This is the accepted manuscript made available via CHORUS. The article has been published as:

Design and implementation of the ABRACADABRA-10 cm axion dark matter search

Jonathan L. Ouellet, Chiara P. Salemi, Joshua W. Foster, Reyco Henning, Zachary Bogorad, Janet M. Conrad, Joseph A. Formaggio, Yonatan Kahn, Joe Minervini, Alexey Radovinsky, Nicholas L. Rodd, Benjamin R. Safdi, Jesse Thaler, Daniel Winklehner, and Lindley Winslow

Phys. Rev. D **99**, 052012 — Published 29 March 2019

DOI: 10.1103/PhysRevD.99.052012

Design and Implementation of the ABRACADABRA-10 cm Axion Dark Matter Search

Jonathan L. Ouellet,[1,*] Chiara P. Salemi,[1] Joshua W. Foster,[2] Reyco Henning,[3,4] Zachary Bogorad,[1] Janet M. Conrad,[1] Joseph A. Formaggio,[1] Yonatan Kahn,[5,6] Joe Minervini,[7] Alexey Radovinsky,[7] Nicholas L. Rodd,[8,9] Benjamin R. Safdi,[2] Jesse Thaler,[10] Daniel Winklehner,[1] and Lindley Winslow[1,†]

[1] *Laboratory for Nuclear Science, Massachusetts Institute of Technology, Cambridge, MA 02139, U.S.A.*
[2] *Leinweber Center for Theoretical Physics, Department of Physics, University of Michigan, Ann Arbor, MI 48109, U.S.A.*
[3] *University of North Carolina, Chapel Hill, NC 27599, U.S.A.*
[4] *Triangle Universities Nuclear Laboratory, Durham, NC 27708, U.S.A.*
[5] *Princeton University, Princeton, NJ 08544, U.S.A.*
[6] *Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL 60637, U.S.A.*
[7] *Plasma Science and Fusion Center, Massachusetts Institute of Technology, Cambridge, MA 02139, U.S.A.*
[8] *Berkeley Center for Theoretical Physics, University of California, Berkeley, CA 94720, U.S.A.*
[9] *Theoretical Physics Group, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, U.S.A.*
[10] *Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, U.S.A.*

(Dated: March 12, 2019)

The past few years have seen a renewed interest in the search for light particle dark matter. ABRACADABRA is a new experimental program to search for axion dark matter over a broad range of masses, $10^{-12} \lesssim m_a \lesssim 10^{-6}\,\mathrm{eV}$. ABRACADABRA-10 cm is a small-scale prototype for a future detector that could be sensitive to QCD axion couplings. In this paper, we present the details of the design, construction, and data analysis for the first axion dark matter search with the ABRACADABRA-10 cm detector. We include a detailed discussion of the statistical techniques used to extract the limit from the first result with an emphasis on creating a robust statistical footing for interpreting those limits.

I. INTRODUCTION

In recent years, the absence of a compelling direct detection of dark matter (DM) in accelerator and Weakly Interacting Massive Particle (WIMP) searches has reignited the search for Axion-like DM (ADM). The coincidence of new developments in quantum sensors and quantum information technology has stoked this reawakened interest, and the past few years have seen a wealth of new experimental ideas and approaches that are beginning to revolutionize the field [1, 2]. While most ADM searches have traditionally focused on a narrow mass range from $10 \lesssim m_a \lesssim 100\,\mu\mathrm{eV}$, recent theoretical work has made a compelling case for ADM in the mass range $m_a \lesssim 1\,\mu\mathrm{eV}$ [3–15].

The ABRACADABRA-10 cm experiment has recently released results of the first direct detection search for ADM below $1\,\mu\mathrm{eV}$ [16]. The design of the experiment was motivated by the proposal of [17–19], and is based on measuring the coupling of ADM to electromagnetism – similar to experiments probing different mass regimes like the long-running ADMX [20, 21] and HAYSTAC [22]. In this lower mass range, the axion field a, in the presence of a large magnetic field can be thought of as an induced effective current

$$\mathbf{J}_{\mathrm{eff}} = g_{a\gamma\gamma}\frac{\partial a}{\partial t}\mathbf{B}\,, \tag{1}$$

where $g_{a\gamma\gamma}$ is the axion-photon coupling. This current sources a small AC magnetic field that can be measured with a sensitive enough magnetometer.

ABRACADABRA-10 cm is a prototype detector for a new ADM search program with the ultimate goal of probing down to the QCD axion couplings. The design and implementation of ABRACADABRA-10 cm contains novel elements that have not previously been used in ultralight dark matter searches:

- A toroidal magnet geometry, with the detection element placed in the near-zero-field region;
- A broadband readout mode involving continuous-stream data-taking for $\sim 10^6$ seconds (roughly 1 month), and several compression techniques to mitigate the total data storage requirements while preserving the desired signal bandwidth;
- A calibration technique where a signal is injected through current in a calibration loop in a similar geometry as the expected axion signal;
- A data analysis pipeline tailored to the expected statistics of the axion field in the quasistatic regime, where the signal is best described as a flux power spectral density rather than photon-counting with the added constraint that "rescanning" is prohibitively time-consuming.

In this paper, we provide context and additional details for each of these novel elements and their specific implementation in ABRACADABRA-10 cm. In Section II,

* ouelletj@mit.edu
† lwinslow@mit.edu

we describe the design and construction of the toroidal ABRACADABRA-10 cm detector. We describe the data collection approach used for the broadband readout in Sec III, and describe the calibration of the detector in Sec. IV. In Sec. V, we describe the data analysis and limit extraction approach used for our broadband search.

II. DETECTOR DESIGN AND CONSTRUCTION

The ABRACADABRA-10 cm detector and setup can be split into six separate systems: the toroidal magnet, the magnet support infrastructure and shielding, the pickup loop circuit, the SQUID electronics, the calibration circuit, and finally the cryostat and detector support infrastructure. In this section, we discuss the design and construction of each.

The expected signal-to-noise ratio (SNR) of ABRACADABRA-10 cm can be written approximately as

$$\mathrm{SNR} = g_{a\gamma\gamma}\sqrt{\rho_{\mathrm{DM}}}\mathcal{G}VB_{\mathrm{max}}\left(\frac{M_{\mathrm{in}}}{L_T}\right)\frac{(\tau t)^{\frac{1}{4}}}{S_{\Phi\Phi}^{1/2}}\,, \qquad (2)$$

where V is the volume of the toroid, $\mathcal{G}$ is a geometric factor, B_{max} is the max field inside the toroid, M_{in} is the inductive coupling of the SQUID, and L_T is the total inductance of the readout circuit. Here, we assume that the integration time t, exceeds the axion coherence time, τ. The final parameter of importance is the flux noise level, $S_{\Phi\Phi}^{1/2}$, typically measured in $\mu\Phi_0/\sqrt{\mathrm{Hz}}$. The relevant parameters are summarized in Table I.

A. The Toroidal Magnet

The magnet structure is built around 80 identical Delrin wedges, (see Fig. 1a). When glued together, they create a toroidal frame with an inner radius of 3 cm and an outer radius of 6 cm, with a total height of 12 cm. The total volume of magnetic field is $V \approx 890\,\mathrm{cm}^3$.

The magnet current is carried by a NbTi(CuNi) wire which is wound 1,280 times around the magnet. Between each pair of wedges is a groove that has 16 winds of wire, laid down in pairs 8 layers deep. The wire is held in place with epoxy. To cancel the azimuthal current, the toroid is counterwound.

The toroid was wound by Superconducting Systems, Inc [23] in three separate pieces, with three separate lengths of NbTi(CuNi) wire. The pieces are then glued together and the wires are connected together with two superconducting crimps. These crimps are then attached to the outside of the toroid. These joints could create small stray fields which contribute to the backgrounds for the axion search, but unfortunately could not be avoided in the construction.

The toroid is mounted in a dilution refrigerator (described below) and cooled to $\lesssim 1\,\mathrm{K}$. The NbTi(CuNi) wire superconducts below $\lesssim 9\,\mathrm{K}$. We charge the magnet by injecting a 121 A current into the toroid. Once charged, we turn off heat to a superconducting switch (located away from the magnet) which then locks the current into the magnet. The current source is disconnected from the charging leads on top of the refrigerator.

When fully charged, the maximum field in the magnet is $B_{\mathrm{max}} = 1\,\mathrm{T}$. This was confirmed with a Hall probe to a precision of $\sim 1\%$, with the uncertainty coming from uncertainty in position of the probe in the field. Once in persistent mode, we observed no decay in the field to a precision of $\lesssim 0.1\%$ on the scale of 1 week. The Hall sensor was removed before normal data taking.

B. Support Infrastructure

The toroid is mechanically supported by a G10 frame held together with nylon bolts (see Fig. 1c). The goal of this structure was to rigidly mount the toroid in place, while minimizing the amount of non-superconducting metal near the magnet. The one exception to this is the copper straps which wrap around the outside of the toroid that provide the required thermalization to cool the magnet. These straps undoubtedly contribute some level of noise for our axion search, though in the current setup it is not the dominant noise source. In the future we will search for alternative thermalization approaches.

The entire toroid and support structure are mounted inside the external shield (see Fig. 1d). The shield consists of a spun copper can that has been coated inside and out with a 25-75 μm layer of tin, for a total thickness of $\approx$1 mm. The copper provides good thermal conductivity to minimize thermal gradients across the shield. It also provides the thermal conductivity to the copper straps which cool the magnet. Once below 3.7 K, the tin becomes superconducting and expels environmental magnetic fields and acts as a shield against electromagnetic interference. Optimizing and characterizing this external shielding will be the subject of future work.

The external shield is built in two hemispheres (top and bottom) which have $\approx$ 12 mm of vertical overlap when assembled. There is a small gap in one location between the inner and outer shield through which the magnet wires, pickup loop wires and calibration loop wires pass as three sets of twisted pairs. The shield halves are connected with a layer of solder and epoxy to ensure both electrical and mechanical connection.

A 12 mm thick aluminum top plate is epoxied to the top of the top shield and acts as the contact point for the thermalization to the rest of the cryostat and mechanical mounting point to the vibration isolation system (see below). The aluminum plate is electrically isolated from the shield to minimize grounding loops. A thermometer is epoxied to the outside of the bottom shield which monitors the temperature of the farthest point from the

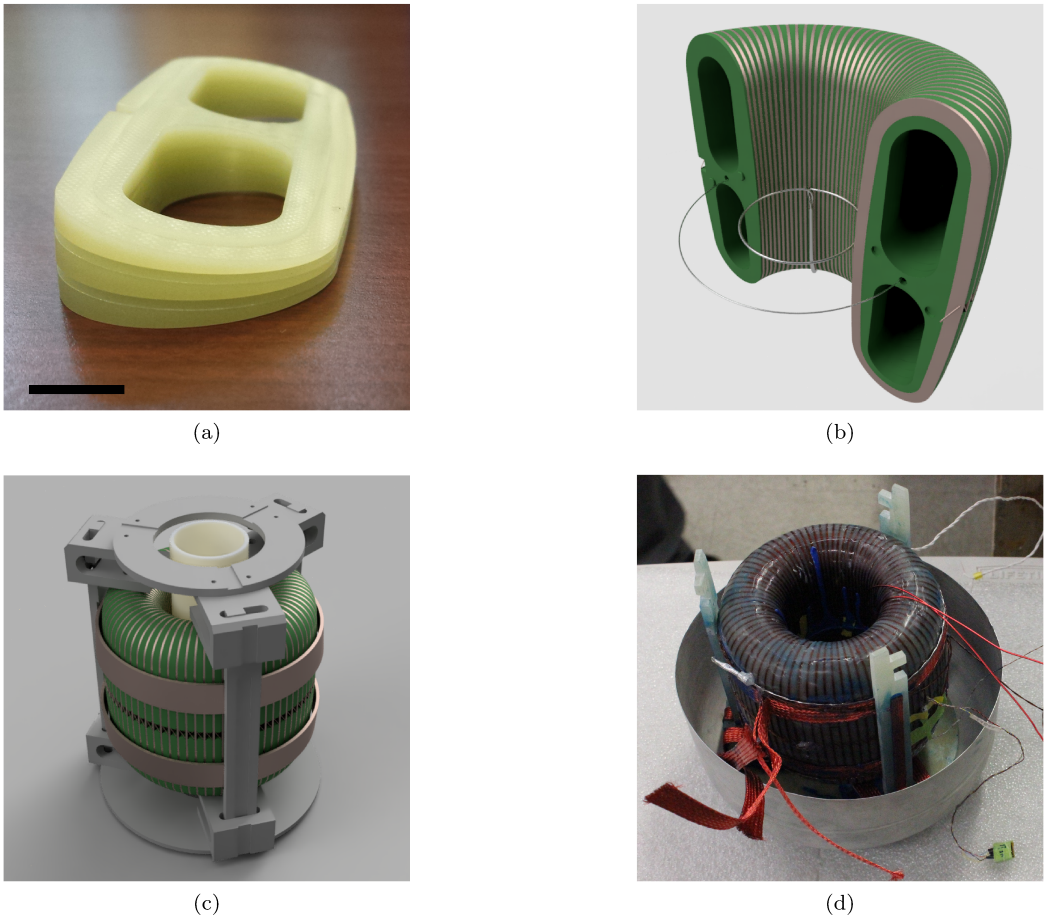

(a) (b) (c) (d)

FIG. 1. (a) Three of the 80 Delrin wedges that form the toroid structure stacked together. The black bar indicates a $\approx 1\,\mathrm{cm}$ scale. (b) Cutaway rendering of the toroid with the 1 mm diameter wire pickup loop in the center. A 0.5 mm diameter wire runs through the center of field region to form the calibration loop. Toroid height is $\approx 12\,\mathrm{cm}$. (c) Rendering of the ABRACADABRA-10 cm support structure. The pickup loop is supported by a PTFE (white) tube through the center. The magnet is supported by an outer G10 support structure and thermalized with two copper bands. (d) Photo of the assembled ABRACADABRA-10 cm, with the top of the superconducting shield and support structure removed.

TABLE I. Summary of the ABRACADABRA-10 cm detector design parameters.

Pickup Loop Radius	R_p	20.1 mm
Pickup Loop Wire Diameter	r_p	1.0 mm
Magnet Inner Radius	R_{in}	30 mm
Magnet Outer Radius	R_{out}	60 mm
Magnet Height	h	120 mm
Magnet Max Field	B_{max}	1.0 T
Geometric Factor	$\mathcal{G}_V$	0.027
Pickup Loop Inductance	L_p	95.5 nH
SQUID Input Inductance	L_{in}	150 nH
SQUID Inductive Coupling	M_{in}	2.5 nH

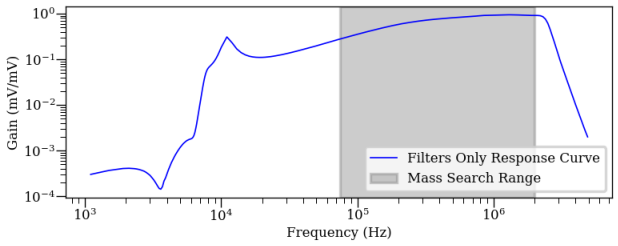

FIG. 2. Gain of the combined high-pass and anti-aliasing filters. All spectra are corrected for this response function – unless otherwise noted. Measured in-situ, using injected signals at different frequencies. This also defines the usable range of data. For our search we use the range 75 kHz – 2 MHz.

thermalization. However, during data taking this thermometer is not active.

C. Pickup Loop Circuit

The pickup loop measures the magnetic field in the center of the toroid – a region that should have zero field in the absence of an axion signal. The time-averaged magnitude of the flux through the pickup loop due to $\mathbf{J}_{\rm eff}$ can be written as

$$|\Phi_a|^2 = g_{a\gamma\gamma}^2 \rho_{\rm DM} V^2 \mathcal{G}_V^2 B_{\rm max}^2 \equiv A\,, \tag{3}$$

where $\rho_{\rm DM}$ is the DM density, V is the volume of the magnet, $B_{\rm max}$ is the maximum field in the magnet and $\mathcal{G}$ is a geometric factor. The pickup loop itself consists of a 1 mm diameter solid NbTi wire wrapped around the outside of a 5.5 cm diameter polytetrafluoroethylene (PTFE) tube that is 18.1 cm tall. The geometric factor $\mathcal{G}_V$ weights the effective current in Eqn. 1, by the contribution to the flux through the pickup loop. This can be written as

$$\mathcal{G}_V \equiv \frac{1}{B_{\rm max} V}\left|\int_{\rm Loop} dA \int_{\rm Toroid} dV' \,\frac{\mathbf{B}(\mathbf{r}')\times(\mathbf{r}'-\mathbf{r})}{|\mathbf{r}'-\mathbf{r}|^3}\cdot\hat{\mathbf{n}}\right| \tag{4}$$

where $\hat{\mathbf{n}}$ is the normal to the plane of the pickup loop, and the integrals are taken over the area enclosed by the pickup loop and the volume of the toroid. The integrand is reminiscent of the Biot-Savart law, with the current taken to be the axion-induced effective current $\mathbf{J}_{\rm eff}$ which follows lines of $\mathbf{B}$ [19].[1] For the ABRACADABRA-10 cm geometry, we calculate this using a COMSOL [24] simulation to be $\mathcal{G}_V = 0.027$. The two wire leads from the pickup loop, which consists of the same wire as the loop, are twisted into a twisted pair configuration and run out under the bottom of the toroid through the gap in the shield. Once outside of the shield, the wires run $\approx$ 15 cm inside a stainless steel mesh sleeve. At this point, the 1 mm wires are joined to 75 μm twisted-pair PFA-insulated wire with superconducting crimped solder. The 75 μm wires run for $\approx$1 m inside hollow superconducting solder capillaries [25] to the input of the SQUIDs mounted on the 700 mK (Still) stage of the cryostat. The SQUIDs have an input inductance of 150 nH to match the calculated inductance of the pickup loop of $L_p = 95.5$ nH. Including the inductance of the wires, the total design inductance of the pickup loop circuit is $\approx$550 nH. However, measuring the inductance of the circuit yielded a value closer to $L_T \approx 3.3\,\mu$H, we discuss this further in Sec. IV. The data presented in [16] was taken in a broadband readout configuration with no resonant amplifier.

D. SQUID Setup And Readout

The first stage was readout with a Magnicon two stage SQUID current sensor [26, 27]. The SQUID is operated at a temperature of 870 mK and has typical flux noise floor of $0.6\,\mu\Phi_0^2$/Hz. The inductive coupling between the input coil of the SQUID and the SQUID is $M_{\rm in} = 2.52$ nH. The SQUID is operated with the Magnicon XXF-1 electronics in flux-lock feedback loop (FLL) mode with a SQUID flux to voltage conversion of $\partial V/\partial\Phi_S = 1.29\,{\rm V}/\Phi_0$. In FLL mode, the response of the SQUID is linear over the dynamic range of the amplifier ± 11 V or $\approx \pm 8.5\,\Phi_0$, however, this comes at the cost of limiting the bandwidth of the system to $\approx$ 6 MHz.

The output voltage of the SQUID was recorded with an AlazarTech 9870 8-bit digitizer [28]. To achieve the needed voltage precision we use the smallest available input range of $\pm$40 mV which leads to a typical digitizer noise floor of $3.5\times10^{-9}\,{\rm mV}^2$/Hz. However, due to the large background below $\sim$ 20 kHz (see Fig. 4), we must first pass the signal through a 10 kHz high-pass filter to prevent railing the digitizer. Additionally, we use a 2.5 MHz anti-aliasing filter. The frequency response of these two filters is shown in Fig. 2. These filters define the usable range of data for our axion search of 50 kHz – 3 MHz. It is worth pointing out that even though the gain is less than unity over the majority of our search range, both the signal and dominant noise is scaled by this gain, so the SNR is unchanged.

The digitizer is clocked to a Stanford Research System FS725 Rb frequency standard, with a ten-second Allan variance of $< 10^{-11}$.

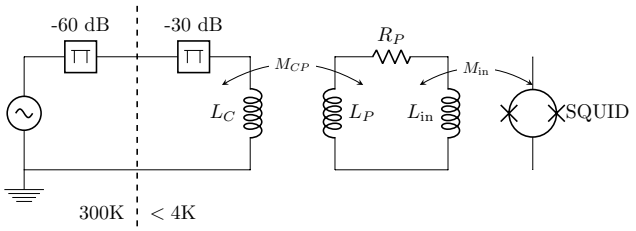

FIG. 3. A conceptual diagram of the ABRACADABRA-10 cm calibration circuit. The calibration loop, $L_C \approx 300$ nH, is concentric with the pickup loop, $L_P = 95.5$ nH. The circuit is plugged into the SQUID with input inductance $L_{\rm in} \approx$ 150 nH. The parasitic resistance in the circuit is measured as $R_P \approx 13\,\mu\Omega$.

E. Calibration Circuit

We measure the end-to-end gain using a calibration system. It consists of the 0.5 mm diameter NbTi wire passing through the body of the toroid (i.e. in the magnetic field region), creating a 9 cm diameter loop concentric with and in the same plane as the pickup loop (see Fig. 1b). This wire runs out of the detector shield as a twisted pair and then into an RG196 coaxial cable. This cable is connected to a 30 dB attenuator at the 4 K stage and then continues up to the top of the cryostat and through a BNC feed-through out of the vacuum region.

[1] In the notation of [19], $\mathcal{G}_V = V_B/V$.

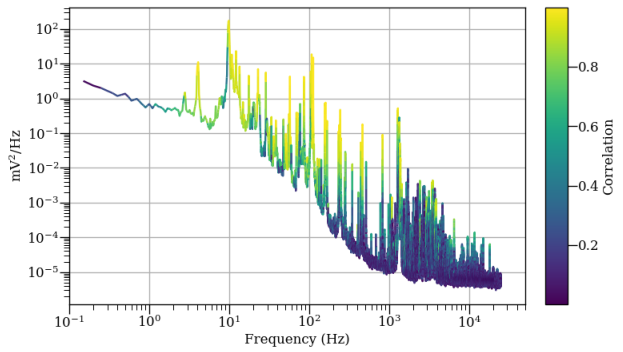

FIG. 4. Low frequency SQUID spectra from ABRACADABRA-10 cm taken with an accelerometer attached the 300 K plate. The spectrum is that of the SQUID output, with the degree of correlation with the accelerometer indicated by color (i.e. the correlation coefficient). The accelerometer begins to lose sensitivity above a few kHz, so it is not clear from this measurement how far up the correlation continues. These data were taken with a larger dynamic range on the digitizer, so have a relatively high ADC noise floor of $\sim 5 \times 10^{-6}\,\mathrm{mV}^2/\mathrm{Hz}$. (Data taken without signal shaping filters.)

During data taking, this BNC is left unplugged, and the attenuator contributes noise from a 50 Ω resistor at 4 K, which is well below our current noise level. We calculate the mutual inductance between the calibration and pickup loops both with an analytic calculation based on the geometry, as well as with a COMSOL [24] simulation. These values agree and predict a mutual inductance of $M_{CP} = 19.3\,\mathrm{nH}$.

While calibrating, we typically add an additional 60 dB of warm attenuation for a total of 90 dB of attenuation to get the signal to reasonable size. All attenuators are impedance matched at 50 Ω. However, the output of the final attenuator is shorted by the calibration loop, which has an inductance of ≈300 nH; for frequencies below ≈ 30 MHz this causes it to behave as a current source driving a current through the calibration loop with amplitude independent of frequency. A wiring schematic of the calibration circuit is shown in Fig. 3.

F. Cryostat and Detector Suspension

The ABRACADABRA-10 cm detector is mounted inside an Oxford Instruments Triton 400 dilution refrigerator. It is mechanically supported by the detector suspension system. This consists of a 0.038" Kevlar thread which attaches to a vented bolt screwed into the center hole of the top aluminum mounting plate on ABRACADABRA-10 cm. The thread runs ∼ 1.5 m up through the various cold stages of the cryostat to a steel spring which supports the weight of the detector. The steel spring has a spring constant of $k \approx 20.4\,\mathrm{N/m}$ and connects to a hook which is mounted about 1 m above the 300 K plate of the cryostat. The hook is at the top of a 1 m long vacuum tube which is rigidly connected to the 300 K plate. Due to the poor thermal conductivity of Kevlar, the spring and top end of the thread stay at 300 K while the bottom of the thread is cooled to $\lesssim$1 K with the detector.

The detector is thermalized to the coldest stage of the cryostat through four 10 mm wide 75 μm thick copper ribbons. Specifically, they are attached to the Mixing Chamber plate of the cryostat and then to the aluminum top plate of the detector. They are mounted with significant slack to minimize vibration through these ribbons. They are electrically isolated from the top of the detector by using thin Kapton pads between the copper ribbon and the aluminum plate and connected with Nylon bolts.

The detector suspension system is designed to act like a pendulum which rolls off lateral vibration above frequencies of $f \approx 0.4\,\mathrm{Hz}$ and in the vertical direction above frequencies of $f \approx 0.3\,\mathrm{Hz}$. The operating frequency of the pulse tube is 1.4 Hz and creates one of the main vibrational noise sources in the ABRACADABRA-10 cm data. It is clear from Fig. 4, that even with this suspension system, vibrational noise still plays a significant role in the ABRACADABRA-10 cm backgrounds, and is a future path for potential improvement.

To improve the magnetic shielding of the detector, we wrapped the cryostat in Mu-Metal shielding. As Mu-Metal is not typically suited to cryogenic temperatures, we only wrapped the outermost vacuum vessel. The vertical walls of the vessel were lined both inside and out with a 200 μm thick layer. The bottom of the vessel was covered with a single layer on the inside. The top of the vacuum vessel and cryostat were not covered with MuMetal due to all the instrumentation and cryostat infrastructure. We measured the DC magnetic field attenuation ex-situ to be a factor of $\sim 5 - 10$.

III. DATA COLLECTION PROCEDURE

A. Axion Search Data

In [19], the original proposal for a broadband search involved collecting time series data at a high sampling frequency continuously for months to years. However, this runs into practical disk space limitations. For example, one month sampled at 10 MS/s would fill ≈26 TB of disk space. This is maneagable, but would not scale well to a 1 GS/s sampling rate for a full year. However, this sort of sampling is not necessary for resolving ADM signals, where the expected signal width is given by $\Delta f/f \sim 10^{-6}$. Instead, we take an approach that maintains the required spectral resolution, while minimizing the required disk space.

For ABRACADABRA-10 cm, we sample continuously at 10 MS/s. Once samples are pulled from the digitizer, the data follows two processing paths: transforming and downsampling. First, the samples are accumulated into

a 10 s buffer (of 10^8 samples), which is then transformed via discrete Fourier transform (DFT) [29] into a power spectral density (PSD). Once the next 10 s is available, it is transformed and its PSD is then averaged with the first, and so on. This builds up an averaged PSD, called $\bar{\mathcal{F}}_{10\mathrm{M}}$, which has Nyquist frequency of 5 MHz and frequency resolution of 100 mHz. This spectrum would be able to resolve axion signals down to $m_a \sim 100$ kHz with at least one bin width. After 80 averages (or 800 s), the average spectrum is written to disk and the averaging is reset. This level of averaging was chosen as a balance between storage space and being able to resolve time variation of background noise.

In parallel with this, the 10 MS/s time series is decimated by a factor of 10, to a 1 MS/s time series. This data is then accumulated into a 100 s buffer – again of 10^8 samples – then transformed with a DFT and converted into a PSD. In this way, we build up a second averaged PSD called, $\bar{\mathcal{F}}_{1\mathrm{M}}$, with a Nyquist frequency of 500 kHz and a frequency resolution of 10 mHz. This spectrum would be able to resolve signals down to $m_a \sim 10$ kHz with at least one bin width. After 16 averages (or 1600 s), the average spectrum is written to disk and the averaging is reset.

The data are then decimated by another factor of 10 and written directly to disk at a sampling rate of 100 kS/s. Offline, we take the time series data and transform it as one 2.45×10^{11} sample long DFT to form a final spectrum, $\mathcal{F}_{100\mathrm{k}}$. Unlike the other spectra, $\mathcal{F}_{100\mathrm{k}}$ is not averaged over multiple integration periods, but is instead a single PSD with Nyquist frequency 50 kHz and frequency resolution of ≈ 408 nHz. This spectrum could be used for searches for axion signals down to below 1 Hz, however, it is not used in the present analysis.

Each decimation step is done by first applying a top-hat filter with a 10-bin width, and then down-sampling by keeping every 10^{th} filtered sample. This approach was chosen because it is fast computationally, though it is not quite optimal. We collected data from July 16, 2018 through August 14, 2018, accumulating a total exposure of $T = 2.45 \times 10^6$ s or 24.5×10^{12} samples. In total, the data consist of 3065 independent $\bar{\mathcal{F}}_{10\mathrm{M}}$ spectra and 1532 $\bar{\mathcal{F}}_{1\mathrm{M}}$ spectra as well as a $2.45{\times}10^{11}$ continuous samples of 100 kS/s data. The total data footprint was about 3.8 TB for an average write rate of 12.4 Mbps – both of which are easily handled by a desktop PC.

B. Magnet Off and Digitizer Noise Data

We also perform a Magnet Off measurement to understand backgrounds that are not correlated with the magnet. This data was collected with the exact same procedure and hardware configuration as the Magnet On data. Neither the cryostat, nor the SQUIDs were stopped in between measurements. We started collecting Magnet Off data within a few days of stopping the Magnet On run. We collected Magnet Off data from August 18, 2018 through August 27, 2018, for a total of $8.00{\times}10^5$ s of data.

We also collected $\approx$ 16 h of digitizer noise data, with nothing plugged in to measure the noise level inherent to the ADC and computer.

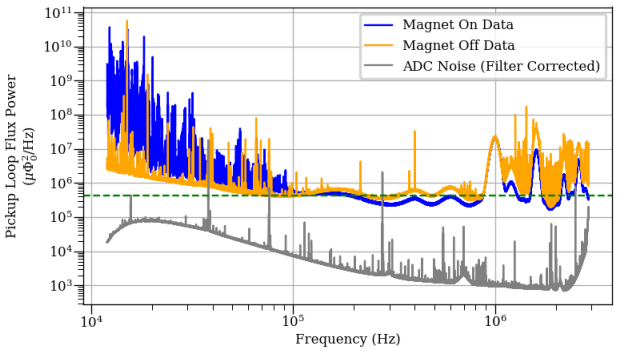

FIG. 5. Example $\bar{\mathcal{F}}_{10\mathrm{M}}$ SQUID spectra with magnet on (*blue*) and off (*orange*), along with the digitizer noise floor (*gray*). SQUID spectra are averaged over ≈9 h, digitizer data averaged over ≈16 h. The typical SQUID noise floor is shown in green dashed line. Note: The spectra were collected at different times and some of the transient noise peaks are not seen in all spectra.

C. Data Quality

Figure 5 shows examples of spectra for magnet on and magnet off data. There are a few features of these spectra worth discussing. The region below 100 kHz shows large noise spikes and a baseline increasing towards lower frequency. These spikes are generally too broad to be identified as ADM, but instead are incoherent noise backgrounds. We also observe that this noise is significantly reduced when the magnet is turned off. We interpret this noise as due to vibration of the detector. In particular, this appears to be the high frequency tail of the noise in Fig. 4. The fact that it is reduced when the magnet is turned off implies that stray fields from the toroidal magnet are being seen by the pickup loop. We see that for $f > 100$ kHz this noise becomes sub-dominant; however, it is likely that it continues to higher frequency. This will present a major challenge for future detector configurations, including those with resonator readouts, that hope to lower the noise floor by many orders of magnitude.

In the region from 100 kHz $\lesssim f \lesssim$ 850 kHz, the noise is mostly flat with a few small broad bumps and is approximately consistent with the expected SQUID flux noise floor. We see a slow variation in this noise level over the month of data taking, associated with variations in the noise level of the SQUIDs.

The region above $\sim$ 850 kHz shows two effects: broad bumps with $\sim$ 100kHz widths and a forest of very narrow transient peaks. The broad peaks are due to an unknown and incoherent noise source that decreases our sensitivity in that frequency range. The origin of this noise will

be the subject of future investigation, but for now we tolerate the decreased sensitivity. The forest of narrow transient peaks, on the other hand, present a larger problem. These peaks are $\lesssim$100 mHz wide and actually narrower than we expect for an ADM signal in this range. They are transient in time and appear to be correlated with working in the lab and with working hours. They were present for a portion of the time that we collected Magnet On data and all the time that we collected Magnet Off data. The transient nature and narrow width of these peaks imply that their source is likely from digital electronics turning on and off somewhere in, or near, the lab.

This transient noise was observed to be either present as a forest of many lines or completely absent. For our ADM search, the easiest approach was to use this fact to tag and eliminate the effected periods of time. This could be done reliably by eye, but we use a more quantitative approach described in Sec. V B. Though the lines only appeared at frequencies above $\sim$ 850 kHz, we excluded all data from the tagged time intervals. In total, these cuts eliminated $\sim$30% of the exposure.

IV. CALIBRATION

We quantify our detector response to a potential ADM signal, by performing a set of calibration measurements. Each measurement involves injecting a series of AC signals with known amplitude and frequency into the calibration system described in Sec. II E. We compare the power measured by our readout circuit to the power expected from the flux through the pickup loop generated by the calibration loop.

The input signal is generated by a Stanford Research Systems SG380 signal generator, locked to the same Rb frequency standard as the digitizer. The SG380 has very low phase noise and is able to output a tone with very long coherence time (longer than our measurement time), such that the resulting peak in the PSD was less than one frequency bin wide. For each amplitude and frequency, we perform a similar data collection to our axion search. We collect, transform and average 1 s buffers to form an averaged PSD. A zoomed example calibration line is shown in Fig. 6a. The resulting peaks are typically only one bin wide. We measure the power in each calibration peak and compare this to the expected flux power generated by the calibration loop.

We perform this procedure for between 120 and 200 frequencies from 10 kHz to 3 MHz, and for four different input amplitudes: 10 mVpp, 20 mVpp, 100 mVpp and 200 mVpp. The resulting gain spectrum is shown in Fig. 6b, and shows good agreement between the different input amplitudes. We perform the calibration measurement before the ADM search run, between the Magnet On run and the Magnet Off run, and again after the Magnet Off run. The resulting calibration curves were very consistent in time and did not depend on whether the magnet was on or off. For the present search, we determine our final calibration by interpolating the 200 mVpp data taken with the magnet on (red curve in Fig. 6b).

From Fig. 6b, we see that our measured gain is a factor of $\approx$ 6.5 below what is expected based on the calculated circuit inductance. We tested each element of the calibration circuit and determined that the discrepancy came from the pickup loop side of the circuit and had a frequency dependence consistent with an additional inductance. Based on this, we suspect that the factor of $\approx$6.5 most likely comes from parasitic inductance in the readout circuit. This degrades the overall sensitivity of our axion search, and is the focus of future upgrades.

V. AXION SEARCH AND LIMIT EXTRACTION

For the present analysis, we restrict our axion search to the frequency range 75 kHz $< f <$ 2 MHz or axion mass range $0.31 < m_a < 8.3$ neV. We therefore do not include the $\mathcal{F}_{100\mathrm{k}}$ spectrum in the rest of this analysis, as it has a Nyquist frequency of 50 kHz. We use the $\bar{\mathcal{F}}_{10\mathrm{M}}$ spectra to search the range from 500 kHz $< f <$ 2 MHz, and the $\bar{\mathcal{F}}_{1\mathrm{M}}$ spectra to search the range 75 kHz $< f <$ 500 kHz. In this way, a potential signal would be covered by at least 10 frequency bins at all frequencies. We further average the averaged spectra $\bar{\mathcal{F}}_{10\mathrm{M}}$ and $\bar{\mathcal{F}}_{1\mathrm{M}}$ to contain 3200 and 480 averages respectively. This decreases our ability to resolve time variations in our background noise to $\approx$9 h and $\approx$18 h, respectively. This step is not necessary for our analysis and is purely to decrease the computational resources required by a factor of $\approx$ 40. After this, our 1 month of data collection is spanned by 75 $\bar{\mathcal{F}}_{10\mathrm{M}}$ spectra and 37 $\bar{\mathcal{F}}_{1\mathrm{M}}$ spectra.[2]

Our data analysis procedure closely follows the method introduced in [30]. Our expected signal is a narrow peak in the pickup loop PSD above the noise background, with a width $\Delta f/f \sim 10^{-6}$ arising from the ADM velocity dispersion. The challenge in a broadband search such as this is that we are scanning a large number of mass points without the benefit of being able to efficiently 'rescan' mass points with possible signal detections. As such, we need to be thorough with our statistical modeling, as at least some points are likely to populate the tails of any distribution. In this section, we describe the statistical modeling of our expected signal and background, as well as a data quality cut for tagging the periods of time when transient noise causes the data to look neither like signal nor background.

[2] We recycle the notation because we have only changed the number of spectra contributing to the average, but otherwise, they are conceptually equivalent.

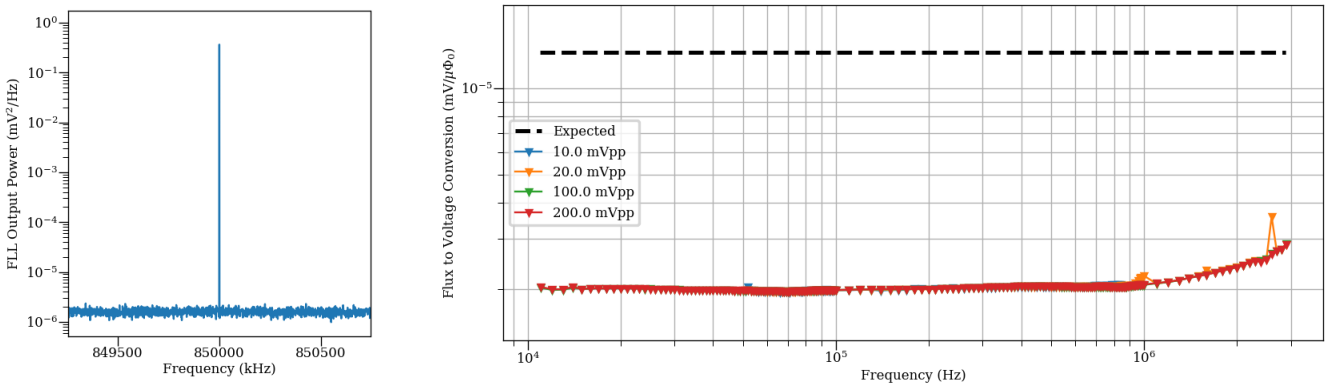

(a) (b)

FIG. 6. (a) Example calibration peak at 850 kHz with 10 mVpp excitation and 90 dB of attenuation. Bin width is 1 Hz wide and all power is contained within a single bin. Output voltage is measured at the output of the amplifier electronics. (b) Measured detector response for four different input amplitudes taken with the magnet on. The measured gain is a factor of ≈ 6.5 below the expected response (dashed line at the top). The outlier in the 20 mVpp spectrum is the result of a background line contributing power to the measured peak.

A. Likelihood Analysis

The local ADM field can be thought of as arising from a partially coherent sum over a very large number of individual axion particles, where the phases of each particle are randomly distributed. As a result, the expected signal power in each frequency bin is drawn from an exponential distribution. When averaged over $N_{\rm avg}$ independent PSDs, the signal in each frequency bin k will follow an Erlang distribution. When combined with background noise that is incoherent and Gaussian distributed in the time domain, the resulting PSD data is still Erlang-distributed [30]. Accordingly, for a single averaged PSD, our combined signal-plus-background model prediction in each bin is an Erlang distribution,

$$P(\bar{\mathcal{F}}_k; N_{\rm avg}, \lambda_k) = \frac{N_{\rm avg}^{N_{\rm avg}}}{(N_{\rm avg}-1)!}\frac{(\bar{\mathcal{F}}_k)^{N_{\rm avg}-1}}{\lambda_k^{N_{\rm avg}}} e^{-\frac{N_{\rm avg}\bar{\mathcal{F}}_k}{\lambda_k}}, \tag{5}$$

with shape parameter $N_{\rm avg}$ and mean $\lambda_k = s_k + b$, where

$$s_k = \begin{cases} A\,\frac{\pi f(v)}{m_a v}\Big|_{v=\sqrt{4\pi f_k/m_a - 2}} & f_k > m_a/2\pi\,, \\ 0 & f_k \le m_a/2\pi\,, \end{cases} \tag{6}$$

and b is the expected background power. Here, A denotes the combination of parameters that control the signal strength, defined in Eqn. 3. We assume $f(v)$ is given by the Standard Halo Model (SHM), with velocity dispersion $v_0 = 220\,$km/s, and $v_{\rm obs} = 232\,$km/s the DM velocity in the Earth frame [31] and $\rho_{\rm DM} = 0.4\,{\rm GeV/cm^3}$ [32, 33].

We build an analysis over a set of $\mathcal{N}$ averaged spectra $\bar{\mathcal{F}}_j$, each one an average over $N_{\rm avg}$ individual PSD. For example, the analysis of the $\bar{\mathcal{F}}_{\rm 10M}$ spectra, we have $\mathcal{N} = 75$ averaged spectra, where each averaged spectrum $\bar{\mathcal{F}}_j$ is an average over $N_{\rm avg} = 3200$ PSDs (with the possible exception of the final averaged spectrum which usually has fewer PSDs contributing). We search for an axion signal at mass $m_a = f/(2\pi)$, by restricting our search to a window containing frequency bins from $f_{k_i(m_a)} = m_a/(2\pi)$ to $f_{k_f(m_a)} = (1+4v_0^2)f_{k_i(m_a)}$ – approximately 8 times the width of the expected signal. Since incoherent background noise varies on frequency scales much larger than this, we can approximate the background noise level in this window as independent of frequency. We tested that our final results were insensitive to the precise choice of this window width. We account for long term variability in our noise levels by allowing the expected background level to vary from one averaged spectrum $\bar{\mathcal{F}}_j$ to the next; we denote $\mathbf{b} = \{b_1, b_2, \ldots, b_{\mathcal{N}}\}$ to be these background values, which we treat as nuisance parameters. The expected axion signal strength A is constant across our data taking period and thus is the same for each $\bar{\mathcal{F}}_j$. For a given axion mass point m_a, signal strength A and background values $\mathbf{b}$, we calculate the likelihood of our data $\mathbf{d}$:

$$\mathcal{L}(\mathbf{d}_{m_a}|A,\mathbf{b}) = \prod_{j=1}^{\mathcal{N}}\prod_{k=k_i(m_a)}^{k_f(m_a)} P(\bar{\mathcal{F}}_{j,k}; N_{{\rm avg},j}, \lambda_{j,k}), \tag{7}$$

where k indexes the (windowed) frequency bins and j indexes the different spectra. We allow for the generic possibility that each spectrum $\bar{\mathcal{F}}_j$ has a different number of averages, $N_{{\rm avg},j}$. This accounts for the spectra collected at the very end of the data taking period which have a different number of averages.

With the likelihood in Eq. (7), we perform a likelihood ratio test to search for a possible axion signal. To claim a detection, we place a 5σ threshold on the profiled likelihood ratio between the signal-plus-background and background-only hypotheses. We define a test-statistic

(TS) for discovery as

$$\mathrm{TS}(m_a) = 2\ln\left[\frac{\mathcal{L}(\mathbf{d}_{m_a}|\hat{A},\hat{\mathbf{b}})}{\mathcal{L}(\mathbf{d}_{m_a}|A=0,\hat{\mathbf{b}}_{A=0})}\right], \tag{8}$$

where $\hat{A}$ and $\hat{\mathbf{b}}$ are the values of A and $\mathbf{b}$ which achieve the global maximum of the likelihood, and $\hat{\mathbf{b}}_{A=0}$ is the value which achieves the constrained maximization with $A = 0$. The maximization of A is performed over a range including positive and negative parameter values, accommodating that a negative parameter value may provide the optimal fit to the data. If $\hat{A} < 0$, it is understood that the corresponding best-fit axion coupling is 0, as no value of $g_{a\gamma\gamma}$ could lead to negative-valued $\hat{A}$. The 5σ condition for discovery at a given m_a is $\mathrm{TS}(m_a) > \mathrm{TS}_{\mathrm{thresh}}$, where

$$\mathrm{TS}_{\mathrm{thresh}} = \left[\Phi^{-1}\left(1 - \frac{2.87\times10^{-7}}{N_{m_a}}\right)\right]^2 \tag{9}$$

accounts for the local significance as well as the look-elsewhere effect (LEE) for the N_{m_a} independent masses in the analysis (here Φ is the cumulative distribution function for the normal distribution with zero mean and unit variance) [30]. For this analysis, $N_{m_a} \approx 8.1\times10^6$ between 75 kHz and 2 MHz (see below), and $\mathrm{TS}_{\mathrm{thresh}} = 56.1$.

Where we do not see a detection, we set a 95% C.L. limit, $A_{95\%}$, with a similar profiled likelihood ratio. To do so, we use the following test statistic for upper limits

$$t(m_a, A) = \begin{cases} 2\ln\left[\frac{\mathcal{L}(\mathbf{d}_{m_a}|\hat{A},\hat{\mathbf{b}})}{\mathcal{L}(\mathbf{d}_{m_a}|A,\hat{\mathbf{b}}_A)}\right] & A \geq \hat{A}\,, \\ 0 & \text{Otherwise}\,. \end{cases} \tag{10}$$

Here, $\hat{\mathbf{b}}_A$ is the background values that maximizes the likelihood for a given A. Using t, we can establish the 95% C.L. limit $A_{95\%}$ where $t(m_a, A_{95\%}) = 2.71$. In this limit setting procedure, it is necessary that $\hat{A}$ was allowed to be negative-valued if this provided the best fit in order to make an accurate calculation of $A_{95\%}$. In addition, we implement one-sided power-constrained limits [34], which in practice means that we do not allow ourselves to set a limit stronger than the 1σ lower level of the expected sensitivity band. This ensure that our constraints are statistically conservative while also addressing the possibility that $A_{95\%}$ is negative-valued.

Finally, we discuss the set of mass points over which we scan. In principle, we can search for an axion signal at any value between $75\,\mathrm{kHz} < m_a/(2\pi) < 2\,\mathrm{MHz}$. In our data, this range is spanned by 57.5×10^6 frequency bins and it might seem natural to search for an axion signal centered at each frequency bin. However, since each axion signal model is resolved by between 10 and 100 frequency bins, neighboring frequency bins would produce very strongly correlated results. Alternatively, a log-spaced set of N masses such that $m_a^{i+1}/m_a^i = 1 + 4v_0^2$, would achieve a minimal coverage such that every frequency bin belongs to exactly one fit window. However, the spacing that achieves a set of statistically independent axion mass points is given by $m_a^{i+1}/m_a^i \approx 1 + 3v_0^2/4$ [30], which yields $N_{m_a} \approx 8.1\times10^6$ independent axion masses within our frequency range. For our search procedure, we therefore increase the granularity of the search and produce a set of log-spaced masses that obey $m_a^{i+1}/m_a^i = 1 + v_0^2/2$. This eight-fold enhancement in the resolution of our tested masses, as compared to the minimal coverage set, results in overlapping signal windows of masses studied in our analysis and allows us to over-resolve a potential axion signal by a factor of two. This yields a total of 13.0×10^6 mass points to test, which is appropriately larger than the estimated number of independent mass points. We emphasize that this choice of mass points is not a fundamental limit on our mass resolution but is instead imposed merely by computational resources. In the event of an observed excess, we could fit a region around it with the mass floating in the fit.

For each mass studied, we also compute the expected sensitivity bands from the null-hypothesis models using the Asimov dataset procedure [35], following implementation outlined in [30].

B. Quality Cuts

We can also use the analysis infrastructure described in the previous section to veto mass points where the condition of Gaussian-distributed incoherent noise does not hold as well as to create a quantitative data quality cut to identify periods of time with excess transient noise.

In order to tag periods of time with increased transient noise, we leverage the fact that the transient noise does not appear as a single peak, but instead as a forest of many correlated peaks. Under the null hypothesis of a flat background, the survival function for the test statistic t is given by

$$S(t) = 2\left(1 - \Phi\left(\sqrt{t}\right)\right). \tag{11}$$

The presence of a true axion signal, would yield a small number of mass points that deviate from this distribution. But a violation of the null hypothesis of a flat background – e.g. due to a forest of correlated transient noise – would lead to a much larger deviation from this distribution.

For each $\bar{\mathcal{F}}_j$, we calculate the number of mass points with at least a 3σ excess within the time period covered by that PSD, accounting for the LEE. We find that the number of 3σ excesses, follows a clear bi-modal distribution with an obvious time correlation, see Fig. 7a. This allows us to place a quantitative cut by requiring that an averaged PSD, $\bar{\mathcal{F}}_j$, have fewer than 30 mass points with a 3σ excess. This effectively eliminates periods of time with transient noise. We emphasize two points here: first, by placing this cut on a statistic which is calculated across a broad ranges of frequency, we do not produce a

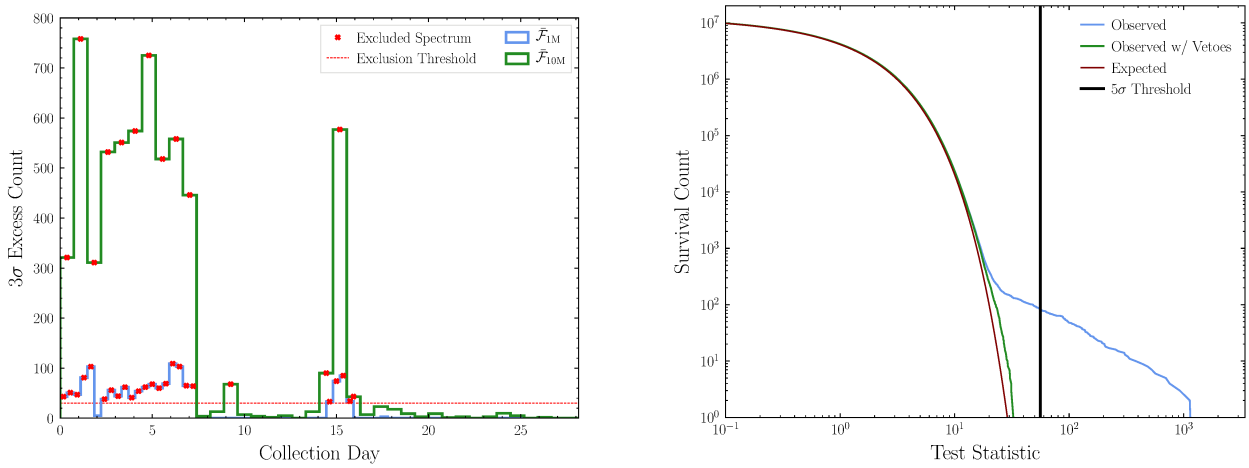

(a) (b)

FIG. 7. (a) The number of 3σ excesses accounting for the look-elsewhere effect in each spectrum after vetoing the excesses that are present in the corresponding Magnet Off data. (b) The distribution of local TS values in the full month of analyzed data after removing periods of transient noise. In blue is the observed distribution of local TS values prior to vetoing the Magnet Off excesses. In green, the observed distribution of TS values after the Magnet Off veto. In red, the expected distribution under the null hypothesis. We see that after applying vetoes, there is excellent agreement down to very low survival counts, with no remaining 5σ excesses.

bias at any one mass point or range. Second, as we describe in the next section, a single axion signal would not be expected to create 30 mass points with 3σ excesses or larger. So, while exotic models with multiple axion could be affected by this cut, it would not present a problem for our baseline model of a single axion. This cut removes $\approx 30\%$ of our exposure.

Once we have removed periods of time with high transient noise, we remove individual mass points that have non-transient noise peaks – or are in other ways inconsistent with our null hypothesis of a flat background. We perform our axion discovery analysis on the Magnet Off data – where we expect no axion signal to be present. Any mass points that show LEE-corrected excess beyond 5σ in the Magnet Off data are vetoed. We consider these mass points to have poorly understood backgrounds where we do not have sensitivity to ADM. Out of 13.0×10^6 mass points, this requirement vetoed 18,733(6,651) points in the range $500\,\text{kHz} < f < 2\,\text{MHz}$ ($75\,\text{kHz} < f < 500\,\text{kHz}$) and implies a decrease in our signal recovery efficiency of 0.2%.

The axion search data, collected with the magnet on, showed 83 excesses with LEE corrected significance $\geq 5\sigma$, however they were all vetoed by cutting against the Magnet Off data. It is worth pointing out that the number of 5σ excesses in the axion search data was much smaller than in the Magnet Off data, due to the lower transient noise levels seen during that run. In Fig. 7b, we show the distribution of TS values before and after the Magnet Off veto, compared against the distribution expected under the null hypothesis. The strong agreement between the theoretical expectation and the observed distribution after vetoes are applied demonstrates that the experimental backgrounds are well-modeled by the null hypothesis and ABRACADABRA-10 cm has strong discovery power under deviations from this theoretical expectation. In particular, in the results presented in [16], we find no significant excesses after vetoes are applied.

C. Recovering an Injected Signal

As a final test of our analysis procedures, we test that we are able to recover an injected signal and discover an axion at the claimed significance. This is crucial because as axion searches achieve greater sensitivity, there will be an inevitable trade-off between broadband and narrowband coverage, and a claimed exclusion at a given $g_{a\gamma\gamma}$ will be used as justification to avoid re-scanning parameter space that has already been tested.

To test the discovery power of our analysis procedure, we generate Monte Carlo (MC) spectra characterized by a mean background level $\hat{b}$ and $\{N_{\text{avg},j}\}$ averaged PSDs identical to that of the observed data. We then add an artificial axion signal with signal template set by the SHM, for a range of axion coupling strengths A. We perform our quality cuts on the individual spectra, then perform the joint analysis on the surviving spectra. We then evaluate the best-fit axion coupling and the 95^{th} percentile limit on that coupling as a function of the "true" axion coupling of the injected signal for each MC realization. Figure 8 shows the resulting distributions of reconstructed $g_{a\gamma\gamma}$ and t-values for six axion masses.

Critically, the ability of our analysis procedure to accurately recover the correct axion parameters when allowable by the background level is unaffected by the qual-

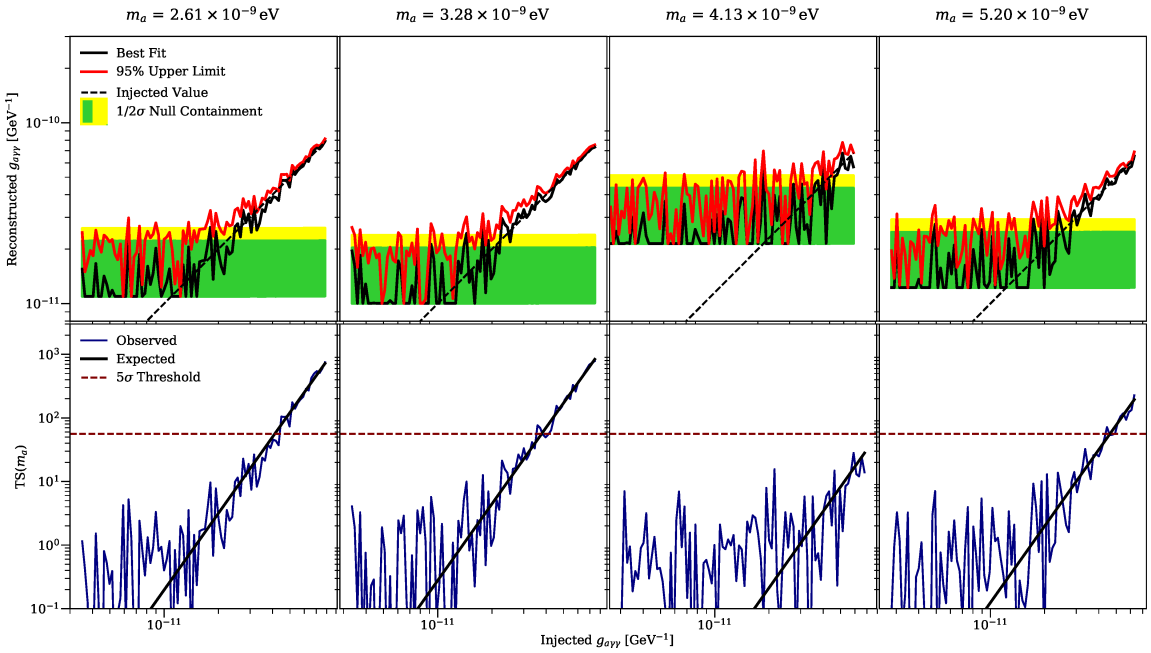

FIG. 8. (*Top row*) The recovered signal parameters as a function of the injected signal parameters in four Monte Carlo realizations with identical mean background levels. Green and yellow bands indicate the expected 1 and 2σ containment for the upper 95% limit on the axion coupling under the hypothesis of no axion signal. (*Bottom row*) The observed and expected test statistic for discovery as a function of the injected signal strength. The dashed red line indicates the threshold for a discovery at 5σ significance accounting for the LEE.

ity cuts. This is most clearly seen in the bottom panels of Fig. 8, which show the t-value as a function of the injected signal strength, as compared to the expected t-value as a function of the injected signal. We see strong agreement between the expected and observed test statistic when the signal is strong enough that we expect to be able to discover it (i.e., when the injected coupling lies above the null model containment bands). These examples also demonstrate that when our signal is not significant enough to be discovered, our limit-setting procedure is unaffected by the quality cuts.

VI. CONCLUSIONS

The successful run of the ABRACADABRA-10 cm experiment [16] introduced and validated several new techniques useful for constraining axion dark matter, including a toroidal magnet geometry sensitive to ADM at $m_a \lesssim 1\,\mu$eV, a broadband readout technique capable of handling the data-load required to study millions of axion masses simultaneously, a signal injection through a calibration loop to characterize this type of detector, and the first implementation of the broadband data analysis technique proposed in [30]. In this paper, we have described in detail the implementation and validation of these techniques, which lend additional to support to the results presented in [16].

Of greatest practical concern for the first results is the identification of the mismatch between expected and measured end-to-end gain, which we aim to rectify in the next data-taking run, and the mitigation of vibrational noise. We have also emphasized the statistical analysis employed to extract the first results. The goal of this analysis is to establish a sure footing for the presented statistical limits with a robust understanding of the exclusion limits. This is important as next generation experiments reach for ever higher sensitivities and re-scanning regions of parameter space becomes prohibitively time-consuming. In addition, the success of the statistical description of our data paves the way for the use of a blind analysis pipeline, which we hope to use in future runs.

ABRACADABRA-10 cm represents the first step in an experimental search program, which aims to ultimately be sensitive to ADM in the coupling range preferred by the QCD axion. Future phases of ABRACADABRA will require larger magnets with higher fields, improved shielding, and strong mitigation of mechanical vibration. Augmenting the techniques described here with a resonant amplification readout and scan strategy will also greatly improve the sensitivity of a future full scale ABRACADABRA detector [19, 36, 37]. We have already begun engineering studies towards designing and build-

ing such a detector and ABRACADABRA-10 cm creates a strong foundation for this ongoing work.

ACKNOWLEDGMENTS

The authors would like to acknowledge the useful conversations and advice from Henry Barthelmess of Magnicon Inc. This research was supported by the National Science Foundation under grant numbers NSF-PHY-1658693, NSF-PHY-1806440, NSF-PHY-1505858 and NSF-PHY-1122374. This work was also supported by DOE Early Career Grant number DE-SC0019225, by the Kavli Institute for Cosmological Physics at the University of Chicago, by the Miller Institute for Basic Research in Science at the University of California, Berkeley, by the MIT Undergraduate Research Opportunities in Physics program, by the Leinweber Graduate Fellowship Program and by the Simons Foundation through a Simons Fellowship in Theoretical Physics. This research was supported in part through computational resources and services provided by Advanced Research Computing at the University of Michigan, Ann Arbor. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics under Award Numbers DE-FG02-97ER41041 and DEFG02-97ER41033 and by the Office of High Energy Physics under grant DE-SC-0012567. We would like to thank the University of North Carolina at Chapel Hill and the Research Computing group for providing computational resources and support that have contributed to these research results.

[1] D. J. E. Marsh, in *Proceedings, 13th Patras Workshop on Axions, WIMPs and WISPs, (PATRAS 2017): Thessaloniki, Greece, 15 May 2017 - 19, 2017* (2018) pp. 59–74, arXiv:1712.03018 [hep-ph].
[2] M. Battaglieri *et al.*, in *U.S. Cosmic Visions: New Ideas in Dark Matter College Park, MD, USA, March 23-25, 2017* (2017) arXiv:1707.04591 [hep-ph].
[3] M. Tegmark, A. Aguirre, M. Rees, and F. Wilczek, Phys. Rev. **D73**, 023505 (2006), arXiv:astro-ph/0511774 [astro-ph].
[4] H. Davoudiasl, D. Hooper, and S. D. McDermott, Phys. Rev. Lett. **116**, 031303 (2016), arXiv:1507.08660 [hep-ph].
[5] P. W. Graham and A. Scherlis, Phys. Rev. **D98**, 035017 (2018), arXiv:1805.07362 [hep-ph].
[6] F. Takahashi, W. Yin, and A. H. Guth, Phys. Rev. **D98**, 015042 (2018), arXiv:1805.08763 [hep-ph].
[7] P. Agrawal, G. Marques-Tavares, and W. Xue, JHEP **03**, 049 (2018), arXiv:1708.05008 [hep-ph].
[8] L. Visinelli and P. Gondolo, Phys. Rev. D **81**, 063508 (2010), arXiv:0912.0015 [astro-ph.CO].
[9] R. T. Co, F. D'Eramo, and L. J. Hall, Phys. Rev. **D94**, 075001 (2016), arXiv:1603.04439 [hep-ph].
[10] L. Di Luzio, A. Ringwald, and C. Tamarit, Phys. Rev. **D98**, 095011 (2018), arXiv:1807.09769 [hep-ph].
[11] A. Arvanitaki, S. Dimopoulos, S. Dubovsky, N. Kaloper, and J. March-Russell, Phys. Rev. **D81**, 123530 (2010), arXiv:0905.4720 [hep-th].
[12] P. Svrcek and E. Witten, JHEP **06**, 051 (2006), arXiv:hep-th/0605206 [hep-th].
[13] P. Agrawal, J. Fan, M. Reece, and L.-T. Wang, JHEP **02**, 006 (2018), arXiv:1709.06085 [hep-ph].
[14] M. Farina, D. Pappadopulo, F. Rompineve, and A. Tesi, JHEP **01**, 095 (2017), arXiv:1611.09855 [hep-ph].
[15] R. T. Co, E. Gonzalez, and K. Harigaya, (2018), arXiv:1812.11186 [hep-ph].
[16] J. L. Ouellet *et al.*, (2018), arXiv:1810.12257 [hep-ex].
[17] P. Sikivie, Phys. Rev. Lett. **51**, 1415 (1983), [Erratum: Phys. Rev. Lett. 52, 695 (1984)].
[18] P. Sikivie, N. Sullivan, and D. B. Tanner, Phys. Rev. Lett. **112**, 131301 (2014), arXiv:1310.8545 [hep-ph].
[19] Y. Kahn, B. R. Safdi, and J. Thaler, Phys. Rev. Lett. **117**, 141801 (2016).
[20] S. J. Asztalos *et al.* (ADMX), Phys. Rev. **D64**, 092003 (2001).
[21] N. Du *et al.* (ADMX Collaboration), Phys. Rev. Lett. **120**, 151301 (2018).
[22] L. Zhong *et al.*, Phys. Rev. D **97**, 092001 (2018).
[23] "Super conducting systems inc." `http://www.superconductingsystems.com`.
[24] COMSOL Multiphysics v. 5.3a. www.comsol.com. COMSOL AB, Stockholm, Sweden.
[25] S. Henry *et al.*, Nucl. Instrum. Method Phys. Res A **763**, 155 (2014).
[26] http://www.magnicon.com.
[27] D. Drung, C. Abmann, J. Beyer, A. Kirste, M. Peters, F. Ruede, and T. Schurig, IEEE Transactions on Applied Superconductivity **17**, 699 (2007).
[28] https://www.alazartech.com.
[29] M. Frigo and S. G. Johnson, Proceedings of the IEEE **93**, 216 (2005), special issue on "Program Generation, Optimization, and Platform Adaptation".
[30] J. W. Foster, N. L. Rodd, and B. R. Safdi, Phys. Rev. D **97**, 123006 (2018).
[31] P. J. McMillan and J. J. Binney, Mon. Not. Roy. Astron. Soc. **402**, 934 (2010), arXiv:0907.4685 [astro-ph.GA].
[32] R. Catena and P. Ullio, JCAP **1008**, 004 (2010), arXiv:0907.0018 [astro-ph.CO].
[33] F. Iocco, M. Pato, G. Bertone, and P. Jetzer, JCAP **1111**, 029 (2011), arXiv:1107.5810 [astro-ph.GA].
[34] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, (2011), arXiv:1105.3166 [physics.data-an].
[35] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, Eur. Phys. J. **C71**, 1554 (2011), [Erratum: Eur. Phys. J.C73,2501(2013)], arXiv:1007.1727 [physics.data-an].
[36] M. Silva-Feaver *et al.*, IEEE Trans. Appl. Supercond. **27**, 1400204 (2017), arXiv:1610.09344 [astro-ph.IM].
[37] S. Chaudhuri, K. Irwin, P. W. Graham, and J. Mardon, (2018), arXiv:1803.01627 [hep-ph].