

## New results from HAYSTAC's phase II operation with a squeezed state receiver

M. J. Jewell<sup>1,2</sup>, A. F. Leder,<sup>3</sup> K. M. Backes,<sup>1,2,\*</sup> Xiran Bai,<sup>1,2</sup> K. van Bibber,<sup>3</sup> B. M. Brubaker,<sup>4</sup> S. B. Cahn,<sup>1,2</sup> A. Droster,<sup>3</sup> Maryam H. Esmat,<sup>5</sup> Sumita Ghosh,<sup>6,2</sup> Eleanor Graham,<sup>1,2</sup> Gene C. Hilton,<sup>7</sup> H. Jackson,<sup>3</sup> Claire Laffan,<sup>1,2</sup> S. K. Lamoreaux,<sup>1,2</sup> K. W. Lehnert,<sup>4</sup> S. M. Lewis,<sup>3,†</sup> M. Malnou,<sup>4</sup> R. H. Maruyama,<sup>1,2</sup> D. A. Palken,<sup>4,‡</sup> N. M. Rapidis,<sup>3,§</sup> E. P. Ruddy,<sup>4</sup> M. Simanovskaia,<sup>3,§</sup> Sukhman Singh,<sup>1,2</sup> D. H. Speller,<sup>5</sup> Leila R. Vale,<sup>7</sup> H. Wang,<sup>1,2,§</sup> and Yuqi Zhu<sup>1,2</sup>

(HAYSTAC Collaboration)

<sup>1</sup>Department of Physics, Yale University, New Haven, Connecticut 06520, USA

<sup>2</sup>Wright Laboratory, Department of Physics, Yale University, New Haven, Connecticut 06520, USA

<sup>3</sup>Department of Nuclear Engineering, University of California Berkeley, California 94720, USA

<sup>4</sup>Department of Physics, University of Colorado, Boulder, Colorado 80309, USA

<sup>5</sup>Department of Physics and Astronomy, The Johns Hopkins University,  
3400 North Charles Street, Baltimore, Maryland 21211, USA

<sup>6</sup>Department of Applied Physics, Yale University, New Haven, Connecticut 06520, USA

<sup>7</sup>National Institute of Standards and Technology, Boulder, Colorado 80305, USA



(Received 26 January 2023; accepted 29 March 2023; published 28 April 2023)

A search for dark matter axions with masses  $> 10 \mu\text{eV}/c^2$  has been performed using the HAYSTAC experiment's squeezed state receiver to achieve subquantum limited noise. This work includes details of the design and operation of the experiment previously used to search for axions in the mass ranges 16.96–17.12 and 17.14–17.28  $\mu\text{eV}/c^2$  (4.100–4.140 GHz and 4.145–4.178 GHz) as well as upgrades to facilitate an extended search at higher masses. These upgrades include improvements to the data acquisition routine which have reduced the effective dead time by a factor of 5, allowing for the new region to be scanned  $\sim 1.6$  times faster with comparable sensitivity. No statistically significant evidence of an axion signal is found in the range 18.44–18.71  $\mu\text{eV}/c^2$  (4.459–4.523 GHz), leading to an aggregate upper limit exclusion at the 90% level on the axion-photon coupling of  $2.06 \times g_\gamma^{\text{KSVZ}}$ .

DOI: 10.1103/PhysRevD.107.072007

### I. INTRODUCTION

Astrophysical and cosmological observations indicate that more than a quarter of the universe's mass energy exists in the form of dark matter [1,2]. While there are a number of hypotheses for the nature of this dark matter, no candidate has been detected directly. One such candidate, the axion, was originally proposed as a solution to the strong *CP* problem [3,4] but also has the necessary characteristics to act as the dominant source of dark matter in the universe [5–9]. Models of the axion do not explicitly predict its mass or

coupling strength to matter, but for QCD axions all of its couplings,  $g_{aii}$ , are proportional to its mass,  $m_a$ . This includes its coupling to radiation,  $g_{\alpha\gamma\gamma}$ , which is relevant to this experiment. The resulting parameter space of the allowed masses and couplings is typically benchmarked by couplings from the Kim, Shifman, Vainshtein, and Zakharov (KSVZ) [10,11] and the Dine, Fischler, Srednicki, and Zhitnitsky (DFSZ) [12,13] models based on heavy quarks and multi-Higgs interactions, respectively. The ability to solve two outstanding questions in physics, coupled with null results from searches for weakly interacting massive particles (WIMPs), situates the axion as one of the most promising candidates for a new fundamental particle [14].

Although the mass of the QCD axion is largely unknown, axions in the mass range  $\sim 1$ –100  $\mu\text{eV}/c^2$  are favored by models in which the Peccei-Quinn (PQ) symmetry breaking occurs after inflation [15–17]. To date, the most sensitive probes for dark matter axions in this range are those using axion haloscopes [18–24]. These experiments exploit the axion's coupling to the pseudoscalar electromagnetic product  $\mathbf{E} \cdot \mathbf{B}$  by employing a large static magnetic field  $\mathbf{B}$  to

\*Present address: The MITRE Corporation.

†Present address: FNAL.

‡Present address: U.S. Senate Office of Senator Hickenlooper.

§Present address: Stanford.

convert the oscillating axion field into a small electric field  $\mathbf{E}$  which oscillates at frequency  $\nu_a \approx m_a c^2/h$  for axions of mass  $m_a$  [25,26]. To resonantly enhance the relatively small coupling strength  $g_{a\gamma\gamma}$  expected for this interaction, the induced field is enclosed inside a high quality factor cavity with a spatially overlapping cavity mode whose bandwidth encompasses  $\nu_a$ . When this resonant condition is met, a small power excess at  $\nu_a$  with spectral width given by the axion's virial velocity  $v \sim 270$  km/s would be present. Through the use of the haloscope technique, experiments have achieved ratios of the expected axion-induced signal power to the total experimental noise power (SNR), which allow for probes of dark matter axions sensitive to QCD couplings.

The main challenge of such an experiment arises from the fact that the neither the axion's mass nor coupling strength are *a priori* known, requiring that a successful search maintain this high sensitivity for axions of many different masses. Because the width of the cavity's resonant mode is small compared to the range of possible axion masses, the defining metric of a haloscope is not the SNR for a single axion mass but rather the rate  $R$  (in Hz/s) at which axions of different masses can be searched for at a given SNR. While modern haloscopes have achieved sensitivity to QCD axions over large mass ranges below  $\sim 10 \mu\text{eV}/c^2$  [18,19,24], extending these searches to higher frequencies becomes increasingly challenging as a number of experimental parameters, most notably the cavity volume, scale poorly with frequency. In order to maintain sensitivity to QCD axions at masses  $\gtrsim 10 \mu\text{eV}/c^2$ , improvements are needed to counter this drop in  $R$ .

One such path is the reduction of the total noise of the system, which can in part be achieved through the use of low-noise amplifiers operating at cryogenic temperatures, but it is traditionally limited by quantum noise from zero-point vacuum fluctuations. These fluctuations impose a standard quantum limit (SQL) to the noise a haloscope can achieve with phase insensitive linear amplifiers. In recent years it has been demonstrated that circumvention of this barrier is possible by coupling the cavity to a squeezed state receiver (SSR) [27] consisting of two Josephson parametric amplifiers (JPAs) [28]. Operation of these JPAs as phase sensitive amplifiers allows for the manipulation of squeezed vacuum states, in which the measurement uncertainty is moved into one quadrature of the field, resulting in a reduction of noise below the SQL. This allows for the cavity's measurement bandwidth to be substantially increased by strongly overcoupling the readout antenna, which in turns gives an enhancement to the scan rate. This was first demonstrated in an axion search by the Haloscope At Yale Sensitive To Axion Cold Dark Matter (HAYSTAC) [24,29,30].

In this work, details from HAYSTAC's quantum enhanced search for dark matter axions are presented in the context of the second phase of data taking with a focus

on new results from continued operation of the SSR since the initial demonstration published in Ref. [24]. A detailed overview of the experimental configuration is presented in Sec. II, highlighting both changes made to the detector in order to integrate the SSR and improvements made since the initial SSR run to further improve the scan rate during the acquisition of new data. Following this, results from HAYSTAC's dark matter search are presented in Sec. III with a focus on new data taken with the upgraded detector.

## II. EXPERIMENTAL OVERVIEW

The HAYSTAC experiment began operation in 2016 and is now in its second phase of axion sensitive measurements. In Phase I which ran from January 2016 to July 2017, HAYSTAC took data with a single Josephson parametric amplifier (JPA) operated in phase-insensitive mode to scan the mass range of  $23.15\text{--}24.0 \mu\text{eV}/c^2$  (5.6–5.8 GHz) [31,32] achieving nearly quantum-limited noise for the first time in a haloscope experiment. In order to push beyond this quantum limit, the receiver chain was modified in August 2018 to incorporate a SSR, and HAYSTAC began a second phase of data taking referred to as Phase II. The first iteration with this setup, Phase IIa, operated from September 2019 to April 2020 and searched for axions with masses between  $16.96\text{--}17.12$  and  $17.14\text{--}17.28 \mu\text{eV}/c^2$  (4.100–4.140 GHz and 4.145–4.178 GHz) achieving noise below the quantum limit for the first time in an axion experiment [24]. Following the conclusion of this search, a new set of JPAs optimized for frequencies  $> 4.2$  GHz were installed. This allowed for a second search, Phase IIb, to be performed in the higher mass range between  $18.44\text{--}18.71 \mu\text{eV}/c^2$  (4.459–4.523 GHz) between July 2021 and November 2021. This also included upgrades to the data acquisition routine, increasing the effective live time and in turn the scan rate of the experiment by a factor of  $\sim 1.6$ .

### A. Principle of squeezed state receiver enhancement

The benefit of a SSR to a haloscope experiment was first described in detail in Refs. [27,33]. The basic design of the SSR enhanced haloscope is shown in Fig. 1 where a microwave cavity converts axions into a detectable signal which is measured by a receiver chain consisting of two JPAs. In such an experiment, the presence or absence of an axion is determined through precise measurements of the cavity's field, which is governed by the Hamiltonian

$$\hat{H} = \frac{\hbar\nu_c}{2}(\hat{X}^2 + \hat{Y}^2), \quad (1)$$

where  $\hat{X}$  and  $\hat{Y}$  are the noncommuting quadratures of the cavity field. The precision by which the magnitude of this field can be measured is ultimately limited by the total noise of the system. If this noise is evenly distributed along both quadratures, any observation which attempts to measure the

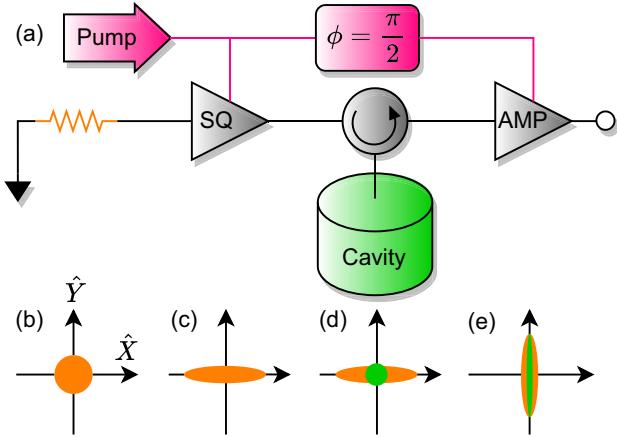


FIG. 1. (a) Simplified illustration of the cavity coupled to the SSR outlining the main components and (b)–(e) how they transform the state in quadrature phase space. (b) An initially unsqueezed vacuum state (orange) sourced by a  $50\ \Omega$  terminator is (c) sent to the SQ JPA, which squeezes the  $\hat{Y}$  quadrature. The state is then (d) reflected off of the cavity, thereby adding cavity noise (green), which would include any axion-induced signal. (e) The state is then routed to the AMP JPA, which is operated  $\pi/2$  out of phase of the SQ JPA such that it amplifies the previously squeezed  $\hat{Y}$  quadrature, which contains both cavity noise and squeezed vacuum noise.

magnitude of both quadratures simultaneously is subject to the SQL for haloscopes which imposes a minimum required measurement uncertainty [34]. Circumventing this limit is possible with the use of phase sensitive amplifiers, such as JPAs, which allow the distribution of noise to be manipulated in the quadrature phase space. In the case of the SSR this is achieved by operating two JPAs in line with each other with the cavity coupled between them. The first JPA, named the squeezer (SQ), allows for the noise incident on the cavity to be prepared in a squeezed state, with the majority of the uncertainty pushed into a single quadrature. The state is then reflected off of the cavity where it picks up noise from the cavity, such as an axion signal, which gets evenly distributed along both quadratures. A measurement of the state is then made with the second JPA, the amplifier (AMP), which noiselessly amplifies the state along the previously squeezed quadrature, resulting in a reduction of the total noise of the system ( $N_{\text{sys}}$ ).

In this picture, the  $\hat{X}$  and  $\hat{Y}$  variables are functions of the frequency difference or detuning [ $\delta_\nu = (\nu_a - \nu_c)$ ] between the axion's frequency ( $\nu_a$ ) and the cavity's frequency ( $\nu_c$ ), equivalent to half the JPA's pump frequency ( $\nu_p$ ), and are comprised of linear combinations of phasor components spaced equally around  $\nu_p$  at  $\pm\delta_\nu$ . The reduction to the total noise is understood by examining the three contributions to  $N_{\text{sys}}$  in terms of these phasor variables:

$$N_{\text{sys}}(\delta_\nu) = N_c(\delta_\nu) + N_r(\delta_\nu) + N_A(\delta_\nu) \quad (2)$$

expressed in units of power spectral density and assuming that the noise sources are statistically independent. The first term represents the Johnson-Nyquist noise sourced from the cavity's internal loss,  $N_c$ , which peaks on cavity resonance, with the same Lorentzian line shape as the cavity's transmission response. This noise would also contain the axion signal and is not impacted by the SSR. The main benefit from squeezing comes from the reduction of noise away from the resonance, where the main contributions are the noise of the total amplification chain referred to the input of the AMP,  $N_A$ , and the Johnson-Nyquist noise incident on and reflected off of the cavity,  $N_r$ , which is modified by the cavity's reflection response. Operation of the SSR creates correlations in the amplifier added noise and reflected noise, and an appropriate choice of the amplifier pump phase yields a partial cancellation of the noise contributions, resulting in an improvement to the experimental sensitivity away from the cavity's resonance where these sources are dominant.

Because the resulting noise reduction is a function of  $\delta_\nu$ , the improvement to an axion search is quantified by comparing  $N_{\text{sys}}$  to the expected power of an axion signal over a range of  $\delta_\nu$ 's to capture improvements both on and off resonant with the cavity. The detected power is given by [25,26]

$$P_{\text{ax}}(\delta_\nu) = \left( \frac{2\pi\rho_a\hbar^3c^3}{\mu_0} \right) \left( \frac{g_{a\gamma\gamma}^2}{m_a^2} \right) \nu_c B_0^2 V C_{nml} \times Q_L \frac{\beta}{1 + \beta} \frac{1}{1 + (2\delta_\nu/\Delta\nu_c)^2}, \quad (3)$$

where  $\rho_a$  is the local dark matter density  $0.45\ \text{GeV}/\text{cm}^3$  [35],  $B_0$  is the amplitude of the magnetic field along the longitudinal axis of the cylindrical cavity,  $V$  is unfilled volume of the cavity, and  $C_{nml}$  is the cavity form factor for the mode indexed by  $nml$ . The cavity's linewidth  $\Delta\nu_c$  is given by the loaded quality factor as  $Q_L = \nu_c/\Delta\nu_c$  which is reduced by a factor of  $(1 + \beta)$  from its unloaded value  $Q_0$ , where the coupling of the antenna is parametrized by  $\beta$ . As with  $N_c$ , this signal is defined as a function of  $\delta_\nu$  and is filtered by the cavity's transmission profile represented by the last term in parentheses. While this power peaks exactly at  $\delta_\nu = 0$ , there is still detectable power at  $\delta_\nu > 0$  proportional to the cavity's linewidth. This results in each cavity tuning step having sensitivity to a range of axion masses. This window, referred to as the measurement bandwidth, is defined as the region where the SNR,

$$\Sigma(\delta_\nu) = \frac{P_{\text{ax}}(\delta_\nu)}{h\nu N_{\text{sys}}(\delta_\nu)} \sqrt{\frac{\tau}{\Delta\nu}}, \quad (4)$$

exceeds a desired threshold. The scaling under the square root represents the improvement with the number of independent power measurement over a bandwidth  $\Delta\nu$  taken over an integration period of  $\tau$ .

A typical haloscope operates in a series of discrete tuning steps where subsequent spectra, whose bandwidths overlap, sum together to improve the total SNR. Because of this, there is a benefit to increasing the bandwidth of each step by increasing the coupling  $\beta$ . The resulting reduction of  $Q_L$  decreases the experiment's sensitivity near the cavity's resonance but can be compensated for by an increase in off-resonant sensitivity if the SNRs at these frequencies are above the desired threshold. For a given haloscope, the optimal overcoupling is defined as the  $\beta$  which maximizes the scan rate

$$R \propto \int_{-\infty}^{+\infty} \Sigma^2(\delta_\nu) d\delta_\nu. \quad (5)$$

Without squeezing, the optimal overcoupling is found to be  $\sim 2$  [36], with operation beyond this point resulting in a reduction of the scan rate as the noise from  $N_r$  quickly washes out signals for all frequencies. By making use of the SSR, the noise contribution from  $N_r$  is reduced over a broad range of frequencies and improves the SNR at  $\delta_\nu > 0$ . This allows the cavity's coupling coefficient to be increased beyond the optimal value for an unsqueezed measurement, resulting in a larger measurement bandwidth for each tuning step, and in turn gives an enhancement to the scan rate of the experiment.

For an experiment which is able to deliver squeezing  $S$ , defined as the reduction in the variance of the system output  $S = \sigma_{\text{on}}^2 / \sigma_{\text{off}}^2$ , the optimal overcoupling is approximately  $\beta = 2S$ . In the case of an ideal system, with no transmission loss, this can be used to arbitrarily increase the scan rate, which approximately scales as  $R \propto S$ . However, the deliverable squeezing in a realistic system is limited by transmission loss between the SQ and AMP JPAs, which results in part of the squeezed state being replaced by unsqueezed vacuum. For a system with transmission efficiency  $\eta$  between the JPAs, the theoretical limit to the delivered squeezing, assuming noiseless amplification, is given by

$$S = \frac{\eta}{G_{\text{SQ}}} + 1 - \eta, \quad (6)$$

where  $G_{\text{SQ}}$  is the one quadrature gain of the SQ. During HAYSTAC's first operation using the SSR in Phase IIa, a transmission efficiency of 0.63 was achieved with a SQ gain of 13 dB. This resulted in a delivered squeezing of 4 dB, allowing for the scan rate to be roughly doubled relative to operation without squeezing [24].

## B. Detector design and performance

A model of the experimental setup showing the cavity installed in the magnet's bore is shown in Fig. 2. The main cavity is composed of a cylinder of copper-plated stainless steel which is 25.4 cm long and has a 10.2 cm inner diameter. The resonant frequency of the cavity is tuned by

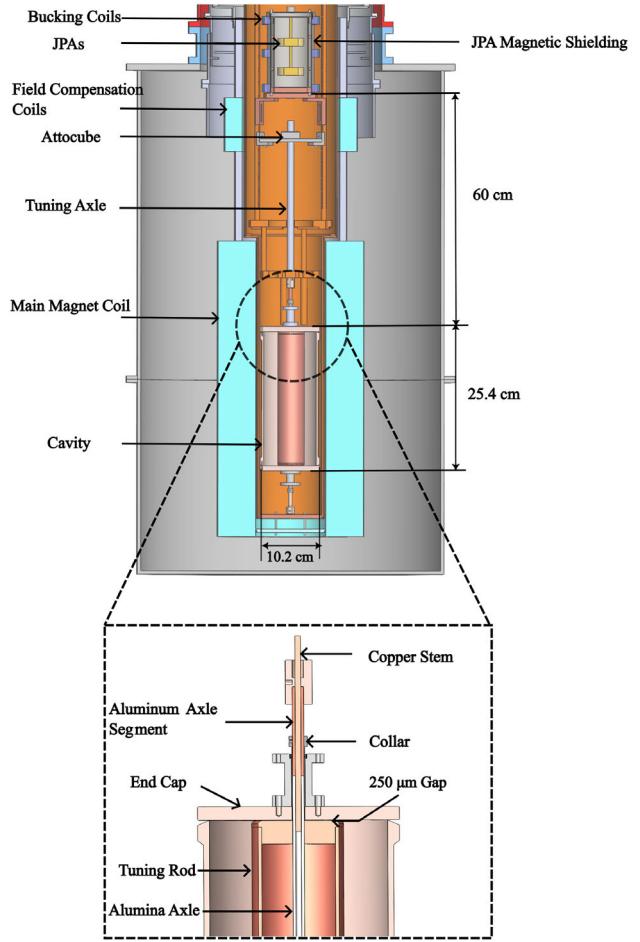


FIG. 2. Top: CAD model of the main cavity installed inside the magnet's bore. Bottom: a detailed view of the top of the cavity highlighting the tuning rod design. The axle used to tune the cavity is segmented such that only the small sections outside of either end cap are conducting, with the sections inside the cavity made of alumina. The tuning rod is further cooled by inserting a copper stem through the hollow axle, from the top, bottom, or both ends. Also shown is the JPA housing located  $\sim 60$  cm above the main cavity. The magnet's 70 K and 4 K shields, and the fridge's 4 K and 1 K shields are omitted for simplicity.

rotating a 5.1 cm diameter tuning rod, also made of copper-plated stainless steel, which pivots around an off-axis axle. The rod takes up  $\sim 25\%$  of the cavity volume, leaving an unfilled volume of  $V = 1.545$  l. Rotary motion of the axle is driven by an Attocube ANR240 precision piezoelectric motor [37] designed to operate at cryogenic temperatures. The cavity mode of interest for the axion search is the  $\text{TM}_{010}$ , as this mode has the highest overlap with the external magnetic field as represented by the form factor  $C_{010}$ . Varying the position of the rod by  $180^\circ$  allows for tuning over the range of 3.6–5.8 GHz for the  $\text{TM}_{010}$  mode. The experiment is operated inside an LD250 BlueFors dilution refrigerator [38] which achieves a base temperature of 61 mK at the mixing chamber plate. The cavity sits inside the 12.7 cm bore of a solenoidal magnet from

Cryomagnetics, Inc., used to generate an 8 T magnetic field within the cavity volume.

To prevent degradation of the quality factor due to coaxial loss, the segment of the axle inside the cavity and extending through the tuning rod is a 6.35 mm diameter hollow alumina tube, which transitions to a short aluminum segment outside of the cavity on either end cap, as shown in Fig. 2 for the top end cap. The axle is fixed in place with collars on either side, which rest on the inner race of their respective bearings. This results in a  $\sim 250$   $\mu\text{m}$  gap between the rod and either end cap to allow for smooth tuning of the rod. Due to the low thermal conductivity of alumina at cryogenic temperatures [39,40], a copper stem with diameter 3.175 mm is inserted into the axle to improve the thermal link between the tuning rod and the cavity with an insertion depth chosen to maximize the thermal contact area while simultaneously minimizing a potential decrease in the quality factor. This was found to be necessary during early operations of Phase I which observed excess thermal noise from the cavity, consistent with a component of the cavity failing to reach the base temperature [32,36]. Including the stem reduced the effective noise temperature from 600 mK to 250 mK while still maintaining an unloaded quality factor of  $Q_0 \approx 45 \times 10^3$  at cryogenic temperatures. Phase IIa used separate stems inserted from both the top and bottom of the cavity, while Phase IIb used a single stem from just the bottom. Both configurations achieved similar performance. While the cavity's effective noise temperature still appears elevated above the base temperature of 61 mK, the exact cause is unknown, with recent studies hinting at possible rf leakage through one of the cavity ports. Eliminating the remaining excess cavity noise is an ongoing effort by the HAYSTAC Collaboration.

Located on the cavity's top end cap are two antenna ports. The first is a weakly coupled port used for inputting microwave tones for calibration of the detector response, while the second is a strongly coupled test port used to read signals out of the cavity. The depth of the strongly coupled antenna is adjusted with a mechanically actuated stepper motor [41] which is attached to the antenna via a Kevlar line. This allows for the antenna's coupling coefficient to be maintained near the optimal value of  $\beta \sim 7.1$  needed for the axion search to maximally benefit from the SSR [24]. This is slightly larger than the estimate in Ref. [27] due to excess cavity noise, as found during the noise calibrations described in Appendix A 2.

To monitor the cavity's field, the two ports described above are coupled to the external receiver chain, as shown in Fig. 3. This chain allows for power emitted by the cavity to be read out as well as for the injection of signals needed to calibrate the detector response. During operation the main input to the receiver chain comes from the thermal noise of a  $50 \Omega$  load which can be toggled by a switch between a cold load held at the base temperature of the mixing chamber and a hot load held at an elevated temperature by a variable

temperature stage (VTS) operating at temperatures between 290 mK and 1 K. The temperature of the VTS is measured by an MFFT-1 Magnicon temperature sensor [42]. During Phase IIa (Phase IIb) this temperature was maintained at 333 mK (290 mK) for use in the calibration measurements described in Appendix A 2. The VTS was incorporated prior to the start of Phase II, allowing a finer temperature control of the hot load, previously limited to the temperature of the still plate,  $\sim 775$  mK.

Signals in the receiver chain are routed using 0.085" coaxial cables with the outer conductors of the cables thermalized at each temperature stage with gold-plated copper clamps which also serve to block stray light coming from higher temperature stages. The cable segments above the 4 K stage use stainless steel cables, while segments below this use superconducting NbTi/NbTi cables. The extremely small losses for the superconducting coax are experimentally shown to be independent of applied fields up to 7 T, due to shielding of the inner surfaces from the applied field by the superconducting jacket. This low-loss behavior can be expected up to  $H_{c2} = 14$  T. In addition, dc electrical signals, used to tune the JPAs and power the HEMT, are sent into the fridge through API 56F715-005 EMI filters mounted on DB15 feedthroughs. Inside the fridge these dc signals are routed to the still plate using 28 AWG phosphor-bronze wire before transitioning to superconducting wire.

Upon exiting the fridge, the output signal from the cavity is split into two paths, with one path leading directly to the return port of an Agilent E5071C vector network analyzer (VNA) [43], while the other leads to a PC with a 14-bit GaGe Oscar CSE4344 PCIe digitizer [44] used to sample the cavity field. Prior to reaching the digitizer the signal is sent to an IQ mixer (Marki Microwave IQ0307LXP [45]) where it is split into the I/Q quadratures and down-converted into an intermediate frequency (IF) band centered at the cavity's resonant frequency or, equivalently, half the JPA pump frequency. To precisely match these frequencies, the same signal generator is used as both the pump for the JPAs and the local oscillator (LO) for the mixer. A frequency divider before the mixer is used to convert the tone down to the cavity frequency. Finally, before reaching the digitizer the signals from both quadratures pass through low-pass filters which cut off the IF window at 1.9 MHz in order to remove high frequency noise outside the cavity's usable bandwidth. All connections between these IF components are made with double-shielded BNC cables, and the signal generators, VNA, and digitizer are all clocked to the same 10 MHz rubidium reference source (SRS FS725m [46]).

The cavity's performance is characterized with a VNA, which injects stimuli into either of the two ports described above. Using a pair of VNA measurements, the cavity's  $\text{TM}_{010}$  mode frequency  $\nu_c$ , coupling coefficient  $\beta$ , and

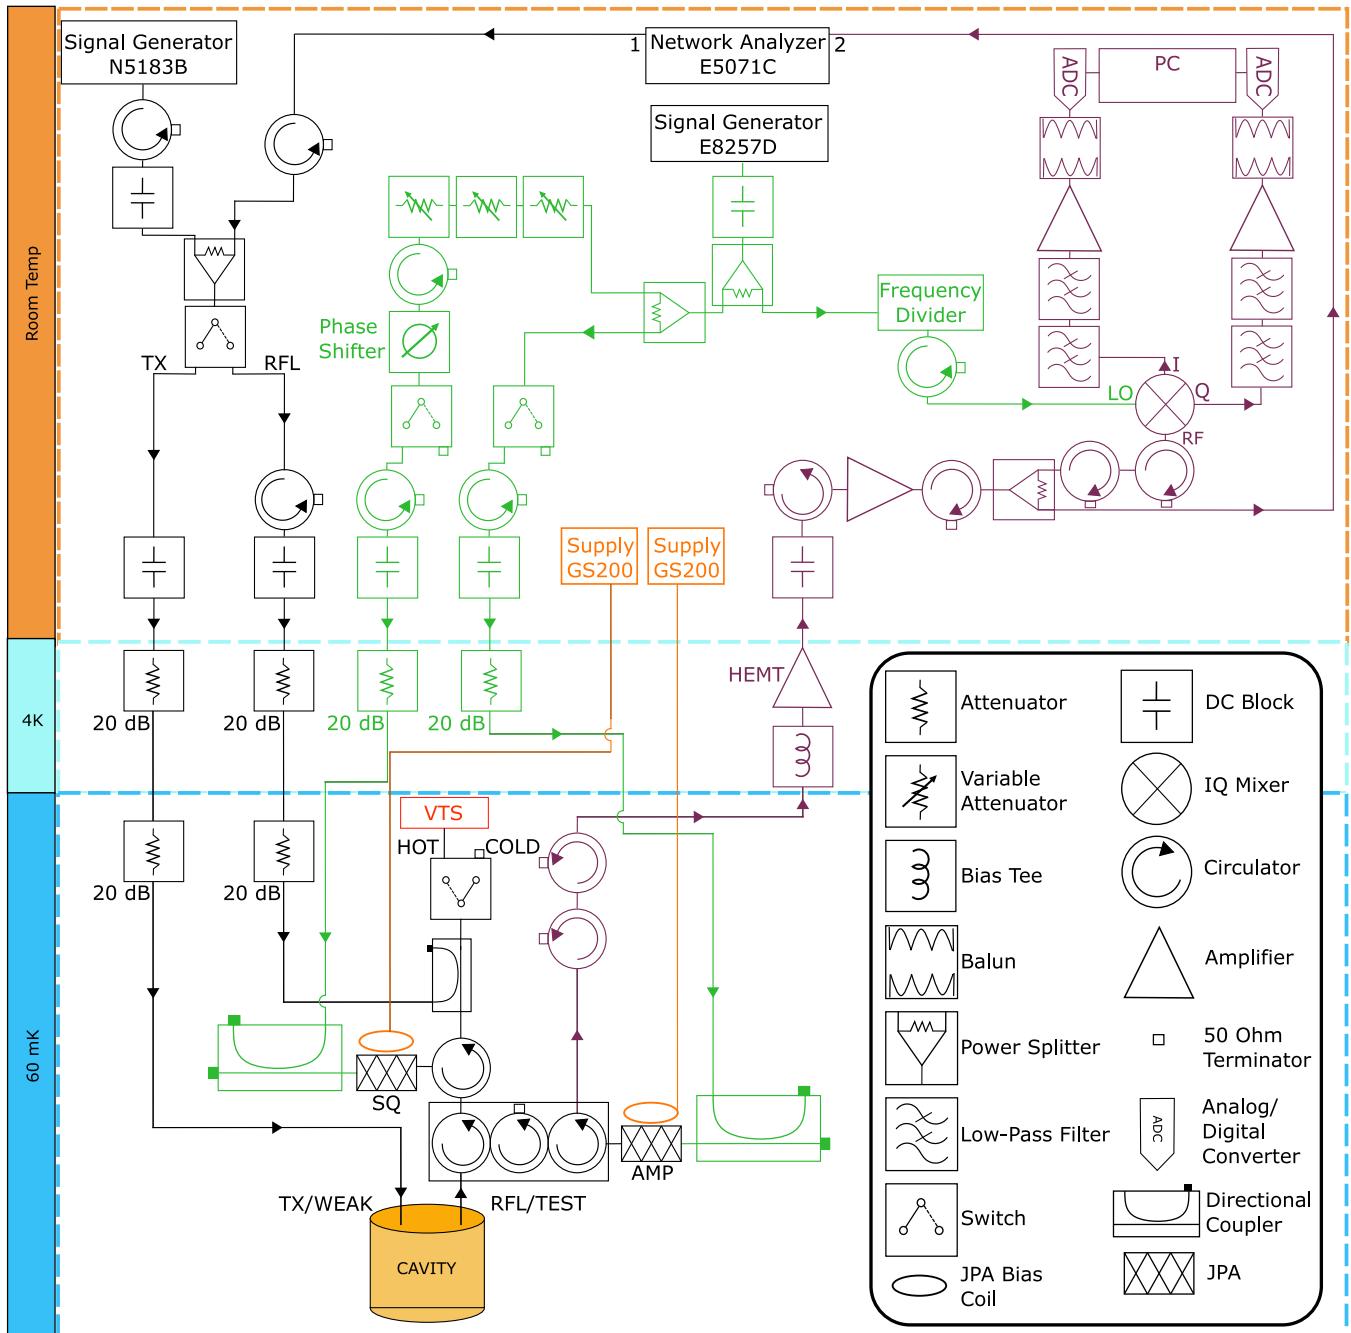


FIG. 3. Schematic of HAYSTAC's electronics including both the IF and rf components. The electronics are divided by their temperature stage, with the top stage held at room temperature, the middle at 4 K, and the bottom at 61 mK. Input lines into the system (black) are used to send signals for calibration and characterization of the system, while output lines (purple) are used to transfer signals from the output of the cavity to the DAQ to be recorded. The lines providing the pump tone to the JPAs and mixer (green) and the lines for the JPA bias currents (orange) are also shown.

loaded quality factor  $Q_L = Q_0/(1 + \beta)$  can be extracted. In both cases the observed VNA response is fit to the expected functional form of the power

$$P(\nu) = \frac{P_0}{(2Q_L(\nu/\nu_c - 1))^2 + 1} + (A\nu + B), \quad (7)$$

where the first term is a Lorentzian with peak amplitude  $P_0$  and the second term in parentheses represents a linear baseline. The first is a transmission measurement (“tx”), taken by sending a stimulus to the weak port and reading it out through the test port via the cavity; it is used to extract both  $Q_L$  and  $\nu_c$ . Following this, a reflection measurement (“rfl”) is taken by routing the stimulus signal to the strongly

coupled port where it is reflected off of the cavity and back to the VNA. This is used to extract  $\beta$  as

$$\beta = \frac{1 + \sqrt{P(\nu_c)/(A\nu_c + B)}}{1 - \sqrt{P(\nu_c)/(A\nu_c + B)}}, \quad (8)$$

when  $\beta > 1$  [47]. Example measurements taken during Phase IIb, along with their corresponding fits, are shown in Fig. 4, and the resulting values of  $Q_L$  and  $\beta$  at each frequency scanned in Phase II are shown in Fig. 5. Their values vary as the cavity frequency changes, with the largest variations coming from discontinuous jumps corresponding to adjustments of the antenna position to maintain the coupling coefficient near the value which optimizes the scan rate. The average values over Phase IIa (Phase IIb) are  $Q_L = 6180(5915)$  and  $\beta = 6.9(6.5)$ .

During Phase IIa operation the  $TM_{010}$  mode exhibited residual drift, taking  $\sim 10$  min to settle to its new resonance each time the rod position was adjusted. To avoid degrading the detector's sensitivity, a 15 min delay was imposed after each tuning step during which no data were taken, slightly decreasing the detector live time. This effect worsened over time such that during the acquisition of rescan data, the tuning between each target candidate became inconsistent and typically required manual intervention. While neither of these issues had a major impact on the results of Phase IIa, the cavity was reassembled and the tuning rod realigned to increase its stability prior to the start of Phase IIb. As a result, the cavity tuning was more reliable in Phase IIb, and the observed mode drift decreased substantially. This allowed the wait time to be reduced from 15 min to 5 min, chosen

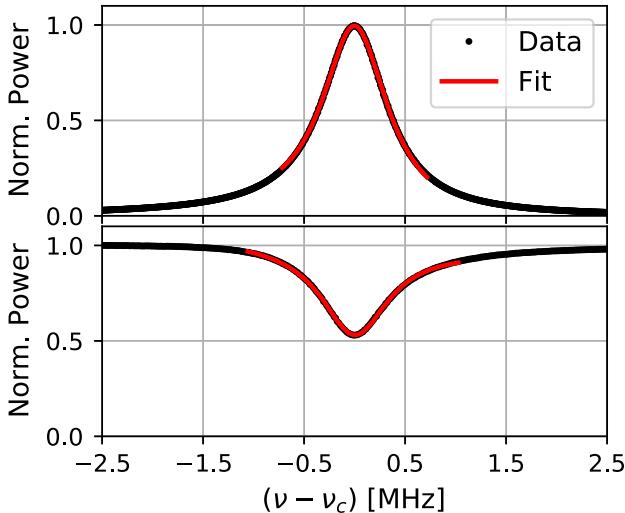


FIG. 4. Example transmission (top) and reflection (bottom) measurements (black) taken with the VNA in Phase IIb along with the fit (red) to Eq. (7) used to extract the relevant cavity parameters. The profile for both is normalized, with the transmission profile normalized to have unit peak power and the reflection normalized to unit power off resonance.

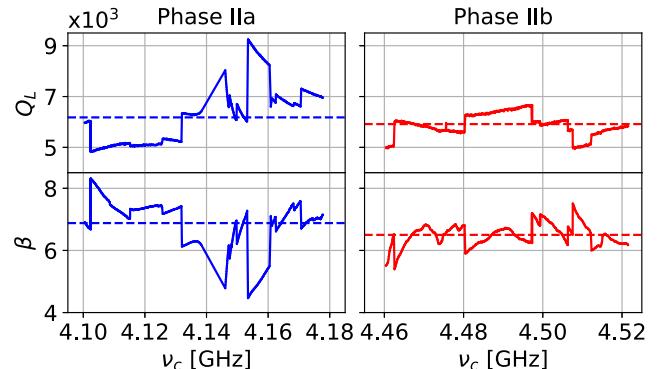


FIG. 5. Measured values of  $Q_L$  and  $\beta$  as a function of  $\nu_c$  for every run taken during the original scans in both Phase IIa (left) and Phase IIb (right) with the average value over each phase shown as a dashed line. The discontinuous jumps in both values correspond to adjustments of the antenna position to maintain  $\beta$  near the optimal value.

conservatively to ensure that heat produced from tuning was dissipated.

### C. Design and performance of the squeezed state receiver

The SSR used in this work follows a similar design to that used in [27] and is composed of two JPAs operated in the phase-sensitive mode which are coupled to the main cavity via the test port, as shown in Fig. 3. Each JPA acts as a nonlinear  $LC$  resonator, where the nonlinear inductance is achieved by an array of superconducting quantum interference devices (SQUIDs). The JPAs used in this work are flux-pumped JPAs which use an on-chip bias line to modulate the magnetic flux threading the SQUID loop at twice the resonance frequency [28]. The resonance of each JPA is then set by a static flux provided by a pair of superconducting bias coils on which each JPA is mounted. The flux through each bias coil is set separately by supplying a current from one of two identical current sources (Yokogawa GS200 [48]), allowing the resonance frequency of each JPA to be adjusted semi-independently. To determine the operating range for each JPA, a flux tuning curve that correlates the JPA's resonant frequency with the coil current is created by sending stimulus signals from the VNA to the JPA. By measuring the phase response of the transmitted signal, a tuning curve is extracted by finding the point at which the phase changes by  $\pi$ . Example tuning curves for the two JPAs used in Phase IIb are shown in Fig. 6. Comparing the two curves, a common region of gain is observed between 4.2 and 4.8 GHz. While an ideal JPA would be periodic in its response, regions of nonideal behavior are observed in Fig. 6. Although the exact cause of these regions is unknown, they are likely the result of flux trapped in the chip's superconducting circuitry or by other sources of nonuniformity in the magnetic flux penetrating each JPA. To achieve stable gain, these regions are avoided

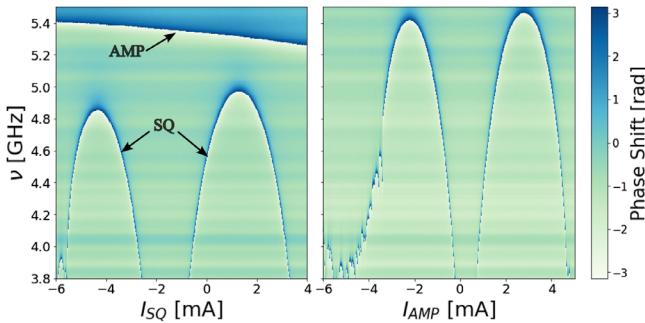


FIG. 6. Example flux tuning curves of the SQ JPA (left) and the AMP JPA (right) produced from measuring the phase response as a function of frequency  $\nu$  with the VNA while varying the current  $I_{\text{JPA}}$  in one of the JPA's bias coils. The resonance of the JPA at a given  $I_{\text{JPA}}$  is the frequency at which the phase undergoes a  $\pi$  phase shift. In the tuning curve for the SQ, the AMP's curve is also present at  $\sim 5.3$  GHz but stays roughly constant due to reverse biasing the AMP's current. The observed residual variation is due to an initially imperfect compensation and was corrected before axion data were taken. The SQ curve is also present in the AMP's curve, but it is below the lowest frequency plotted here. Additionally, the nonsmooth regions in the curves are likely due to nonuniformity of the magnetic flux penetrating each JPA.

when operating the JPAs. For the JPAs used in Phase IIa this operating range only extended to  $\sim 4.2$  GHz by design; to extend the search to higher frequencies, a new set of JPAs were obtained and installed prior to the start of Phase IIb.

Once the flux bias needed to obtain the desired resonant frequency was found, amplification was achieved by pumping each JPA at twice this frequency with a single low noise signal generator (PSG E8257D [49]) shared by both JPAs as shown in Fig. 3. The absolute amplification gain is controlled by the power of the pump tone sent to each JPA. An automated procedure, which scans over a range of coil currents and pump powers to find the optimal parameters for the desired gain, is used to individually bias each JPA during operation of the experiment.

Complications to this approach not present in the test setup from Ref. [27], which used an injected signal and did not require a magnetic field, arise from the proximity of the JPAs to the large magnetic field required to convert axions into a detectable signal. To minimize the impact, the JPAs are held  $\sim 60$  cm above the cavity as shown in Fig. 2. Here the magnetic field is reduced to  $\sim 50$  G by a pair of NbTi field compensation coils, oppositely oriented and in series with the magnet. Further shielding from the remaining stray field is achieved in a multistaged shielding configuration. First, a set of three persistent superconducting NbTi coils are mounted around the amplifier canister to counter changes to the magnetic flux. Inside these coils, the JPAs are housed in a four-layered shielding can. From inside out, this consists of a layer of Amumetal 4 K, aluminum, Amumetal 4 K, and niobium. In this way the Amumetal provides shielding as the JPAs are cooled to cryogenic temperatures, at which

point the aluminum and niobium layers become superconducting and provide additional shielding. Together, this reduced the stray flux to  $\ll 1$  flux quantum over the area of the JPAs' SQUID loops, negligible to the operation of the JPAs as confirmed by measurements of the tuning curves with and without the field which show no measurable change. Due to limitations of space inside the fridge, the two JPAs are housed inside the same shielding can and are located  $\sim 10$  cm apart with no shielding between them. As a result of their close proximity, changing the current of one coil affects the flux in both JPAs, and biasing either JPA independent of the other is not possible. To decouple them, the residual flux induced in the secondary JPA (the JPA that is not actively being tuned) is empirically measured as a fraction  $\epsilon$  of the change of the primary JPA (the JPA that is actively being tuned) current. This allows for the secondary JPA's frequency to be held constant by applying a compensating current  $\epsilon\Delta I$  each time the primary JPA's current is changed. This phenomenon is visible in Fig. 6 (left), where the AMP's curve at  $\sim 5.3$  GHz is slowly tuned as the bias for the SQ is changed due to an initially imperfect compensation.

As described in Sec. II A, vacuum fluctuations in the experiment are squeezed by operating the two JPAs at the same resonant frequency, with a  $\pi/2$  phase shift between the two pump tones such that their amplified and squeezed quadratures are perpendicular. This is achieved by using a single signal generator to generate the tone used by both JPAs, with a variable phase shifter located before the squeezer used to apply a phase shift proportional to the voltage  $V_\theta$ . In addition, a variable attenuator is added to the tone sent to the SQ JPA to allow for its gain to be set independently from the AMP JPA by varying the applied voltage  $V_A$ . This is required as the optimal SQ gain is typically  $\sim 15$  dB lower than that of the AMP to avoid potential saturation [27,33]. As a result, operating the SSR requires finding the pair of  $V_A$  and  $V_\theta$  for which the squeezing is optimal. This is done by scanning over a range of  $V_A$  and  $V_\theta$  and collecting short (10 ms) time traces with data acquisition software, first with the SQ on and then with it off. Squeezing is then quantified by the ratio of the variances of the two time traces ( $S = \sigma_{\text{on}}^2/\sigma_{\text{off}}^2$ ). This results in a map of squeezing  $S$  over a range of  $V_A$  and  $V_\theta$  as shown in Fig. 7. This allows for the set of parameters which result in the optimal squeezing being found. The example shown here is from a coarse scan used to approximately find the location of the optimal parameters and to initialize the finer scan performed during acquisition of axion data. This finer scan is done in a small region of the full parameter space using an automated tuning algorithm that searches in the vicinity of the initial guess in decreasing voltage step sizes in order to narrow down both  $V_A$  and  $V_\theta$  until the optimal value of  $S$  is reached. Because the tuning step size of the experiment ( $\sim 100$  kHz) is small relative to the change in these variables with frequency, the most

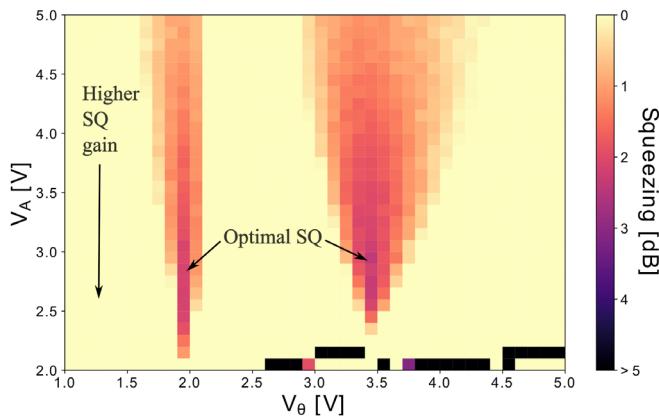


FIG. 7. Example plot of delivered squeezing mapped over both the voltage of the variable phase shifter  $V_\theta$  and the variable attenuator  $V_A$ . This mapping allows for the identification of the  $V_\theta$  and  $V_A$  which maximize squeezing, where the optimal  $V_\theta$  corresponds to a  $\pi/2 + n\pi$  phase difference between the SQ and AMP, and the optimal  $V_A$  corresponds to the largest SQ gain before saturation. Additionally, beyond a certain gain this fully saturates the chain, causing regions of fake squeezing seen here as the disconnected points at  $V_A < 2.1$  V.

recent set of optimal parameters are used to initialize the algorithm each time the cavity is tuned, allowing for efficient tuning of the SSR.

Following this optimization, the gain along the amplified quadrature as well as its frequency dependence are measured. In order to measure the gain in just one quadrature (1Q), a calibration tone is generated by a signal generator (N5183B MXG [50]) and injected into the system through the “rfl” path at a detuning of 200 kHz relative to the JPA’s resonance. The ratio of the observed power with and without the JPA of interest being pumped, as measured by the digitizer, gives a measurement of the gain along the amplified quadrature at this one detuning. Because the 1Q gain measurement is time-consuming, the full frequency dependence of the gain at all detunings is mapped out by using the VNA, which sweeps a stimulus signal through many frequencies along the same path. This results in a measurement of the two-quadrature (2Q) gain profile, an example of which is shown from Phase IIb for both the AMP and SQ in Fig. 8. In Phase II the target 2Q gains were  $\sim 22$  dB for the AMP and  $\sim 5$  dB for the squeezer. The AMP gain is chosen such that it is large enough to overwhelm the added noise of the high-electron-mobility transistor (HEMT) which follows the AMP in the receiver chain. The target gain of the SQ JPA is kept low enough to prevent saturation of the AMP JPA.

For the SSR used in HAYSTAC the delivered squeezing off resonant from the cavity is measured prior to the start of each run to be 3.7 dB in Phase IIb and 4 dB in Phase IIa. This level of squeezing is consistent with the expectation in Eq. (6) given the measured transmission efficiency  $\eta$  of 0.60 in Phase IIb and 0.63 in Phase IIa. This efficiency is largely

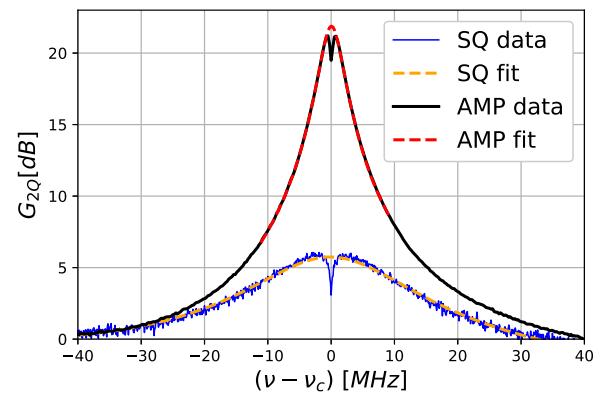


FIG. 8. Example 2Q gain profiles of the AMP and the SQ as measured by the VNA (solid lines) and their fits to a Lorentzian (dashed lines). Because the data are taken through rfl and the JPA is centered on the cavity’s resonant frequency, the cavity’s reflection profile appears as a dip at 0 MHz. When fitting the profile the  $\sim 2$  MHz region where this effect is most prominent is ignored so as not to distort the fit.

dominated by losses in the custom triple-junction circulator (QuinStar QCE-070100CM30 [51]) located between the SQ and AMP JPA in Fig. 3 which is used to isolate the JPAs from unwanted noise reflecting backwards through the system. The difference between the transmission in the two phases is likely due to the frequency dependence of the circulator’s loss. Assuming this level of squeezing is maintained over the full frequency range, the scan rate is expected to increase by a factor of 1.9 in Phase IIb and 2.0 in Phase IIa relative to an equivalent measurement taken in the same configuration but without squeezing. To precisely calibrate the noise of the system, the performance of the SSR and its delivered squeezing are monitored periodically during data taking as described in Appendix A 2.

#### D. Data acquisition

Once the signal reaches the digitizer, the voltages of both IF channels are sampled at a rate  $f_s$ , and the sampled data are transferred to the RAM of the data acquisition (DAQ) PC. The rate at which the data are sampled is chosen such that the Nyquist frequency is comfortably above both the 1.9 MHz cutoff of the IF filters and the cavity’s sensitive bandwidth, roughly given by  $2\Delta\nu_c \sim 1400$  kHz. In Phase IIa a sampling rate of 10 MHz is used but is lowered to 5 MHz in Phase IIb in an effort to reduce the total amount of data that is transferred to RAM.

Before saving the data to disk, they are first processed *in situ*. The reason for this is that at these sampling rates the raw data from a single 1 hour observation would require roughly  $\sim 60$  GB of disk space. Rather than save all of these data to disk, the data are processed such that only a fraction of the original data is stored. The processing begins by sampling the voltage fluctuations  $V_{I,Q}(t)$  of both quadratures’ output by the mixer in 5 s segments. Because these

quadratures are coming from the mixer, they are not necessarily aligned to the amplified and squeezed quadratures of the SSR and must be rotated by an angle  $\theta_{\text{AS}}$  by the linear combinations

$$V_{\text{AMP}}(t) = V_I(t) \sin \theta_{\text{AS}} + V_Q(t) \cos \theta_{\text{AS}} \quad (9)$$

for the amplified quadrature and

$$V_{\text{SQZ}}(t) = V_I(t) \cos \theta_{\text{AS}} - V_Q(t) \sin \theta_{\text{AS}} \quad (10)$$

for the squeezed quadrature. This angle is empirically determined at each tuning step as the angle which maximizes the variance of  $V_{\text{AMP}}$ , indicating alignment with the amplified quadrature.

The resulting voltage fluctuations of the amplified and squeezed quadratures are further divided into subsegments consisting of  $N_s$  samples, and the power spectrum of each is computed from the FFT. For Phase IIa a subsegment size of  $10^5$  is chosen, but this is reduced to  $2^{15}$  in Phase IIb. In both cases this results in a frequency resolution ( $\Delta\nu_b = f_s/N_s$ ) which is smaller than the  $\sim 5$  kHz expected axion linewidth  $\Delta\nu_a$ , but the reduction in Phase IIb speeds up the FFT computation. The power spectrum of each subsegment is then combined to produce an average power spectrum for the entire 5 s period. These steps are repeated until the total observation time reaches the desired level, chosen to be 1 h for both Phase IIa and Phase IIb, with the average power spectrum from each 5 s period being further combined to produce a single power spectrum representing the average over the full observation time.

As a result of this processing, the data storage requirement is reduced by a factor of  $\sim 10^5$  but requires usage of PC resources also needed by the digitizer. This introduces dead time during which voltage fluctuations from the cavity are not being recorded, and it reduces the effective scan rate of the experiment. The dead time from the *in situ* processing, defined as the fraction of the total observation time spent not recording voltage fluctuations from the cavity, is 40% in Phase IIa. In addition, time spent tuning the cavity and biasing the JPAs results in an additional 12% dead time in Phase IIa, resulting in a total dead time of 52%.

Several improvements were implemented for Phase IIb, resulting in a reduction of the total dead time by a factor of  $\sim 1.6$ . Two improvements were made to speed up the data acquisition. The first was to parallelize the data transfer with the use of the two separate buffers on the digitizer card. By swapping between these buffers, data can be transferred to disk while they are simultaneously being sampled, resulting in a reduction in the dead time. The second improvement was to decrease the time needed to compute the power spectrum by parallelizing the FFT algorithm using an NVIDIA Titan V GPU [52] while also reducing the sampling frequency as described above. While this reduced both the frequency resolution and the width of the FFT,

both still comfortably fit the requirements for an axion search previously described. In total, the dead time associated with the data acquisition was reduced by a factor of  $\sim 10$ . In addition to improvements to the *in situ* processing, the reassembly of the cavity prior to the start of Phase IIb described in Sec. II B resulted in a significant reduction to the mode drift of the cavity. As a result, the time needed for the cavity to settle after each tuning step was reduced by a factor of 3. Altogether, the improvements to the tuning and DAQ allowed for the total dead time to be reduced to 18% resulting in an effective scan rate enhancement of  $\sim 1.6$  in Phase IIb.

## E. Data processing and analysis

Following acquisition of the average power spectrum at each tuning step, data are processed and analyzed in order to search for potential candidates, identified as power excesses in the spectrum above a chosen threshold. We give a brief outline of this procedure. A more detailed procedure is found in [53], with the main difference arising from the homodyne configuration which centers the cavity around the JPA's resonant frequency.

The first step of this procedure is to combine the average power spectrum observed at each tuning step. While each of these spectra individually has sensitivity to axion candidates, the benefit of having a larger bandwidth comes from the overlap between subsequent spectra, which allows for the sensitivity at a given frequency to continue increasing even as it is detuned from the cavity. In order to take advantage of this benefit, the power observed at each tuning step must be combined to form a single power spectrum which optimizes the SNR at each frequency.

Prior to their addition, each spectrum must be filtered to remove spectral features that are not indicative of the axion's presence. This procedure starts by removing such structure shared between all power spectra in the IF. To facilitate this, spectra from different tuning steps are first aligned along the IF band and summed to produce an average IF response. Because the IF band is, by definition, centered around the resonant frequency of the cavity and JPA, this response largely captures the spectral features of the cavity noise and the JPA gain. Signals in the rf, like those resulting from an axion, would be washed out as they would move in the IF band relative to the mixing frequency and effectively average away. By applying a Savitzky-Golay (SG) filter [54] (width = 50 kHz, order = 10) to the IF response, the large scale spectral shape can be extracted and is subsequently removed from each spectrum by dividing out the filter's output. Additionally, contaminated IF frequencies can be identified by comparing the average response before and after the filter. Because noise from rf signals would average away in this procedure, any spikes in the average response which deviate from the filtered average must be due to IF noise, which would add coherently. In addition, each spectrum is truncated into the analysis band defined by a

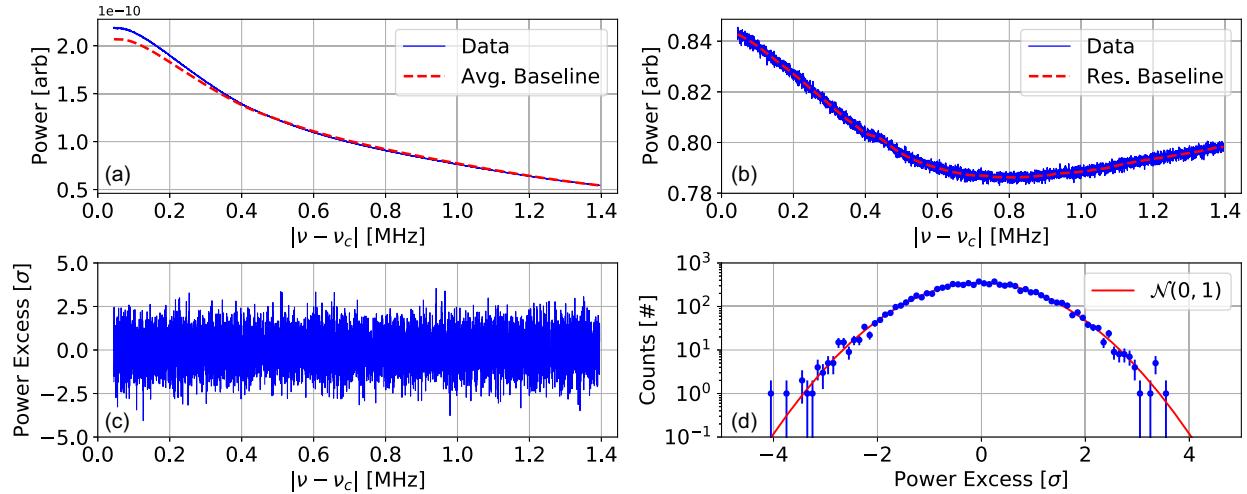


FIG. 9. Example steps of the analysis procedure described in the text for a single spectrum taken during Phase IIb. (a) Raw data (solid blue) for a single tuning step compared to the average baseline (dashed red) calculated from the full data set. The shape of both follow from the cavity's transmission profile, with differences between the two arising from variations of experimental properties between each tuning step. (b) Same spectrum after dividing out the average baseline (solid blue) and removing IF contaminated bins, along with the calculated residual baseline (dashed red). (c) Same spectrum after removing the residual baseline and normalized such that each bin is approximately sampled from a normal distribution with  $\mu = 0$  and  $\sigma = 1$ . (d) Distribution of the observed power in each frequency bin (blue points) from (c), showing good agreement with  $\mathcal{N}(0, 1)$  (solid red).

lower bound at 45 kHz to remove low frequency 1/f noise and an upper bound at 1345 kHz  $\sim 2\Delta\nu_c$  which covers the majority of the cavity's measurement bandwidth while staying below both the Nyquist and low pass cutoff frequencies.

While this effectively removes a majority of the non-axionic spectral features, variations of the frequency-dependent profiles of both the JPA and cavity result in residual structure unique to each spectrum remaining, as seen in Fig. 9. Because both the cavity and JPA have bandwidths which are  $\gg \Delta\nu_a$ , this residual structure can be effectively removed with a minimal effect to the axion signal by applying a lower-order SG filter (width = 50 kHz, order = 4) to each spectrum and dividing out the filtered response from the original spectrum. With these parameters the filter has a 3 dB point around  $\sim 60$  kHz, which is far enough above the  $\sim 5$  kHz axion linewidth to avoid significant attenuation of its signal. The resulting attenuation to an axion signal is found to be 0.91 using a simulation similar to that described in Ref. [53]. Following this, the resulting spectra consist of Gaussian-distributed noise with a mean of 0, as shown in Fig. 9.

The final step before summing the spectra from each tuning step is to weight the contribution of each IF bin of individual spectra by their observed SNR, so as to produce a sum spectrum which optimizes the total SNR. These weights are calculated separately for each spectrum, with parameters extracted from the noise and signal characterizations described in the Appendix. Following this scaling, the spectra are aligned in the rf and summed to produce a single combined spectrum. Because the data come from a

homodyne measurement in which the JPA's pump tone is also centered on the cavity's resonance, each frequency bin in the output spectrum contains contributions from both positive and negative detunings from cavity resonance. The double-sided profile, shown in Fig. 10, is recovered by mirroring each spectrum and its associated SNR around

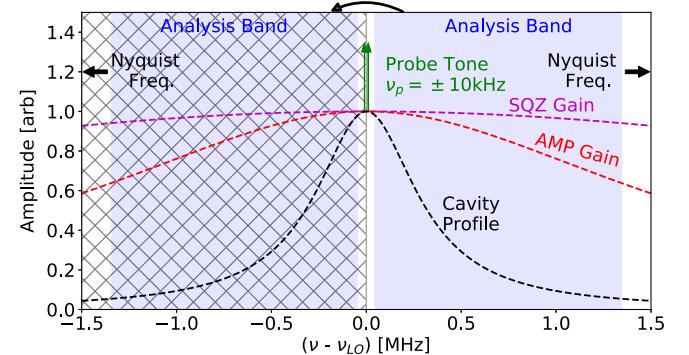


FIG. 10. Diagram of the relevant frequencies for HAYSTAC's data taking and analysis, showing the relative shape of the cavity (black), AMP (red), and SQ (purple) profiles as a function of detuning and normalized to the same peak height. The measurement is taken in a homodyne configuration, with all the profiles centered at 0 MHz. As a result, each bin contains power from both positive and negative detunings (Appendix A 2), and the double-sided profile is recovered by mirroring (gray cross-hatches) the spectrum around 0 MHz. Also shown is the analysis band (blue shaded region) and the probe tone (green) at 10 kHz used to monitor gain fluctuations over the course of each measurement.

0 kHz. This allows for observations from both the positive and negative detuned frequencies to contribute to the combined spectrum with the appropriate weighting.

In the absence of any rf tones, this procedure results in a final combined spectrum which is composed of normally distributed noise fluctuations, with any axionlike signal appearing as a power excess with shape given by the expected virialized axion line shape in the terrestrial lab frame [55]. Because the resulting spectrum still has frequency resolution given by the DAQ and is intentionally set much smaller than the axion linewidth, an axion-induced signal would appear divided across multiple bins. To account for this, the spectral densities from adjacent bins are combined by convolving the combined spectrum with the expected axion line shape. The grindspectrum, which is optimally sensitive to an axion, is produced by normalizing each bin of the previously generated spectrum by its total SNR such that the power fluctuation in each bin is standard Gaussian distributed with  $\mu = 0$  and  $\sigma = 1$ .

Frequencies in the resulting spectrum at which the observed power exceeds a chosen threshold are then recorded. The threshold chosen for this analysis was  $3.468\sigma$ , which, for a target significance of  $5.1\sigma$ , would correspond to a 10% two-scan false negative rate in a frequentist analysis. While the frequentist framework is no longer used, the threshold is chosen to match previous analyses for convenience. These candidates can be the result of real rf noise sources as well as statistical fluctuations in the noise. For each such candidate an additional observation or rescan is made to determine if the excess persists, and the resulting data are processed and analyzed following the same procedure as for the initial data but with a higher order SG filter (width = 30 kHz, order = 6) to account for the lack of tuning, as described in Ref. [53], which gives a slightly larger attenuation of 0.76 to the expected axion signal. Any candidate which persists above threshold in the rescan can be further interrogated with the magnetic field at 0 T. Because axions require a magnetic field to convert into a detectable signal, this would result in any real axion-induced signal vanishing. Any signal which passes this final test would be manually studied in follow-up scans to determine if the signal has the necessary characteristics of an axion, such as the spectral shape, scaling with field, and scaling with cavity profile. Any signal which persists at zero field would be the result of rf noise from a different source. While this noise could be from a known rf signal, it could also be connected to new physics such as dark photons [56,57]. Determination of its origin would require the manual study of the behavior relative to that predicted by the respective model.

### III. AXION SEARCH RESULTS

Data taken during Phase IIa were previously analyzed to search for the presence of axion dark matter, and the results were published in Ref. [24] and are shown in Fig. 11. No

signal was observed in the mass ranges 16.96–17.12 and 17.14–17.28  $\mu\text{eV}/\text{c}^2$  resulting in an aggregate exclusion over this region of  $g_\gamma \geq 1.95 \times g_\gamma^{\text{KSVZ}}$  at the 90% level. This result differs by a factor of  $\sqrt{2}$  from that originally published in [24]. The correction was needed to properly handle the noise in the IF, which previously incorrectly assumed a 1/4 quantum of noise as the quantum-limited value of the input referred noise in the IF spectrum rather than the correct value of 1/2. This resulted in the noise being underestimated by exactly a factor of 2. The analysis has since been corrected to properly treat the IF noise, and the results presented here for both Phase IIa and Phase IIb reflect the correct analysis. Below are details and results from the analysis of data taken during Phase IIb.

Initial data were taken over a ~53 day period between July 1 and August 20, 2021 covering a 64 MHz region in the range 4.459–4.523 GHz and consisting of 801 total spectra spaced ~80 kHz apart. Using the processing routine described in Sec. II E, these spectra were optimally summed to produce a single grindspectrum showing the observed power at each probed frequency, and axion candidates were identified with a  $3.468\sigma$  threshold. Using two independent analyses of the data, 37 potential axion candidates were identified, with 19 of the candidates being common to both. Each of the remaining candidates was found to be near threshold in both, with small differences in the analyses causing the discrepancy. While many of these candidates are likely the result of statistical fluctuations in the noise, both parts of Phase II data taking also observed a set of nonaxionic rf noise candidates which presented as large as ~40 kHz wide power deficits in the grindspectrum. This unique structure made these candidates easy to identify, resulting in five candidates in Phase IIa and one in Phase IIb. Because these regions have poor sensitivity to an axion, a rf cut was applied to remove data in the 200 kHz region around each candidate in the initial scans. The source of this noise is currently unknown, but so far it has been observed to vanish during the rescans, indicating this noise is either time dependent or related to the different configuration used during rescans. This allowed for the resulting gaps in data to be filled with additional observations at these frequencies taken during rescans.

Rescans of each of the observed power excesses were performed between September 20 and November 18, 2021. Because these data are taken with the intention of probing individual frequencies, rather than multiple frequencies simultaneously, the optimal search configuration is slightly different than for the initial scans. Instead, the cavity is operated near critical coupling ( $\beta = 1$ ) in order to maximize the peak sensitivity at the frequency of interest with the squeezer turned off. Analysis of the rescan spectrum found that no excess persisted above the target threshold, ruling out all candidates as potential axion signals.

While no candidate was observed to persist, an exclusion on  $g_\gamma$  over the mass range scanned can be set by applying the

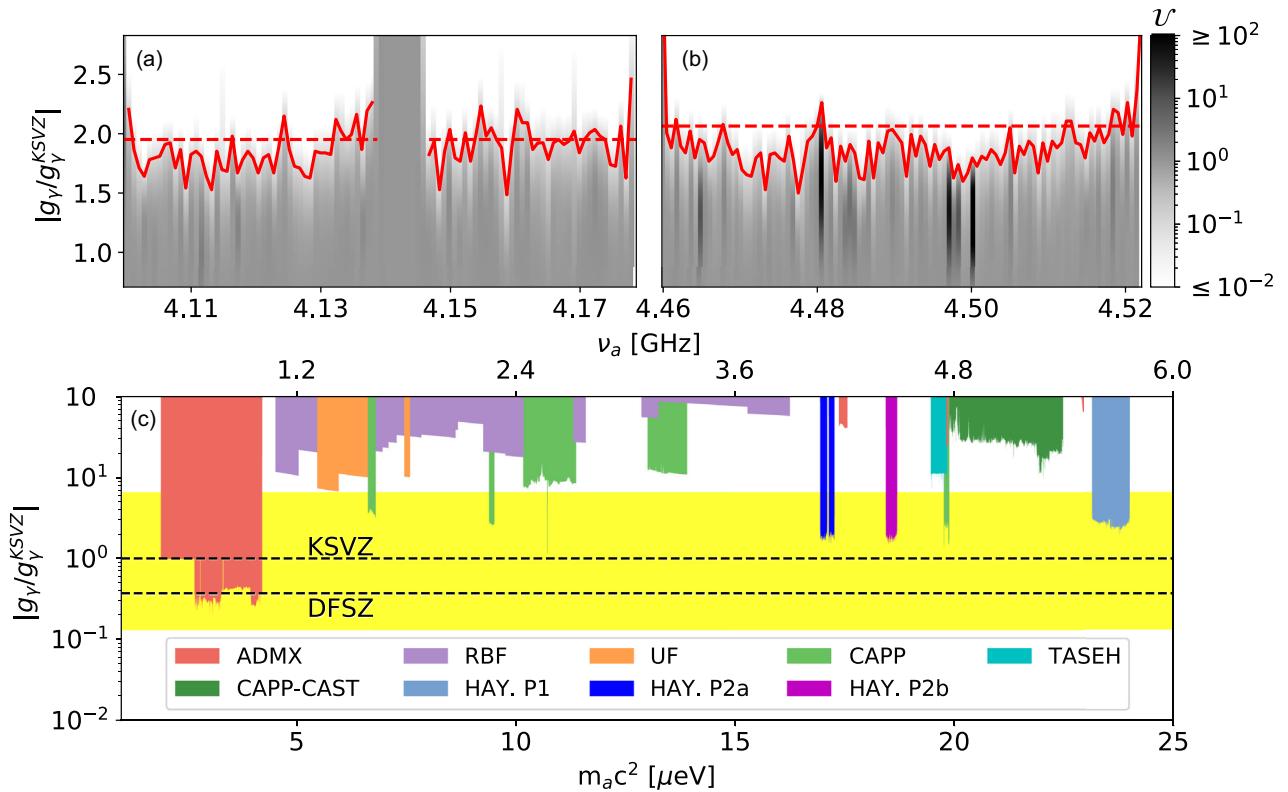


FIG. 11. Axion exclusions showing the two-dimensional prior update  $\mathcal{U}$  in grayscale (top panel) for (a) Phase IIa [24] and (b) Phase IIb. This includes the 10% prior update contour (solid red) as well as the 90% aggregate exclusion (dashed red). (c) Results of this work shown alongside previous exclusion results from HAYSTAC Phase I [31,32] as well as from other haloscopes RBF [58], UF [59], ADMX [19,60–62], CAPP [18,63–66], CAST-CAPP [67], and TASEH [68]. The QCD axion model band representing the most natural KSVZ and DFSZ models is shown in yellow [69], with the benchmark KSVZ [10,11] and DFSZ [12,13] model lines shown as black dashed lines.

Bayesian framework outlined in Ref. [70]. This framework is a straightforward application of Bayes' theorem and gives the change in probability of the axion existing at each frequency which was probed. These prior updates ( $\mathcal{U}$ ) are determined by both the experimental sensitivity and actual observed power at each frequency, allowing for a “maximally informative” result. By combining the prior update from both the initial data and the rescan data, a 90% aggregate exclusion of  $g_\gamma \geq 2.06 \times g_\gamma^{KSVZ}$  is found over the entire mass range  $18.44$ – $18.71 \mu\text{eV}/c^2$ . This result is summarized in Fig. 11, which shows both the aggregate exclusion and the full two-dimensional space of  $\mathcal{U}$  with the 10% prior update contour in each frequency bin. While no candidate in Phase IIb was observed to persist above threshold during rescans, a handful of candidates are shown to have large ( $\geq 100$ ) prior updates in the Bayesian framework. Each of these points in the prior update is the result of large rf noise spikes during the initial scan which are found to fall below threshold during the rescan and thus are rejected as an axion candidate. Due to the nature of the Bayesian analysis, the prior update calculated from the observed excess in both scans can result in a large combined prior update at couplings lower than HAYSTAC’s

90% confidence limit in certain instances, such as when the initial scan has a larger than average upward fluctuation, possibly due to time-dependent rf noise. The results were validated with semi-independent analyses carried out by two groups, one based at Yale and the other at Berkeley. The two analyses agreed on average to within 5% in both the aggregate exclusion and the 10% prior update contour. This level of agreement is within the 8.3% systematic uncertainty in the exclusion arising from parameter estimation, as described in the Appendix. In addition, Fig. 11 shows the results from both Phase I and Phase II of HAYSTAC alongside results from other axion haloscopes in the mass region  $1$ – $25 \mu\text{eV}$ .

#### IV. CONCLUSION

The results presented here detail the design and performance of HAYSTAC’s Phase II data taking campaign and include new results from the axion dark matter search taken during Phase IIb. This expands on the previously published results from Phase IIa [24,29,30] and includes specific focus on the upgrades made prior to Phase II to incorporate the SSR needed to achieve subquantum limited noise.

In addition, results from Phase IIb are presented for the first time along with a detailed outline of updates made to facilitate a search for axions above the range covered in Phase IIa. No evidence of an axion-induced power excess is observed in this new region, resulting in an aggregate exclusion on the axion photon coupling at the 90% level of  $2.06 \times g_\gamma^{\text{KSVZ}}$  in the mass range  $18.44\text{--}18.71 \mu\text{eV}/c^2$ . Combined with the aggregate exclusion of  $1.95 \times g_\gamma^{\text{KSVZ}}$  from Phase IIa over the axion mass ranges  $16.96\text{--}17.12$  and  $17.14\text{--}17.28 \mu\text{eV}/c^2$ , a total of 137 MHz of axion masses has now been scanned using the quantum enhanced haloscope pioneered by the HAYSTAC experiment. This demonstrates the ability to systematically operate a sub-quantum limited dark matter axion search over large swaths of parameter space. Continued operation with the current JPAs will allow an additional  $\sim 500$  MHz of yet unexplored parameter space to be covered with comparable sensitivity using the current cavity. In addition, extension to higher frequencies can be facilitated with higher frequency JPAs as well as new cavity designs currently being characterized with a dedicated research and development campaign [71].

## ACKNOWLEDGMENTS

HAYSTAC is supported by the National Science Foundation under Grants No. PHY-1701396, No. PHY-1607223, No. PHY-1734006, and No. PHY-1914199, and the Heising-Simons Foundation under Grants No. 2014-0904 and No. 2016-044. We thank Kyle Thatcher and Calvin Schwadron for work on the design and fabrication of the SSR mechanical components, Felix Vietmeyer for his work on the room-temperature electronics, Steven Burrows for graphical design work, and Cady van Assendelft for help in assembling the experiment. We thank Vincent Bernardo and the J. W. Gibbs Professional Shop as well as Craig Miller and Dave Johnson for their assistance with fabricating the system's mechanical components. We also wish to thank the Cory Hall Machine Shop at University of California, Berkeley, and acknowledge Sergio Velazquez for fabricating several of the prototype components tested here. We thank Dr. Matthias Buehler of Low-T Solutions for cryogenics advice. Finally, we thank the Wright Laboratory for housing the experiment and providing computing and facilities support. K. M. B.'s affiliation with The MITRE Corporation is provided for identification purposes only and is not intended to convey or imply MITRE's concurrence with, or support for, the positions, opinions, or viewpoints expressed by the author.

## APPENDIX: DETECTOR CHARACTERIZATION AND PARAMETER UNCERTAINTY

To optimally combine each spectrum and determine the sensitivity to an axion with coupling strength  $g_{a\gamma\gamma}$ , the expected SNR as a function of  $\delta_\nu$  must be precisely estimated for each spectrum. A set of characterization

measurements are used to extract the relevant parameters needed to determine the expected axion conversion power and the total system noise. Details about this characterization procedure, along with estimates of the uncertainty associated with each extracted parameter, are presented below.

### 1. Characterization of expected signal

The power expected from an axion-induced signal comes from Eq. (3),

$$P_{\text{ax}}(\delta_\nu) = \frac{1}{2} \left( \frac{g_{a\gamma\gamma}^2 2\pi \rho_a \hbar^3 c^3}{m_a^2 \mu_0} \right) (B_0^2 V C_{010}) \times \left( \nu_c Q_L \frac{\beta}{1 + \beta} \frac{1}{1 + (2\delta_\nu/\Delta\nu_c)^2} \right) \quad (\text{A1})$$

reorganized here into three groups of parameters by parentheses, with a factor of 1/2 added to account for the SSR only measuring the signal in a single quadrature. The parameters in the first set of parentheses capture the dependence on both fundamental constants as well as astrophysical parameters of axion cold dark matter. To be consistent with the literature and other axion dark matter searches, the axion density of  $\rho_a = 0.45 \text{ GeV}/\text{cm}^3$  is used. The axion coupling is parametrized by the dimensionless coupling  $g_\gamma$  related to  $g_{a\gamma\gamma}$  by

$$g_{a\gamma\gamma} = g_\gamma \frac{m_a \alpha_e}{\pi \Lambda^2}, \quad (\text{A2})$$

where  $\alpha_e$  is the fine-structure constant and  $\Lambda$  is a parameter which encodes the dependence of the axion mass on hadronic physics, taken to be 77.6 MeV [72]. For KSVZ axions  $|g_\gamma| = 0.97$  [10,11] while for a DFSZ axion  $|g_\gamma| = 0.36$  [12,13]. Because these parameters are generic to all axion experiments and are not experimentally measured here, they are not considered in the estimate of the total uncertainty.

Contained in the second set of parentheses of Eq. (A1) are system parameters which are assumed to be static over the course of a run and determined outside of the normal characterization measurements. The first of these is the volume of the cavity,  $V$ , excluding regions occupied by the tuning rod and estimated as 1.545 l from a geometric model. Uncertainty on the volume comes from the  $\sim 125 \mu\text{m}$  machining precision required for the alignment of the tuning rod, which approximately translates to a  $\sim 0.3\%$  uncertainty. Second, the magnetic field  $B$  is set to 8 T, as measured by the superconducting current in the solenoidal magnet. Operating the magnet in the persistent mode allows for long-term stability, and the dominant uncertainty comes from the calibration of the current readout estimated with a Hall probe to be  $\leq 1\%$  by the manufacturer [73]. Finally,  $C_{010}$  is calculated by modeling the spatial profile of the  $\text{TM}_{010}$  electric field as a function of the rod angle with the

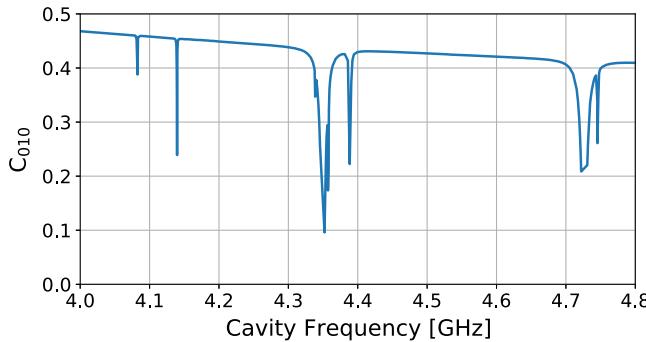


FIG. 12. Plot of the cavity's form factor for the  $TM_{010}$  mode over the band of interest in this paper. The regions with sharp drops in  $C_{010}$  are due to mode crossings and are avoided during the axion search.

use of the CST Microwave Studio [74]. Results of these simulations showing  $C_{010}$  as a function of the  $TM_{010}$  mode frequency can be found in Fig. 12. Extensive validation of this model has been summarized in Ref. [75], showing good agreement between the model and data. In addition, these validations show that a tilt of the tuning rod up to 125  $\mu\text{m}$  can cause  $C_{010}$  to vary by  $\leq 2.6\%$  relative to the baseline configuration. Because the actual misalignment is unknown, this is taken as a conservative estimate on the uncertainty.

The last set of parentheses in Eq. (A1) contains parameters defining the cavity response which are subject to change over the course of a run and must be empirically measured each time the tuning rod is adjusted. As described in Sec. II B this is achieved with a pair of VNA measurements used to measure  $Q_L$ ,  $\nu_c$ , and  $\beta$ . The uncertainty in each of these measurements derives from two main sources. First, residual vibrations and other instabilities in the tuning system can cause the cavity mode to drift over the course of a single measurement which can cause time-dependent fluctuations in each of these parameters. This is most important for  $\nu_c$ , which needs to be held near the JPA pump frequency in order to maintain high squeezing. In order to limit drift to within  $\Delta\nu_c \sim 60$  kHz, conservatively used as a data quality cut on drift, data taking is paused after each tuning step to allow the system time to settle as described in Sec. II B. The residual drift is then measured by taking VNA measurements both before and after each tuning step. The average drifts over the course of Phase IIa and Phase IIb are measured to be 18.5 kHz and 1.5 kHz, respectively. This same method is used to measure the average variation in  $Q_L$  and  $\beta$  due to cavity drift, with the average uncertainty on  $Q_L(\beta)$  found to be 0.4% (0.3%) in Phase IIa and 0.6% (0.3%) in Phase IIb. In addition to drift, uncertainty in these fits can arise from bias introduced by an incomplete or inaccurate fit model. A study using different fit windows and background models for Eq. (7) found an average variation of 6 kHz on  $\nu_c$ , 1.9% on  $Q_L$ , and 2.9% on  $\beta$ . Because the cavity operates at GHz frequencies, the absolute uncertainty on  $\nu_c$  is small but can impact the estimate on  $P_{\text{ax}}$  by distorting the expected shape.

Given the observed variations in  $\nu_c$  described above, the impact on  $P_{\text{ax}}$  was estimated to be  $\leq 1\%$  in both Phase IIa and Phase IIb.

## 2. System noise calibration

In order to determine the total system noise, an approximate model of the system, shown in Fig. 13, is used to express the expected noise spectral density coming from each part of the system. This model is designed to capture the main components of the system without introducing unnecessary complexity. The input of this model is sourced from the Johnson-Nyquist noise of a  $50\Omega$  load held at a temperature  $T$  which has noise spectral density given by

$$S(T) = \frac{1}{2} \left( \frac{1}{2} + \frac{1}{e^{h\nu/k_b T} - 1} \right), \quad (\text{A3})$$

where the overall factor of 1/2 accounts for measuring the noise along just a single quadrature. The temperature of this load can be switched between a cold load held at the mixing chamber's base temperature  $T_{\text{mc}}$  or a hot load held at an elevated temperature  $T_{\text{VTS}}$  by the VTS described in Sec. II B. Noise produced by either load is routed through the remaining components, which perform discrete operations on the spectral density. The multiplicative components of the model are the frequency-dependent gains  $G_S$ ,  $G_A$ , and  $G_H$  and the fractional transmissivities  $\lambda$ ,  $\rho$ , and  $\alpha$ . Each of these loss elements results in a fraction of the state being replaced by Johnson-Nyquist noise at the physical temperature of the component [76], all of which are assumed to be thermalized with the mixing chamber and to be frequency independent. The cavity is modeled as an additional frequency-dependent loss defined by its susceptibility  $|\chi_{\text{rfi}}|^2$ , given by the cavities' reflection profile measured with the VNA as in Sec. II B. This results in a noise contribution proportional to its apparent physical

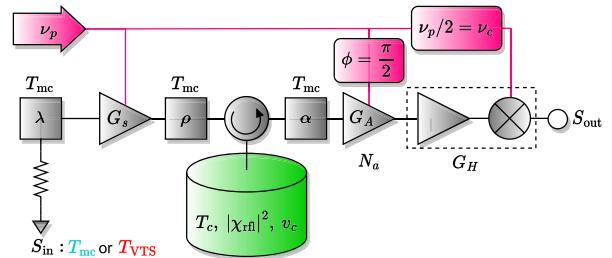


FIG. 13. Model of the HAYSTAC detector used for calibration measurements described in the text. Noise is sourced from a  $50\Omega$  load which can be switched between a cold ( $T_{\text{mc}}$ ) and hot ( $T_{\text{VTS}}$ ) load. This noise is transmitted through a series of discrete components, including loss elements ( $\lambda$ ,  $\rho$  and  $\alpha$ ), gain elements ( $G_S$ ,  $G_A$ , and  $G_H$ ), and the cavity, which adds noise proportional to  $|\chi_{\text{rfi}}|^2$  at its effective temperature  $T_c$ . The signal is then mixed down into the IF by half the JPA pump frequency ( $\nu_p$ ), equivalent to the cavity's frequency ( $\nu_c$ ).

temperature  $T_c$ , which, as described in Sec II B, is observed to be at an elevated temperature relative to  $T_{mc}$ . The final gain stage encompasses gain from the HEMT as well as all subsequent stages, which includes the gain from the mixer used to mix the signal down into the IF with a LO pumped at half the JPA pump frequency  $\nu_p$ , equivalent to the cavity frequency  $\nu_c$ . Additional noise added by the amplification chain, also assumed to be frequency dependent, is represented by  $N_a$ . For an ideal parametric amplifier this noise would vanish but is, in general, nonzero due to imperfections causing deviations from this ideal behavior.

Cascading all of these elements together allows for the expected output spectral density in the IF to be expressed as

$$S_{\text{out}}(|\delta_\nu|) = G_A G_H (N_a + 2[\alpha(1 - \rho)|\chi_{\text{rf}}|^2 S_{mc} + \rho\alpha G_S |\chi_{\text{rf}}|^2 ((1 - \lambda)S_{mc} + \lambda S_{in})] + 2[\alpha(1 - |\chi_{\text{rf}}|^2)S_c + (1 - \alpha)S_{mc}]), \quad (\text{A4})$$

where the explicit frequency dependence of each variable has been suppressed and  $S_{mc}$ ,  $S_c$ , and  $S_{in}$  are the noise spectral densities of the mixing chamber, cavity, and input load, respectively, as given by Eq. (A3). The terms in Eq. (A4) can be arranged to resemble the noise components in Eq. (2), but they now describe the variance of the  $\hat{X}$  quadrature rather than its phasor elements. In this case the first term in brackets comes from the noise reflected off of the cavity,  $N_r$ , and the second term in brackets comes from the excess noise of the cavity,  $N_c$ , allowing  $S_{\text{out}}$  to be written as

$$S_{\text{out}} = G_A G_H (N_a + 2(N_r + N_c)) \quad (\text{A5})$$

with the terms in parentheses being  $N_{\text{sys}}$  in the IF. The factor of 2 which appears in front of  $N_c$  and  $N_r$  represents the noise present in the IF spectrum from positive and negative detunings in the rf,

$$S_{\text{out}}^{\text{IF}}(|\delta_\nu|) = S_{\text{out}}^{rf}(+\delta_\nu) + S_{\text{out}}^{rf}(-\delta_\nu), \quad (\text{A6})$$

added prior to the addition of  $N_a$ . Because both noise sources are symmetric about  $\nu_p$ , as illustrated in Fig. 10, the contribution of both is exactly doubled in each IF bin at the output of the JPA. As a result, the quantum limited noise in one quadrature for any frequency bin in the IF is  $S_{\text{out}} = 1/2$  when operating at 0 K without squeezing. Calibration of the total noise requires measuring each of the parameters in Eq. (A4).

This procedure starts with a measurement of the three transmissivities, performed prior to the start of each run and assumed to be constant over the course of data taking. The first measurement follows the procedure outlined in Ref. [27] by which the total transmission efficiency between the squeezer and amplifier ( $\eta$ ) is extracted through an observation of the change in noise power when swapping

the preamplification between the amplifier and squeezer while detuned from the cavity. This procedure results in a measured transmission efficiency of  $0.63 \pm 0.03$  ( $0.60 \pm 0.03$ ) in Phase IIa (Phase IIb). These measurements agree with the observed squeezing in the system, which is fundamentally limited by this total loss as described by Eq. (6). Following this, the loss before the squeezer is extracted through a measurement of the total system transmission efficiency ( $\zeta = \alpha\rho\lambda$ ) and by dividing out the already determined  $\eta$ . This measurement is taken by observing the dependence of noise power on both the input load and mixing chamber temperatures, the slope of which gives the total system transmissivity observed to be  $\zeta = 0.50 \pm 0.07$ . The uncertainties on these two measurements come from the statistical variations over repeated measurements. Unfortunately, the combinations of measurements do not distinguish between losses before and after the cavity; instead, it is approximately assumed that these losses are evenly divided between the two components, giving  $\alpha = \rho = \sqrt{\eta}$ . Because  $\rho$  only appears in Eq. (A4) as the total transmissivity  $\eta$ , this assumption only has an impact on  $S_{\text{out}}$  through  $\alpha$ . An estimate on this uncertainty can be made by comparing previous measurements of  $\alpha$  from Phase I [36] which show a  $\sim 10\%$  relative change in the transmission efficiency for Phase II. While this drop is likely due to the addition of the more lossy triple junction circulator, a conservative estimate of  $\sim 10\%$  is associated with  $\alpha$  to account for this assumption.

Extraction of the remaining parameters is done by periodically collecting a set of five power spectrum measurements taken with the DAQ described in Sec. II D to record the average power spectrum over an integration time  $\tau$ . Each measurement is taken with certain components removed, allowing for the dependence of each parameter to be isolated with a comparison of the relative change in the spectral shape defined by the ratio

$$Y^{ij} = \frac{S_{\text{out}}^i G_A^j}{S_{\text{out}}^j G_A^i}, \quad (\text{A7})$$

where  $i$  and  $j$  represent the calibration measurement ID defined in Table I. While each measurement is short relative

TABLE I. Summary of the different configurations of the system during the periodic calibration measurements used for the calculation of  $N_a$ ,  $G_s$ , and  $S_c$ .

Calibration ID	SQZ	$(\nu_p/2 - \nu_c)$	Load		
			Temperature	$\tau_{\text{cal}}$	$G_A$
<i>g</i>	ON	0 MHz	$T_{mc}$	10 s	$G_A^g$
<i>b</i>	OFF	0 MHz	$T_{mc}$	10 s	
<i>a</i>	OFF	-3 MHz	$T_{mc}$	10 s	$G_A^c$
Cold ( <i>c</i> )	OFF	-3 MHz	$T_{mc}$	900 s	
Hot ( <i>h</i> )	OFF	-3 MHz	$T_{VTS}$	900 s	$G_A^h$

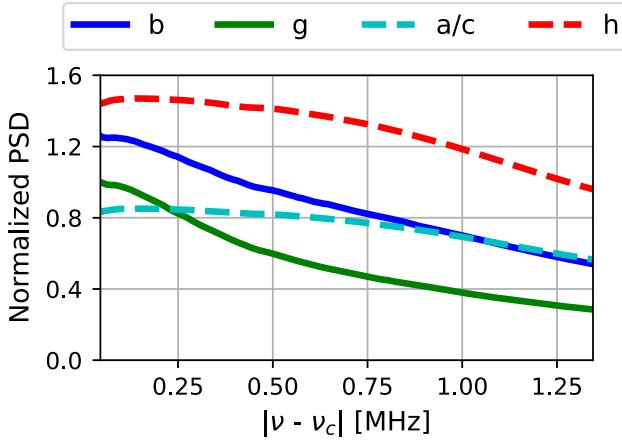


FIG. 14. Example set of calibration spectra taken periodically to extract parameters for determining the system noise, shown for both on-resonant (solid lines) and off-resonant (dashed lines) measurements. The spectra are smoothed to remove expected statistical fluctuations and normalized relative to the on-resonant power of the “g” measurement, which is taken in the same configuration as that used to search for axions.

to the integration time of the axion sensitive noise measurement taken at each tuning step, the total time required for the full set of measurements is  $\sim 1$  hr and is largely dominated by time spent waiting for the system to reach equilibrium after switching between the cold and hot loads. For this reason, these measurements are roughly taken once every nine spectra so as to minimally impact the dead time discussed in Sec. II D. An example set of power spectrum taken during the calibration routine are shown in Fig. 14 where the spectra are normalized to the on-resonant power of the “g” measurement.

In order to limit uncertainty in the parameter extraction, each measurement would ideally be taken such that only the intentionally modified parameters would change between measurements. While this is possible for most parameters in the model, it is not possible for the amplifier gain which is subject to change when rebiased as is required for the off-resonant and hot-load measurements. Direct measurements of the single quadrature gain are made each time the amplifier bias is changed, but these measurements have large uncertainties which can result in deviations of the calibration parameters outside of their physically realistic regions if not corrected. Because the gains enter the calculation in Eq. (A7) only as ratios, realistic estimates of each parameter are extracted by allowing the gain ratios  $G_A^{bg}/G_A^{ac}$  and  $G_A^{ac}/G_A^h$  to deviate from their approximately measured value. In order to avoid overestimating the detector sensitivity, the set of gains which produce physically realizable values are scanned, and the pair which give the highest total noise are conservatively used as the final estimate.

The first two measurements are taken with  $S_{in}$  sourced by a load at  $T_{mc}$  and the amplifier tuned to match the cavity

frequency, the same as for axion data. Isolation of  $G_S$  is achieved by switching the JPA used to prepare the squeezed state on for measurement “g” and off in measurement “b,” allowing for  $G_S$  to be found as

$$G_S = \frac{1}{2\alpha\rho|\chi_{rfl}|^2 Y^{bg} S_{mc}} [2N_c^b(1 - Y^{bg}) + 2N_c^b + N_a(1 - Y^{bg}) - 2\alpha(1 - \rho)|\chi_{rfl}|^2 Y^{bg} S_{mc}]. \quad (A8)$$

The parameter scan restricts values of  $G_S$  to the range  $0 < G_S < 1$ , where the upper bound asserts that squeezing corresponds to a gain of less than 1 and the lower bound is a physical requirement. As can be seen in Eq. (A8), both  $S_c$  and  $N_a$  still appear in the extraction of  $G_S$ . As a result, the uncertainty on  $G_S$  is relatively high, and an estimate of its value is derived by comparing the extracted  $G_S$  to the approximate target gain of 13 dB needed to achieve optimal squeezing. This results in a  $\sim 50\%$  uncertainty on the extraction of  $G_S$  in both phases. While this uncertainty is high, the relative contribution to the total noise and, in turn, the exclusion is small.

Following this a measurement “a” is taken with the squeezer still off but the amplifier tuned 3 MHz below the cavity’s resonant frequency. Comparing this to measurement “b,” the cavity’s excess temperature is found as

$$S_c = \frac{(Y^{ba} - 1)(N_a + 2S_{mc}) + 2\alpha S_{mc}(1 - |\chi_{rfl}|^2)}{2\alpha(1 - |\chi_{rfl}|^2)} \quad (A9)$$

Bounds on  $S_c$  are given by an absolute measurement of the cavity’s physical temperature taken at the start of the run. This is performed by raising the mixing chamber temperature until the excess cavity noise vanishes at  $\sim 225$  mK, indicating the system is in thermal equilibrium. While this would imply  $S_c \approx 0.6$ , measurements of the cavity’s exterior temperature indicate that only an internal component is at an elevated temperature. Currently this is assumed to be the tuning rod, which, as described in Sec. II B, has only a weak thermal link to the mixing chamber. As a result it is assumed that the effective cavity noise temperature  $T_C < T_{rod}$  and a bound of  $0.38 < S_c < 0.44$  is derived from the rod’s geometric orientation relative to the antenna. In addition, an approximate uncertainty of 20% is assigned to  $S_c$  based on the initial estimate of  $T_{rod}$  taken before the start of the run.

Finally, extraction of  $N_a$  is performed with a pair of measurements where  $S_{in}$  is switched between a cold load ( $T_{mc}$ ) and a hot load ( $T_{VTS}$ ) with the amplifier tuned 3 MHz below the cavity’s resonant frequency and the squeezer switched off. This technique is similar to the Y-factor measurement presented in Refs. [31,36], but it makes use of the VTS which allows for control of the load temperature below 775 mK. The ratio of the two measurements then gives  $N_a$  as

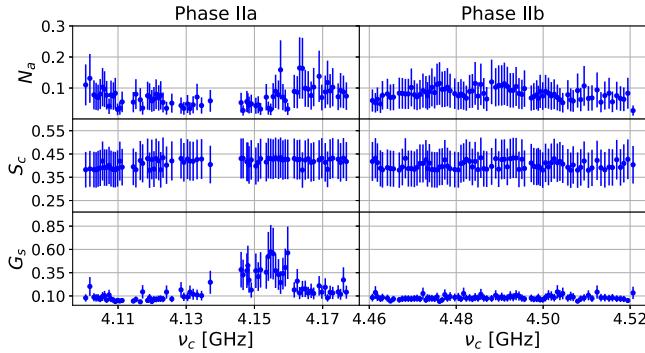


FIG. 15. Average value over the analysis band of the three model parameters extracted from the periodic calibrations used to determine the system noise, shown for all of Phase II as a function of the cavity frequency. Error bars are the result of the systematic uncertainty on each parameter shown in Table II.

$$N_a = \frac{2(N_c^h + N_r^h - Y^{hc}(N_c^c + N_r^c))}{Y^{hc} - 1}. \quad (\text{A10})$$

Because neither the calculation of  $S_c$  nor  $N_a$  is fully decoupled from the other, their calculation is done in two iterations. First,  $N_a$  is calculated by ignoring the cavity's contribution to Eq. (A10). The resulting estimate is used in Eq. (A9) to calculate  $S_c$ , which is then fed back to correct for the cavity's contribution to the determination of  $N_a$ . A lower bound of  $N_a > 0.008$  is set based on the HEMT's added noise spec (LNF-LNC4 8A [77]), and no upper limit is applied. Uncertainty on  $N_a$  is largely due to uncertainty in the temperature of the VTS, which is known to within 5 mK from the Magnicon temperature sensor described in Sec. II B, as well as to the uncertainty in the total transmission efficiency  $\zeta$  which gives a total uncertainty of 63% in both phases.

The extracted values for  $G_s$ ,  $S_c$ , and  $N_a$  for each set of calibration measurements, reported as the average over the analysis band, are shown in Fig. 15 for both data taking phases. Using the extracted value for each parameter, the total system noise can be calculated as a function of detuning using Eq. (A4), as shown in Fig. 16 in the IF using the average parameters from Phase IIb. As expected, the SSR reduces the added noise from  $N_r$  over a large range of  $\delta_\nu$ , allowing for an improved measurement bandwidth when operating strongly overcoupled to the cavity.

### 3. Uncertainty on exclusion

As many of the parameters used to calculate  $\Sigma$  vary with  $\delta_\nu$ , their relative contribution to the total uncertainty also varies. To account for this, the total  $\Sigma$  is calculated by generating a set of fake spectra matching the frequencies in real data. For each spectrum, the individual  $\Sigma$  is calculated using the average parameters from the phase and shifted by the estimated uncertainty on each parameter. The total  $\Sigma$ ,

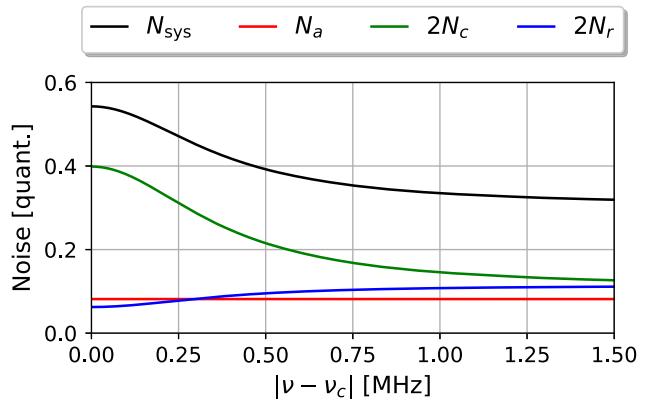


FIG. 16. Contributions to  $N_{\text{sys}}$  in the IF, plotted in units of power spectral density as a function of detuning using the average calibration parameters from Phase IIb shown in Table II.

resulting from the squared sum of overlapping spectra, is calculated and used to estimate the overall bias in the exclusion on  $g_\gamma$ . The results of this study are summarized in Table II, which also includes a summary of the individual uncertainty contribution from each parameter. The total uncertainty on  $g_\gamma$  is 8.3%, with the largest contributions coming from uncertainty in  $S_c$ ,  $\alpha$ , and  $N_a$ . Currently this uncertainty is not propagated to the final exclusion and is instead included here as an approximate estimate on the reliability of the final exclusion. A full treatment of this uncertainty is needed to account for correlations in the extracted parameters introduced by the calibration procedure.

TABLE II. Breakdown of each experimentally determined parameter which is needed to determine the expected  $\Sigma$  for an axion-induced signal. For parameters which vary between tuning steps, the average value over the entire phase is used. Details about the quantification of each parameter and the associated uncertainty are given in the text.

Parameter	Value		Fractional Uncertainty	
	Phase IIa	Phase IIb	$\sigma_{\text{par}} (\%)$	$\sigma_{g_\gamma} (\%)$
$\nu_c$	4.14 GHz	4.49 GHz		$\leq 0.4$
$B_0$	8 T	8 T	1	1
$V$	1.5451	1.5451	0.3	0.15
$C_{010}$	0.43	0.43	2.6	1.3
$Q_L$	6180	5915	2.0	0.6
$\beta$	6.9	6.5	2.9	0.1
$\alpha\rho$	0.63	0.60	2	1
$\alpha$	0.794	0.775	10	4.5
$G_s$	0.11	0.08	50	1.5
$S_c$	0.41	0.40	20	3.5
$N_a$	0.08	0.08	63	5.5
$g_\gamma$	$1.95 \times g_\gamma^{\text{KSVZ}}$	$2.06 \times g_\gamma^{\text{KSVZ}}$		8.3

[1] N. Aghanim *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A6 (2020); **652**, C4(E) (2021).

[2] M. Markevitch, A. H. Gonzalez, D. Clowe, A. Vikhlinin, L. David, W. Forman, C. Jones, S. Murray, and W. Tucker, *Astrophys. J.* **606**, 819 (2004).

[3] R. D. Peccei and H. R. Quinn, *Phys. Rev. Lett.* **38**, 1440 (1977).

[4] R. D. Peccei and H. R. Quinn, *Phys. Rev. D* **16**, 1791 (1977).

[5] S. Weinberg, *Phys. Rev. Lett.* **40**, 223 (1978).

[6] F. Wilczek, *Phys. Rev. Lett.* **40**, 279 (1978).

[7] J. Preskill, M. B. Wise, and F. Wilczek, *Phys. Lett.* **120B**, 127 (1983).

[8] L. Abbott and P. Sikivie, *Phys. Lett.* **120B**, 133 (1983).

[9] M. Dine and W. Fischler, *Phys. Lett.* **120B**, 137 (1983).

[10] J. E. Kim, *Phys. Rev. Lett.* **43**, 103 (1979).

[11] M. A. Shifman, A. I. Vainshtein, and V. I. Zakharov, *Nucl. Phys.* **B166**, 493 (1980).

[12] M. Dine, W. Fischler, and M. Srednicki, *Phys. Lett.* **104B**, 199 (1981).

[13] A. R. Zhitnitsky, *Sov. J. Nucl. Phys.* **31**, 260 (1980).

[14] G. Bertone and T. M. P. Tait, *Nature (London)* **562**, 51 (2018).

[15] M. Gorgetto and G. Villadoro, *J. High Energy Phys.* 03 (2019) 033.

[16] V. B. Klaer and G. D. Moore, *J. Cosmol. Astropart. Phys.* 11 (2017) 049.

[17] M. Buschmann, J. W. Foster, and B. R. Safdi, *Phys. Rev. Lett.* **124**, 161103 (2020).

[18] J. Kim *et al.*, *Phys. Rev. Lett.* **130**, 091602 (2023).

[19] C. Bartram *et al.* (ADMX Collaboration), *Phys. Rev. Lett.* **127**, 261803 (2021).

[20] W. Wuensch, S. De Panfilis-Wuensch, Y. K. Semertzidis, J. T. Rogers, A. C. Melissinos, H. J. Halama, B. E. Moskowitz, A. G. Prodell, W. B. Fowler, and F. A. Nezrick, *Phys. Rev. D* **40**, 3153 (1989).

[21] C. Hagmann, P. Sikivie, N. S. Sullivan, and D. B. Tanner, *Phys. Rev. D* **42**, 1297 (1990).

[22] D. Alesini *et al.*, *Phys. Rev. D* **103**, 102004 (2021).

[23] A. P. Quiskamp, B. T. McAllister, P. Altin, E. N. Ivanov, M. Goryachev, and M. E. Tobar, *Sci. Adv.* **8**, abq3765 (2022).

[24] K. M. Backes *et al.* (HAYSTAC Collaboration), *Nature (London)* **590**, 238 (2021).

[25] P. Sikivie, *Phys. Rev. Lett.* **51**, 1415 (1983); **52**, 695(E) (1984).

[26] P. Sikivie, *Phys. Rev. D* **32**, 2988 (1985); **36**, 974(E) (1987).

[27] M. Malnou, D. A. Palken, B. M. Brubaker, L. R. Vale, G. C. Hilton, and K. W. Lehnert, *Phys. Rev. X* **9**, 021023 (2019); **10**, 039902(E) (2020).

[28] T. Yamamoto, K. Inomata, M. Watanabe, K. Matsuba, T. Miyazaki, W. D. Oliver, Y. Nakamura, and J. S. Tsai, *Appl. Phys. Lett.* **93**, 042510 (2008).

[29] K. M. Backes, A quantum-enhanced search for dark matter axions, Ph.D. thesis, Yale University, 2021, [https://elischolar.library.yale.edu/gsas\\_dissertations/294/](https://elischolar.library.yale.edu/gsas_dissertations/294/).

[30] D. A. Palken, Enhancing the scan rate for axion dark matter: Quantum noise evasion and maximally informative analysis, Ph.D. thesis, University of Colorado Boulder, 2020, <https://jila.colorado.edu/bibcite/reference/12122>.

[31] B. M. Brubaker *et al.*, *Phys. Rev. Lett.* **118**, 061302 (2017).

[32] L. Zhong *et al.* (HAYSTAC Collaboration), *Phys. Rev. D* **97**, 092001 (2018).

[33] M. Malnou, D. A. Palken, L. R. Vale, G. C. Hilton, and K. W. Lehnert, *Phys. Rev. Appl.* **9**, 044023 (2018).

[34] S. K. Lamoreaux, K. A. van Bibber, K. W. Lehnert, and G. Carosi, *Phys. Rev. D* **88**, 035020 (2013).

[35] J. I. Read, *J. Phys. G* **41**, 063101 (2014).

[36] S. Al Kenany *et al.*, *Nucl. Instrum. Methods Phys. Res., Sect. A* **854**, 11 (2017).

[37] Attocube, ANR240, <https://www.attocube.com>.

[38] Bluefors Cryogenics, BF-LD250, <https://bluefors.com>.

[39] Z. Xie, W. Xue, H. Chen, and Y. Huang, *Ceram. Int.* **37**, 2165 (2011).

[40] A. Drobizhev, J. Reiten, V. Singh, and Y. G. Kolomensky, *Cryogenics* **85**, 63 (2017).

[41] Applied Motion, HT23-595, <https://www.applied-motion.com>.

[42] MAGNICON, MFFT-1 Data-Sheet, <http://www.magnicon.com>.

[43] Keysight, E5071C ENA Vector Network Analyzer, <https://www.keysight.com>.

[44] GaGe, GaGe Oscar CSE4344, <https://www.gage-applied.com>.

[45] Marki Microwave, IQ-0307 IQ Mixer, <https://www.markimicrowave.com>.

[46] Stanford Research Systems, Rubidium Frequency Std., <https://www.thinksrs.com>.

[47] D. M. Pozar, *Microwave Engineering* (Wiley, Hoboken, NJ, 2011), 4th ed.

[48] Yokogawa, GS200 DC Voltage/Current Source, <https://tmi.yokogawa.com>.

[49] Keysight, E8257D PSG Analog Signal Generator, <https://www.keysight.com>.

[50] Keysight, N5183B MXG X-Series Microwave Analog Signal Generator, <https://www.keysight.com>.

[51] QuinStar, Cryogenic Circulators and Isolators, <https://quinstar.com>.

[52] NVIDIA, NVIDIA TITAN V, <https://www.nvidia.com>.

[53] B. M. Brubaker, L. Zhong, S. K. Lamoreaux, K. W. Lehnert, and K. A. van Bibber, *Phys. Rev. D* **96**, 123008 (2017).

[54] A. Savitzky and M. J. E. Golay, *Anal. Chem.* **36**, 1627 (1964).

[55] M. S. Turner, *Phys. Rev. D* **42**, 3572 (1990).

[56] A. Caputo, A. J. Millar, C. A. J. O'Hare, and E. Vitagliano, *Phys. Rev. D* **104**, 095029 (2021).

[57] S. Ghosh, E. P. Ruddy, M. J. Jewell, A. F. Leder, and R. H. Maruyama, *Phys. Rev. D* **104**, 092016 (2021).

[58] S. De Panfilis, A. C. Melissinos, B. E. Moskowitz, J. T. Rogers, Y. K. Semertzidis, W. Wuensch, H. J. Halama, A. G. Prodell, W. B. Fowler, and F. A. Nezrick, *Phys. Rev. Lett.* **59**, 839 (1987).

[59] C. Hagmann, P. Sikivie, N. S. Sullivan, and D. B. Tanner, *Phys. Rev. D* **42**, 1297 (1990).

[60] C. Boutan *et al.* (ADMX Collaboration), *Phys. Rev. Lett.* **121**, 261302 (2018).

[61] C. Bartram *et al.* (ADMX Collaboration), *Phys. Rev. D* **103**, 032002 (2021).

[62] C. Bartram *et al.* (ADMX Collaboration), [arXiv:2110.10262](https://arxiv.org/abs/2110.10262).

[63] Y. Lee, B. Yang, H. Yoon, M. Ahn, H. Park, B. Min, D. Kim, and J. Yoo, *Phys. Rev. Lett.* **128**, 241805 (2022).

[64] J. Jeong, S. Youn, S. Bae, J. Kim, T. Seong, J. E. Kim, and Y. K. Semertzidis, *Phys. Rev. Lett.* **125**, 221302 (2020).

[65] O. Kwon *et al.* (CAPP Collaboration), *Phys. Rev. Lett.* **126**, 191802 (2021).

[66] S. Lee, S. Ahn, J. Choi, B. R. Ko, and Y. K. Semertzidis, *Phys. Rev. Lett.* **124**, 101802 (2020).

[67] C. M. Adair *et al.*, *Nat. Commun.* **13**, 6180 (2022).

[68] H. Chang *et al.* (TASEH Collaboration), *Phys. Rev. Lett.* **129**, 111802 (2022).

[69] L. Di Luzio, M. Giannotti, E. Nardi, and L. Visinelli, *Phys. Rep.* **870**, 1 (2020).

[70] D. A. Palken *et al.*, *Phys. Rev. D* **101**, 123011 (2020).

[71] M. Simanovskaia, A. Droster, H. Jackson, I. Urdinaran, and K. van Bibber, *Rev. Sci. Instrum.* **92**, 033305 (2021).

[72] C. Patrignani *et al.* (Particle Data Group), *Chin. Phys. C* **40**, 100001 (2016).

[73] Cryomagnetics, private communication, <https://www.cryomagnetics.com> (2022).

[74] Dassault Systemes, <https://www.3ds.com>.

[75] N. M. Rapidis, S. M. Lewis, and K. A. van Bibber, *Rev. Sci. Instrum.* **90**, 024706 (2019).

[76] H. B. Callen and T. A. Welton, *Phys. Rev.* **83**, 34 (1951).

[77] Low Noise Factory, LNF-LNC4 8A, <https://lownoisefactory.com>.