quarks are arbitrary complex numbers and hence $\arg \det m_q$ and $\bar{\theta}$ are expected to be of order one. The puzzle of why $\bar{\theta} < 10^{-9}$ is called the ‘‘strong CP problem.’’

Peccei and Quinn [2] proposed a solution that postulates the existence of a global $U_{\text{PQ}}(1)$ quasisymmetry. $U_{\text{PQ}}(1)$ must be a symmetry of the theory at the classical (i.e., at the Lagrangian) level, it must be broken explicitly by those nonperturbative effects that make the physics of QCD depend upon θ , and finally it must be spontaneously broken. The

axion [3] is the associated quasi-Nambu-Goldstone boson. One can show that if a $U_{\text{PQ}}(1)$ quasi-symmetry is present, then

$$\bar{\theta} = \theta - \arg(m_1, \dots, m_n) - \frac{a(x)}{f_a}, \quad (3)$$

where $a(x)$ is the axion field and f_a , called the axion decay constant, is of order the vacuum expectation value (VEV) which spontaneously breaks $U_{\text{PQ}}(1)$. It can further be shown [12] that the nonperturbative effects that make QCD depend upon $\bar{\theta}$ produce an effective potential $V(\bar{\theta})$ whose minimum is at $\bar{\theta}=0$. Thus $\bar{\theta}$ is allowed to relax to zero dynamically and the strong CP problem is solved.

The existence of an axion is the signature of the PQ solution to the strong CP problem. Its properties can be derived using the methods of current algebra [13]. The axion mass is given in terms of f_a by

$$m_a \approx 6 \mu\text{eV} \frac{10^{12} \text{ GeV}}{f_a}. \quad (4)$$

All the axion couplings are inversely proportional to f_a . Of particular interest here is the axion coupling to two photons:

$$\mathcal{L}_{a\gamma\gamma} = g_\gamma \frac{\alpha}{\pi} \frac{a(x)}{f_a} \vec{E} \cdot \vec{B}, \quad (5)$$

where \vec{E} and \vec{B} are the electric and magnetic fields, α is the fine structure constant, and g_γ is a model-dependent coefficient of order one. $g_\gamma = 0.36$ in the Dine-Fischler-Srednicki-Zhitnitskii (DFSZ) model [14] whereas $g_\gamma = -0.97$ in the Kim-Shifman-Vainshtein-Zakharov (KSVZ) model [15]. *A priori*, the value of f_a , and hence that of m_a , is arbitrary. However, negative searches for the axion in high energy and nuclear physics experiments [4], combined with astrophysical constraints [16], rule out $m_a \gtrsim 10^{-3}$ eV. In addition, as is discussed below, cosmology places a lower limit on m_a of order 10^{-6} eV from the requirement that axions do not overclose the universe.

2. Production of cold relic axions

The implications of an axion for the history of the early universe may be briefly described as follows. At a temperature of order f_a , a phase transition occurs in which the $U_{\text{PQ}}(1)$ symmetry becomes spontaneously broken. This is called the PQ phase transition. At these temperatures, the nonperturbative QCD effects, which produce the effective potential $V(\bar{\theta})$ are suppressed [17], the axion is massless, and all values of $\langle a(x) \rangle$ are equally likely. Axion strings appear as topological defects. Subsequently, one must distinguish two cases: (1) inflation occurring with reheat temperature less than the PQ transition temperature or (2) inflation occurring with reheat temperature higher than the PQ transition temperature (equivalently, for our purposes, inflation does not occur at all). In case (1) the axion field gets homogenized by inflation and the axion strings are diluted away, whereas in case (2) axion strings are present from the PQ transition to the QCD epoch.

When the temperature approaches the QCD scale, the potential $V(\bar{\theta})$ turns on and the axion acquires mass. There is a critical time, defined by $m_a(t_1)t_1=1$, when the axion field starts to oscillate in response to the axion mass turn on [18]. The corresponding temperature $T_1 \approx 1$ GeV. In case (1), where the axion field has been homogenized by inflation, the initial amplitude of this oscillation depends on how far from zero the axion field is at t_1 . The axion field oscillations do not dissipate into other forms of energy and hence contribute to the cosmological energy density today [18]. This contribution, called “vacuum realignment,” is the only contribution in case (1). In terms of the critical density $\rho_c = (3H_0^2)/(8\pi G)$, the energy density in axions is

$$\Omega_a = \frac{\rho_a(t_0)}{\rho_c} \approx \frac{1}{6} \alpha^2(t_1) \left(\frac{f_a}{10^{12} \text{ GeV}} \right)^{7/6} \left(\frac{0.7}{h} \right)^2 \quad (\text{case 1}), \quad (6)$$

where h parametrizes the present Hubble rate $H_0 = h \times 100$ km/s Mpc and $\alpha(t_1) = a(t_1)/f_a$ is the initial misalignment angle. Note that Ω_a may be accidentally suppressed in case (1), if the homogenized axion field happens to lie close to zero.

In case (2) the axion strings radiate axions [19,20] from the time of the PQ transition until t_1 . At t_1 , each string becomes the boundary of N domain walls. If $N=1$, the network of walls bounded by strings is unstable [21,22] and decays away. If $N>1$, there is a domain wall problem [23] because axion domain walls end up dominating the energy density, resulting in a universe very different from the one observed today. Henceforth, we assume $N=1$.

There are three contributions to the axion cosmological energy density in case (2). One contribution is from vacuum realignment [18]. The vacuum realignment contribution is, in case (2),

$$\Omega_a^{\text{vac}} \approx \frac{1}{3} \left(\frac{f_a}{10^{12} \text{ GeV}} \right)^{7/6} \left(\frac{0.7}{h} \right)^2. \quad (7)$$

In case (2), the vacuum realignment contribution cannot be accidentally suppressed because it is an average over many horizon volumes at QCD time, each with a causally independent value of the initial misalignment angle $\alpha(t_1)$. Note also that the vacuum realignment contribution is larger, by approximately a factor of 2, in case (2) than it is in case (1) with $\alpha(t_1) \approx 1$. This is because only the zero momentum mode contributes in case (1), whereas in case (2) there are contributions from the zero momentum mode and from higher modes [30]. A second contribution is from axions that were produced in the decay of walls bounded by strings after t_1 [24,28–30]. The contribution from wall decay is [30]

$$\Omega_a^{\text{d.w.}} \approx \frac{2}{\gamma} \left(\frac{f_a}{10^{12} \text{ GeV}} \right)^{7/6} \left(\frac{0.7}{h} \right)^2, \quad (8)$$

where $\gamma \equiv \langle \omega \rangle / m_a$ is the average Lorentz factor of the axions produced in the decay of walls bounded by strings. In simulations, it was found [30] that $\gamma \sim 7$ for $\ln(f_a/m_a) \approx 6$, and that

γ increases approximately linearly with $\ln(f_a/m_a)$. Extrapolation of this behavior to the parameter range of interest, $\ln(f_a/m_a) \approx 60$, yields $\gamma \sim 60$, suggesting that the wall decay contribution is less than the vacuum realignment contribution.

A third contribution [19,20,24–27] is from axions that were radiated by axion strings before t_1 .

The string decay contribution has been the most controversial, with three estimates in the literature. The result of Battye and Shellard [25] is

$$\Omega_a^{\text{str,BS}} \approx 10.7 \left(\frac{\xi}{13} \right) [(1 + \alpha/\kappa)^{3/2} - 1] \frac{1}{h^2} \left(\frac{f_a}{10^{12} \text{ GeV}} \right)^{7/6}, \quad (9)$$

where we have put in the dependence on ξ explicitly. ξ parametrizes the average distance between axion strings at time t as ξt . The ratio of parameters α/κ is expected [25] to be in the range $0.1 < \alpha/\kappa < 1.0$. Battye and Shellard obtain the estimate $\xi \approx 13$ from their simulations of local string networks in an expanding universe. However, axion strings are global strings, which are different from local strings. Global strings have most of their energy in the Nambu-Goldstone field surrounding the string core, whereas local strings have most of their energy inside the string core. For this reason, global string networks may behave differently from local string networks (see below). Regardless, for the expected range of α/κ ,

$$\Omega_a^{\text{str,BS}} = (3.4 \text{ to } 40) \left(\frac{\xi}{13} \right) \left(\frac{f_a}{10^{12} \text{ GeV}} \right)^{7/6} \left(\frac{0.7}{h} \right)^2. \quad (10)$$

Yamaguchi *et al.* [26] carried out simulations of *global* string networks in an expanding universe and find $\xi = 1.00 \pm 0.08$. Their result for the string decay contribution to the cosmological axion energy density is

$$\Omega_a^{\text{str,YKY}} = (1.4 \pm 0.94) \left(\frac{f_a}{10^{12} \text{ GeV}} \right)^{7/6} \left(\frac{0.7}{h} \right)^2. \quad (11)$$

Hagmann *et al.* [24,27] carried out simulations of the motion and decay of single string loops and of oscillating bent strings, focusing their effort on obtaining a reliable estimate of the spectrum of axions radiated by strings. They assume $\xi \approx 1$. Their result is

$$\Omega_a^{\text{str,HCS}} \approx 0.27 \ 2^{\pm 1} \left(\frac{f_a}{10^{12} \text{ GeV}} \right)^{7/6} \left(\frac{0.7}{h} \right)^2. \quad (12)$$

The last two results are consistent with one another. The result of Battye and Shellard is compatible with the other two if one sets $\xi \approx 1$ instead of $\xi \approx 13$. Since axion strings are global rather than local, the global string network simulations of Yamaguchi *et al.* could be considered more reliable in determining ξ . With $\xi = 1$, all estimates are consistent and suggest

$$\Omega_a = \Omega^{\text{vac}} + \Omega^{\text{d.w.}} + \Omega^{\text{str}} = (0.5 \text{ to } 3.0) \times \left(\frac{f_a}{10^{12} \text{ GeV}} \right)^{7/6} \left(\frac{0.7}{h} \right)^2 \quad (\text{case 2}). \quad (13)$$

Equations (4), (7), and (13) indicate that the mass range we plan to cover in our experiment, $1.3 \mu\text{eV} < m_a < 13 \mu\text{eV}$, is prime hunting ground for dark-matter axions, with the high end favored in case (2), and the low end in case (1).

From the point of view of experimental design, because we do not know which cosmological scenario pertains, we must be prepared to search from the lowest estimated over-closure bound ($\sim 10^{-6}$ eV) to the upper bound set by SN 1987a ($\sim 10^{-3}$ eV). We note that in any cosmological scenario, Ω_a increases as m_a decreases, which motivates a search beginning from the lowest possible mass and proceeding upwards.

The axions produced when the axion mass turns on at the QCD phase transition, whether from string decay, vacuum realignment, or wall decay, have momenta $p_a \sim 1/t_1 \sim 10^{-8}$ eV/c when the surrounding plasma has temperature $T_1 \approx 1$ GeV. They are nonrelativistic from the moment of their first appearance, and because of their negligible couplings to normal matter and radiation are therefore a form of cold dark matter (CDM). Studies of large-scale structure formation support the view that the dominant fraction of dark matter is CDM. Moreover, any form of CDM necessarily contributes to galactic halos by falling into the gravitational wells of galaxies. Hence, there is a special opportunity to search for axions by direct detection on Earth.

3. The Sikivie microwave cavity dark-matter detector

Axions can be detected by stimulating their conversion to photons in a strong magnetic field [31]. The relevant coupling is given in Eq. (5). In particular, an electromagnetic cavity permeated by a strong static magnetic field can detect galactic halo axions. These halo axions have velocities β of order 10^{-3} and hence their energies $E_a = m_a + \frac{1}{2} m_a \beta^2$ have a spread of order 10^{-6} above the axion mass. When the frequency $\omega = 2\pi f$ of a cavity mode equals m_a , galactic halo axions convert resonantly into quanta of excitation (photons) of that cavity mode. The power from axion-to-photon conversion on resonance is found to be [31,32]

$$\begin{aligned} P &= \left(\frac{\alpha}{\pi} \frac{g_\gamma}{f_a} \right)^2 V B_0^2 \rho_a C \frac{1}{m_a} \min(Q_L, Q_a) \\ &= 0.5 \times 10^{-26} \text{ W} \left(\frac{V}{5001} \right) \left(\frac{B_0}{7 \text{ T}} \right)^2 C \left(\frac{g_\gamma}{0.36} \right)^2 \\ &\quad \times \left(\frac{\rho_a}{1/2 \times 10^{-24} \text{ g/cm}^3} \right) \\ &\quad \times \left(\frac{m_a}{2\pi(\text{GHz})} \right) \min(Q_L, Q_a), \end{aligned} \quad (14)$$

where V is the volume of the cavity, B_0 is the magnetic-field strength, Q_L is its loaded quality factor, $Q_a = 10^6$ is the

“quality factor” of the galactic halo axion signal (i.e., the ratio of their energy to their energy spread), ρ_a is the density of galactic halo axions on Earth, and C is a mode dependent form factor given by

$$C = \frac{\left| \int_V d^3x \vec{E}_\omega \cdot \vec{B}_0 \right|^2}{B_0^2 V \int_V d^3x \epsilon |\vec{E}_\omega|^2}, \quad (15)$$

where $\vec{B}_0(\vec{x})$ is the static magnetic field, $\vec{E}_\omega(\vec{x})e^{i\omega t}$ is the oscillating electric field of the mode in question, and ϵ is the dielectric constant. For a cylindrical cavity and a homogeneous longitudinal magnetic field, $C=0.69$ for the TM_{010} mode. The form factors of the other TM_{0n0} modes are much smaller, and form factors for pure transverse electric (TE), TEM, and the remaining transverse magnetic (TM) modes are zero.

Because the axion mass is only known in order of magnitude at best, the cavity must be tunable and a large range of frequencies must be explored in seeking a signal. The cavity can be tuned by moving a dielectric rod or metal post inside it. Using Eq. (14), one finds that to perform a search with signal-to-noise ratio (SNR), the maximum scanning rate is

$$\begin{aligned} \frac{df}{dt} &= \frac{12 \text{ GHz}}{\text{year}} \left(\frac{4}{\text{SNR}} \right)^2 \left(\frac{V}{5001} \right)^2 \left(\frac{B_0}{7 \text{ T}} \right)^4 C^2 \left(\frac{g_\gamma}{0.36} \right)^4 \\ &\quad \times \left(\frac{\rho_a}{1/2 \times 10^{-24} \text{ g/cm}^3} \right)^2 \left(\frac{3K}{T_n} \right)^2 \left(\frac{f}{\text{GHz}} \right)^2 \frac{Q_L}{Q_a}, \end{aligned} \quad (16)$$

where T_n is the sum of the physical temperature of the cavity plus the noise temperature of the microwave receiver that detects the photons from axion-to-photon conversion.

Two earlier experiments [33], based on rf cavity axion detection, reported limits on halo axion couplings in the axion mass window. However, these earlier experiments lacked the sensitivity, by factors of 100 or more, required to detect axions with benchmark couplings. A recent effort [34], using Rydberg atom single-photon detection [35], is still in development, but holds the potential for greatly improved sensitivity should the Rydberg atom technique be sufficiently advanced.

4. The phase space structure of galactic halos

If a signal is found in our detector, it will be possible to measure the energy spectrum of cold dark-matter particles on Earth with great precision and resolution. Note that this experiment measures the total energy of the axion, e.g., its rest mass plus kinetic energy. The spread ($\Delta f/f = \Delta E_a/m_a \approx 10^{-6}$) of the axion signal is due to the kinetic energy of motion of the axions through the galactic halo. Our experiment is equipped with a high-resolution spectrometer that resolves $\Delta f/f \sim 10^{-11}$ and hence divides the axion signal width into 10^5 bins. We argue below that the spectrum will have narrow lines associated with the late infall of dark mat-

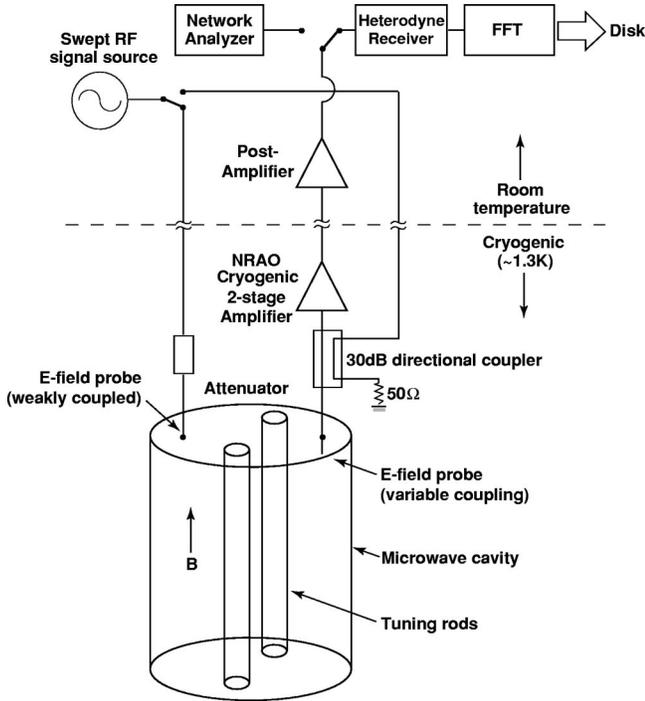


FIG. 1. Schematic of the apparatus.

ter onto our galaxy. Even a few percent of the axion signal in any one of these lines would greatly increase the signal-to-noise ratio and thus the experiment's discovery potential. Further, the lines are rich in information on the history of the formation of the galaxy.

In many past discussions of dark-matter detection on Earth, it has been assumed that the dark-matter particles have an isothermal distribution. Thermalization has been argued to be the result of a period of "violent relaxation" following the collapse of the protogalaxy [36]. However, substantial deviations from an isothermal distribution are expected [37] due to late infall. If the dark matter is collisionless, as axions are, our galaxy must be surrounded by a sea of dark-matter particles. The particles fall into the gravitational well of the galaxy, forming a continual "rain." The flows of late-infalling dark matter do not get thermalized, even on time scales of the order of the age of the universe [37]. Hence there is discrete set of dark-matter flows with well-defined velocity vectors at any location in the halo, in addition to a thermalized component. One flow is comprised of particles falling onto the galaxy for the first time, a second flow is comprised of particles falling out of the galaxy for the first time, a third flow is of particles falling into the galaxy for the second time, etc. Each of the flows could contain a few percent of the local halo density [38].

The intrinsic width of the peaks due to the most recent infall, is estimated [37] to be of order 10^{-17} , although each peak may be fragmented into subpeaks of that approximate width. Because of the Earth's rotation ($v_{\text{rot}} \approx 10^{-6} c$) and its revolution around the Sun ($v_{\text{rev}} \approx 10^{-4} c$), the peaks frequencies ω_n are modulated by the relative amount

$$\frac{\delta\omega_n}{\omega_n} = \frac{\vec{v}_n \cdot (\vec{v}_{\text{rot}} + \vec{v}_{\text{rev}})}{c^2} + \mathcal{O}\left(\frac{v_{\text{rot}}^2, v_{\text{rev}}^2}{c^2}\right), \quad (17)$$

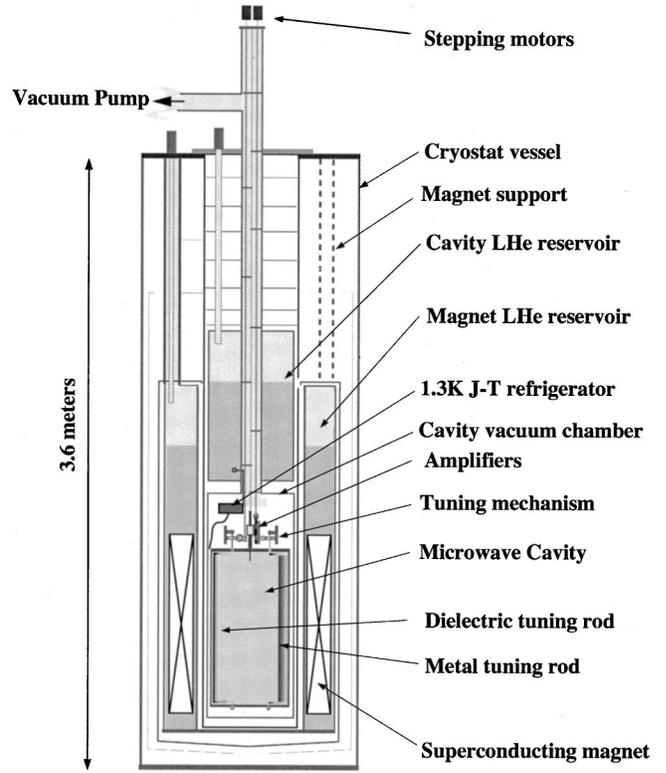


FIG. 2. Sketch of the experiment insert within the magnet system. The experiment insert can be withdrawn without warming the magnet.

where $v_n \approx 10^{-3} c$ is the velocity of the n th phase-space sheet relative to the Sun. Observation of the diurnal frequency modulation of a peak would allow us to determine the corresponding velocity vector \vec{v}_n completely.

II. THE U.S. rf CAVITY AXION SEARCH: DESCRIPTION AND PHYSICS RESULTS

A. Hardware

1. Hardware overview

Figure 1 is a diagram of the apparatus. The heart of the axion detector is a cylindrical rf cavity containing two moveable tuning rods positioned parallel to a strong magnetic field produced by a solenoid surrounding the cavity [39]. The cavity electromagnetic field is coupled through a small adjustable electric-field probe to ultra-low-noise receiver electronics. These components are sketched in Fig. 2.

The cavity is cooled to 1.3 K to reduce its blackbody noise. A 7.6 T static magnetic field, parallel to the cylinder axis, threads the cavity volume. We use the TM_{010} cavity mode since it has the largest form factor for axion-to-photon coupling. The unloaded quality factor Q of the TM_{010} resonance is typically $7-20 \times 10^4$. The frequency step between adjacent power spectra in the data stream is typically 2 kHz, about 1/15th of the width of the cavity resonance width ω_0/Q_L . Hence, many power spectra overlap each cavity fre-

quency in the search range. The frequency range covered in the initial data set is 701–800 MHz, corresponding to the axion mass range 2.9–3.3 μeV .

The cavity electromagnetic fields are coupled to the receiver chain by an electric dipole field probe with adjustable insertion depth (i.e., with adjustable cavity loading). During normal running, the insertion depth is occasionally adjusted to maintain near-critical coupling of the cavity to the cryogenic amplifier input. A directional coupler between the cavity and cryogenic amplifier allows for application of test and calibration signals.

The voltage across the electric-field probe is applied to cryogenic RG-402 coaxial cable [40], then amplified by two cryogenic heterojunction field-effect transistor (HFET) amplifiers in series, built by the National Radio Astronomy Observatory (NRAO) [41]. The overall power gain for the two cryogenic amplifiers is approximately 34 dB. Each amplifier is thermally tied to the cavity. The first set of amplifiers used in data taking had a noise temperature of about 4.3 K. The more recent amplifiers have noise temperatures better than 2 K. Figure 3 (upper) shows the noise temperature and gain versus frequency of one of these first amplifiers. Figure 3 (lower) shows noise and gain for one of the later amplifiers. The output of the cryogenic amplifiers is passed out of the cryostat along cryogenic RG-401 coaxial cable [40], then further amplified by a low-noise room-temperature amplifier with about 38 dB of gain mounted in an rf shielded enclosure directly on the room-temperature top flange of the detector. The amplified microwave signal is applied to 25 m of RG-213 flexible coaxial cable leading to the mixing, i.f., and audio stages of the receiver [42,43].

The axion receiver has a double-heterodyne design strongly influenced by experience gained in the earlier Florida and Rochester-Brookhaven-Fermilab experiments [33]. The first downconversion stage, by image rejection mixing, generates a 10.7-MHz i.f. An eight-pole crystal bandpass filter in the 10.7 MHz i.f. rejects power outside a 35-kHz frequency window centered on the cavity resonant frequency. A second conventional mixing stage downconverts the 10.7-MHz i.f. to near audio frequency (AF) centered at 35 kHz. The audio signal is applied to commercial fast Fourier transform (FFT) electronics [44] which computes a 50-kHz bandwidth power spectrum centered at 35 kHz.

The normal data-taking sequence starts by moving the tuning rods incrementally to establish a new cavity resonant TM_{010} frequency. Then, the electromagnetic power spectrum about the cavity TM_{010} mode is determined by the receiver electronics and commercial FFT instrumentation. The 50-kHz wide FFT power spectrum consists of 400 bins, each 125 Hz wide. A uniformly weighted average of 10 000 power spectra is saved together with other parameters of the experiment including the measured cavity Q and resonant frequency. We refer to this 10 000 spectra average as a “trace.” A total of about 4.5×10^5 traces were recorded over the 2.9–3.3 μeV mass range.

The averaged power spectrum at each cavity setting, together with other experimental parameters (e.g., cavity Q and resonant frequency) constitutes the raw data set. The ther-

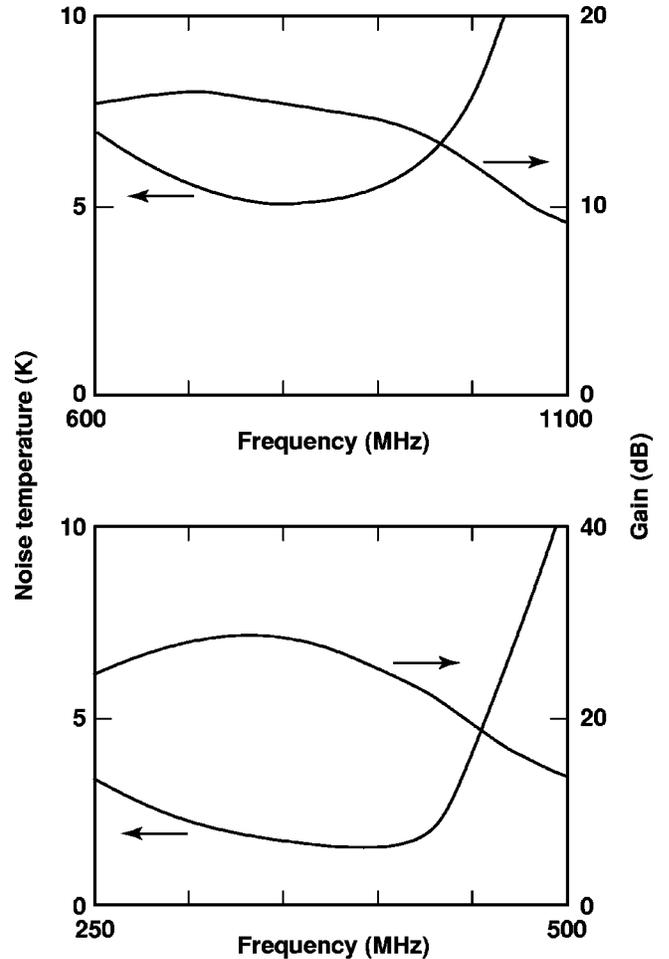


FIG. 3. (Upper) The noise temperature and gain vs frequency of one of the first amplifiers used in the experiment. (Lower) The noise and gain of one of our best cryogenic amplifiers. The data shown were taken at a bath temperature of about 12 K. The noise temperature improves slightly (5%–10%) by lowering the bath temperature to 4.2 K. These latter amplifiers have noise temperatures of about 1.5 K.

malized axion signature is excess power above the background power spectrum concentrated in a peak of bandwidth of the order of 1 kHz (about 6 frequency bins wide). The dominant background sources are relatively broadband cavity blackbody noise and broadband electronic noise in the first cryogenic amplification stage.

2. The resonant cavity

The resonant cavity consists of a copper-plated stainless-steel right-circular cylinder 1 m long and 50 cm diameter with two endplates. The cavity volume is ~ 200 l. (The first such cavity had the lower endplate welded in place and the upper plate removable. This configuration was difficult to electroplate, so the design was changed to employ two removable endplates.) Each removable endplate is firmly seated on a knife-edged lip to ensure low resistance to the TM_{010} wall currents. The cavity components are plated on all major surfaces with high-purity oxygen-free copper, then annealed. Both metal and dielectric tuning rods were designed



FIG. 4. The resonant cavity viewed from above with the top flange removed. The cavity is a right circular cylinder of diameter 50 cm and depth 1 m. An alumina tuning rod is at the upper right, a copper tuning rod is at the lower left.

to be used in the experiment. The metal rods consist of copper cylinders, capped at both ends, plated with oxygen-free high-purity copper and then annealed. The dielectric rods consist of low-microwave-loss alumina. The rods were mounted at the end of alumina swing arms, which pivot about alumina axles that penetrate the top and bottom end plates of the cavity. By rotating these axles, the tuning rods may be swung in circular arcs from close to the cavity side wall to close to the center. Metal rods increase the frequency of the cavity, whereas dielectric rods decrease it. Figure 4 is a photograph of two tuning rods within the cavity. The axles and arms used to move the tuning rods are visible. In this picture, the upper-right tuning rod is made of alumina. Our first data-taking run used copper tubes for tuning rods, each similar to the lower-left rod in the picture, 8.25 cm in diameter. The tuning rods are driven by stepper motors mounted on top of the cryostat. The stepper motors drive G10 rods into a two-stage antibacklash gear reduction mounted on the cavity endplate. The reduction, besides allowing small tuning steps, also reduces the effects of play in the G10 rods. There are 8.4 million motor steps per full revolution, with a single step corresponding to a change in the angle of the tuning rod by 0.15 arcseconds. The tuning precision is about 1 kHz in 100 MHz.

a. Quality factor measurement. The Q of the TM_{010} mode is determined from the transmission response. A swept rf signal is applied through an electric dipole probe to the very weakly coupled rf calibration port on the top plate of the cavity, and the transmitted signal from the major port is directed to a scalar network analyzer. Figure 5 is the swept transmission response across the TM_{010} resonance for one particular tuning-rod configuration. The width of the resonance is around 7 or 8 kHz (at resonant frequencies near 750 MHz) with the cavity critically coupled to the receiver chain, implying an unloaded Q of $\sim 2 \times 10^5$.

b. Obtaining critical coupling. The scan rate at constant SNR depends on the coupling of the external amplifier to the cavity; the rate has a broad maximum at slightly over critical coupling. Hence, during normal data taking, the insertion depth of the major port electric-field probe was periodically adjusted to maintain critical coupling between the cavity

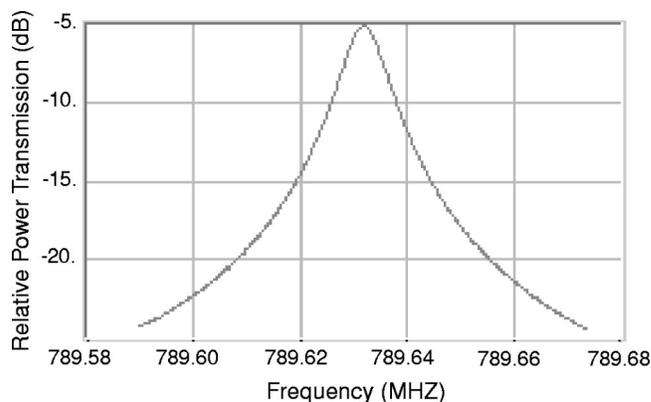


FIG. 5. The cavity transmission resonance taken from the on-line data acquisition display. The vertical axis is relative transmitted power after amplification. The horizontal axis is the swept input frequency.

TM_{010} mode and the receiver chain. At critical coupling, the power incident on the cavity from the major port is absorbed without reflection. To set the coupling, a swept rf signal is injected into the field probe through the weak-coupled side of the directional coupler (see Fig. 1). Almost all the signal power is absorbed in the cold 50Ω terminator. The remaining power (around 1 part in 1000) is directed towards the cavity. A network analyzer and cryogenic directional coupler are used to measure the power reflected from the cavity; when the probe is critically coupled essentially no power is reflected at the cavity resonant frequency. Figure 6 shows the reflected power in the neighborhood of the TM_{010} mode after the critical coupling procedure. The reflected power is measured after the rf amplifier chain using a scalar network analyzer. The bottom of the absorption dip is the noise floor of the electronics, the actual reflected power is considerably less. We considered a ~ 30 dB absorption dip critical coupling.

c. The cavity mode structure and form factor. There is a complicated mode structure within the cavity. There are TE and TM modes associated with the smooth-wall cylindrical geometry. TEM modes are introduced by the metallic tuning

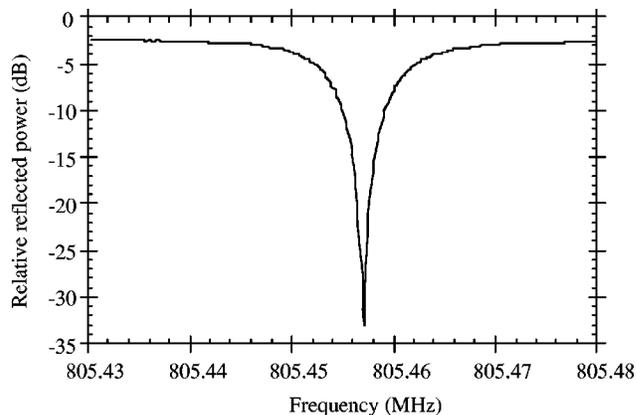


FIG. 6. The relative reflected power off the cavity major port in the neighborhood of the TM_{010} mode after the critical coupling procedure. The resonant frequency is 805.457 MHz.

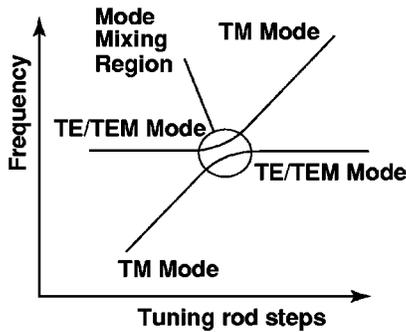
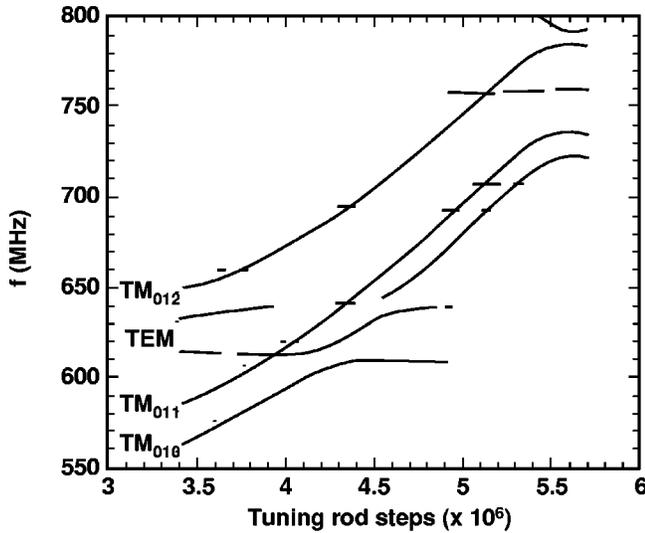


FIG. 7. (Upper) The cavity mode structure with one rod fixed near the wall and the other rotating towards the cavity center. The horizontal axis is the rod position in units of stepper motor steps. (Lower) Sketch of an avoided mode crossing.

rods. There are, in addition, hybrid modes associated with deviations from a uniform cross section. The cavity itself is not a perfect circular cylinder, and the insertions and tuning rods also distort the cylindrical geometry, thus there is considerable mixing among modes. Consequently, there is no sufficiently accurate closed analytic form for the cavity mode structure, so the form factor C is therefore estimated by numerical simulation.

With a cavity containing two metal tuning rods, the TM_{010} mode is lowest in frequency when both copper rods are nearest the wall. As one or both rods approach the cavity center, the frequency increases. Figure 7 (upper) shows the frequencies of various cavity modes versus tuning rod positions for the case of one rod fixed near the wall and the other moving towards the center. The vertical axis is cavity frequency f_0 and the horizontal axis is number of stepper motor steps from a rod position near the wall. These mode maps were obtained by feeding a swept rf signal into the cavity through the weak-coupled rf port and measuring the transmitted power. The TE and TEM modes of the cavity are only weakly excited by the rf probes oriented perpendicular to the TE and TEM electric fields.

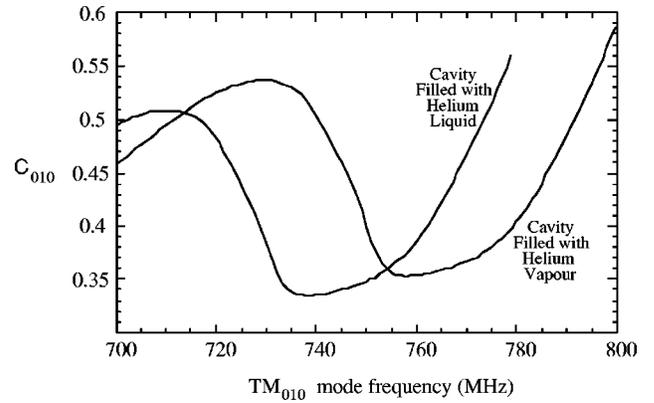


FIG. 8. The calculated cavity form factor C_{010} as a function of the cavity resonant frequency. One curve is for the cavity containing a small amount of helium vapor, the other for the cavity filled with superfluid helium.

Notice that there are mode crossings (mixings) in the various TM modes. These crossings occur when the frequencies of two modes are nearly degenerate and the modes mix. Figure 7 (lower) is a sketch of mode crossing of a TM with a TE or TEM mode. The mode crossings introduce frequency gaps that cannot be covered with a pure TM_{010} mode. These gaps were scanned by filling the cavity with liquid helium at 1.3 K. Liquid helium has a relative dielectric permittivity of 1.055 and hence alters the microwave index of refraction by a factor of 1.027. The frequencies of all the modes of the cavity and the mode crossings are decreased by 2.7%: 19 MHz at 700 MHz. By combining data taken when the cavity is filled with low-pressure gas, as is normally the case with some data taken when the cavity is filled with helium, we cover the entire frequency range without gaps. There is a slight penalty in filling the cavity with helium as the reduced electric field gives a lower axion conversion rate.

The form factor C_{010} of the TM_{010} mode, defined in Eq. (15), is calculated from a numerical solution of the wave equation for the axial component of the electric field E_z in the cavity, including tuning rods. A relaxation code using the Gauss-Seidel method [45] is used on a 110×110 lattice for the simplified case of a constant two-dimensional cavity cross section and tuning rods. (The effects of the approximately 1-cm gap at each end between the tuning rods and the cavity end plates were neglected; we estimate this simplification results in an error on the form factor of the order of 1%.) The resultant axial electric field is combined with the calculated magnetic-field shape (calculated from the actual magnet winding current density) to yield Fig. 8, a plot of the form factor versus TM_{010} resonant frequency for copper rods with dimensions the same as in our experiment. The two lines represent the form factor with the cavity filled with low-density gas and the cavity filled with liquid ^4He . The form factor for different tuning rod materials (metallic and dielectric) and rod diameters is shown in Fig. 9; here the horizontal axis is the cavity resonant frequency, and the vertical axis is the form factor C_{010} .

3. The receiver electronics

a. The cryogenic amplifiers. The characteristics of the cryogenic amplifiers are particularly important because their

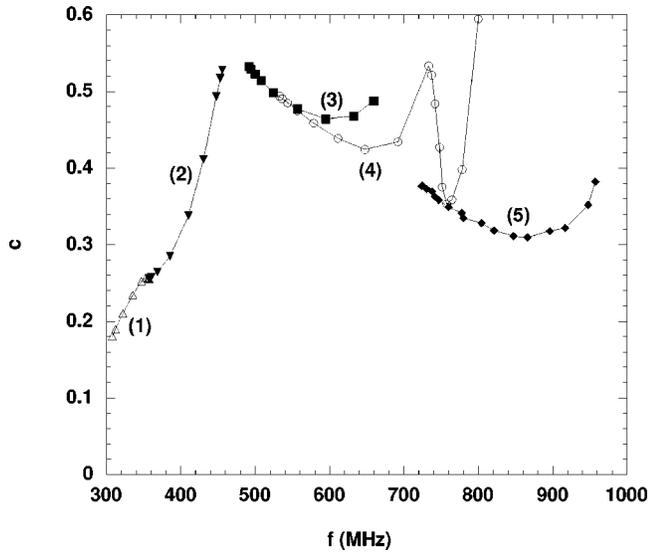


FIG. 9. The frequency dependence of the form factor. The horizontal axis is the cavity resonant frequency, the vertical axis is the form factor C_{010} . The curves show the form factor for different tuning rod materials and rod geometries: (1) two alumina rods ($\epsilon = 9.5$), ratio of rod diameter (r) to cavity diameter (R) $r/R = 0.13$, one rod fixed at cavity center; (2) two alumina rods, $r/R = 0.13$, one rod fixed at cavity center; (3) one metal rod, $r/R = 0.16$; (4) two metal rods, $r/R = 0.16$, one rod fixed at wall (for the 500–700 MHz range) or fixed at the cavity center (for the 700–800 MHz range); (5) two metal rods, the $r/R = 0.16$ rod fixed at the cavity center, the $r/R = 0.3$ rod moving.

noise dominates the system noise. Two cryogenic amplifiers, built by NRAO [41], are cascaded so as to have sufficient gain to render negligible further noise contributions. Each cryogenic amplifier is a balanced design; the input signal is split into two paths by a 90° hybrid (one path phase shifted by 90° with respect to the other). Each path is then amplified and recombined by another 90° hybrid (reversing the initial 90° phase shift). Any reflection from the input from one path is canceled by the phase-shifted reflection from the other path. The reflection instead is directed to a termination resistor at the phase-shifted input of the 90° hybrid [46]. Such a balanced amplifier provides good matching without isolators or other nonreciprocal devices that are difficult to use in the large ambient magnetic field. This balanced design is realized by a pair of single stage HFET amplifiers, the two a matched pair with nearly identical characteristics, with single-pole matching networks at the input. The amplifier package is cooled to near liquid-helium temperatures.

The input impedance of our balanced amplifiers are typically well matched to 50Ω over 500-MHz bandwidths; the input power reflection coefficient is smaller than -18 dB everywhere between 600 and 1300 MHz, limited by the characteristics of the hybrids.

b. Cryogenic amplifier characteristics in high ambient magnetic fields. We were careful to ensure that the gain, noise, and bandwidth of the cryogenic amplifiers are not degraded by the experiment's ambient magnetic field. *In situ* measurements of the gain of the cryogenic amplifier were compared with bench-test data taken by NRAO.

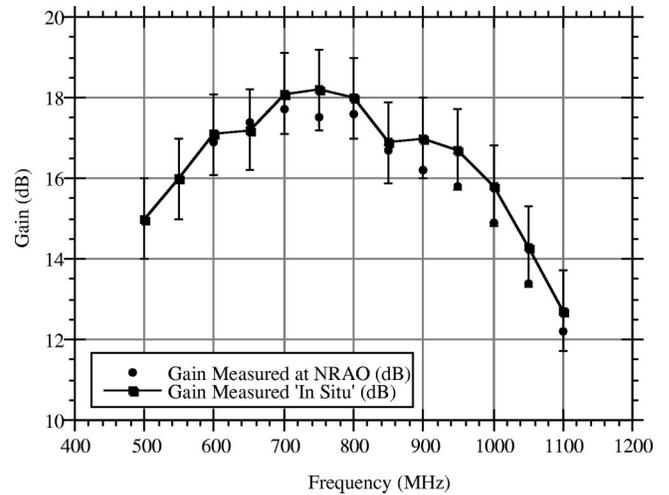


FIG. 10. The power gain of a typical cryogenic amplifier input measured *in situ*. The error bars are due to variations in cable losses due to temperature changes and flexing. Our results are in good agreement with independent measurements of the same devices by NRAO, which are also shown.

For the *in situ* gain measurements, a swept rf signal is applied through the cryogenic directional coupler to the cryogenic amplifier input. With the experiment insert at room temperature and the magnetic field off, a scalar network analyzer at the output of the two cryogenic amplifiers allows for determining the cable and connector attenuation response since the room temperature gain and other amplifier characteristics are known. The insert is then cooled and the nominal magnetic field restored (the magnetic field in the neighborhood of the amplifier is around 4 T). Keeping the swept power the same, another power measurement at the output of the two amplifiers yields the cryogenic amplifier gain in high ambient magnetic fields. Studies of the passive components showed that temperature-dependent losses are negligible. The *in situ* gain is shown in Fig. 10 and is in good agreement with independent measurements by NRAO.

The noise temperature of the cryogenic amplifier is measured with the amplifier installed for data taking. Ultimately, our calibration of amplifier noise temperature is referenced to the temperature of the cavity as a Johnson noise source.

We first critically coupled the amplifier to the cavity via the reflection minimization method. On resonance, the impedance looking towards the cavity from the amplifier is 50Ω ; the cavity is then effectively a $50\text{-}\Omega$ resistor whose Nyquist noise temperature is the physical temperature of the cavity. The measurements begin with the cavity temperature at 1.3 K, and the liquid-helium reservoir below the cavity empty. A small amount of helium liquid is then released into the cavity space, forming a vapor of a few Torr and ensuring good thermal contact between the tuning rods and the rest of the cavity. Heaters on the top and bottom of the cavity then warm the cavity to about 10 K over 1 h. During warming, a swept signal is injected through the weak port and the transmitted spectral density determined with a scalar network analyzer (the swept signal is used to correct for any slight thermal drift, of order 1 dB or less, of the amplifier gain over

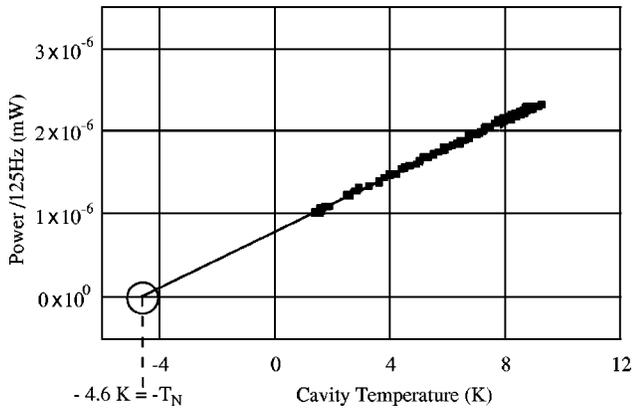


FIG. 11. Results of an early *in situ* noise measurement on a typical cryogenic amplifier. The vertical axis is rf spectral density recorded by the receiver chain, the horizontal axis is the cavity temperature. The straight line is a least-squares fit.

the measurement temperature swing). Between sweeps, a 90-s power spectrum of 30-KHz bandwidth about the cavity resonance is recorded by the receiver. Figure 11 shows the gain-corrected power in a 125 Hz wide bandwidth at the center of the power spectrum (on resonance) of an early cryogenic amplifier as a function of the physical temperature of the cavity tuned to 700 MHz. The straight line is a least-squares fit. The noise temperature is the absolute value of the intercept of the fitted line with the cavity temperature axis; here the inferred value of the noise temperature is 4.6 K, a value consistent with the noise temperature of that amplifier measured independently by NRAO. We ascribe a 10% measurement uncertainty to our *in situ* noise temperature result. The amplifiers currently taking data are considerably better, with measured noise temperatures below 2 K, as determined independently by NRAO and us. Figure 3 (lower) shows the noise and gain for one of these recent amplifiers at a bath temperature of around 12 K; the noise is slightly reduced (by 5–10%) at the operating temperature of the amplifier in the detector.

The orientation of the amplifier during the production data taking is such that the magnetic field is parallel the HFET channel electron flow. During commissioning we observed that if the cryogenic amplifier was rotated so the magnetic field is perpendicular to the electron flow in the HFET junctions, the amplifier noise temperature rises from 4.6 K to around 8 K. A simple model is that the Lorentz force on the electrons in the HFET junctions distorts the electron paths so that their trajectory across the gate region lengthens. In the simplest version of the model, the HFET amplifier noise temperature is proportional to the length of the electron path

through the gate. The model predicts an increase in noise temperature consistent with that seen in [47].

c. The room temperature receiver electronics. The output of the cryogenic amplifiers is carried by a cryogenic RG-401 coaxial cable [40] to a low noise room-temperature postamplifier mounted in a room-temperature rf-shielded enclosure on top of the cryostat. The power gain of the postamplifier is about 35 dB in the frequency range 300 MHz–1 GHz with noise temperature around 90 K. The power gain of the two cryogenic amplifiers in series is 34 dB, hence the contribution of the postamplifier to the system noise temperature is negligible, less than 0.03 K. The signal lines terminate in a shielded enclosure and mate with type-*N* bulkhead connectors. The connectors then mate with flexible RG-213 cables leading to the room-temperature receiver electronics in the analysis hut. Overall, there are three rf lines from the detector to the analysis hut: one for the cryogenic amplifier output, one for the weakly coupled port, and one for the directional coupler. The total power gain from cavity probe to the post-amplifier output is 69 dB (8×10^6); with this gain, the noise power from the cavity and amplifier at the postamplifier output in a 125-Hz bandwidth is 8.3×10^{-14} W.

The amplified rf signal from the detector then enters the double-heterodyne section of the receiver. Figure 12 is a schematic of the receiver i.f. and af processing components. The first component in this section is an image rejection mixer; it shifts rf power from the TM_{010} resonant frequency to an i.f. frequency range centered at 10.7 MHz, while rejecting image noise power from the i.f. The image rejection mixer is a MITEQ IRM045-070-10.7 [42] with insertion loss of 6 dB, and image rejection better than 20 dB. During normal running, the local oscillator frequency is maintained at 10.7 MHz below the measured cavity resonant frequency so that the cavity resonance is centered at 10.7 MHz in the i.f. strip.

After the first mixer is an adjustable attenuator (maximum 63 dB attenuation). Attenuation is needed to avoid saturating the receiver during tests with the cryogenic components at room temperature. The i.f. signal is then amplified by about 20 dB and passes through weakly coupled signal samplers used for monitoring the internal receiver power levels.

The i.f. section has a bandpass filter to suppress noise outside a 30-kHz bandwidth centered on the 10.7 MHz i.f. Out-of-band power would otherwise have been aliased in the af output by the next mixing stage. On account of the strong dependence of the receiver transfer function to the filter’s response, the filter is enclosed in a temperature controlled oven maintained at $40 \pm 0.5^\circ$ C. The filter has eight poles and ripple from this pole structure appears across the receiver passband. After another 20 dB of amplification, the i.f.

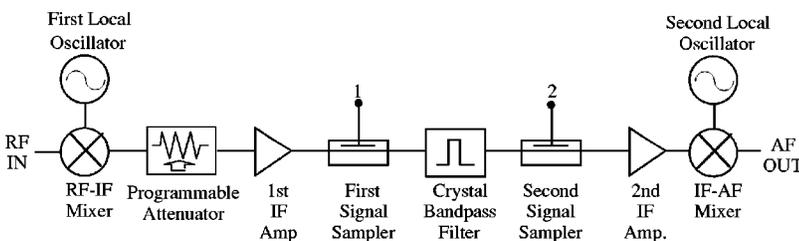


FIG. 12. Simplified schematic of the intermediate (IF) and near-audio frequency (AF) components in the receiver chain.

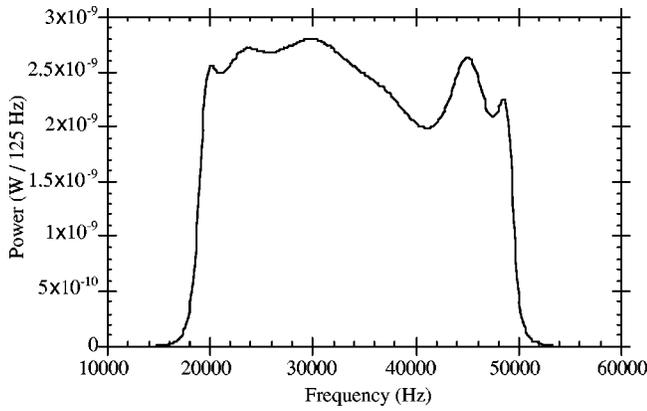


FIG. 13. The spectral output density at the output of the af section with a spectrally flat noise source at the i.f. input. The structure shows mainly the response of the crystal i.f. filter.

signal is mixed down again to near af in the range 10–60 kHz. The output of the second mixer leads to a commercial FFT spectrum analyzer [44], and the resulting power spectra are saved to disk as our raw data. The frequency of the local oscillators is stabilized to drift less than ± 0.005 Hz by a cesium frequency reference. The frequency stability of the receiver chain is established *in situ* by noting that occasional large external radio peaks drift in frequency by less than ± 0.02 Hz.

The receiver response, including the filter, is calibrated by measuring the receiver output spectrum with a white-noise source at the input. Noise generated by a Noise/Com 346B broadband noise source [43], amplified by 35 dB, is injected at the rf input of the image rejection mixer. At the output of the receiver, the FFT spectrum analyzer takes and averages

400-point (125 Hz per point) power spectra of the filter pass-band response. After a total integration time of 2.2 h, we have the calibration of the receiver response shown in Fig. 13. The vertical axis is power per 125 Hz, the horizontal axis is frequency at the af output of the receiver.

The search for axionlike signals in our data is affected by even very small drifts in the receiver response. A typical integration time for the acquisition of a single averaged FFT power spectrum (containing 10^4 individual spectra, described in the next section) is 80 s. Hence fluctuations in the power spectrum are expected to be $1/\sqrt{10000}$ of the raw noise power. A receiver response change of only 1% or less noticeably affects the averaged spectrum. Changes of approximately 0.2% in the receiver response over the first 1.5 years of operation were seen; the treatment of these drifts is discussed in Sec. II B.

d. Signal and noise power through the receiver chain. Table I shows the approximate power at the output of the major components of the receiver electronics. We list power over two different bandwidths: the first over a 125-Hz bandwidth (the width of the frequency bins in the FFT spectrum analyzer), the second over the full bandwidth of the specific component (this to ensure the power output of a component over its full bandwidth is not so large that the next component in the receiver chain is saturated). The input power over a 125-Hz bandwidth incident on the first cryogenic amplifier, is around 10^{-20} W (-170 dBm). Notice the i.f. crystal filter significantly reduces the broadband power in decreasing the bandwidth from about 1 GHz to 35 kHz.

e. The FFT spectrum analyzer and a typical power spectrum. The FFT spectrum analyzer operates here at an effective sampling rate of 100 kHz. Every 8 ms, an individual single-sided spectrum (one in which the negative and posi-

TABLE I. Estimated power levels at selected points in the receiver electronics chain. We take the bandwidth of the cryogenic amplifiers as 1 GHz and the cavity plus amplifier noise temperature of 6 K.

Component	Gain (dB)	Output Power per 125 Hz (dBm)	Output power over full bandwidth (dBm)
Cavity		-170	-101
Cryogenic amplifiers	34	-136	-67 (over 1 GHz)
Room temperature postamplifier	35	-101	-32 (over 1 GHz)
Flexible cable to analysis hut	-6	-107	-38 (over 1 GHz)
Image reject mixer	-7	-114	-45 (over 1 GHz)
First i.f. amplifier	30	-84	-15 (over 1 GHz)
Crystal filter	-3	-87	-60 (over 30 kHz)
Second i.f. amplifier	30	-57	-30 (over 30 kHz)
i.f.-af mixer	-7	-64	-37 (over 30 kHz)

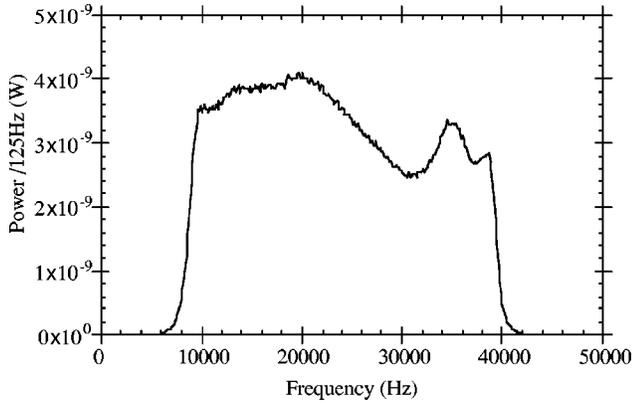


FIG. 14. A typical trace from the raw data. The vertical axis is spectral density at the af output. The mixed-down cavity resonant frequency is near the center of the frequency range.

tive frequency components are folded on top of each other) is taken with uniform (flat) windowing consisting of 400 bins, each 125 Hz wide, spanning the frequency range 10–60 kHz. Over 80 s of integration time per frequency setting, 10 000 such individual spectra are taken and averaged and it is this averaged spectrum that is saved as raw data. Figure 14 shows a typical averaged power spectrum taken during production running.

The most obvious feature of this spectrum is the rapid falloff in power at the edges of the 20–50 kHz frequency range at the skirts of the i.f. crystal filter response. There are slight differences between Figs. 13 and 14 due to additional structure in the latter introduced by standing waves in the cavity, transmission line, and amplifier system. There is also a slight difference in noise power levels on- and off-resonance; the Johnson noise power on resonance (where the cavity noise source dominates) is less than the Johnson noise power off resonance (where internal noise sources in the amplifier dominate). Understanding the underlying structure of the raw data spectra is important as the expected axion signal appears as additional structure above noise background across six adjacent bins in the power spectrum.

f. Extremely narrow bin electronics. There is, in addition, a special electronics channel optimized for the detection of extremely narrow lines from late-infall axions. For this case, the af cavity signal (centered at 35 kHz) is fed into a passive LC filter [48] (passband ≈ 6.5 kHz), amplified, and mixed down to a 5-kHz center frequency. The signal is applied to an ADC/DSP board [49] in a PC host computer, where it is oversampled at 20 kHz by a 16-bit ADC and stored in DSP memory. The PC receives the start command along with cavity Q and frequency data from the main DAQ computer immediately after the main FFT channel begins collecting data. A single spectrum is obtained by acquiring 2^{20} points in about 53 s, performing the FFT calculation in about 8 s, and uploading the 2^{19} -point power spectrum to the PC host for analysis. The entire process is finished before the medium resolution spectrum is done and the cavity frequency returned.

About 3.4×10^5 points lie in the 6.5 kHz passband with a frequency spacing of 19 mHz. In order to get higher sensi-

tivity for lines of intermediate width, we averaged nonoverlapping segments of 8 and 64 points, resulting in two additional power spectra of resolution 152 mHz and 1.2 Hz, respectively.

4. The cryogenic hardware

a. The magnet. The magnetic field is produced by a superconducting solenoid [50]. The magnet consists of niobium-titanium windings immersed in a liquid ^4He cryostat. The helium consumption of the magnet is approximately 55 l per day. The axial field at the center of the solenoid is 7.6 T, falling to about 70% of this value at the cavity end plates.

b. The insert and helium system. Figure 2 shows a sketch of the insert and the magnet showing the cryogenic system. The cavity is at the bottom of the insert. The entire insert can be withdrawn from the magnet bore without warming the magnet.

The cavity is surrounded by a stainless-steel can. Liquid helium is released into the bottom of this can from a reservoir directly above the cavity space through a thin capillary line and an adjustable inlet valve. During normal operation the liquid helium forms a pool several cm's deep in the space under the cavity. A level gauge monitors the depth of the helium pool and controls the release of liquid helium from the valve.

A 4-in diameter stainless-steel column rises from the top of the can enclosing the resonant cavity through the center of the liquid-helium reservoir, through radiation baffles, and continues to a manifold protruding from the top of the detector. Through this column, a large Roots blower pumps on the helium pool below the cavity, thus cooling the cavity to a physical temperature of 1.3 K. The pressure of the helium gas in the insert is typically less than 1 Torr. While scanning near mode crossing regions, the liquid-helium pool is allowed to rise until the entire cavity is filled. (This liquid helium is eventually boiled off by heaters attached to the bottom of the cavity.) The reservoir supplying liquid helium to the insert is refilled once every two weeks. Between the insert and the inner wall of the cylindrical magnet dewar is an insulating space at 10^{-6} Torr vacuum. Heat losses from the insert are slight as it is surrounded on three sides by the magnet dewar and on the fourth side by layers of vapor cooled baffles.

B. Analysis method

Our analysis is a search for the signature of axions in the spectra from the detector [51]. The analysis has two paths. The first path (the “single-bin” search) is motivated by the possibility that at least some halo axions may not have thermalized and therefore have a negligible velocity dispersion; these axions would deposit their energy into a single power spectrum bin. The second path (the “six-bin” search) makes no assumption about the halo axion energy distribution other than its velocity dispersion is $\mathcal{O}(10^{-3}c)$ or less. (Axions with velocity larger than about 2×10^{-3} escape from the halo.) The results from the six-bin search are the most conservative since they are valid whether or not the halo axions

have thermalized, completely or in part. Should the halo be nonthermal because of late infall, the single-bin analysis would be more sensitive. The axion search using an extremely narrow line channel is outlined at the end of this section. Both search paths incorporate a simple power excess in the search statistic.

1. Overview of the single- and six-bin search methodology

Both 1- and 6-bin searches use the same raw data, which covers without gaps 701–800 MHz. These raw data consist of a set of overlapping set of averaged power spectra. We refer to this data set as the “run 1 raw data.” These spectra are combined into a single power spectrum over the entire 701–800 MHz range with a bin width of 125 Hz.

a. The single-bin search. For the single-bin search, individual frequency bins exceeding a power level threshold are selected from the combined power spectrum. The threshold is chosen relatively low so as to select a considerable number of candidates. For just the selected candidate frequencies, a second, independent, set of raw data are taken with the same integration time. This new data set is combined with the first, producing a single combined power spectrum with higher signal-to-noise ratio at the selected candidate frequencies. The selection process is repeated; individual frequency bins exceeding a power level threshold are selected from the new combined power spectrum, and so on. The few surviving candidates are carefully checked as to whether they are identified with known sources of external interference. If all candidates were to be identified with external interference, then no candidates would survive and the excluded axion couplings are computed from the near-Gaussian statistics of the single-bin data.

b. The 6-bin search. For the 6-bin search, all six adjacent frequency bins exceeding a power level threshold are selected from the combined power spectrum. The threshold is again chosen relatively low so as to accept a considerable number of candidates. (These candidates are correlated with candidates from the 1-bin search.) For just the selected candidate frequencies, a second, independent, set of raw data is taken with the same integration time. These new data are combined with the first, producing a single combined power spectrum with higher signal-to-noise ratio at the selected candidate frequencies. Again, repeating the selection process, any six adjacent frequency bins exceeding a power level threshold are selected from the new combined power spectrum. For just these remaining candidate frequencies, yet a third, independent set of raw data are taken with the same integration time as the initial data set, then combined with the previously taken data. Six adjacent frequency bins exceeding a certain power level threshold are selected from the new combined power spectrum. After these three steps of candidate selection there are few surviving candidates, and again we carefully check whether they are identified with known sources of external interference. Should all candidates be identified with external interference, then no candidates survive and the excluded axion couplings are computed by Monte Carlo techniques.

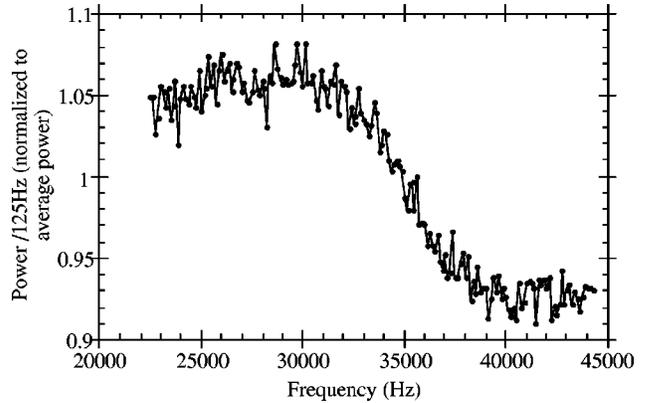


FIG. 15. A raw trace after the first 100 and last 125 bins (out of 400 bins) have been removed, then divided bin-by-bin by the receiver response. The cavity resonant frequency was mixed-down to near 34 kHz af.

2. The data combining algorithm

a. Preliminary treatment of raw traces. Figure 14 is a typical “trace” from raw data. (This trace is an 80-s single-sided power spectrum centered on the cavity resonant frequency. It contains 400 frequency bins, each bin 125 Hz wide.) The first stage in the processing of a raw trace is to remove frequency bins outside the filter passband. From the 400 bins in each trace, the first 100 and the last 125 bins are discarded. Slightly more bins at the higher-frequency side of the passband are removed due to a narrow pole in crystal filter response at this side of the filter passband with a width of about 10 bins, which is somewhat close to the expected axion signal width (approximately 6 bins). We refer to the portion of each trace remaining after the end bins are discarded as the “cropped trace.”

Figure 13 shows the receiver passband response calibration. Comparing this with Fig. 14, much of the slowly varying structure in the power spectrum is due to the receiver passband response. The next operation in data analysis therefore divides the cropped trace bins by the corresponding bins in the receiver passband response. The result of performing this normalization operation on the trace in Fig. 14 is shown in Fig. 15.

We refer to the traces normalized by the receiver passband response as “corrected traces.” Notice there is additional smooth structure in the corrected traces. Furthermore, this structure is not symmetric about the center bin (corresponding to the resonant frequency of the cavity).

b. Origin of asymmetry in the receiver corrected traces. If all the structure in the raw traces were attributable to the frequency-dependent effects in components downstream of the receiver rf amplifier, the corrected traces would consist of nearly Gaussian fluctuations on a background, symmetric about the cavity resonant frequency. It is clear from Fig. 15 that, in fact the power spectra are not symmetric. Furthermore, the structure of the corrected traces depends on the cavity TM_{010} resonant frequency. Figure 16 shows corrected traces at two different resonant frequencies. Such frequency dependence implies that the structure of the corrected traces comes from an interaction between the rf cavity and the first

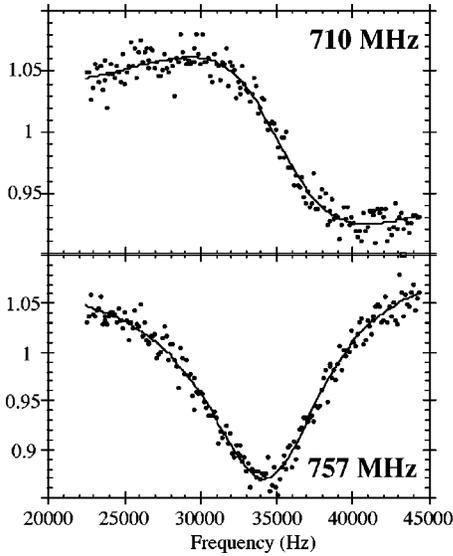


FIG. 16. Typical corrected traces at the af output for two different cavity resonant frequencies. Points are measured relative power. The line is from an equivalent circuit model of the amplifier, transmission line, and cavity interaction.

cryogenic amplifier. This interaction is complicated; there are noise sources both in the cavity and in the amplifier, a transmission line connecting the two, and complex impedances looking into the cavity and amplifier. In order to understand the asymmetric structure, we developed an equivalent circuit model of the amplifier coupled to the cavity.

c. Equivalent circuit model for the resonant cavity coupled to the first cryogenic amplifier. Our equivalent circuit for the resonant cavity coupled to the amplifier is shown in Fig. 17. This equivalent circuit is approximate in that we have assumed that the amplifier (which is a relatively complicated balanced design) can be represented as a perfectly matched 50- Ω input impedance containing current and voltage noise sources. With P the power at the amplifier output and Δ the displacement from the center frequency of the power spectrum (in units of 125 Hz) the power spectrum from the equivalent circuit is

$$\frac{dP}{df}(\Delta) = \frac{a_1 + 8a_3(\Delta - a_5/a_2)^2 + 4a_4(\Delta - a_5/a_2)}{1 + 4(\Delta - a_5/a_2)^2}, \quad (18)$$

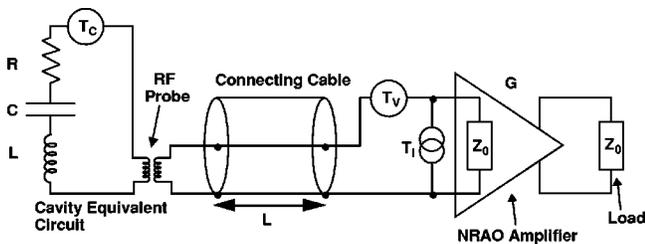


FIG. 17. An equivalent circuit model of the amplifier, transmission line, and cavity interaction. T_C is the noise contributed from cavity Johnson noise. T_V and T_I are the voltage and current noise contributed by the amplifier.

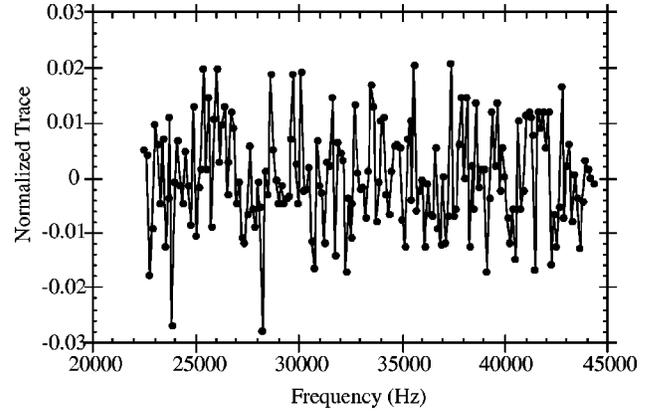


FIG. 18. A single trace after correcting for the receiver response by the equivalent circuit model of the amplifier, transmission line, and cavity interaction.

where $a_1 = T_c + T_I + T_V$, $a_2 = f_0/Q$, $a_3 = T_I + T_V + (T_I - T_V)\cos 2kL$, and $a_4 = (T_I - T_V)\sin 2kL$, with k the wave number corresponding to frequency $f = f_0 + 125 \text{ Hz } \Delta$, and the remaining parameters (the electrical length L from cavity to amplifier, and the amplifier voltage and current noise contributions T_I and T_V) are as shown in Fig. 17. The last coefficient a_5 is actually not an equivalent circuit parameter and requires some explanation: The receiver downconverts power from about the measured cavity resonant frequency to a band centered at 35 kHz. Therefore, in principle the 100th bin of the corrected spectrum should be the cavity resonant frequency. Actually, the cavity resonant frequency is centered by the DAQ system to about ± 500 Hz. The final parameter a_5 is therefore the displacement of the cavity resonant frequency (in 125 Hz units) from the 100th bin in the corrected trace.

The lines overlaid on the corrected traces of Fig. 16 show the result of the equivalent circuit model. The equivalent circuit description describes the smooth background well [52].

d. Fit parameters and the radiometer equation. The equivalent circuit description of the corrected traces gives an estimate of the average power for each. The residual bin-to-bin fluctuations about this average power level should be approximately Gaussian-distributed with rms deviations consistent with the radiometer equation [53]. We shall now check that this holds.

Recall that each trace is divided by its corresponding fitted average power; the resulting trace consists of fluctuations about one. To simplify later processing, we subtract one from each bin, leaving Gaussian-distributed fluctuations about zero. Figure 18 is the corrected trace from the top left corner of Fig. 16 normalized again from the five parameter equivalent circuit model.

We refer to the dimensionless fluctuations about zero in these traces as “deltas.” An important quantity for our analysis is the rms of these deltas for a single trace

$$\sigma = \sqrt{\frac{n}{i} \frac{\delta_i^2}{n}}, \quad (19)$$

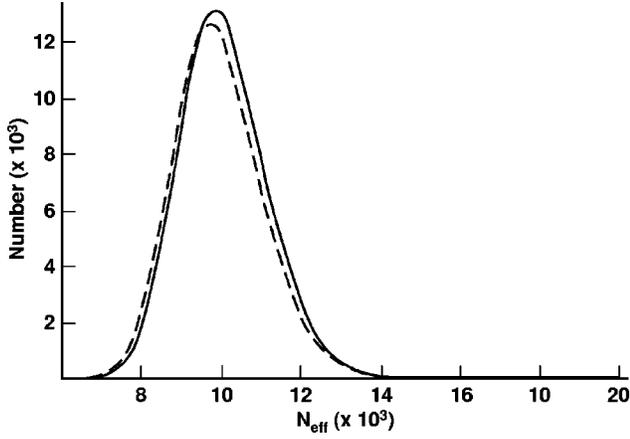


FIG. 19. Histogram of the inferred number of effective averages N_{eff} from the rms spread in the 175 bins in each trace. The line is the expectation for Gaussian noise.

where n is the number of bins and δ_i is the delta from the i th bin. There is a simple relationship between σ and N , the number of spectra in the linear average taken by the FFT to form the trace. This relationship comes from the radiometer equation [53]: $\sigma = 1/\sqrt{N}$, providing a simple relationship between the rms deviations in a normalized trace and the number of spectra averaged together to form the trace. Let $N_{\text{eff}} = 1/\sigma^2$ be the estimate of the number of averages from any one trace. Directly from the χ^2 probability distribution for σ^2 , we have the probability distribution for N_{eff} given by

$$P(N_{\text{eff}}) = \left(\frac{n}{2}\right)^{n/2} \frac{N_{\text{eff}}^{-(1+n/2)} e^{-n/2 N_{\text{eff}} \sigma^2}}{\sigma^n \Gamma(n/2)}, \quad (20)$$

where $n=175$ for the raw traces. Figure 19 is the measured distribution of N_{eff} from the rms deviations for all the normalized traces (approximately 250 000 in number) in the first run data. The measured distribution is centered near 10 000 averages, which is also the number of spectra averaged for each trace in the first run data. The curve is the expected N_{eff} distribution. The reasonable agreement suggests we do not have a large contribution from spurious non-Poisson or non-stationary noise sources.

Figure 20 is the same normalized trace shown in Fig. 18 except that, in software, we have simulated an axion candidate at 32.5 kHz. The corresponding power from this candidate peak in the normalized trace of height H_s is $P_s = H_s k_b T_N B$. Notice the height H_s of a signal seen in the normalized trace is inversely proportional to the system noise temperature T_N . Software peak injection of known power into the analysis chain enables us to establish confidence limits.

e. Factors affecting the relative weights of data from different raw traces. The change in TM_{010} mode frequency between adjacent traces is typically 1/15 of the receiver passband. As a result, any one 125-Hz frequency bin appears in many traces, and these data points are combined so as to maximize sensitivity. Four factors determine the weight assigned to a bin in combining the data. First, there is the system noise temperature T_N for the trace. Second, the num-

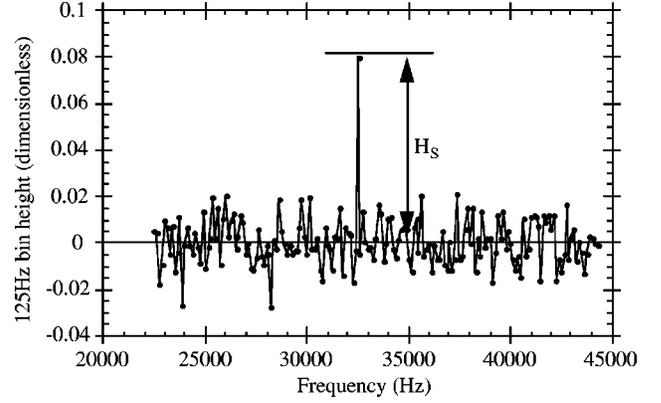


FIG. 20. A single trace from the raw data with an overlaid artificial single bin axion peak.

ber of spectra N averaged to form the trace. Third, the difference between the bin frequency and the resonant frequency of the TM_{010} mode. Finally, the power from axion-to-photon conversion, which may vary from trace to trace due to changes in the magnetic field, cavity form factor, etc.

We now consider the effect of the position of the TM_{010} mode on the weighting applied to a bin from a trace, that is, properly accounting for the frequency offset in bins from the peak of the cavity resonance curve. We define h as the ratio of the signal height in a bin from axion-to-photon conversion to the signal height seen if the bin were at the peak of the TM_{010} mode resonance. We have $h=1$ for a bin exactly on resonance, $h=0$ for a bin far off resonance, and between these extremes h is described by a Lorentzian shape. Recall that in operation of the experiment, the first local oscillator is set to center the cavity resonance in the middle of the trace. In terms of n , the number of 125 Hz wide bins to the center of the resonance, and Γ , the Lorentzian resonance width h is given by

$$h(n) = \frac{1}{1 + 4(125)^2 n^2 / \Gamma^2}. \quad (21)$$

Our weighting strategy in combining overlapping traces is as follows: First, consider just two overlapping traces; let the bin from the first (second) have power excess δ_1 (δ_2), Lorentzian factor h_1 (h_2), rms power fluctuations σ_1^W (σ_2^W), and hypothetical signal power P_1 (P_2). The weightings w_1 and w_2 that maximize signal-to-noise ratio (SNR) are [54]

$$\delta_{WS} = w_1 \delta_1 + w_2 \delta_2 = \frac{h_1 P_1}{(\sigma_1^W)^2} \delta_1 + \frac{h_2 P_2}{(\sigma_2^W)^2} \delta_2. \quad (22)$$

The signal-to-noise ratio in any one trace is $S = hP/\sigma$. The signal-to-noise ratio for a single bin in two overlapping spectra is

$$S_{WS} = \left\{ \frac{h_1 P_1}{\sigma_1} + \frac{h_2 P_2}{\sigma_2} \right\}^{1/2} = \sqrt{S_1^2 + S_2^2}. \quad (23)$$

The extension to more than two overlapping traces is straightforward.

3. Useful quantities from the combined data

In the previous section we described the weighting scheme. The weighted sum of δ_{WS} in a single bin of the combined data, is given by

$$\delta_{WS} = \sum_i \frac{h_i P_i \delta_i}{(\sigma_i^W)^2}, \quad (24)$$

where the index i refers to the i th spectrum contributing to the single bin in the combined data. The standard deviation for this weighted sum σ_{WS} is

$$\sigma_{WS} = \left\{ \sum_i \frac{h_i P_i}{(\sigma_i^W)^2} \right\}^{1/2}. \quad (25)$$

The data array in which we search for peaks is the ratio

$$\frac{\delta_{WS}}{\sigma_{WS}} = \frac{\sum_i h_i P_i \delta_i / (\sigma_i^W)^2}{\left\{ \sum_i h_i P_i / (\sigma_i^W)^2 \right\}^{1/2}}. \quad (26)$$

Notice that neither δ_{WS} nor σ_{WS} has units of watts; they are dimensionless as a result of multiplying the excess power δ_i contributing to a bin by a weighting factor w_i , having dimensions of $1/W$. To express the weighted-sum deviations in units of power, divide by the summed weights

$$\sum_i w_i = \sum_i \frac{h_i P_i}{(\sigma_i^W)^2}. \quad (27)$$

Using the summed weights to normalize the weighted sum of the deltas, we obtain the following expression for fluctuation about the mean power in watts, $\delta_j(W)$, in the j th bin of the combined data

$$\delta_j(W) = \frac{\sum_i h_i P_i \delta_i / (\sigma_i^W)^2}{\sum_i h_i P_i / (\sigma_i^W)^2}. \quad (28)$$

Similarly, the standard deviation in watts, $\sigma_j(W)$, in the j th bin of the combined data, is

$$\sigma_j(W) = \frac{\left\{ \sum_i h_i^2 P_i^2 / (\sigma_i^W)^2 \right\}^{1/2}}{\sum_i h_i P_i / (\sigma_i^W)^2}. \quad (29)$$

The expected signal power in watts P_S^W in the combined data is similarly

$$P_S^W = \frac{\sum_i h_i^2 P_i^2 / (\sigma_i^W)^2}{\sum_i h_i P_i / (\sigma_i^W)^2}. \quad (30)$$

Finally, the SNR in the combined data is obtained by adding in quadrature the SNR's from the contributing bins in the raw traces

$$\text{SNR} = \left\{ \sum_i \left(\frac{h_i P_i}{\sigma_i^W} \right)^2 \right\}^{1/2}. \quad (31)$$

4. The run 1 raw data

Production data were taken over the frequency range 701–800 MHz corresponding to a range of axion masses 2.90–3.31 μeV . This data consists of 4.2×10^5 raw traces saved to disk.

a. Cuts on the raw data. Three different cuts were used that together removed 1.4% of the raw traces. The first cut removed traces taken when the pressure in the resonant cavity exceeded 0.8 Torr. These relatively high pressures almost always developed as sudden rises, resulting in rapid variations in the cavity resonant frequency.

The second cut removed data that occurred when the change in cavity resonant frequency between successive spectra exceeded 7 kHz. Failing this cut was correlated with failing the first cut, as sudden changes in cavity resonant frequency were often caused by sudden changes in pressure in the cavity. Data taken during rapid frequency changes would suffer from poorly measured cavity Q and resonant frequency.

The third cut on the raw data eliminated traces written when the cavity temperature exceeded 2 K. Such a relatively high cavity temperature indicated a failure in the vacuum system or the detector warming because the liquid-helium supply system had failed.

b. The run 1 data. Rather than having a single continuous sweep over the frequency range, we scanned the frequency interval with several sweeps. This allows us to identify background from transient rf peaks (e.g., if a high-power peak appears at some frequency, but fails to reoccur at the same frequency weeks or months later, it is unlikely to be due to an axion).

Run 1 occurred in three distinct stages. First, three sweeps of the entire frequency span were made. Second, a fourth, less continuous sweep of parts of the span was made to raise the SNR across the frequency search range to a nominal value. Third, gaps left in the first four sweeps due to mode crossings were filled in with the cavity flooded with liquid ^4He .

c. Signal-to-noise ratio in the run 1 combined data. Figure 21 shows the expected power from KSVZ axions versus frequency and the corresponding rms noise. Figure 22 shows the signal-to-noise ratio versus frequency assuming an axion energy density $0.45 \text{ GeV}/\text{cm}^3$.

This signal-to-noise ratio assumes that all the signal power appears in a single 125-Hz bin. In the case where the signal power is thermalized over 6 bins, the rms noise level increases by a factor of $\sqrt{6}$. This implies that the run 1 ratio of the KSVZ signal power to the rms noise in a 750-Hz (6 bin) bandwidth is at least 4.7 everywhere in our frequency range.

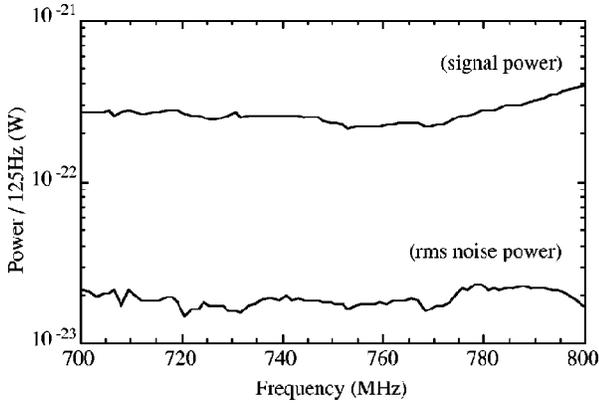


FIG. 21. Signal (upper line) and noise (lower line) power spectral density vs frequency, assuming that all the signal power appears in a single 125 Hz bin. If, as expected, the signal power is spread over 6 bins the noise power increases by a factor of $\sqrt{6}$.

d. The statistics of the run 1 combined data. The raw data from run 1 was combined according to the algorithm described above. In this combined spectrum, the bin contents deviate slightly from a mean-zero Gaussian distribution. Although the shape is Gaussian, the mean of the deviations is not zero, so we have not entirely removed the non-Gaussian components of the power spectra. This result is to be expected, because knowledge of the underlying physical parameters governing the shape of the power spectra is limited. For instance, there are many noise sources in the HFET amplifier, yet our equivalent circuit model uses only two. Furthermore, the model does not include the effects of the internal 90° hybrids and the directional coupler between the cavity and the amplifier input. Introducing more fit parameters is undesirable because axionlike signals in the spectra would be then unacceptably diluted in the fit. A second example of our imperfect knowledge of the contributions to power spectra shapes is the crystal filter. We have assumed that the transfer function of the receiver electronics is well known, and can be applied as an errorless correction to each raw trace. The temperature of the crystal filter was regulated to better than 0.5° C, yet it is possible that some crystal aging caused the transfer function to change over time. Alternatively, the gain of one of the other components in the

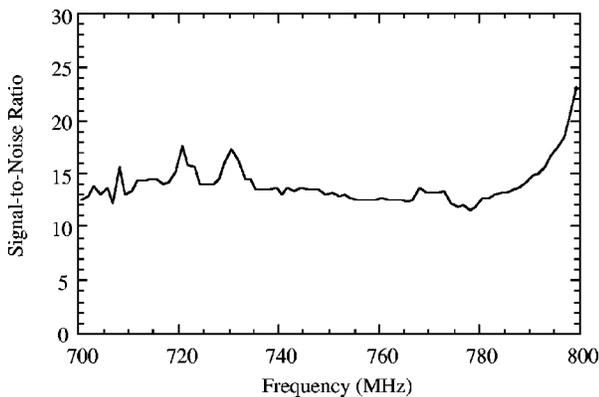


FIG. 22. Signal-to-noise ratio vs frequency, assuming that all the signal power appears in a single 125 Hz bin.

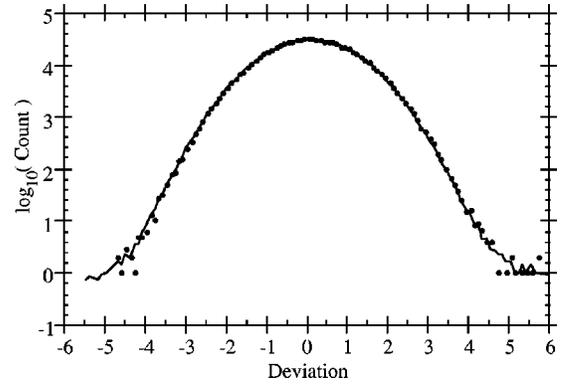


FIG. 23. Dispersion of single 125 Hz bin power about the mean for the run 1 data. The line is the prediction from Monte Carlo simulation.

receiver chain could change in time. Our receiver normalization corrections would then be inaccurate.

We found that the non-Gaussian component of the power spectra in the run 1 production data originates in the correction made to each power spectrum for the receiver response, likely due to a sudden aging of the filter crystals between calibrating the response during commissioning and run 1 data. The response function used throughout run 1 is shown in Fig. 13. Fortunately, throughout run 1 we found this effect had stabilized. We can therefore generate an effective receiver calibration by taking a linear average of all the traces from run 1 (power spectra that have been normalized both to the receiver response and the five parameter fit). That is, bin j of this average of N spectra is $(1/N)\sum_i^N P_{i,j}$. We refer to this average as the “run 1 fit residual.”

We hypothesize that for any single receiver normalized power spectrum, the distribution of power in 125-Hz frequency bins about the run 1 fit residual is Gaussian, and we test this hypothesis by simulation. We take each raw trace in the run 1 data and normalize the trace with the receiver response of Fig. 13 and fit the normalized trace to the five parameter cavity-amplifier model. We then keep the fitted curve and discard the real trace. We then build an artificial trace to replace the real one as follows: We start with the run 1 fit residual curve described above. We add to it computer generated Gaussian distributed noise of rms deviation 0.01 (as would be expected for a trace from an average of 10 000 power spectra). We then multiply this curve by the fit curve to the real trace. The resulting artificial trace is then analyzed exactly like a real trace from the data. It is fitted and normalized with the five parameter model, and the resulting normalized trace is added to the fake combined traces exactly as the real trace would have been.

This procedure is repeated for every raw trace until a complete set of artificial combined traces has been produced. To improve the statistics of the resulting data set, the procedure is repeated ten times, so the artificial set of combined traces, resulting from the simulation, has roughly ten times the statistics of the combined data. Figure 23 shows the histogram of the real combined data (points) overlaid by the simulation (line) adjusted for the different normalizations.

TABLE II. Persistent candidates in the 1-bin search channel. The first column gives the candidate frequency. The second and third columns give the heights of the candidate signals in run 1 and the rescan. The final column gives the results of investigation of the persistent candidate by manual scanning.

Frequency (MHz)	Run 1 Pk height	Run R Pk height	Comment and ID method
729.910 750	3.5655	17.958	T, A
748.944 125	3.7367	3.4041	N
750.014 000	3.5827	3.5450	N
771.218 500	28.893	3.6456	N
771.218 625	28.893	3.6456	N
771.234 375	23.576	5.3310	N
771.265 750	27.365	9.5139	N
780.299 625	3.5043	3.3814	N

5. The single-bin peak search

The single-bin peak search is most sensitive to signals with a bandwidth less than 125 Hz. The search statistic is power excess in any single 125-Hz frequency bin of the combined data. We assume that the distribution of single-bin power excesses in single bins of the combined data is purely Gaussian, so that the sensitivity to axions in this channel can be calculated analytically. The search proceeds in several steps. There is first a single-bin search in the run 1 data, followed by the single-bin rescan (“run R”) of candidates. Finally we carefully study any candidate frequency passing the cut in both scans to see whether it was due to axion-to-photon conversion.

a. Single-bin peak search in the run 1 data. The data sample is the array of power excess in units of standard deviation in each of the 792 000 125-Hz-wide frequency bins between 701 and 800 MHz. Single-bin candidates are selected by requiring the single-bin power exceed a fixed number of standard deviations above the mean power. Frequency bins passing this cut are then rescanned to see if the power excess is persistent. In the single-bin channel the same candidate selection criterion is applied to run 1 and run R.

In the single-bin search channel we use a threshold of 3.2σ power excess. As discussed above, the effect of drifts in the receiver passband response on the single-bin data distribution is to shift the mean of the Gaussian distribution upwards by 0.1σ . Therefore for a 3.2σ cut on the Gaussian data centered at $+0.1\sigma$, the threshold must be at $+3.3\sigma$ in the combined data set. This resulted in 538 single-bin candidates. The expected number of candidates n_c in a sample of N bins from a Gaussian distribution of unit width with threshold at p sigma is

$$n_c = \frac{N}{2} \left(1 - \operatorname{erf} \frac{p}{\sqrt{2}} \right) \quad (32)$$

or 544 candidates above $+3.2\sigma$.

The frequency of each of the 538 candidates in the 1-bin channel was rescanned, with enough traces acquired at each frequency to equal or exceed the run 1 SNR. The rescan data was combined according to the same algorithm as that used in run 1, and the same 3.2σ candidate threshold applied in

the combined run R data. Of the 538 candidates from the run 1 search, 8 reappeared as rescan candidates.

b. Investigation of persistent 1-bin candidates. Table II is a list of candidates surviving both the run 1 and the run R selection. The fourth column contains one or more code letters indicating how the persistent candidate was eliminated as a possible axion signal. Below is an explanation of each code letter:

A: Elimination by detection with an antenna. Here, the cable connecting the output of the room-temperature post amplifier to the input of the axion receiver is disconnected, and a stub antenna connected to the receiver input in its place. A trace is taken to see if the candidate corresponds to a background signal in the laboratory. If it does, this is a strong indication, but not proof, that the peak in the detector is due to rf contamination from outside the detector.

T: To establish that a peak is due to rf contamination, the two main sources of rf leakage into the cavity, the weakly coupled port and the cable connected to the directional coupler weak port, are terminated at the rf feedthroughs on top of the cryostat. The candidate is finally discarded if the peak disappears. More specifically, a candidate is discarded if it does not exceed the 3.2σ candidate threshold at 99.5% confidence following termination of test ports.

N: Nonrecurring peak. On-line software accumulates an average of at least 100 traces centered at the frequency of a persistent candidate. If the Gaussian probability is less than 0.5% that the bin height we measured is a power greater than 3.2σ in run 1 plus a large negative statistical fluctuation consistent with no peak observed in the 100 averaged traces, then we conclude that the peak seen in run 1 was due to a transient rf source and the candidate is rejected.

As can be seen from Table II, all of the candidates surviving run 1 and run R were rejected. Figure 24 shows the resulting axion-to-two-photon couplings excluded as a function of axion mass at 90% confidence assuming the axion line is significantly narrower than 125 Hz and a halo density of 0.45 GeV/cm^3 .

6. The 6-bin peak search

The 6-bin (750 Hz) search is well matched to the line shape of thermalized galactic halo axions. It is based on a

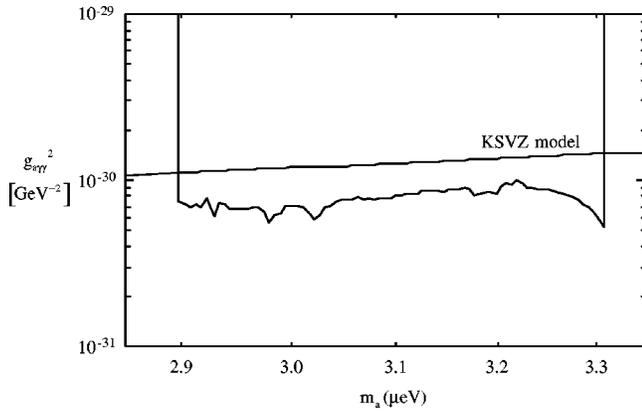


FIG. 24. Axion-photon couplings excluded at 90% confidence, assuming the axion line is significantly narrower than 125 Hz.

power excess in the sum of any 6 adjacent frequency bins in the combined data (the 6 bins covering a total bandwidth of 750 Hz). There are three factors which, while not affecting the single-bin search, are important for the 6-bin search.

First, because the tail of the assumed Maxwellian halo axion velocity distribution extends the axion line shape to high frequency, we do not expect all axions would convert into photons within a 750-Hz search bandwidth; 20% of the power is in the tail of the distribution above a frequency $m_a + 750$ Hz. At most, 80% of power from axion-to-photon conversion is detected in any adjacent 6 bins. Second, if a 750-Hz bandwidth axion signal is present in a raw trace, the five parameter fit to the trace baseline will deviate from the shape it would have in the absence of the signal. For KSVZ axion signals, the resultant dilution of the signal strength between a raw trace and the resulting normalized trace is between 10% and 15%. The third effect concerns the statistic used to search for the 6-bin power excess. The simplest procedure is to separate n 125-Hz-wide bins into $n/6$ nonoverlapping 750-Hz-wide bins, each of which is the sum of 6 underlying 125-Hz-wide bins. Unfortunately, an axion signal would be split between two adjacent 750-Hz bins. Depending on the position of the peak with respect to the chosen bins, the power excess in a single 750-Hz bin is degraded by up to 50%.

a. Use of overlapping frequency bins. To avoid the third problem above, we take overlapping 6-bin segments of combined data. The i th 6-bin segment is the sum of the i th to the $(i+5)$ th 125-Hz bins, the $(i+1)$ th 6-bin segment is the sum of the $(i+1)$ th to the $(i+6)$ th 6-bin segment, and so on. Hence a set of n 125-Hz bins yields $(n-5)$ 6-bin segments. This algorithm ensures that the bulk of the power in the axion line will be deposited in at least one of the 6-bin segments. The penalty of this gain in SNR over the previous method is that there are now $\sim n/6$ 750 Hz wide bins compared to the $n/6$ from the above nonoverlapping method. On the other hand, adjacent bins in the array are highly correlated, so that a thermalized Maxwellian axion signal will yield power excesses in several neighboring 6-bin sets. These adjacent candidates are rescanned together, without incurring extra rescan time. We refer to this procedure of

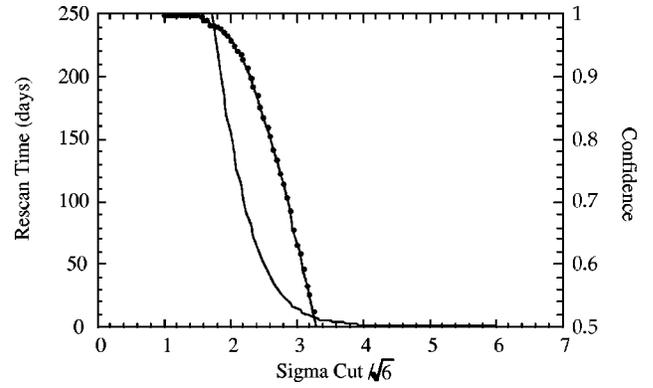


FIG. 25. (Right axis) The fraction of injected artificial detected peaks vs the candidate threshold, and (left axis) the time needed to rescan noise candidates passing the candidate threshold. The candidate threshold applied in this analysis yields a confidence greater than 90%.

constructing overlapping 6-bin segments as “coadding” the data.

The search starts by setting a candidate threshold at a high value ($6\sqrt{6}\sigma$ was used in our search). If any bin exceeds the threshold, then the frequency of that bin is added to a rescan list. This bin and the surrounding 16 bins are flagged so that none of the nearby correlated bins exceeding the threshold are duplicated on the rescan list. The candidate threshold level is then reduced by a small amount ($0.1\sqrt{6}\sigma$) and the process repeated until some target threshold is reached. We refer to this technique as a “running sigma cut.”

One might guess that if the same threshold level were set in the $n/6$ uncorrelated 6-bin segments generated using a simple 6-bin cut, and in the n correlated 6-bin segments generated using a running sigma cut, there would be the same number of candidates in the two cases since the running sigma cut prevents two correlated bins from both exceeding the threshold. However, in simulation studies, a search on n correlated bins with a running sigma cut consistently yielded 2–3 times as many candidates as a search on $n/6$ uncorrelated bins at the same threshold level.

Owing to its greater sensitivity, the correlated bin search is adopted in our 6-bin search channel. The data after passing through the running sigma cut filter is non-Gaussian, and the sensitivity to axion signals using this search method must be determined from simulation.

b. Simulation of the 6-bin peak search in the run 1 data. The actual candidate threshold level required in the run 1 data was determined from Monte Carlo simulation. The simulation generates a set of combined traces with artificially injected KSVZ axionlike signals injected at a random set of frequencies. The array of power excesses from the combined traces is coadded to make correlated simulated 6-bin data, and a running sigma cut peak search is implemented on this data set.

Figure 25 is a plot of the fraction of the injected peaks located by the peak search on the Monte Carlo generated combined traces versus cut level. From this simulation, a candidate threshold in the 6-bin peak search on the run 1 data

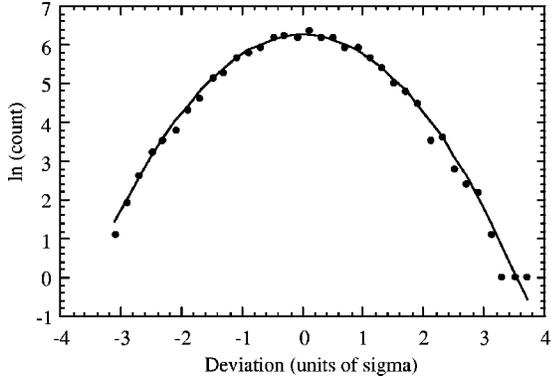


FIG. 26. Dispersion of single 125 Hz bin power about the mean (points) for the (1+R) combined data. The curve shows the expectation from Monte Carlo simulation.

was set at 2.25σ . There were 6535 candidate frequencies exceeding the threshold.

c. Re-scanning the 6-bin candidates. Each candidate was subsequently rescanned in run R. The run R data was combined using the same data combining algorithm as in run 1. Next, the run 1 and run R data were combined with each other into a data set called the “1+R combined data.”

The combined data from run R differs from that in run 1 in one important aspect: the effect of the receiver drift on the run R data is more severe than on the run 1 data. This is because in run R, the cavity resonant frequency is held at a fixed value while a single candidate frequency is rescanned, hence, the same systematic shift is applied to every new trace that is added to the combined data. If the systematic shift is δ_{SYS} and the rms spread in Gaussian fluctuations in a single trace is σ , then the ratio R_s of the amplitude of the systematic shift to the rms spread in the Gaussian noise in the sum of n traces weighted equally is $R_s = (\delta_{\text{SYS}}\sqrt{n})/\sigma$.

In order to correct the run R combined data for the receiver drift, we exploit the fact that in these data the rescan frequency is always close to the center of the raw traces, and that the systematic shift in spectra taken within a few weeks of each other is nearly the same. We take the set of run R combined data points near rescan frequencies, consisting of 6535 data points.

Associated with each of these data points is the average time when the underlying traces contributing to each point were taken. These data points were then rearranged into increasing average time order. For each of these points, a receiver drift correction was computed from the average of 300 nearby points.

With R_s a good approximation and the raw traces contributing to each candidate point taken within a short enough time so that the evolution of the systematic shift is small, the statistics of the corrected combined data will be Gaussian. Figure 26 is the histogram of the dispersion at candidate frequencies in the combined data from run R; it agrees with the Gaussian expectation.

d. The algorithm for combining the run 1 with the run R data. The run 1 combined data and the run R combined data near each rescan frequency are combined using a weighting scheme similar to that used in combining the raw traces. The

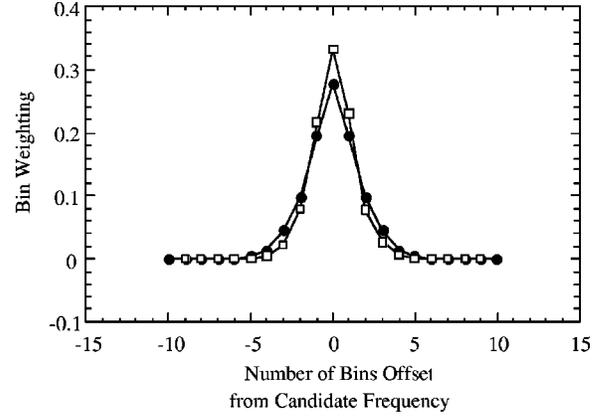


FIG. 27. Neighboring bin weighting functions for (circles) the (1+R) combined data and (squares) the (1+R+P) combined data. These functions represent the probability, due to noise, that an axion signal at a particular bin appears instead in a nearby bin. The circles show the initial bin relocation after run 1, and the squares the relocation after run (1+R).

41 nearest 125-Hz frequency bins surrounding each rescan frequency are taken from each of the two data sets. Matching bins from the two data sets are combined as follows. With δ_1 (δ_R) the contents of one of the 41 run 1 (run R) bins, the weighted sum δ_{1+R} of a bin from run 1 with the corresponding bin from run R is given by

$$\delta_{1+R} = \frac{\delta_1 S_1 + \delta_R S_R}{\sqrt{S_1^2 + S_R^2}}. \quad (33)$$

The SNR in the combined data set is the quadrature sum of the component SNR’s.

The 6-bin search is more subtle. In constructing the 6-bin search statistic, we need to consider that the axion peak is many 125 Hz bins wide, and each bin also contains Gaussian noise. A peak may pass the run 1 cut due to a power excess in the i th bin; if in the rescan the i th bin happens to contain a large negative noise fluctuation while the j th bin (where $|j-i| < \text{axion linewidth}$) exceeds the threshold, then by only re-examining the same bin where we first found the power excess, we would miss the axion. It is most likely that the same bin will pass the cut in the two data sets ($i=j$), with the probability of two bins separated by $(i-j)$ passing cuts in runs 1 and R decreasing as $i-j$ increases.

We therefore developed a candidate search in the run (1+R) combined data as follows. First the arrays of 41 bins about each rescan frequency are coadded, generating overlapping 6-bin sets. The 41 element array is then multiplied by the function in Fig. 27, with the peak of the function centered on the rescan frequency. This function represents the probability, due to noise, that a simulated axion signal injected in a particular bin appears instead in a nearby bin. The two curves show the initial bin relocation after run 1, and the second relocation after run (1+R). We refer to the product of the overlapping six bin sets near each rescan frequency and the weighting function as the weighted (1+R) distribution. Candidates are those frequencies in the upper

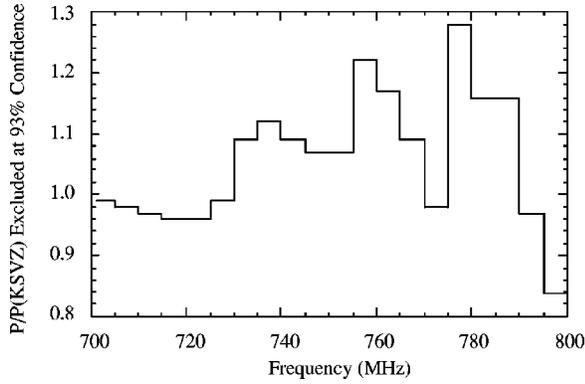


FIG. 28. Axion power (in units of KSVZ axion power) excluded at 93% confidence after serially applying the (1) and (1+R) candidate thresholds.

10% probability in the weighted (1+R) distribution; the corresponding candidate threshold level in the weighted (1+R) distribution is $2.06\sqrt{6}\sigma$. The number of surviving 6-bin candidates is 654. These candidates are scanned for a third time in “run P.”

e. 6-bin search sensitivity in the run (1+R) data. The sensitivity of the experiment when the axion signal exceeds both the $2.25\sqrt{6}\sigma$ threshold in the run 1 data and the $2.06\sqrt{6}\sigma$ threshold in the run (1+R) data is determined from Monte Carlo simulation. In this simulation, combined traces are created that replicate the run 1 and the run R data, except that axionlike signals are injected at a random subset of the rescan candidate frequencies (approximately 46 per 10 MHz of bandwidth) passing the run 1 cut. (We inject the axionlike signals at this stage because in the run R data, only frequencies on the run 1 candidate list have a good signal-to-noise ratio.) The peak search procedure is the same as that applied to the data. First a $2.25\sqrt{6}\sigma$ threshold is set for candidates from the fake combined traces corresponding to the run 1 data. A list is made of the frequencies that exceed the threshold and are correlated with the frequencies at which artificial peaks were injected. Then the simulated combined traces, corresponding to runs 1 and R, are combined using the same algorithm as for combining the run 1 and R data. A $2.06\sqrt{6}\sigma$ candidate threshold is set for these simulated combined traces. Finally, a list is made of the frequencies of injected signals that exceeded both thresholds.

This simulation is repeated for various power levels of the injected signals. The full frequency range 701–800 MHz is split up into 20 ranges. The result of the simulation is the fraction of peaks injected in each of the 20 ranges exceeded both thresholds as a function of the injected signal power. Inverting this result gives the signal power detectable at any confidence level as a function of frequency.

The Monte Carlo simulation was implemented in the C programming language and ran on the MIT Laboratory for Nuclear Science “Abacus” computer facility. Figure 28 is a plot of the power level that exceeds both thresholds in the Monte Carlo simulation at 93% confidence as a function of frequency. The power level is in units of power expected from KSVZ axions assuming a halo density of 0.45 GeV/cm^3 . Referring to Fig. 22, the regions where we

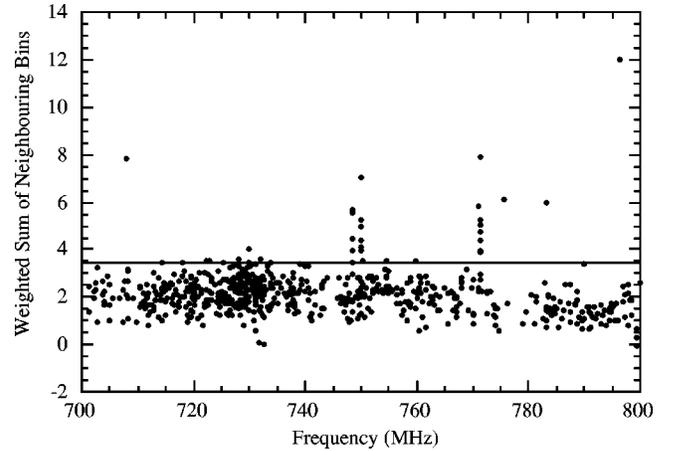


FIG. 29. The weighted sum of bins in the neighborhood of persistent candidate frequencies. The line is the candidate threshold.

are most sensitive corresponds to the regions of highest SNR in the run 1 combined data.

f. The run P combined data. The 654 6-bin candidate frequencies passing the $2.06\sqrt{6}\sigma$ threshold in run R were scanned again. We refer to this scan as run P. Three distinct phases occurred in the run P data taking. First persistent candidates between 778 MHz and 800 MHz were scanned. Next, the cavity was filled with liquid helium and persistent candidates near mode crossings were scanned. Finally, the rest of the run P candidates were scanned. There resulted 28 traces, each of 10 000 power spectra, for each candidate frequency.

To correct for the receiver drift in the run P data, the receiver passband response was re-measured every month, and the corresponding measurement used to normalize traces taken within ± 2 weeks. The run P combined data set is made using a minor extension of the procedure used in making the run 1+R combined data. With δ_1 , the deviation from mean power in one of the 41 bins near a rescan frequency in the run 1 combined data and with δ_R (δ_P) the corresponding power excess in the run R (P) combined data, and the signal-to-noise ratios at that frequency by S_1 , S_R , and S_P , respectively, the weighted sum δ_{1+R+P} of the three data streams is

$$\delta_{1+R+P} = \frac{\delta_1 S_1 + \delta_R S_R + \delta_P S_P}{\sqrt{S_1^2 + S_R^2 + S_P^2}}. \quad (34)$$

The SNR of the combined data is again the quadrature sum of the component SNR’s.

The 41 bin regions about each persistent candidate frequency in the run (1+R+P) data were coadded and again weighted according to a weighted sum of these bins, and shown in Fig. 27 as the weighting function. The weighting function used for the run (1+R+P) data is narrower than the one used to process the run (1+R) data. This is because the SNR in the run (1+R+P) data is higher, hence it is less likely that fluctuations result in a bin far away in frequency from the bin with the maximum signal power.

Figure 29 is the weighted sum for each rescan frequency in the run (1+R+P) data. Also shown is the candidate

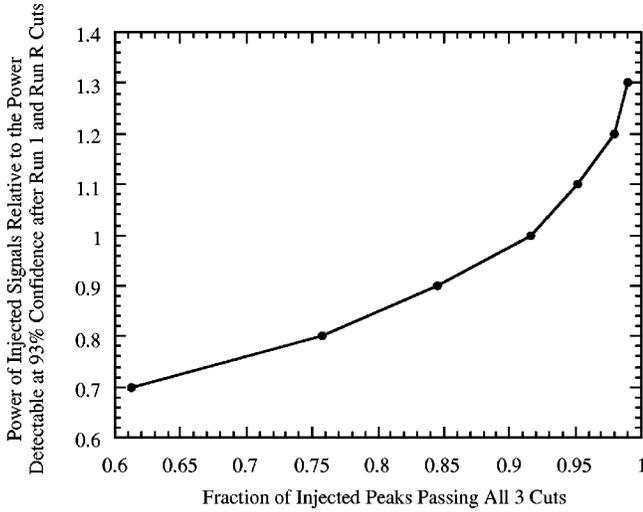


FIG. 30. Power of injected artificial peaks (in units of KSVZ axion power) vs the fraction of peaks passing all three candidate thresholds. The KSVZ axion would be detected at over 90% confidence.

threshold that determines which frequencies to re-examine. These candidates were carefully examined individually to determine whether the peaks at those frequencies were consistent with axion-to-photon conversion.

The effect of the run P candidate selection procedure is determined using a simulation similar to that of the run (1+R) combined data. Axion signals were added to a set of Monte Carlo generated raw traces similar to the run 1, run R, and run P data. The power of injected axion signal depended

on frequency; the injected power at a frequency f was set equal to the signal power detectable at 93% confidence in runs 1 and R, multiplied by a scale factor. The simulation was run at various values of the scale factor. Figure 30 is a plot of the fraction of signals passing the cuts in the run 1, run (1+R), and the run (1+R+P) combined data as a function of the scale factor. The 6-bin search channel is sensitive to KSVZ axions at over 90% confidence.

g. Examination of surviving 6-bin candidates. Table III shows all the 6-bin candidate frequencies surviving the cuts in runs 1, R, and P. The second, third, and fourth columns in the table are the candidate peak heights in the three data sets. The fifth column encodes the method used to eliminate that candidate. For examining these few remaining candidates, an on-line trace combining software is used. Many raw traces are acquired at a candidate frequency and the average of these traces is computed in real time. Then, overlapping sets of 6 bins are coadded for the 6-bin peak search.

“R, T”: The axion receiver is disconnected from the long cable connecting it to the room-temperature post amplifier. A stub antenna is attached and one or more traces are acquired at the candidate frequency. The flag R means that a radio peak was detected in the lab at the same frequency as the candidate. This hints there may be leakage of a radio signal into the axion detector. The flag T means that for such candidates, the weakly coupled port and the line to the cryogenic directional coupler weak port are terminated on top of the cryostat. If the peak then disappears, it was eliminated as a candidate.

“N”: During this examination no peak was seen at this frequency. To establish that there is no peak at high confi-

TABLE III. A list of candidate frequencies surviving cuts in runs 1, R, and P. The last column gives the method by which the candidate frequency was eliminated. Candidates marked with an asterisk were below the run P candidate threshold (3.5σ), but were examined anyway.

Frequency (MHz)	Run 1 Pk height	Run R Pk height	Run P Pk height	Method of elimination
707.792 250	2.2633	7.0965	7.8610	R, T
714.362 250*	2.2589	2.2014	3.4795*	N
718.078 125*	2.2712	2.6417	3.4418*	N
722.143 875	3.2541	3.1852	3.5557	N
729.489 750*	2.4225	2.1543	2.9880*	N
729.910 750	3.3062	2.6648	4.0381	R, T
731.953 750	2.6281	2.5627	3.6342	N
740.692 500*	2.7180	2.0656	3.2812*	N
748.426 625	6.4740	6.4740	5.5094	N
748.434 250	6.9001	5.3010	5.6607	N
750.427 125	6.2973	4.2480	3.5190	R, T
754.498 125	3.4347	3.4347	3.5136	N
759.703 000	3.2476	2.7648	3.5429	N
775.749 500	8.6955	4.8389	6.1293	R, T
771.249 375	15.559	9.1932	5.2231	R, T
783.249 375	8.5138	5.6073	5.9795	R, T
790.000 000*	2.4320	2.2307	3.3705*	N
796.315 000	3.1884	2.8133	12.017	N
799.976 250*	5.4750	3.1005	2.5656*	R, T

dence, the coadded power spectrum is calibrated by calculating the power level corresponding to 1σ rms in the Gaussian noise background from the amplifier and cavity noise temperature. Figure 21 is used to determine the signal power. The actual height of the candidate frequency bin is measured, and if the probability is less than 0.5% that the bin height seen is from a signal of the power level from Fig. 21, then the candidate is eliminated.

After this examination, none of the 6-bin persistent candidates remained.

h. The exclusion limit in the 6-bin search. The sensitivity to axions in the 6-bin search is deduced from the results of the Monte Carlo simulation of Fig. 28. A slight dilution in candidate selection efficiency from the run P cut at 3.5σ reduces the confidence level to just above 90% (see Fig. 30 and the preceding discussion). Recall that the signal power in our cavity is proportional to the square of the axion photon coupling constant $g_{A\gamma\gamma}^2$, and together with the relation between axion mass and PQ symmetry breaking scale, we obtain the following expression for the limit on $g_{A\gamma\gamma}^2$ imposed by this experiment:

$$g_{A\gamma\gamma}^2(\text{excluded})^2 = (3.6 \times 10^{-7} \text{ GeV}^{-1})^2 \left(\frac{m_A}{1 \text{ eV}} \right)^2 \frac{2P_{\text{excl90}}}{P_{\text{KSVZ}}}, \quad (35)$$

where P_{KSVZ} is the power expected in the KSVZ axion model and P_{excl90} is the signal power excluded at 90% confidence.

Figure 31 shows the range of $g_{A\gamma\gamma}^2$ excluded at 90% confidence by the 6-bin search. On the main figure, the independent variable has been divided by m_A^2 to flatten the exclusion line. The inset shows the results from the two previous pilot experiments [33]. Also shown are the KSVZ and DFSZ model predictions. This experiment has achieved an im-

provement in sensitivity of two orders of magnitude over previous experiments, and is the first such search to probe realistic axion models.

7. The search for extremely narrow lines

According to the radiometer equation, the sensitivity of our search can be increased if strong and narrow spectral lines exist.

For our nominal experimental conditions, the integration time for each tuning rod setting is around 60 s. The fractional Doppler shift due to the Earth's rotation in this time is $\sim 10^{-12}$ and amounts to \sim mHz (less than the inverse integration time).

Besides the highest-resolution spectrum of bandwidth 19 mHz, two more high-resolution spectra (of Nyquist resolutions 152 and 1200 mHz) are formed by combining adjacent bins of the highest-resolution spectrum. The three high-resolution spectra were examined for peaks on-line by calculating the mean and standard deviation of each 100 bin wide segment. Bins with an excess of 11σ in the unaveraged spectrum, and excesses of 5σ and 4σ in the two lower resolution spectra (152-mHz and 1.2-Hz binwidth, respectively) were flagged and saved for further analysis.

After completion of all frequency scans, the frequency coverage was fairly uniform with every frequency sampled several times. Any undersampled regions were rescanned separately to achieve a minimum of four samples at each frequency.

In order to reduce the number of candidate peaks, a coincidence search was performed. Since the contributing spectra were generally taken at different times, the annual frequency modulation of the narrow axion lines must be considered. We required the peaks to fall within a variable frequency window of size $\Delta f/f = 4 \times 10^{-7} |\sin(\pi\Delta t/365 \text{ days})|$, where Δt is the time separation.

To reduce the number of candidates for subsequent rescans, the thresholds for coincidence pairs at the three resolutions was set to 15σ , 8σ , and 5σ , respectively. This pro-

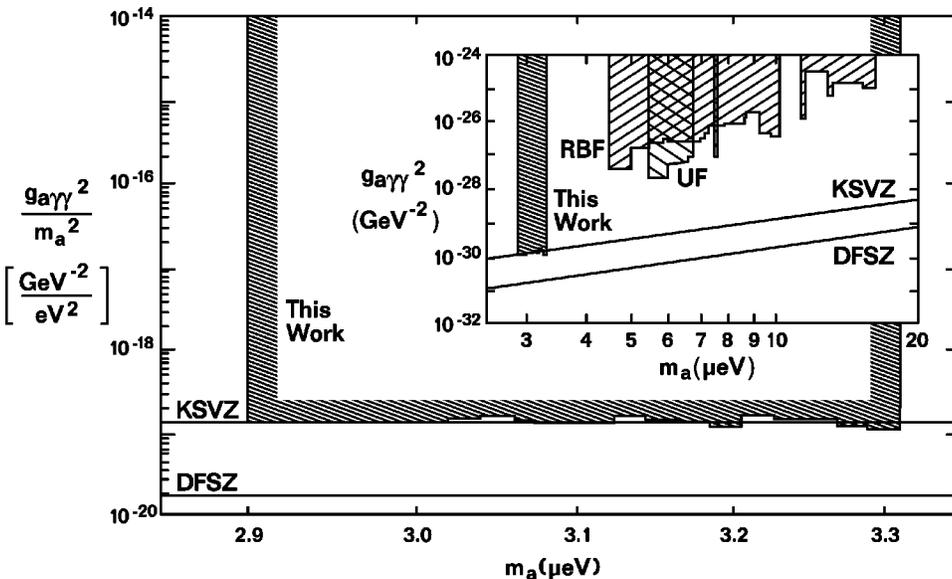


FIG. 31. The range of $g_{A\gamma\gamma}^2$ excluded at 90% confidence by the 6-bin search. On the larger figure, the independent variable has been divided by m_A^2 to flatten the exclusion line. The inset shows the results from the two previous pilot experiments. Also shown are the KSVZ and DFSZ model predictions.

duced a total of several hundred candidate frequencies for the three resolutions. We later examined each peak by taking six additional spectra. If the peak persisted, further tests were performed, similar to the 6- and 1-bin searches, by disconnecting or attenuating various diagnostic coaxial cables leading into the experiment dewar. Some spurious peaks disappeared with this procedure, while the remainder were matched to strong narrow band signals in the environment.

The minimum detectable power in a bin of bandwidth B is given by $P_{\min} = k_B T_s B N_\sigma$, where N_σ is the number of standard deviations above the mean. Because the signal can fall between two bins, the detectable power is smaller than P_{\min} . A simulation shows that at an approximate 90% confidence and with the signal frequency falling randomly between two bin frequencies, the least power in the strongest of the two bins for two out of four trials is $0.7 \times P_{\min}$. Thus, the minimum detectable power in for the three different resolutions are approximately $P_{\min} \approx 3 \times 10^{-23}$ W, 5×10^{-23} W, and 9×10^{-23} W, respectively, assuming a system noise temperature of $T_s \approx 5.5$ K. This represents sensitivity to approximately 10%, 17%, and 30% of the total KSVZ axion signal power.

III. THE U.S. rf CAVITY AXION SEARCH: HIGHER SENSITIVITY UPGRADE

A. dc SQUID amplifiers in the axion search

For a given axion-photon coupling, the scan rate improves with decreasing system noise temperature as T_s^{-2} , and for a given scan rate, the power sensitivity, which can be reached, is proportional to $\rho_A g_A^2 \gamma \gamma$, and goes as T_s . We are therefore planning to improve the scan rate and sensitivity by lowering the system noise temperature $T_s = T_n + T$, i.e., the electronic noise temperature plus the cavity blackbody temperature. In the present experiment, $T_n \sim T \sim 1.5$ K, and thus $T_s \sim 3$ K at best. Recently, however, the Berkeley group has developed dc superconducting quantum interference device (SQUID) amplifiers in the 100–1000 MHz range specifically for the axion experiment. This section is a brief description of these SQUID amplifiers, and the new results.

The essence of the dc SQUID is shown in Fig. 32(a). Two junctions are connected in parallel, breaking a superconducting loop of inductance L . Each junction has a capacitance C and is shunted by a resistance R . The I – V characteristics are shown in Fig. 32(b) for $\Phi = n\Phi_0$ and $(n + 1/2)\Phi_0$, where Φ is the external flux applied to the loop, n is an integer, and $\Phi_0 = h/2e \approx 2.07 \times 10^{-15}$ Wb is the flux quantum. When the SQUID is biased with a constant current (greater than $2I_0$, where I_0 is the critical current of each junction), the voltage across the SQUID oscillates with period Φ_0 as Φ is increased [Fig. 32(c)]. To maximize the performance of the amplifier, the SQUID is biased at the steep portion of the V – Φ curve, where the flux-to-voltage transfer coefficient $V_\Phi \equiv |\partial V / \partial \Phi|_I$, is a maximum. Thus, the SQUID produces an output voltage in response to a small input flux $\delta\Phi$ ($\ll \Phi_0$), and is effectively a flux-to-voltage transducer.

The signal and noise properties of the SQUID are obtained by solving numerically the equations relating the volt-

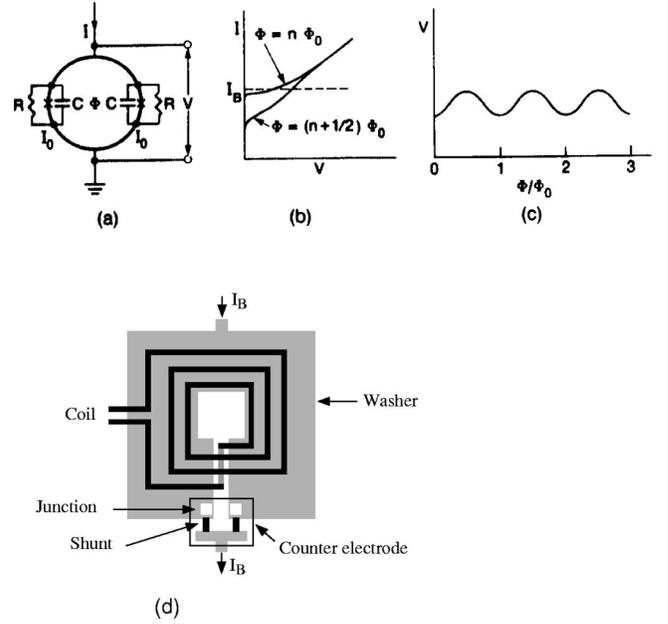


FIG. 32. Left: (a) The dc SQUID; (b) I – V characteristics; (c) V vs Φ/Φ_0 at constant bias current I_B ; (d) square washer dc SQUID.

age across the SQUID to the bias current and the applied flux [55]. The noise arises from Nyquist current in the shunt resistors, which has a current spectral density $4kT/R$. One finds that V_Φ peaks smoothly as a function of bias current with a shallow maximum around $(2n + 1)\Phi_0/4$. Optimum performance is found for $\beta_L = 2LI_0/\Phi_0 = 1$.

For a typical SQUID at 4.2 K with $\beta_L = 1$ and $2\pi kT/I_0 \Phi_0 = 0.05$, one finds the flux-to-voltage transfer coefficient is given approximately by $V_\Phi \approx R/L$. The corresponding white voltage noise across the SQUID has a spectral density $S_v(f) \approx 16kTR$, approximately 8 times the Nyquist noise in a resistance $R/2$. The spectral density of the equivalent flux noise in the SQUID is thus $S_\Phi(f) = S_v(f)/V_\Phi^2 \approx 16kTL^2/R$. It is often convenient to introduce the flux noise energy:

$$\mathcal{E}(f) = S_\Phi(f)/2L \approx 9kTL/R. \quad (36)$$

These expressions for $S_v(f)$, $S_\Phi(f)$, and $\mathcal{E}(f)$ show that the noise power of the SQUID is expected to scale with the temperature T . For a representative SQUID at 4.2 K with $L = 100$ pH, $I_0 = 10$ μ A, and $R = 4\Omega$, one finds $V_\Phi = 80$ μ V Φ_0^{-1} and $\mathcal{E}(f) \approx 1.3 \times 10^{-32}$ J/Hz.

The most widely used configuration for the dc SQUID [56] is shown schematically in Fig. 32(d). The loop of the SQUID consists of a square washer of niobium with a slit through one side. The two Josephson junctions are grown on each side of the slit near the outer edge, and connected with a second niobium film to form a closed loop. The bias current is passed between the SQUID washer and the counter-electrode of the junctions. The magnetic flux is applied to the SQUID by passing a current through a spiral coil deposited over the square washer, with an intervening insulating layer.

To form a tuned amplifier, one connects the voltage source in series with a capacitor and the input coil of the SQUID to form a tuned circuit. The lowest noise temperature

at the resonant frequency of the tank circuit $\omega/2\pi$ is achieved for a source resistance $\omega M_i^2 V_\Phi / R$, where M_i is the mutual inductance between the input coil and the SQUID loop. Under these conditions, one finds an optimized noise temperature

$$T_n \approx 7\omega T / V_\Phi \approx 7T/G, \quad (37)$$

where $G \approx V_\Phi / \omega$ is the corresponding power gain. Experimental results in good agreement with these predictions have been attained at frequencies up to about 100 MHz [57]. Unfortunately, at higher frequencies, the effects of parasitic capacitance between the input coil and the SQUID washer cause a substantial roll-off in the amplifier gain [58]. Although attempts have been made to circumvent this problem by placing the input coil inside the hole in the SQUID [59], because of the decreased coupling between the coil and the SQUID, the resultant gain is quite small.

An alternative way of achieving high gains and low noise temperatures at high frequencies is to configure the SQUID input resonantly in order to tune out the parasitic loading. Here, the input coil is used as a microstrip resonator [60], the input signal is no longer coupled to the two ends of the input coil, but rather between one end of the coil and the SQUID loop, which acts as a ground plane for the coil. The microstrip resonator is thus formed by the open-ended stripline including the inductance of the input coil and its ground plane and the capacitance between them.

A number of such SQUID amplifiers have been fabricated and operated. These started as conventional square-washer SQUID's, with inner and outer dimensions of 0.2 mm \times 0.2 mm and 1 mm \times 1 mm; the 31-turn input coil had a width of 5 μ m and a length $l \approx 71$ mm. Since the conventional washer SQUID is an asymmetric device (the two Josephson junctions are situated close together rather than on opposite sides of the SQUID loop), one can either ground the washer or ground the counterelectrode close to the Josephson junctions. Using the washer as a ground plane for the input coil suggests one should ground the washer. However, it is also possible to ground the counterelectrode and have the washer at the output potential. In this case, depending on the sign of V_Φ , one can obtain either a negative or positive feedback from the output to the input. For the sign of the transfer function (denoted V_Φ^+) the power gain was 18 dB at 620 MHz and the Q of the input resonator was about 31.

The SQUID's were mostly operated with the counter electrode grounded and the washer at output potential. Figure 33(a) shows the gain achieved for four different lengths of coil on the same SQUID. The gain is typically 18 dB.

The noise temperature of a number of SQUID amplifiers were measured referenced to a heated resistor as a noise source. Figure 33(b) shows the gain and noise temperature for an 11-turn SQUID. On resonance, the gain is 22 dB and the noise temperature $T_n = 0.9$ K \pm 0.3 K, measured at a bath temperature of 4.2 K. This is the overall system noise temperature, which includes a 0.4 K contribution from a room-temperature post amplifier.

In order to reduce this noise contribution, a single-stage HFET amplifier was operated at a bath temperature of 1.8 K.

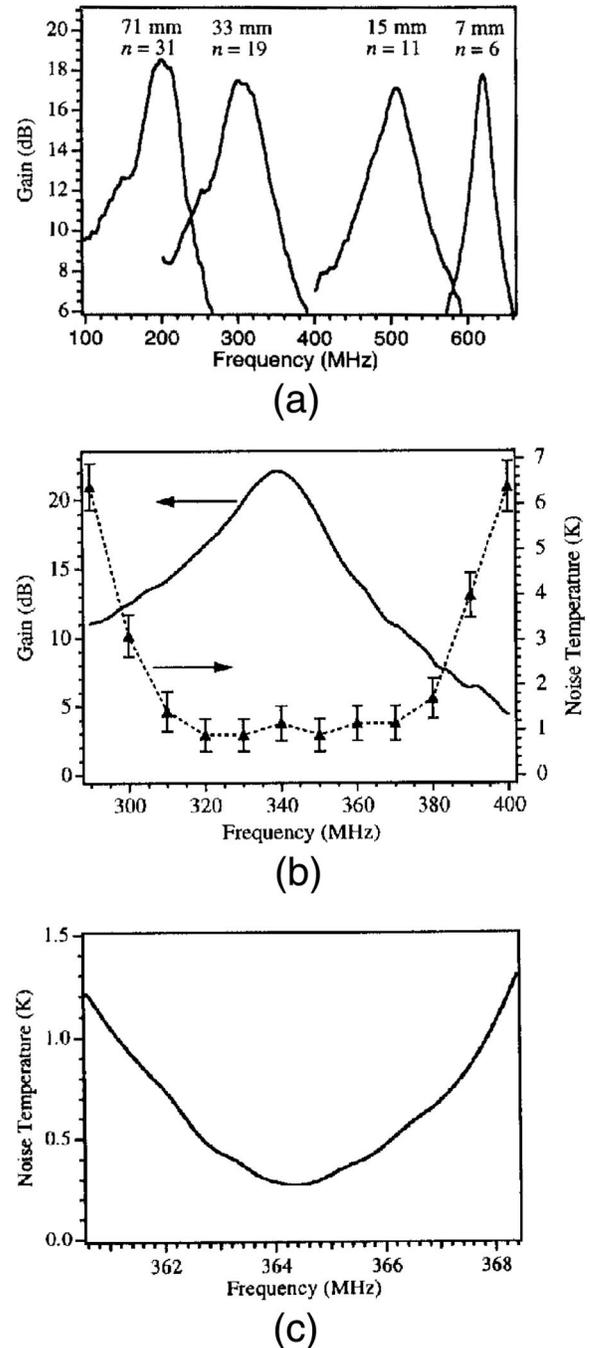


FIG. 33. (a) Gain vs frequency for four coils on the same SQUID, with signal applied to the innermost turn. Data are for counterelectrode grounded and V_Φ^+ . (b) Gain and noise temperature vs frequency for an 11-turn SQUID cooled to 4.2 K. (c) Noise temperature vs frequency for a 29-turn SQUID cooled to 1.8 K, with a HFET postamplifier at the same temperature.

The lowest noise temperature achieved with this amplifier was about 12 K. Two SQUID amplifiers were characterized, one with a resonant frequency of 250 MHz and a second with 365 MHz. When the SQUID's were cooled to about 1.8 K, the gain increased slightly, with a system noise temperature of about 0.3 ± 0.1 K and 0.25 ± 0.1 K, respectively [the latter shown in Fig. 33(c)]. To our knowledge, these are the

lowest noise temperatures achieved in this frequency range. (The apparent narrow bandwidth was set by the resonant input matching network of this specific post amplifier, and does not reflect the bandwidth intrinsic to the SQUID.)

More recently, a microstrip SQUID amplifier was cooled to 0.4 K, and found to have a noise temperature of about 0.1 K. In addition, demonstrations have been made of adjustable microstrip tuning with the stripline loaded by a varactor, SQUID post amplification, and SQUID amplification with a resonant input load. Further studies are planned to characterize the properties of these SQUID's at dilution refrigerator temperatures.

B. Cryogenic upgrade

The upgraded detector will be a considerably more complex instrument than the present one. The upgrade of the experiment to incorporate SQUID amplifiers poses several unique design challenges. Since the system noise temperature $T_s = T(1 + 7\omega/V_\Phi)$, the physical temperature T will be reduced to around 200 mK using a dilution refrigerator. In practice, T_n saturates at about 100 mK because of hot electron effects [61].

The upgraded experimental insert will be compatible with the existing 8 T magnet, i.e., the cavity arrays and refrigerator will fit into a 53-cm diameter clear bore. The cavity array with its vacuum and thermal shields (~ 200 kg) occupies the lower third of the tower, situating it in the middle of the solenoidal magnet. The array is suspended from the cryogenic assembly, comprised of the dilution refrigerator, field compensation coils, LHe reservoir, thermal baffles, support members, etc., occupying the upper two-thirds of the tower. The entire tower is suspended from a 60-cm diameter flange, mating with the top of the magnet cryostat via an O-ring seal. The dilution refrigerator itself will fit within a 17.5-cm bore down the center, which in addition to serving as the vacuum pump-out column, accommodates a helium pre-cool line, microwave signal cables, actuator, and sensor wires. To minimize the thermal load from the outside, all mechanical motion for tuning and coupling will be performed by piezoelectric actuators; these have already been fitted to the present experiment.

The SQUID amplifier requires a field-free environment for operation ($\lesssim 1$ mT at the shields), but is located less than a meter above the magnet coil where the fringe field would normally be on the order of 1 T. Hence, there is a field compensation coil to reduce the field to a few tens of mT surrounding the hermetic niobium can containing the SQUID. The field compensation coil is actually a pair of opposing coils mechanically and electrically locked together, with relative dipole moments such that the net force on the pair is zero, but able to produce a zero field region at the location of the SQUID. Whereas a single bucking coil would have an upward force of several tons, this configuration has only internal forces between the two coils for any excitation of the main field magnet. Figure 34 exhibits key details.

IV. CONCLUSIONS

While much has happened both in particle physics and cosmology over the past two decades, the motivation for the

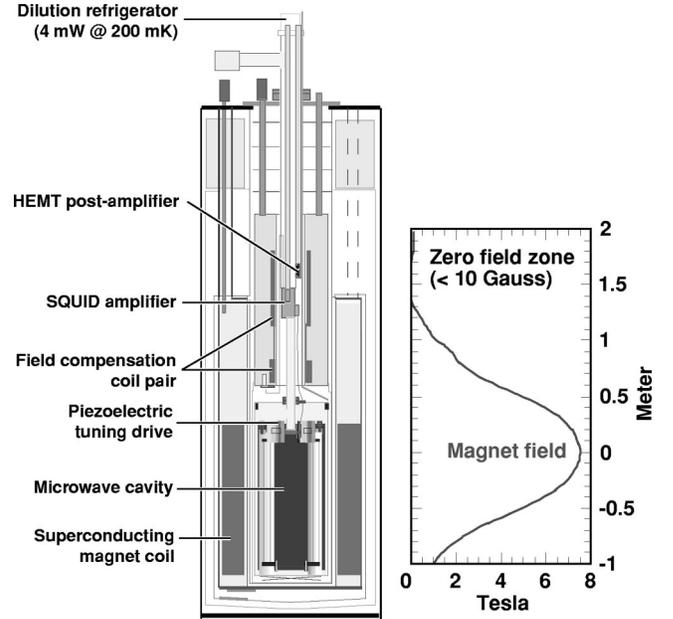


FIG. 34. Sketch of upgraded experiment, and field profile with compensation coil pair.

axion remains strong. It is still the most compelling mechanism to enforce strong- CP conservation, and the sum total of all cosmological observables still point to a large CDM component in the Universe.

Figure 31 shows the rf cavity experimental constraints on the axion axion mass and coupling through the axion-photon coupling, along with two prototypical axion models (KSVZ and DFSZ). It provokes several key points.

The first concerns axion models. While it is encouraging that two very different axion models have axion-photon couplings of the same order of magnitude, one should be cautious in drawing conclusions. Should the ratio of the electromagnetic and color anomaly E/N of the PQ symmetry be far away from that corresponding to exact suppression (~ 1.95) there could be a pleasant surprise for experimentalists; conversely a value close to it could make the axion nearly impossible to see. Kim has evaluated numerous models with very different but plausible PQ charges [62], and nevertheless finds their corresponding $g_{A\gamma\gamma}^2$ clustered in roughly two orders of magnitude, mostly more favorable than the implementation of the benchmark DFSZ model. Second, the open range of axion masses ($m_A = 10^{-6} - 10^{-3}$ eV) is “soft” on both ends, particularly on the low-mass side.

It is also interesting that the axion-photon coupling $g_{A\gamma\gamma}$ has lent itself to the greatest variety of imaginative and practical search techniques. If the axion is someday found, it will likely be through the $\vec{E} \cdot \vec{B}$ interaction. Purely laboratory experiments, such as the precision measurement of vacuum birefringence using lasers and magnets, are elegant and will soon provide beautiful demonstrations of higher-order QED. On the other hand, it strikes us as unlikely those techniques could ever be sensitive to DFSZ axions, although verifying

the astrophysical limits would be an important and significant achievement.

The most encouraging development in the past two years has been our demonstration that the standard axion models could be tested with the microwave cavity experiments. Upgrades to extend the sensitivity and mass reach are already being planned. In the meantime, we plan to operate the experiment indefinitely; the (1–10 μeV) region of mass is arguably the most interesting range to explore.

The microwave cavity experiments are also not without limitations. The local CDM density of our halo is not well determined. Furthermore, we do not see how to easily extend the mass search of the microwave cavity searches to cover the entire range, particularly upward. If $m_A \gg 10^{-4}$ eV, another technique will need to be found as we consider it unrealistic to operate an experiment with hundreds or more

cm-scale cavities. It is impressive to consider that as small as the axion signal is expected to be—requiring radio receivers orders of magnitude more sensitive than anything else on earth—power sensitivity is not likely to be a limitation of the microwave cavity experiments, in view of the very promising work on dc SQUID amplifiers.

ACKNOWLEDGMENTS

This research is supported by the U.S. Department of Energy under Contract Nos. DE-FC02-94ER40818, W-7405-ENG-48, DE-FG02-97ER41029, DE-AC02-76CH0300, DE-DE-FG02-90ER40560, DE-AC-03-76SF00098, and the National Science Foundation under Award No. PHY-9501959. We thank the LLNL Technical Information Department for timely and efficient production of the graphics.

-
- [1] I.S. Altarev *et al.*, Phys. Lett. B **276**, 242 (1992); K.F. Smith *et al.*, *ibid.* **234**, 191 (1990).
- [2] R.D. Peccei and H. Quinn, Phys. Rev. Lett. **38**, 1440 (1977); Phys. Rev. D **16**, 1791 (1977).
- [3] S. Weinberg, Phys. Rev. Lett. **40**, 223 (1978); F. Wilczek, *ibid.* **40**, 279 (1978).
- [4] Axion theory reviews include J.E. Kim, Phys. Rep. **150**, 1 (1987); H.-Y. Cheng, *ibid.* **158**, 1 (1988); R.D. Peccei, in *CP Violation*, edited by C. Jarlskog (World Scientific, Singapore, 1989), pp. 503–551; M.S. Turner, Phys. Rep. **197**, 67 (1990); G.G. Raffelt, *ibid.* **198**, 1 (1990).
- [5] L.J. Rosenberg and K.A. van Bibber, Phys. Rep. **325**, 1 (2000).
- [6] D.J. Gross and F. Wilczek, Phys. Rev. Lett. **30**, 1343 (1973).
- [7] G. 't Hooft, Phys. Rev. Lett. **37**, 8 (1976); Phys. Rev. D **14**, 3432 (1976); R. Jackiw and C. Rebbi, Phys. Rev. Lett. **37**, 172 (1976); C.G. Callan, R.F. Dashen, and D.J. Gross, Phys. Lett. **63B**, 334 (1976).
- [8] A.A. Belavin, A.M. Polyakov, A.S. Shvarts, and Yu.S. Tyupkin, Phys. Lett. **59B**, 85 (1975).
- [9] S. Adler, Phys. Rev. **177**, 2426 (1969); J.S. Bell and R. Jackiw, Nuovo Cimento A **60**, 47 (1969).
- [10] S. Weinberg, Physica D **11**, 3583 (1975).
- [11] M. Kobayashi and K. Maskawa, Prog. Theor. Phys. **49**, 652 (1973).
- [12] C. Vafa and E. Witten, Phys. Rev. Lett. **53**, 535 (1984).
- [13] S. Weinberg [3]; W.A. Bardeen and S.-H.H. Tye, Phys. Lett. **74B**, 229 (1978); J. Ellis and M.K. Gaillard, Nucl. Phys. **B150**, 141 (1979); T.W. Donnelly *et al.*, Phys. Rev. D **18**, 1607 (1978); M. Srednicki, Nucl. Phys. **B260**, 689 (1985); P. Sikivie, in *Cosmology and Particle Physics*, edited by E. Alvarez *et al.* (World Scientific, Singapore, 1987), pp. 143–169.
- [14] M. Dine, W. Fischler, and M. Srednicki, Phys. Lett. **104B**, 199 (1981); A.P. Zhitnitskii, Sov. J. Nucl. Phys. **31**, 260 (1980).
- [15] J. Kim, Phys. Rev. Lett. **43**, 103 (1979); M.A. Shifman, A.I. Vainshtein, and V.I. Zakharov, Nucl. Phys. **B166**, 493 (1980).
- [16] M.S. Turner, Phys. Rep. **197**, 67 (1990); G.G. Raffelt, *ibid.* **198**, 1 (1990).
- [17] D.J. Gross, R.D. Pisarski, and L.G. Yaffe, Rev. Mod. Phys. **53**, 43 (1981).
- [18] J. Preskill, M. Wise, and F. Wilczek, Phys. Lett. **120B**, 127 (1983); L. Abbott and P. Sikivie, *ibid.* **120B**, 133 (1983); M. Dine and W. Fischler, *ibid.* **120B**, 137 (1983); J. Ipser and P. Sikivie, Phys. Rev. Lett. **50**, 925 (1983).
- [19] R. Davis, Phys. Rev. D **32**, 3172 (1985).
- [20] D. Harari and P. Sikivie, Phys. Lett. B **195**, 361 (1987).
- [21] A. Vilenkin and A.E. Everett, Phys. Rev. Lett. **48**, 1867 (1982).
- [22] P. Sikivie, in *Where are the Elementary Particles?*, Proceedings of the 14th Summer School on Particle Physics, Gif-sur-Yvette, 1982, edited by P. Fayet *et al.* (Inst. Nat. Phys. Nucl. Phys. Particules, Paris, 1983).
- [23] P. Sikivie, Phys. Rev. Lett. **48**, 1156 (1982).
- [24] C. Hagmann and P. Sikivie, Nucl. Phys. **B363**, 247 (1991).
- [25] R.A. Battye and E.P.S. Shellard, Nucl. Phys. **B423**, 260 (1994); Phys. Rev. Lett. **73**, 2954 (1994); **76**, 2203 (1996).
- [26] M. Yamaguchi, M. Kawasaki, and J. Yokoyama, Phys. Rev. Lett. **82**, 4578 (1999).
- [27] C. Hagmann, S. Chang, and P. Sikivie, Nucl. Phys. B (Proc. Suppl.) **72**, 81 (1999); Phys. Rev. D **63**, 125018 (2001).
- [28] D. Lyth, Phys. Lett. B **275**, 279 (1992).
- [29] M. Nagasawa and M. Kawasaki, Phys. Rev. D **50**, 4821 (1994).
- [30] S. Chang, C. Hagmann, and P. Sikivie, Phys. Rev. D **59**, 023505 (1999).
- [31] P. Sikivie, Phys. Rev. Lett. **51**, 1415 (1983); Phys. Rev. D **32**, 2988 (1985).
- [32] L. Krauss, J. Moody, F. Wilczek, and D. Morris, Phys. Rev. Lett. **55**, 1797 (1985).
- [33] W. Wuensch *et al.*, Phys. Rev. D **40**, 3153 (1989); S. De Panfilis *et al.*, Phys. Rev. Lett. **59**, 839 (1987); C. Hagmann *et al.*, Phys. Rev. D **42**, 1297 (1990).
- [34] M. Tada *et al.*, Nucl. Phys. B (Proc. Suppl.) **72**, 164 (1999).
- [35] D. Kleppner and T.W. Ducas, Bull. Am. Phys. Soc. **21**, 600 (1976); Thomas F. Gallagher, *Rydberg Atoms* (Cambridge University Press, Cambridge, England, 1994), pp. 66–68.
- [36] D. Lynden-Bell, Mon. Not. R. Astron. Soc. **136**, 101 (1967).
- [37] J.R. Ipser and P. Sikivie, Phys. Lett. B **291**, 288 (1992).
- [38] P. Sikivie, I.I. Tkachev, and Y. Wang, Phys. Rev. Lett. **75**, 2911

- (1995); Phys. Rev. D **56**, 1863 (1997).
- [39] Additional details on the apparatus may be found in H. Peng *et al.*, Nucl. Instrum. Methods Phys. Res. A **444**, 569 (2000).
- [40] Precision Tube, 620 Naylor Rd, Salisbury, MD 21801.
- [41] National Radio Astronomy Observatory, Charlottesville, VA 22903.
- [42] MITEQ, Inc., Hauppauge, NY 11788.
- [43] Noise/Com, Inc., Paramus, NJ 07652.
- [44] Stanford Research Systems, Inc., Sunnyvale, CA 94089.
- [45] V. Vemuri and W.J. Karplus, *Digital Computer Treatment of Partial Differential Equations* (Prentice-Hall, Englewood Cliffs, NJ, 1981).
- [46] See, for example, S. Ramo, J.R. Whinnery, and T. van Duzer, *Fields and Waves in Communication Electronics*, 2nd ed. (Wiley, New York, 1984), Chap. 11.
- [47] E. Daw and R.F. Bradley, J. Appl. Phys. **82**, 1925 (1997).
- [48] Model Q70-35K-1000/1000-720B, TTE, Inc., 1162 West Olympic Blvd., Los Angeles, CA 90064.
- [49] Model "Gamma," Bittware Research Systems, 26 South Main Street, Concord, NH 03301.
- [50] Wang NMR Inc., Livermore, CA.
- [51] C. Hagmann *et al.*, Phys. Rev. Lett. **80**, 2043 (1998).
- [52] For a simpler equivalent circuit model, see B.E. Moskowitz and J. Rogers, Nucl. Instrum. Methods Phys. Res. A **264**, 445 (1988).
- [53] R.H. Dicke, Rev. Sci. Instrum. **17**, 268 (1946).
- [54] E.J. Daw, Ph.D. thesis, Massachusetts Institute of Technology, 1998.
- [55] C.D. Tesche and J. Clarke, J. Low Temp. Phys. **29**, 301 (1977).
- [56] M.B. Ketchen and J.M. Jaycox, Appl. Phys. Lett. **40**, 736 (1982).
- [57] C. Hilbert and J. Clarke, J. Low Temp. Phys. **61**, 237 (1985).
- [58] For references and a discussion, see for example, T. Ryhänen, H. Seppä, R. Ilmoniemi, and J. Knuutila, J. Low Temp. Phys. **76**, 287 (1989).
- [59] M.A. Tarasov, G.V. Prokopenko, V.P. Koshelets, I.L. Lapitskaya, and L.V. Filippenko, IEEE Trans. Appl. Supercond. **AS-5**, 3226 (1995).
- [60] M. Mück, M.-O. André, J. Clarke, J. Gail, and C. Heiden, Appl. Phys. Lett. **72**, 2885 (1998); **75**, 698 (1999); **75**, 3545 (1999).
- [61] F.C. Wellstood, C. Urbina, and J. Clarke, Appl. Phys. Lett. **54**, 2599 (1989).
- [62] J.E. Kim, Phys. Rev. D **58**, 055006 (1998).