

Any Light Particle Searches with ALPS II: Description of the first science campaign

Aaron D. Spector,^{1,*} Daniel C. Brotherton,² Ayman Hallal,³ Henry Frädrieh,¹
 Jacob Egge,¹ Li-Wei Wei,^{1,†} Todd Kozlowski,^{1,‡} Kanioar Karan,^{4,†}
 Aldo Ejlli,³ Katharina-Sophie Isleif,⁵ Hartmut Grote,⁴ Harold Hollis,²
 Guido Mueller,³ David B. Tanner,² Benno Willke,³ and Axel Lindner¹

¹*Deutsches Elektronen-Synchrotron DESY, 22603 Hamburg, Germany*

²*Department of Physics, University of Florida, 32611 Gainesville, Florida, USA*

³*Max-Planck-Institut für Gravitationsphysik (Albert-Einstein-Institut)
 and Leibniz Universität Hannover, 30167 Hannover, Germany*

⁴*School of Physics and Astronomy, Cardiff University, Cardiff CF24 3AA, United Kingdom*

⁵*Helmut-Schmidt-Universität, 22043 Hamburg, Germany*

From February to May of 2024 the Any Light Particle Search II (ALPS II) conducted its first science campaign using the ‘light-shining-through-a-wall’ technique to search for pseudo-Goldstone bosons that lie beyond the Standard Model of particle physics and which are inaccessible by accelerator-based experiments. The experimental setup consists of two strings of superconducting dipole magnets, each more than 100 m long, that are separated by a wall. Laser light is directed through the first magnet string and a heterodyne detection system is used to measure the electromagnetic power that traverses a wall via the conversion to and then from a bosonic field. After the wall, a high-finesse optical cavity resonantly enhances the signal power. Two searches were carried out, one with the laser polarized perpendicular to the magnetic field direction and another with its polarization state aligned parallel to the magnetic field. No evidence for the existence of new bosons was found. In its first science campaign, ALPS II reached photon-boson conversion probability sensitivities of a few 10^{-13} . The ongoing upgrade of the optical system aims to increase this sensitivity by about four orders of magnitude.

CONTENTS		III. Results and Backgrounds	
List of Acronyms	2	A. Alternative Frequency Analysis	30
I. Introduction	3	1. Quantum Noise	31
A. ALPS II Design	4	2. Technical Noise	32
B. Optical System	7	B. Background Estimation	32
C. Heterodyne Sensing	10	C. Statistics of the Background	35
D. Central Optical Bench	15	1. Detection Thresholds	36
II. Closed-Shutter Calibration	17	2. Setting Exclusion Limits	38
A. Potential Systematic Effects	19	D. Summary of Results	38
B. System Locked Flags	21	E. Verifying the Result	38
C. Normalization by HPL Power	21	1. Science-to-Veto Ratio	39
D. Closed to Open Shutter Ratio	23	2. Moving Averaged Data	39
E. Veto Detection System	25	3. Time Evolution of Backgrounds	42
F. Calibration Uncertainty	26	F. Central Optical Bench	
G. Uncertainty in the Results	28	Transmissivity	42
		IV. Performance of the Optical System	44
		A. Regeneration Cavity Parameters	45
		1. Resonant Enhancement	45
		2. Transmissivity	46
		3. Reflectivity	47
		B. Field overlap	48
		1. Longitudinal Overlap	48
		2. Polarization Overlap	49
		3. Open Shutter Power Ratio	51
		4. Spatial Overlap Results	55

* aaron.spector@desy.de

† Now at Max-Planck-Institut für Gravitationsphysik (Albert-Einstein-Institut) and Leibniz Universität Hannover, 30167 Hannover, Germany

‡ Now at Helmut-Schmidt-Universität, 22043 Hamburg, Germany

5. Field Overlap Results	55	VI. Bibliography	59
6. Optimizing the Spatial Overlap	57	A. Appendix	60
V. Conclusion	57	1. Digital Frequency Representation	60
Acknowledgments	59	2. Changes in Open Shutter Data	60
		3. Signal Leakage	61
		4. Double Checking the Calibration	63

LIST OF ACRONYMS

ADC	analog-to-digital converter	LO	local oscillator laser
AL	auxiliary laser	LSW	Light-shining-through-a-wall
ALPS	Any Light Particle Search	MZ	Mach Zehnder
BD	beam dump	NL	north left
BSM	beyond the Standard Model	NR	north right
CH	central hall	PC	production cavity
COB	central optical bench	PD	photodetector
DESY	Deutsches Elektronen-Synchrotron	PDH	Pound-Drever-Hall
EOM	electro-optic modulator	PLL	phase-locked loop
FIR	finite impulse response	PSD	power spectral density
FSR	free spectral range	RC	regeneration cavity
HPL	high-power laser	RMS	root-mean-square
HR	highly reflective	PDF	probability density function
HWP	half-wave-plate	RL	reference laser
		ULE	ultra-low expansion

I. INTRODUCTION

While the necessity of physics beyond the Standard Model (BSM) is increasingly apparent, there remains nonetheless a shortage of experimental evidence to explain phenomena like dark matter, dark energy or the small mass of the Higgs boson. Many theories for physics beyond the Standard Model of particle physics predict the existence of very lightweight and extremely weakly interacting pseudo-Goldstone bosons like the axion, which are out of reach at accelerator-based experiments (see [1] and references therein). This calls for complementary approaches.

Light-shining-through-a-wall (LSW) experiments at optical frequencies offer unique opportunities to search for new physics [2–4]. Here, a laser is directed through a conversion region (typically a beam tube with vacuum or well-controlled gas-pressure, which might be permeated by a magnetic dipole field), generating (potentially) a beam of bosons capable of propagating through a wall that blocks the laser light. The bosonic field then travels through another (re-)conversion region beyond the wall, where some of its energy is reconverted to an electromagnetic field. By measuring the rate at which energy converts between the electromagnetic and bosonic fields, it is possible to deduce their coupling.

Optical cavities can be used to enhance significantly this process. A cavity before the wall will increase the intensity of the electromagnetic field traveling through the conversion region, while a cavity after the wall will resonantly enhance the electromagnetic field reconverted from the interaction between the bosons field [5–7]. LSW experiments in the optical regime generate an enormous photon flux while also being able to accommodate a long string of magnets due to a small beam divergence, making them compelling tools in the hunt for BSM physics. A laser operating at a wavelength of 1064 nm with a high-finesse cavity before the wall can direct up to $\sim 10^{24}$ photons per second through a string of dipole magnets before the wall. Behind the wall, a cavity might boost the re-conversion probability by more than a factor 10,000 and photon rates of 10^{-5} s^{-1} can be sensed [8]. Therefore, photon-boson-photon conversion probabilities below 10^{-33} at O(100 m) long conversion and re-conversion regions can be probed by such

optical LSW-experiments.

The Any Light Particle Searches (ALPSs), an LSW experiment conducted at Deutsches Elektronen-Synchrotron (DESY) from 2007 to 2010, utilized an optical cavity before the wall with a single 5 T HERA dipole magnet to set, at the time, new exclusion limits for purely laboratory based searches for scalar, pseudoscalar bosons, hidden photons and other hypothetical bosons [9]. This was followed by the OSQAR experiment based at CERN, which utilized two 9 T LHC magnets to become the most sensitive LSW experiment for scalar and pseudoscalar bosons of its time [10].

In 2012, a new LSW experiment, ALPS II, was started at DESY and in 2024 ALPS II launched its first science campaign. With two strings of 12 HERA dipole magnets providing 5.3 T magnet fields, a high-power laser (HPL) generating the bosonic field before the wall, and a high-finesse optical cavity after the wall implemented for its first science campaign, ALPS II is now the most sensitive LSW experiment ever built. Over the course of four months, from February to May 2024, two data runs were conducted searching for beyond the Standard Model lightweight bosons.

We would like to emphasize that these searches are sensitive to a rich variety of bosons including vector and tensor bosons in addition to the typically cited scalar and pseudoscalars [1, 11, 12]. Some of the electromagnetic couplings of these fields, such as vector bosons, do not even require the magnets. The results in this manuscript are therefore reported generally in terms of the conversion rate between electromagnetic and BSM fields during a single pass through one of the magnet strings, denoted as $\mathcal{P}_{\gamma \leftrightarrow \phi}$. An interpretation of these results in terms of the coupling between electromagnetic and scalar, pseudoscalar, vector, and tensor bosons is reported in Ref. [1]. Additionally, throughout this paper we refer to lightweight bosons, bosonic fields, and BSM fields interchangeably as potential sources of a signal. A discovery of any such field would represent a fundamental breakthrough in our understanding of the universe as the first evidence of a new energy scale between the electroweak and Planck scales as well as the first direct measurement of a particle residing in the BSM ‘dark sector’.

This document provides an overview of the optical system, calibration procedure, results,

and performance of the instrument during the first science campaign, and is organized as follows. The remainder of Section I discusses the design of the ALPS II optical system as well as the optical system that was implemented during the first science campaign, and then explains the fundamentals of the heterodyne detection system, and the central optical bench. Section II explains how the results are calibrated and the various sources of systematic error in the calibration. The results, expected background rates, and derived exclusion limits are then presented in Section III. The performance of the optical system during the first science campaign is examined in Section IV and our conclusions are presented in Section V.

A. ALPS II Design

As Figure 1 shows, the ALPS II optical system enhances the observable photon flux from photon-boson-photon conversion via the resonant power buildup of optical cavities before and after the wall [8, 13]. A HPL [14] is injected into the optical cavity before the wall, referred to as the production cavity (PC), which increases the laser power propagating through the initial magnetic field. The cavity after the wall, referred to as the regeneration cavity (RC), resonantly enhances the signal power measured at the detection system. The electromagnetic field generated by the reconversion of the BSM field in the RC will be referred to as the regenerated, reconverted, or signal field.

The magnetic fields for ALPS II are generated by superconducting dipole magnets formerly used by the Hadron Electron Ring Accelerator (HERA). With 12 magnets in each of the strings before and after the wall, the products of the magnetic flux density times the lengths of the magnetic fields are $BL = 563.2 \pm 0.1 \text{ T} \cdot \text{m}$ for each string [15, 16]. The magnets are oriented such that their magnetic fields are pointing in the vertical direction with respect to the ground. The polarization state of the HPL can be controlled with a half-wave-plate (HWP) in the injection path to the PC. This allows us to search for different types of BSM fields based on whether the HPL is polarized parallel or perpendicular to the magnetic field.

At sufficiently low masses ($< 1 \text{ meV}$) the BSM field will remain coherent over the length of the

experiment with the electromagnetic field that generates it. Therefore, for the experiment to work properly, it is essential that the light circulating in the PC is resonant with the RC and for the cavities to share the same spatial eigenmode. These conditions are quantified by a parameter that we refer to as the field overlap. The wall is also equipped with a shutter that can be opened to allow light to couple directly from the PC to the RC to check the field overlap by measuring the power in transmission of the RC, referred to as P_{open} , and comparing it to the HPL power incident on the RC. As we explain in more detail in Section II, the open shutter measurements can be used to directly calibrate the results in terms of the conversion rate between electromagnetic and BSM fields $\mathcal{P}_{\gamma \leftrightarrow \phi}$.

To provide specific examples, we will examine the interaction of scalar fields and a laser polarized perpendicular to the external magnetic field, as well as pseudoscalar fields and a laser polarized parallel to the magnetic field. As Ref. [11] explains, with the shutter closed, the power measured after the wall, P_ϕ , generated by the reconversion of the BSM field, is the same in both of these cases and can be approximated with the equation

$$P_\phi = \eta^2 \beta_{\text{RC}} P_{\text{PC}} \mathcal{P}_{\gamma \leftrightarrow \phi}^2 \quad (1)$$

$$\simeq \eta^2 \beta_{\text{RC}} P_{\text{PC}} \left(\frac{g_{\phi\gamma\gamma} BL}{2} \right)^4, \quad (2)$$

for scalar or pseudoscalar fields with a rest mass energy below 0.1 meV . In this equation β_{RC} is the resonant enhancement factor of the RC and P_{PC} is the laser power circulating in the PC [13]. To maximize the total field overlap between the light circulating in the PC and the RC eigenmode η , the frequency of the HPL must be held on a resonance of the RC and coupled to its spatial eigenmode. The field overlap is the product of the spatial, spectral, and polarization components,

$$\eta = \eta_{xy} \eta_{\text{pol}} \eta_z. \quad (3)$$

Here, η_{xy} is the spatial component field overlap between the light circulating in the PC and the RC eigenmode, while η_{pol} is the polarization overlap between the field generating the BSM field and the direction of the magnetic field, and finally η_z , is the spectral component of the total overlap and expresses the degree to which the




FIG. 1: Side view of the design of the full ALPSII experimental system. The HPL beam and PC spatial eigenmode is shown in red, while the BSM field traveling to the right is shown in blue. The dashed purple lines show the projection of the RC spatial eigenmode. A HWP before the PC can be used to configure the polarization state of the HPL with respect to the direction of the magnetic field (shown as the black arrow).

HPL is resonant with the RC. Additional inefficiencies in the detection systems are neglected for now and discussed later.

In this manuscript the direction of the magnetic fields will be referred to as the \hat{y} direction, while the optical axis of the system (and hence both cavities) will be defined as the \hat{z} direction, and the horizontal component with respect to the ground that is also orthogonal to the optical axis will be the \hat{x} direction. These directions are illustrated in the diagram of the optical system in Figure 2, discussed later in Section IB. The wavefronts of fields propagating along the optical axis of the system will therefore be parallel to the xy plane (ignoring their curvature). With this in mind, the magnitude of the spatial component of the field overlap will be denoted as η_{xy} . For two fields that share the same polarization state and have complex spatial field distributions $E_1(x, y, z)$ and $E_2(x, y, z)$, η_{xy} can be defined with the following overlap integral [17],

$$\eta_{xy} = \frac{|\int E_1^* E_2 dx dy|}{\sqrt{\int |E_1|^2 dx dy \int |E_2|^2 dx dy}} \quad (4)$$

$$= \frac{|\int E_1^* E_2 dx dy|}{\sqrt{P_1 P_2}}. \quad (5)$$

In this case P_1 and P_2 give the powers of the fields; however, this equation can also be applied using any complex field distribution including that of a cavity eigenmode. In the case of ALPSII, E_1 refers to the spatial mode of the HPL light propagating through the production magnet string and E_2 refers to the spatial mode of the RC.

The polarization component of η , denoted by η_{pol} , is equal to the ratio of the circulating field

in the PC that is in the ideal polarization state,

$$\eta_{\text{pol},\perp} = \frac{|E_x|}{|\vec{E}|} \quad (6)$$

and

$$\eta_{\text{pol},\parallel} = \frac{|E_y|}{|\vec{E}|}, \quad (7)$$

either perpendicular or parallel to the magnetic field for a scalar or pseudoscalar search, respectively.

The spectral component of the field overlap η_z , also referred to as the longitudinal field overlap, is discussed in more detail Section IV B 1. It is dependent on the resonance condition of the cavity and can be approximated by the cavity Lorentzian:

$$\eta_z \simeq \frac{1}{\sqrt{1 + \left(\frac{\Delta\nu}{\nu_0}\right)^2}}. \quad (8)$$

Here $\Delta\nu$ is the difference between the steady state frequency of the HPL light circulating in the PC and the nearest resonance of the RC, ν_0 is the frequency half-width-half-maximum of the cavity. In principle, the longitudinal field overlap also contains a component related to the frequency noise of the HPL light with respect to the resonance of the RC [13], however, as discussed later in this section, the contribution of this term is insignificant due to the configuration of the optical system for the first science campaign.

While ALPSII was designed to utilize both a PC and an RC, the first science campaign was performed without the PC in place for a number of reasons that will be discussed later in this section. Without the PC, P_{PC} is equivalent to

the input power of the HPL P_i in the calculation of the power measured at the output of the RC:

$$P_\phi = \eta^2 \beta_{\text{RC}} P_i \mathcal{P}_{\gamma \leftrightarrow \phi}^2. \quad (9)$$

A magnetic field length product of $563.2 \text{ T} \cdot \text{m}$ and a coupling of $4 \times 10^{-10} \text{ GeV}^{-1}$ (a factor of 20 above the ALPS II design sensitive with two cavities) would result in $\mathcal{P}_{\gamma \leftrightarrow \phi} \approx 1 \times 10^{-14}$. Therefore, using a laser wavelength of 1064 nm, with an input power of 30 W, a resonant enhancement factor in the RC of 7000, and $\eta^2 = 100\%$ would lead to roughly 1 photon per day¹ at the output of the regeneration cavity.

With the shutter open the HPL power directly after the RC is given by

$$P_{\text{open}} = \eta^2 \beta_{\text{RC}} P_i (T_1 T_2 \dots). \quad (10)$$

In this equation the factor $(T_1 T_2 \dots)$ refers to the combined transmissivity of all components mounted along the optical axis of the system in between the magnet strings and is analogous to $\mathcal{P}_{\gamma \leftrightarrow \phi}$ in the closed shutter case. In fact, under the assumption that the parameters of the optical system η , β_{RC} , and P_i do not change between open and closed shutter measurements, the conversion rate $\mathcal{P}_{\gamma \leftrightarrow \phi}$ can be expressed entirely in terms of $(T_1 T_2 \dots)$ and the ratio P_ϕ/P_{open} :

$$\mathcal{P}_{\gamma \leftrightarrow \phi} = \sqrt{(T_1 T_2 \dots) \frac{P_\phi}{P_{\text{open}}}}. \quad (11)$$

As Section ID explains, the optics mounted between the magnet strings not only generate the interference beatnotes that are used to maintain the resonance condition of the HPL with respect to the RC, but also play a role in the containment of the stray-light background. Therefore the transmissivity of these optics is very low with a measured combined transmissivity of the optics used in ALPS II of $T_1 T_2 \dots = (9.7 \pm 1.2) \times 10^{-23}$ (see Section III F for a detailed description of these measurements). Again assuming an input power $P_i = 30 \text{ W}$, a power build of

$\beta_{\text{RC}} = 7000$, and a field overlap of $\eta = 100\%$, this would produce an open shutter power just after the cavity of $P_{\text{open}} \sim 2 \times 10^{-17} \text{ W}$ or roughly 100 photons per second.

To sense such ultra-weak signals, a heterodyne detection system was used [18]. This technique works by optically mixing the weak signal field with an additional local oscillator laser to generate an interference beatnote that can then be measured with a photodiode, whereas the signal field alone would be much smaller than the photodiode's dark noise. This approach is extremely practical, as even with off-the-shelf components sensing the beatnote, heterodyne interferometry is capable of reaching the shot noise limit corresponding to a background of roughly one photon per measurement time [19].

Searching for signals on the order of one photon per day when the shutter is closed means that the suppression of the stray-light background is critical to the sensitivity of the experiment. Any stray light from the HPL that couples to the detection system can create a background that is indistinguishable from the science signal. This necessitates strict requirements on the 'light tightness' of the optical system. Therefore, the goal is to allow less than a single photon of the 30 W of input HPL power at the exact frequency of the HPL to couple spuriously to the RC over the duration of the measurement (an equivalent power of $2 \times 10^{-25} \text{ W}$ for a measurement of one million seconds). Random backgrounds with a mean value γ_{BR} significantly higher than 1 photon per measurement time will cause a $\sqrt{\gamma_{\text{BR}}}$ loss in sensitivity in terms of $\mathcal{P}_{\gamma \leftrightarrow \phi}$ and a $\sqrt[4]{\gamma_{\text{BR}}}$ loss in sensitivity in terms of $g_{\phi\gamma\gamma}$.

The exceptionally fine frequency resolution of the heterodyne detection system has the advantage of being able to exclude spurious signals from scattered light that experience more than a cycle of phase drift over the period of data taking. Therefore, for the measurement to be shot noise limited, only the power spectral density of the background integrated over the frequency resolution of the detection system must be below one photon per measurement time. Achieving such a sensitivity does not come without its own set of challenges. One of these challenges is that the phase noise between the HPL and the laser used as the local oscillator laser must be held to a fraction of a cycle over the duration of the measurement. This is explained in more detail in Section IC.

¹ Both Watts and 'Photons/time' will be used as a unit of power throughout this paper. To convert power between each set of units the equation $P = h\nu \times (\text{photons/s})$ can be used, with P representing power in Watts, h being Planck's constant, and ν the frequency of the light.

While the PC considerably boosts the sensitivity of the experiment, there are several advantages to characterizing the system without a PC initially in place. Because the PC has a nominal transmissivity² of only 2.5%, operating the optical system without it also allows for a factor of $40\times$ more laser power to be incident on the wall, reducing the integration time needed to find the ‘light leaks’. It is also important to note that the design of the system, even without the PC in place, significantly attenuates the power of the HPL incident on the RC in order to contain the stray light, making it more challenging to measure the field overlap. Its characterization is therefore much easier without the PC in place due to the increase in HPL power at the detection system. Finally, in the two-cavity design, the field overlap is maintained by a control system designed to change the length of PC to track the length changes of the free-running RC. Not using a PC significantly simplifies the control of the optical system. In this case, the RC length can be left free-running, but the feedback control loop can actuate on the frequency of the HPL instead of the position of one of the PC mirrors. Because the resonances in the piezoelectric actuator attached to the laser crystal are at much higher frequencies than the resonances of the mount used to perform the length actuation of one of the PC mirrors it is possible to actuate much faster with this setup. The faster actuation is therefore able to suppress the residual root-mean-square (RMS) frequency noise of the HPL with respect to the RC resonance to a level in which its contribution to the loss in field overlap is less than 1%. This noise is therefore not considered in this analysis because it is responsible for significantly lower losses in the field overlap than the other sources and their corresponding uncertainties.

B. Optical System

The optical system for the first science campaign was designed to achieve five goals: to maximize the (1) HPL power, (2) resonant enhance-

ment in the RC, and (3) coupling between the HPL and RC as much as possible while suppressing (4) the background signals due HPL stray light as well as (5) the relative phase noise between the HPL and the laser used as the local oscillator. In the following sections we explain how this is done.

A diagram of the optical system (top down view) used for the first science campaign is shown in Figure 2. The optical components are located in one of three clean rooms, north left (NL), central hall (CH), and north right (NR), that appear in the diagram in order from left to right. These cleanrooms are connected by the magnet strings, each of which are 116.8 m in physical length (not shown in this diagram). The magnet string on the left side of the diagram will be referred to as the production magnet string, while the regeneration magnet string is on the right side of the figure. The RC is shown on the right side of the figure, its eigenmode propagating through the bore of the regeneration magnet string. The curved mirror of the RC, housed in NR, will be referred to as RC1, while the flat RC mirror mounted to the central optical bench (COB) (discussed later in this section) will be referred to as RC2. This configuration will produce a cavity eigenmode with a waist position at RC2. A detailed description of the RC and its characterization can be found in Ref. [20]. There, a measurement of the cavity mirror g parameters is used to infer (assuming that RC2 is flat) a spatial eigenmode with waist position at RC2 having a radius of 7.2 mm along an axis oriented $11^\circ \pm 5^\circ$ with respect to the horizontal and 6.6 mm along an axis oriented $150^\circ \pm 5^\circ$ with respect to the horizontal. Likewise, the results of calculating the beam radius at RC1 for the same orientations are 9.1 mm and 9.2 mm respectively. The minimum free aperture of the magnets varied between 47 mm and 51 mm and the magnets with the widest aperture were positioned near the end stations of the experiment to ensure that the losses from the cavity circulating fields clipping on the beam tube were as low as possible [16]. In doing this, the free aperture of each of the individual magnets should be at least a factor 2.7 times the projected beam diameter at their position. Scans of the excess optical losses as a function of the RC eigenmode position on its end mirrors presented in Ref. [20] suggest the free aperture of the entire regeneration magnet

² In this document transmissivity (and reflectivity) will refer to the ratio of power transmitted (and reflected) by an optic (or an optical system like a cavity) with respect to the input power.

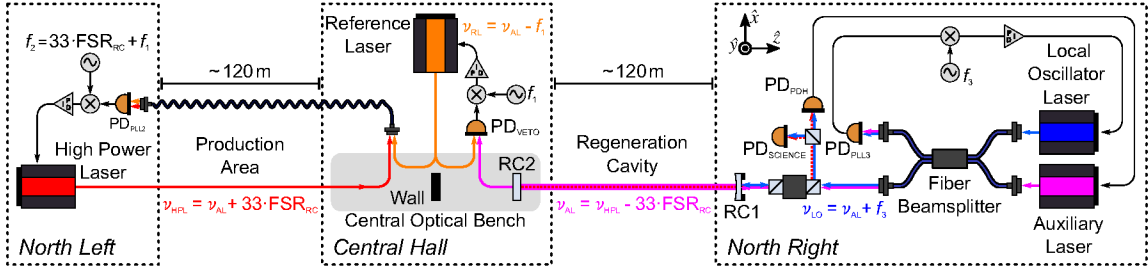


FIG. 2: Top down view of the optical system used during the ALPS II first science campaign.

string may be smaller than this due to an error in the positioning of one of the magnets. Nevertheless, this error is believed to only contribute roughly 16 ppm of extra clipping losses, well below other sources of optical losses such as scattering due to the surface roughness of the cavity mirrors. A summary of the RC parameters during the closed shutter periods of the two science runs is also shown in Table I. A more detailed explanation of these parameters can be found in Section IV A.

The HPL supplies the power that drives the production of the BSM field. Its frequency must therefore be held on a resonance of the RC. This is accomplished with a series of phase-locked loops (PLLs) using two other lasers, referred to as the auxiliary laser (AL) and reference laser (RL). Here, the frequency of the AL is stabilized to a resonance of the RC using the Pound-Drever-Hall (PDH) technique [21–23]. The RL is then interfered with the transmitted AL on the COB and phase locked to it. Finally, the HPL is phase locked to the RL. With this configuration the length changes of the RC will be encoded in the phase of the AL in transmission of the cavity and will transfer to the phase of the RL and then the HPL via the PLLs allowing the HPL to be frequency stabilized to the RC while maintaining some degree of isolation between them. This isolation is critical because the presence of stray HPL light that couples to the RC will create a background that is indistinguishable from a signal created from the regeneration of BSM fields. As Section ID describes, the COB utilizes two ‘light-tight’ boxes (not shown in Figure 2) to help block possible paths that scattered light from the HPL could take to the PC.

For all frequency stabilization loops in the optical system, whether they are PLLs or PDH

crystals provide the fast actuation at frequencies from 100 mHz to 40 kHz, while a Peltier element also mounted to the laser crystal changes the laser crystal temperature, providing the slow, long-range actuation at frequencies below 100 mHz.

A fourth laser on the NR table, referred to as the local oscillator laser (LO), is incident on the RC, but held off resonance. It is used to form the interference beatnote with the regenerated signal for the heterodyne detection system. With all of this, the HPL is held on a resonance of the RC at a frequency 33 free spectral ranges (FSRs) away from the AL.³ The system also maintains the phase stability of the beatnote between the HPL and the LO. This is critical as the regenerated signal is measured via this beatnote, by a heterodyne detection system (see Section IC) at the science photodetector (PD) in NR. A PD located in CH, referred to as the veto PD, is used to monitor the stray light contribution to the light reflected from the RC along with other parameters critical to evaluating the state of the optical system.

The first science campaigns consisted of two measurement runs. In the first measurement run, referred to as S_{\perp} , the polarization states of the lasers traveling through the magnetic fields were aligned in the \hat{x} direction, perpendicular to the magnetic fields ($\vec{E} \perp \vec{B}_{\text{ext}}$), while for the second measurement run, S_{\parallel} , the laser polarization states were aligned in the \hat{y} direction, parallel to the magnetic fields ($\vec{E} \parallel \vec{B}_{\text{ext}}$). The laser polarization states on either end of both magnetic fields could be adjusted using HWPs (not shown in Figure 2).

³ Here the FSR of the cavity is defined as $c/2L_{\text{RC}}$ with c being the speed of light in vacuum and L_{RC} being the length of the cavity

Regeneration cavity parameter	Value (S_{\perp})	Value (S_{\parallel})
Free spectral range (f_0)	$(1222632.34^{+0.18}_{-0.46})$ Hz	$(1222632.41^{+0.79}_{-0.55})$ Hz
Length (L_{RC})	$(122.601230^{+0.000030}_{-0.000017})$ m	$(122.601223^{+0.000055}_{-0.000079})$ m
Storage time (τ)	$(6.850^{+0.079}_{-0.129})$ ms	$(6.753^{+0.076}_{-0.117})$ ms
Total round-trip attenuation (A)	$(238.8^{+2.7}_{-4.6})$ ppm	$(242.2^{+2.3}_{-4.7})$ ppm
Resonant enhancement (β_{RC})	7010^{+190}_{-300}	6820^{+190}_{-270}
Power overlap (η^2)	$0.47^{+0.18}_{-0.14}$	$0.49^{+0.23}_{-0.22}$

TABLE I: Summary of the measured parameters of the RC during S_{\perp} and S_{\parallel} . All parameters in the table are dynamic that fluctuated during the science campaign. These fluctuations were observed between measurements performed during maintenance periods in the run. To assess the performance of the optical system during the closed shutter periods, the nearest set of maintenance parameters was assigned to each point of valid closed shutter data. The mean of these assigned values over the full set of valid closed shutter data is what is reported here, while the uncertainties give the combination of the peak to peak range in combination with the systematic uncertainty for each measurement. For the FSR, length, and storage time of the RC, the dynamic drifts were the dominant source of uncertainty, while for the excess optical losses and the resonant enhancement the uncertainty had roughly equal contributions from the dynamic drifts of the parameters and the systematic uncertainty of the RC mirrors transmissivities. A detailed description of the RC and its characterization can be found in Ref. [20] as well as Section IV A.

Figure 3 shows how the frequencies of the various lasers were configured for the first science campaign. The optical frequencies of the lasers are roughly 282 THz and their frequencies are all referenced to the length of the RC via the frequency of the auxiliary laser which is frequency stabilized to a resonance of the cavity. The LO is then phase locked to the AL with an offset frequency of $f_3 = 54.95$ MHz. This PLL will be referred to as PLL3. As Figure 2 shows, both lasers are coupled to optical fibers and interfered using a fiber beam splitter and therefore share the same spatial eigenmode as they propagate toward the cavity.

The COB, located at the center of the experiment is a critical component of the optical system. It is designed to help sense the relative phase changes between HPL and AL, while also suppressing the stray-light background and housing the wall/shutter system. This is discussed in more detail in Section ID. On the COB, the RL is interfered with the AL light in transmission of the RC. This beatnote is then detected with the veto PD and is used to phase lock RL via PLL1 to AL with an offset frequency of $f_1 = 12.2$ MHz to the AL light in transmission of the RC. The RL light is also interfered with the HPL light exiting the production magnet string on the COB. The combine laser

fields are then coupled to a 140 m polarization-maintaining single mode optical fiber and sent to NL where it is used to phase lock the HPL to the RL at PLL2. Here the PLL2 offset frequency is set to 33 times the FSR of the cavity (the FSR is roughly 1.22263 MHz) plus an additional 12.2 MHz offset to account for the PLL1 frequency. This ensures that the HPL is held at a frequency 33 FSRs away from the AL and thus on resonance with the RC.

The transmissivities of the cavity mirrors are configured such that the majority of the regenerated field will exit the cavity via the mirror in NR and be sensed by the science PD. This light will form an interference beatnote with the LO as it exits the RC, at the heterodyne frequency $f_s = 54.95$ MHz $- 33$ FSR ~ 14.603 MHz. Some of the regenerated light will also exit the cavity via the mirror in CH and be incident on the veto detector. The regenerated signal will interfere with AL inside the RC to form an interference beatnote at the veto heterodyne frequency of $f_v = -33$ FSR.

The phase stability of all of the interference beatnotes is critical to the ALPSII heterodyne detection scheme. Therefore, the oscillators used to demodulate the science and veto signals, as well as those used for the PLLs, are all synchronized to a global clock and left to

run uninterrupted for the duration of a science run. To set the frequencies relative to the RC FSR (referred to from here on as f_0), f_0 is measured and used to fix the beatnote frequencies before the run. Changes to the length of the RC that push its FSR away from the initial value of f_0 can introduce phase offsets to the measurement. This is discussed in more detail in Section IV B 1. Another point discussed in more detail in Appendix A 1, is that because the oscillators are driven by digital clocks, they cannot use any arbitrary frequency, but must instead use a frequency that can be represented by an integer fraction of the clock frequency.

The HPL must remain coupled to the spatial eigenmode of the RC with a frequency corresponding to one of its resonances. This can be checked by opening the shutter in the wall and allowing a small portion of the HPL light to be injected to the RC and measuring this power in transmission of the RC. Due to the highly reflective optics used to contain the stray light on the COB, the open-shutter HPL power (P_{open}) is too weak to be observed directly using a power

meter or photodetector. Instead, the amplitude of the HPL-LO beatnote must be measured at the science detector using the heterodyne detection system. The demodulation and calibration procedure is then applied to the data to give the total coupling of the HPL to the RC, along with information on the HPL's phase coherence with respect to the heterodyne detection system. During these open shutter periods the veto detector can also be used to check the spatial coupling of the HPL to the RC, as well as its phase coherence.

C. Heterodyne Sensing

ALPSII utilizes a heterodyne detection system in which the LO is used as the local oscillator laser and is interfered with the signal field exiting the RC at RC1. The power of the resulting field traveling away from the curved mirror of the cavity toward the science PD can be expressed by

$$P_{\text{out}}(t) = P_{\text{LO}} + P_{\text{s}} + 2\eta_{xy,\text{LO}}\sqrt{P_{\text{LO}}P_{\text{s}}}\cos\{2\pi(\nu_{\text{LO}} - \nu_{\text{s}})t + \theta(t)\} + \dots \quad (12)$$

In this equation, the terms that include the power of AL are not shown. P_{LO} is the power of the LO in reflection of the RC, while P_{s} is the power of the field at the signal frequency. It should be noted here that P_{s} can refer to either the power of the regenerated field with the shutter closed (P_{ϕ} , discussed in Section IA), or the power in transmission of the cavity when the shutter is open (P_{open} , discussed in more detail later in Section II). Because the LO is not resonant with the RC nearly all of its power is reflected from the cavity, while the signal field is generated inside the cavity itself and is transmitted through the RC curved mirror. In reality, there is also some residual power of AL reflected from the cavity; however due to the RC impedance matching, this power is much lower than the LO power and is ignored in this analysis. The spatial overlap between the signal field and LO is given by $\eta_{xy,\text{LO}}$. This is also equivalent to the spatial overlap between the LO and the RC as the signal field is entirely in the spatial mode of the RC. The frequencies of the fields

are ν_{LO} for LO and ν_{s} for the signal field which is identical to the frequency of the HPL. $\theta(t)$ represents the time dependent relative phase offset between the interference beatnote at $\nu_{\text{LO}} - \nu_{\text{s}}$ and the phase of the oscillator used to demodulate the data at the heterodyne frequency. It can be expressed as a static phase plus a dynamic term $\theta(t) = \theta_0 + \delta\theta(t)$. The static phase offset θ_0 is arbitrary; however, the dynamic term $\delta\theta(t)$ must be suppressed as much as possible. Here, a RMS of 0.1 rad would lead to a loss of 1% in the power measured by the heterodyne detection system [19].

Due to several diagnostic pick-off mirrors, as well as the imperfect reflectivity and aperture of the mirrors in the path to the cavity, there are some optical losses as the light travels from the RC to the science PD. The power transmissivity on this path T_{sci} , was measured to be between 80% and 90% throughout the first science campaign. The power incident on the photodetector results in a voltage output by the photodetector that is proportional to its conversion gain G .




FIG. 3: Frequencies of the lasers and PLLs.

The static voltage is removed with an analog high-pass filter, suppressing the terms P_{LO} and

P_S in Equation 12, and the resulting voltage, approximated by

$$V_{ADC}(t) \simeq 2\eta_{xy,LO} T_{sci} G_{AC} \sqrt{P_{LO} P_S} \cos \{2\pi(\nu_{LO} - \nu_S)t + \theta(t)\}, \quad (13)$$

is digitized by the analog-to-digital converter (ADC) with a sampling frequency $f_{m_1} = 500$ MSa/s. Here G_{AC} refers to the total conversion gain of the electronic chain of the science PD and its electronic filters at the heterodyne demodulation frequency (~ 14 MHz). The amplitude of the voltage signal will be referred to as $V_{AC} = 2\eta_{xy,LO} T_{sci} G_{AC} \sqrt{P_{LO} P_S}$, while its frequency will be denoted by $f_S = \nu_{LO} - \nu_S$. While V_{AC} in this equation refers to the amplitude of the deterministic signal measured at the science PD, from here on the measured value of V_{AC} is assumed to include some measurement noise.

The amplitude of this beatnote is the primary measurable for the experiment as it can be used to deduce the signal power P_S . The measurement of V_{AC} is complicated by the fact that the phase of the beatnote (which is assumed to be static) is unknown. Therefore, the time series of

the beatnote must be projected onto both cosine and sine components. Here, the ‘in-phase’ term I_{DM1} is used to refer to demodulation via a cosine function, while the ‘quadrature’ term Q_{DM1} refers to demodulation with a sine function. To avoid offsets being introduced by asymmetries in the demodulation process, a technique referred to as double demodulation is used [19]. This process is shown in the upper gray box of the flow chart in Figure 4. Here, the first stage of demodulation is performed digitally at the nominal frequency difference between the two lasers plus a 2.4 Hz offset: $f_S + 2.4$ Hz. By expressing the data series resulting after the first demodulation $H_{DM1}[m]$, as the complex number $I_{DM1} + iQ_{DM1}$ the result of the first demodulation can be treated as a series of points on the complex plane:

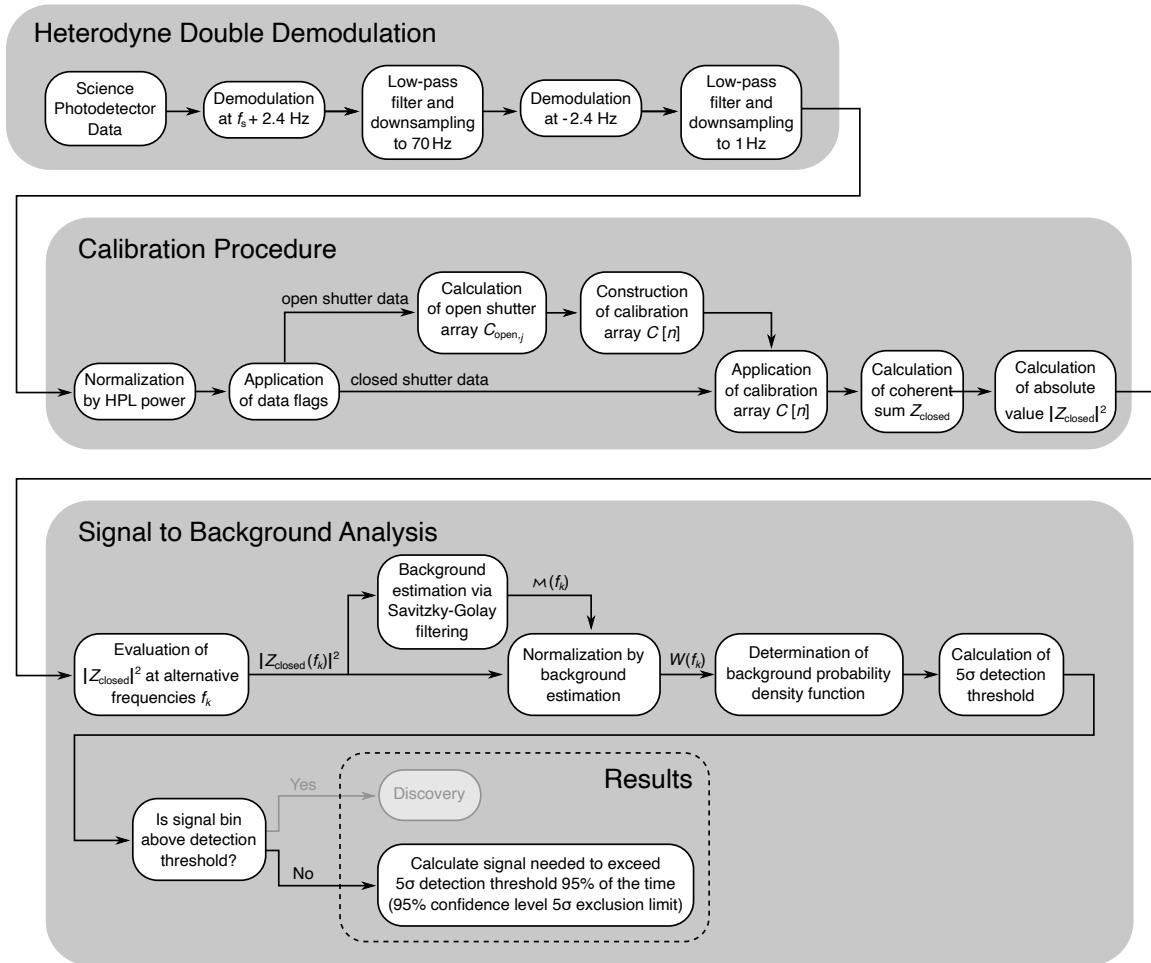


FIG. 4: Flow chart of the analysis for the science PD data. The data first undergo the process of double demodulation by the heterodyne detection system. The calibration procedure is then applied to the resulting data series to calculate the final result at the heterodyne signal frequency. The background is then estimated by performing the same analysis at alternative frequencies and a detection threshold is calculated based on the statistics of the background at these other frequencies. A discovery can be claimed if the signal at the expected frequency is above the detection threshold. If the signal is below the detection threshold, as was the case in the case in both runs of the first science campaign, exclusion limits can then be calculated.

$$H_{\text{DM1}}[m] = \overbrace{V_{\text{ADC}}[m] \cos \left\{ 2\pi(f_s + 2.4 \text{ Hz}) \frac{m}{f_{m1}} \right\}}^{I_{\text{DM1}}} + i \overbrace{V_{\text{ADC}}[m] \sin \left\{ 2\pi(f_s + 2.4 \text{ Hz}) \frac{m}{f_{m1}} \right\}}^{Q_{\text{DM1}}} + \dots \quad (14)$$

$$= V_{\text{AC}} \cos \left\{ 2\pi m \frac{f_s}{f_{m1}} + \theta[m] \right\} e^{i2\pi(f_s + 2.4 \text{ Hz}) \frac{m}{f_{m1}}} + \dots \quad (15)$$

$$= \frac{V_{\text{AC}}}{2} \left(e^{i \left(2\pi m \frac{f_s}{f_{m1}} + \theta[m] \right)} + e^{-i \left(2\pi m \frac{f_s}{f_{m1}} - \theta[m] \right)} \right) e^{i2\pi(f_s + 2.4 \text{ Hz}) \frac{m}{f_{m1}}} + \dots \quad (16)$$

$$= \frac{V_{\text{AC}}}{2} e^{i \left(2\pi m \frac{2.4 \text{ Hz}}{f_{m1}} - \theta[m] \right)} + \dots \quad (17)$$

An finite impulse response (FIR) filter is then used to down-sample I_{DM1} and Q_{DM1} to the desired data rate of $f_m = 70$ samples per second. Terms at higher frequencies, such as those at the doubled frequency $2f_s$, or at the PLL3 frequency, are not relevant as they are eliminated by this process. Also, the result will be the same whether the demodulation is expressed in terms of the arrays $I_{\text{DM1}}[n]$ and $Q_{\text{DM1}}[n]$ or using the complex amplitude and phase $\frac{V_{\text{AC}}}{2} e^{i2\pi(f_s + 2.4 \text{ Hz}) \frac{n}{f_m}}$. To reduce the number of terms to keep track of in the second demodulation we used the formulation on the complex plane. The science signal is thus a complex 2.4 Hz oscillation in this data.⁴

The second demodulation is then applied in post-processing by multiplying the result of the first demodulation with $e^{i2\pi(-2.4 \text{ Hz}) \frac{n}{f_m}}$. Here the sign of the frequency matters because the data are now complex. Performing the additional demodulation in post-processing helps mitigate the effects of static offsets in the result of the first demodulation stage. More details on this process can be found in Ref. [19]. After the data series is demodulated for the second time the data is down-sampled to the final data rate of 1 sample per second using another FIR

filter. An investigation into the down-sampling showed no appreciable difference for measurements derived with the down-sampled data using this process versus the data at 70 samples per second.

The full process of two stages of demodulation, then filtering and down-sampling produces what we will refer to as the ‘ H -function’ or,

$$H[n] = \frac{V_{\text{AC}}}{2} e^{-i\theta[n]}. \quad (18)$$

In this equation $H[n]$ is a data series with an index n that also happens to be n seconds into the run due to the data rate of 1 Hz. While it is not explicitly shown in the equation above, each point in the series will also have a random value drawn from a probability distribution on the complex plane to account for the measurement noise in the system.

The expectation value of $H[n]$ for a given sample is a complex number with an amplitude given by half the amplitude in voltage of the beatnote between the LO and the signal field at that time and a phase equal to the phase difference between the beatnote and the oscillators used to perform both demodulations:

$$\langle H[n] \rangle = \frac{V_{\text{AC}}}{2} e^{-i\theta_0} \quad (19)$$

$$= \eta_{k,y,\text{LO}} T_{\text{sci}} G_{\text{AC}} \sqrt{P_{\text{LO}} P_{\text{S}}} e^{-i\theta_0}. \quad (20)$$

Figure 5a shows a representation of $H[n]$ on the complex plane. Here the ‘cloud’ is the probability distribution representing the measurement noise. As the figure shows, the sum of the signal

⁴ 2.4 Hz is chosen as the offset frequency as it is well within the Nyquist frequency of the downsampled sampling frequency at 70 Hz and also shares a greatest common divisor with this frequency of less than one. This is meant to reduce potential effects from aliasing when the signal is down sampled [19].




FIG. 5: (a) $H[n]$ represented on the complex plane. (b) Z -function represented on the complex plane after summing over 100 coherent points of $H[n]$.

and the cloud is offset from the origin by a blue vector which for $H[n]$ has an amplitude of $\frac{V_{AC}}{2}$ and a phase of $-\theta_0$. As mentioned above, the actual values of θ_0 from the runs are unknown and do not affect the results; however, it is extremely important to understand the dynamic component of the phase $\delta\theta$, as it could lead to decoherence of the appearance of the signal in the data.

The values of $H[n]$ can then be integrated to produce what is referred to as the Z -function as the equation

$$Z = \frac{1}{N} \sum_{n=1}^N H[n] \quad (21)$$

shows. The Z -function is equivalent to calculating the average value of $H[n]$. Thus, as $H[n]$ is summed, the expectation value of Z remains constant. If the measurement noise present in

V_{AC} (and hence $H[n]$) is uncorrelated, the variance of Z decreases with $1/N$ (this is shown experimentally later on in Figure 18). The reason is that even though the variance in the sum $H[n]$ increases linearly with N , when calculating the variance of Z , the variance of the sum is multiplied by the square of the prefactor, or $1/N^2$. Figure 5b shows the result of calculating Z from a series of 100 values of $H[n]$ with an expectation value and uncertainty equivalent to Figure 5a. Here, the area on the complex plane that contains some fraction of the probability distribution in Z is a factor of 100 less than the area that contains the same fraction of the probability distribution for individual values of $H[n]$.

The power of the signal leaving the RC can then be expressed as the absolute square of the Z -function using the equation

$$P_s = \frac{|Z|^2}{\eta_{xy,LO}^2 T_{sci}^2 G_{AC}^2 P_{LO}}, \quad (22)$$

assuming that the relative phase $\theta(t)$ is constant, and ignoring the effects of photon counting statistics and technical noise on Z . If the measurement is limited by photon counting statistics, the expected measured photon rate will be one photon per measurement time. Therefore the sensitivity of the system in terms of power will improve linearly with integration time. If a signal is present above quantum noise,⁵ the signal-to-noise ratio of the measurement of $|Z|^2$ and thus P_s will be equal to the expected number of photons measured over that period of time in the signal field [18].

For the heterodyne detection system to work properly, the phase difference θ between the LO, the signal field, and the oscillators used in the double demodulation scheme must remain stable during the period of time in which the summation takes place. This is one of the most significant tasks of the optical system. As mentioned earlier, the phase coherence between these fields can be directly monitored during the open shutter periods by calculating Z for these periods and examining the stationarity of its phase. When the shutter is open, the signal

⁵ Throughout this manuscript the phrases quantum noise and shot noise will refer to the noise due to photon counting statistics.

field is the HPL light incident on the RC after passing through the COB optics.

With a total run time of more than a month for S_{\parallel} the frequencies of the PLLs and demodulation stages must be known relative to each other with nHz accuracy to ensure that the error in these frequencies does not contribute to significant phase noise in the measurement.⁶ It is therefore important to understand how these signals are actually represented by the digital oscillators producing them. In all cases, the signals are supplied by function generators with a clock rate of 1 GHz and 48 bit precision.⁷ This means that the actual frequency of any signal output by these devices must be some integer multiple of $1 \text{ GHz}/2^{48}$ ($\sim 3.5527 \mu\text{Hz}$). To be clear, this does not prevent the frequencies from being *known* with nHz precision, only that they cannot be set with that level of precision, which is unimportant. This difference is discussed in more detail in Appendix A 1.

D. Central Optical Bench

The COB has three primary functions in the experiment. It houses the wall and shutter system, suppresses the spurious coupling of HPL light to the RC, and supports the transfer of the phase information of the AL light transmitted by the RC to the HPL even when the shutter is closed. Figure 6 shows the design of the COB for the first science campaign. In this diagram the thicker black edge of the mirrors indicate their reflective surface. The main practical differences between this layout and the COB design for the full optical system with a PC is that this layout does not have a PC mirror on it. Thus, up to 30 W of HPL power is available on the COB, when only 1 W would be possible if the PC were in place.

The shutter system is a remotely operable wheel in which different materials can be used

to fill slots and, depending on the position of the wheel, act as the wall. One of the slots must be left unoccupied for the open shutter periods. The upper diagram of the COB in Figure 6 shows the system when the shutter is open, while the lower diagram shows the optical axis of the COB when the shutter is closed. Here, the beam dumps (BDs) BD1 and BD2 contain the HPL stray light when the shutter is closed and open respectively. For S_{\perp} , a silver mirror that directed the residual HPL light to a beam dump was used as the wall. Between the S_{\perp} and S_{\parallel} run, two other shutter materials were implemented, including a piece made from solid aluminum, and another solid aluminum piece covered in foil coated with a black ultra-highly-absorptive layer. Subsequent testing indicated that the aluminum piece with the black foil provided the best suppression of stray light over short test phases and therefore it was used as the wall for S_{\parallel} .⁸ No matter what material was used for the wall no light is expected to directly transmit it as the silver mirror should provide over 25 *orders of magnitude* suppression of the directly transmitted light (in addition to the suppression of the COB optics), while the solid aluminum pieces should provide more than 100,000 *orders* of suppression.

The COB has two light-tight boxes that subdivide it into three areas. The HPL area on the left, the RC area on the right, and the intermediate RL area in between them. The purpose of this design is to contain the stray light from the HPL as well as possible and try to prevent it from coupling to the RC, while also allowing the HPL to be phase locked to the AL. The light-tight boxes therefore cannot be completely sealed in order to allow for the stabilization of the HPL frequency as well as open shutter measurements. To accommodate this they are equipped with the highly reflective (HR) mirrors LT1 and LT2. Because there are also light-tight boxes outside the vacuum system connected to the exit flanges of the light-tight boxes inside the vacuum system, the only way in and out of them is through the magnet bore or LT1 and LT2. Each has a transmissivity of $12.4 \pm 0.4 \text{ ppm}$ for

⁶ While roughly 1,000,000 seconds of valid data were acquired during S_{\parallel} , the experiment ran for over 4,000,000 seconds during this time frame. Therefore, 0.1 rad would translate to roughly a 4 nHz error in the frequencies used.

⁷ It should also be noted here that all devices used in the experiment are synchronized to the same global clock.

⁸ This may also have been related to the fact that only one of the five open slots on the filter wheel was occupied by a shutter for S_{\perp} and all but one was occupied for S_{\parallel} . This is under further investigation.

p -polarized (\hat{x} direction) light with a wavelength of 1064 nm at room temperature using an angle of incidence of 35° . The low transmissivity was chosen to help suppress the stray light signals that pass between the areas of the experiment.

To keep the HPL and the RC as separated as possible, RL is split into two paths to perform the transfer of phase information from the AL to the HPL. This splitting is done by a Mach Zehnder (MZ) interferometer. The beam is initially split by the mirror MZ1. Any changes in the differential path-length between the two paths from MZ1 to MZ4 would lead to a phase offset between the heterodyne local oscillator field and the regenerated signal, as well as the open shutter signal. For this reason the MZ mirrors are mounted to an ultra-low expansion (ULE) glass plate. The MZ alignment was performed with the COB installed in the central vacuum tank while it was vented. This was done using a CCD outside the vacuum system to observe the position of the beam in transmission of one of the MZ paths, while blocking the other path. With this process, it was possible to align the paths to within $\pm 10 \mu\text{rad}$ of each other.

The RL light for the RL-AL interference is reflected by MZ1 ($R_{\text{MZ1}} = 1.51 \pm 0.02\%$) and then passes through LT2 via reflections from two HR mirrors, MZ2 and MZ4. With 200 mW of RL power injected to the COB at the ‘Injection Port’, this leaves roughly 40 nW of power incident on the RC at the mirror RC2, where most of it is reflected. With 10 μW of power in AL in transmission of RC, the interference beatnote formed here has an amplitude on the order of 1 μW .⁹ This light is then directed to the out-of-vacuum veto PD which is used to perform the sensing for PLL1.

To form the interference beatnote with HPL, RL passes through MZ1 and MZ3 ($T_{\text{MZ3}} = 452 \pm 10 \text{ ppm}$) where it interferes with the HPL field transmitted by LT1. At this point the RL power is 90 μW while the HPL power here with 30 W injected to the ‘PC port’ is 370 μW , leading to an interference beatnote with an amplitude on the order of 360 μW . This light is sent out of the vacuum system at the ‘PLL2 port’ and coupled

to the 140 m fiber that transfers the light to the PLL2 photodetector in NL.

Because tests of the backgrounds before the first science campaign did not show any dependence on the HPL polarization, the HPL light was set to be p -polarized upon incidence on LT1 to provide the maximum signals for the interference beatnotes and open shutter periods, as all on axis COB mirrors were more transmissive for p -polarized light. All of the lasers in the setup are therefore tuned to nominally be p -polarized when incident on the 35° angle of incidence mirrors on the COB. For this reason, the transmissivity values quoted in this manuscript are all for p -polarized light unless stated otherwise.

For S_\perp , the HPL light injected to the production magnet string was already p -polarized so no other optics were necessary to achieve this and HWP₀ was not installed. For S_\parallel , the light injected to the production magnet string was s -polarized and HWP₀ was oriented to rotate the polarization of the HPL light by 90 degrees. For S_\perp HWP₂ was tuned to have no effect on the AL light in transmission of the cavity as it was already p -polarized (oriented in the \hat{x} direction). For S_\parallel , because AL was s -polarized (oriented in the \hat{y} direction), HWP₂ was adjusted to rotate the AL light in transmission of the RC to be p -polarized. HWP₁ and HWP₂ were mounted to remotely rotatable stages in order to periodically invert the sign of specific stray light paths and break their coherence, however tests of this procedure showed its impact on the stray light was only marginal and therefore it was not used during the science runs to avoid the risk of introducing additional systematic effects.

When the shutter is open, the HPL must pass through mirrors LT1, MZ3, MZ4 ($T_{\text{MZ4}} = 275 \pm 10 \text{ ppm}$), LT2 before being incident on the RC. The combined transmissivity of these optics, and hence T_{COB} , was measured to be $T_{\text{COB}} = 1.9 \pm 0.1 \times 10^{-17}$. These measurements are discussed in more detail in Section III F. With the shutter open, the HPL incident on the RC was roughly $6 \times 10^{-16} \text{ W}$ when 30 W of HPL power was used. Including mirror RC2, the on-axis optical attenuation between the two magnet strings was $T_{\text{COB}} \cdot T_{\text{RC2}} = 9.7 \pm 1.2 \times 10^{-23}$ when the shutter was open. When considering the transmissivity of the RC and the coupling of the HPL to its eigenmode, on the order of 10 aW of HPL power is expected at the science detector when the shutter is open.

⁹ Assuming unity spatial overlap, the interference beatnote amplitude in power P_{AC} , for two fields with powers P_1 and P_2 can be found using $P_{\text{AC}} = 2\sqrt{P_1 P_2}$.

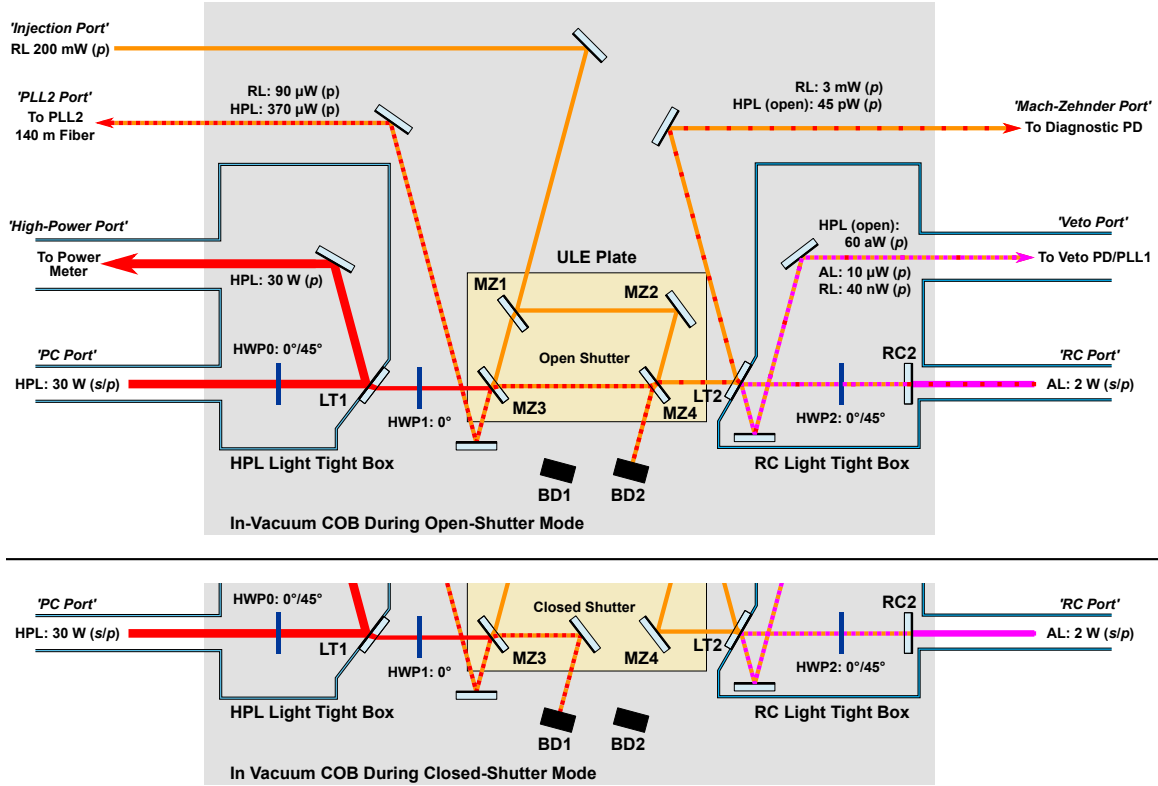


FIG. 6: The COB layout for the first science campaign during open-shutter mode is shown in the upper diagram with the lower diagram showing the optical axis of the COB during closed shutter mode. Here, the lasers powers at different points on the COB are inexact, but represent typical values measured during the science runs.

With RL aligned to AL in transmission of the cavity, RL serves as a reference of the position of the cavity eigenmode along the optical axis after MZ4. Therefore, a diagnostic PD at the output of the ‘Mach-Zehnder Port’ in reflection of LT2 was used to help align the HPL to RL where the power of the HPL was much higher than after LT2.

The HPL light in reflection of LT1 is directed out of vacuum at the ‘High-Power Port’, to a power meter that measured the HPL power continuously during the science runs. Tests over short time periods showed no apparent difference in the background levels when this light was incident on the power meter versus a number of other beam dumps. The power meter is housed in an out-of-vacuum light-tight box mentioned earlier in this section, which connects to the vacuum flange at the High-Power port to try to limit any stray light from scattering into the

central cleanroom.

II. CLOSED-SHUTTER CALIBRATION

As explained in Section IA, the conversion rate of an electromagnetic field passing once through one of the magnet strings and converting to a BSM field (and vice versa) can be expressed using measured parameters as¹⁰

$$\mathcal{P}_{\gamma \leftrightarrow \phi} = \sqrt{\frac{P_{\phi}}{\eta^2 \beta_{\text{RC}} P_i}}. \quad (23)$$

¹⁰ Here the double pass conversion rate, $\mathcal{P}_{\gamma \rightarrow a \rightarrow \gamma}$, of an electromagnetic field converting to a BSM field in the magnet string before the wall and reconverting to an electromagnetic field in the magnet string after the wall is simply $\mathcal{P}_{\gamma \rightarrow a \rightarrow \gamma} = \mathcal{P}_{\gamma \leftrightarrow \phi}^2$.

Assuming there are no effects that alter the power or pointing of the HPL outside of the transmissivity on the optics on the COB, the equation

$$P_{\text{open}} = \eta^2 \beta_{\text{RC}} T_{\text{RC2}} T_{\text{COB}} P_1, \quad (24)$$

gives the HPL power in transmission of the RC when the shutter is open. Here, Equation 10 from Section IA was modified to include the specific transmissivities for the factor ($T_1 T_2 \dots$). In this case T_{COB} is the combined transmissivity of the four mirrors oriented at 35 degrees on the COB (for p -polarized light), while T_{RC2} is the transmissivity of the flat cavity mirror on the COB (RC2 in Figure 2). The factor $\beta_{\text{RC}} T_{\text{RC2}}$ gives the nominal transmissivity of the RC for a laser on resonance with a perfectly spatially coupled beam. The open shutter power is also dependent on the total field overlap η . The spatial and spectral components of η will determine the coupling of the HPL to the RC. Because the COB is designed to be significantly more transmissive for the properly configured polarization state, light in the incorrect polarization state will not pass through the COB optics. Thus, the open shutter power will depend on the polarization component of η .

With this the equation for $\mathcal{P}_{\gamma \leftrightarrow \phi}$ can be further simplified using the measured open shutter power:

$$\mathcal{P}_{\gamma \leftrightarrow \phi} = \sqrt{T_{\text{COB}} T_{\text{RC2}} \frac{P_\phi}{P_{\text{open}}}}. \quad (25)$$

This shows that if the open shutter power could be predicted for a given period of closed shutter data, $\mathcal{P}_{\gamma \leftrightarrow \phi}$ could be calculated in terms of the transmissivity of the optics between the two magnetic field regions and the mean of the ratio of the measured power when the shutter is closed versus the predicted open shutter power. This is critical as the transmissivity of the optics on the COB and the flat cavity mirror are essentially static parameters that do not need to be measured during data taking. Therefore the only dynamic variables which must be measured during the run are P_ϕ and P_{open} .¹¹

In using the ratio P_ϕ/P_{open} to calibrate the results, we rely on the assumption that the COB optics will only attenuate the power of the HPL and not perturb the spatial eigenmode with respect to the eigenmode of the BSM field. There are a number of systematic effects that could invalidate this assumption and, as Section II A explores, several checks were in place specifically to ensure that this was not the case.

The open and closed shutter powers cannot be measured simultaneously. However, because the misalignment of the system due to drifts typically only became significant over several days, it was possible to calibrate the closed shutter data using the nearest neighboring open shutter periods. Open shutter measurements were performed daily during S_\perp and several times per day for S_\parallel , to minimize the drifts between open and closed shutter periods. For S_\perp this corresponded to 19 open shutter periods during the 15 day run, while for S_\parallel 78 open shutter periods were performed over 41 days. The average time between the closed shutter data and its nearest open shutter period was roughly 8 hours for both runs. A critical component of this is ensuring that the system was in a stable state during the periods in which data was acquired for both the times when the shutter was open or closed. The process of identifying these times and flagging the data based on the status of the system is explained in Section II B.

Because it was possible to continuously measure P_{HPL} , the closed and open shutter data was also normalized by the HPL power to further reduce the systematic uncertainty of the measurement, and this is described in Section II C. The rest of the calibration procedure to calculate the result from the heterodyne detection system in terms of Z and then calibrate this in $\mathcal{P}_{\gamma \leftrightarrow \phi}$ is explained in Section II D. A second PD called the veto detector, was implemented on the opposite side of the RC from the science detector, to verify potential signals and we examine this in Section II E. The uncertainty in the calibration of the ratio P_ϕ/P_{open} and the results in terms of $\mathcal{P}_{\gamma \leftrightarrow \phi}$ is then scrutinized in Sections II F and II G respectively.

It should be noted that the purpose of the calibration procedure is to provide the results in terms of the electromagnetic field to BSM field conversion rate $\mathcal{P}_{\gamma \leftrightarrow \phi}$ and therefore this manuscript does not consider the effect of the magnet strings on the coupling of the BSM

¹¹ Other dynamic variables such as the HPL power injected to the production string were measured during the run, but under these assumptions they are not strictly necessary to calibrate the results in terms of $\mathcal{P}_{\gamma \leftrightarrow \phi}$.

fields. Exclusion limits, derived from the results presented in Section III D, on the coupling between electromagnetic fields and pseudoscalar, scalar, vector, and tensor bosons can instead be found in Ref. [1].

A. Potential Systematic Effects

There are several effects that could spoil the calibration of the regenerated signal via the open shutter signal. For example, any wedge angle in the optics between the two magnetic fields will deflect the light that transmits from the production area to the regeneration cavity during the open shutter periods. The BSM fields however, will be unaffected by these wedge angles during the closed shutter periods. Therefore, if the system is aligned by optimizing the open shutter signal, compensating for the wedge angles would lead to a misalignment of the BSM field with respect to the cavity eigenmode. The wedge angles of the COB substrates were measured by the manufacturer to be between 1.7 to 5.5 μrad , producing a deflection angle on the beams between 0.8 to 2.8 μrad .

The total deflection angle of the COB due to the wedge angles of the on-axis optics was measured by two different methods. First, a green He-Ne laser was directed through the beam tube from NR to NL through the COB and its position was recorded on a camera. Then, the COB was removed and the change in the position of the green beams on the camera was measured. This gave a deflection angle of $10 \pm 5 \mu\text{rad}$ in the horizontal direction and $5 \pm 5 \mu\text{rad}$ in the vertical direction for a total deflection angle of $11 \pm 7 \mu\text{rad}$. The deflection angle of the COB optics was also measured using an autocollimator once the COB was removed. In these measurements the autocollimator beam passed through the COB optics, was then normally incident on a mirror, and then once again passed through the COB optics, before being incident on the autocollimator detection system. This procedure was then repeated several times pushing the COB in and out of the setup, but always being normally incident on the mirror. The difference of the angle when the COB was in the setup versus when it was not was measured by the autocollimator to be $9 \pm 2 \mu\text{rad}$ in the horizontal direction and $1 \pm 1 \mu\text{rad}$ in the vertical direction. This would correspond to a total COB deflection

angle of $5 \pm 1 \mu\text{rad}$ and is within the error of the measurements using the He-Ne laser. Considering that the COB optics were mounted with a random azimuthal orientation, this is consistent with the wedge angles measured by the manufacturer. If the system were ideally aligned with the HPL perfectly coupled to the RC during the open shutter periods, such a deflection angle would lead to an additional loss of the spatial power overlap of the regenerated field to the RC eigenmode of 2%. If there were an additional 10 μrad offset in the alignment of the HPL with respect to the RC eigenmode along the same axis as the wedge angle (corresponding to a loss in the spatial overlap of 8%), this would correspond to an additional uncertainty in the spatial power overlap of between 6 and 10%. Therefore, an additional 10% uncertainty has been incorporated into the closed to open shutter power ratio because such misalignments were typical during the first science campaign.

Another effect which could reduce the measured signal power is the loss in phase coherence between the regenerated field and the heterodyne detection system due to optical path length changes on the COB. One cause of this could be temperature changes in the substrates on the COB that the various lasers pass through on their way to forming interference beatnotes for the PLLs. The COB was originally designed to reduce the influence of this effect as the faces of the optics are oriented such that, after being split, the RL will pass through two optics on both paths before it interferes with the HPL and the AL in transmission of the RC.¹² Thus, if these optics experience common changes in their temperature, their optical path lengths will all change by the same amount. PLL1 will adjust the phase of RL to compensate for the changes in the optical path-length through the mirrors LT2 and RC2. This should be nearly canceled by the optical path-length changes in the other RL path through mirrors MZ1 and MZ3.¹³ The HPL only passes through the mirror before interfering with RL. Therefore only its path-

¹² On the path to the point of interference with AL one of these optics is oriented at 0° as opposed to the others which are oriented at 35° and a more rigorous calculation should include this additional difference in path length.

¹³ The word nearly is used here again due the path length difference through the 0° mirror.

length changes are encoded on the phase of the HPL, and hence the BSM field, with respect to the heterodyne detection system.

Although this may appear to be ideal, that is, suppressing effects that lead to any decoherence between the bosonic field and the heterodyne detection system, it must also be considered that the open shutter path passes through all on axis substrates on the COB. Therefore, in this configuration, this effect will cause the PLLs to detune the open shutter phase with respect to the free space phase determined by the path length between the reflective interfaces of mirrors MZ3 and RC2. If the open shutter signal is then used to compensate for the potential phase changes in the system this effect will lead to a loss in coherence of the regenerated signal with respect to the heterodyne detection system. While this effect should produce a measurable change on the phase of the open shutter signals during the run, currently no way has been found to disentangle this effect from other effects that the open shutter calibration technique should compensate for. These other effects could be phase changes induced by misalignments in RL or the alignment of the RL-HPL beatnote to the CH-NL optical fiber. If these signals were present and large enough to cause a decoherence in the open shutter data, this would appear as changes in the phase of the open shutter data unless they were perfectly canceled by other effects. Because the phase changes in the open shutter data were not large enough to produce a significant decoherence in the data, it seems this was not actually a problem for the optical system. Still, it was important to verify that the optical path length changes induced by temperature changes in the on axis COB optics are significantly less than a wavelength.

The optical substrates each have a thickness of 9.5 mm. Assuming an index of refraction for fused silica at a wavelength of 1064 nm of 1.45, each of the 35° substrates have an optical path length of 15.0 mm,¹⁴ while the flat cavity mirror has an optical path length of 13.8 mm. Assuming a thermo-optic coefficient of 8.5 ppm/K for fused silica at room temperature [24], and a coefficient of thermal expansion of 0.5 ppm/K [25],

¹⁴ While the angle of incidence is 35°, the angle the on-axis light propagates through the substrate is only 23.3° due to Snell's law.

the optical path length through the 35° mirrors will change by 95 nm/K, while for the cavity mirror it will change by 87 nm/K. Together, this will produce a 0.35 cycle/K phase change in the open shutter signal with respect to the free space optical path length.

Measurements of the temperature in the central cleanroom showed the RMS deviations from the mean were less than 0.1 K over the duration of the closed shutter periods of S_{\perp} and less than 0.07 K during the closed shutter periods of S_{\parallel} . Therefore, the effects of path length changes in the substrates are expected to be less than 0.035 cycles for S_{\perp} and 0.025 for S_{\parallel} . In terms of the losses in the measured regenerated power, this would correspond to 4.8% for S_{\perp} and 2.4% for S_{\parallel} . Each of these values have been included as an additional uncertainty in the closed to open shutter ratio in their respective runs.

Other effects that could change the phase of the open shutter signal were also investigated. They included misalignments of RL with respect to the HPL as well as alignment errors between these lasers and the long fiber used to transfer the RL-HPL beatnote to NL. Both of these effects were shown to have an influence on the open shutter phase of the HPL measured at the veto detector. Based on the misalignments experienced during the run and tests performed afterward, the alignment noise in these paths could have been responsible for phase changes as high as 0.1 cycles during the run, although the unexplained phase changes in the run were not that large. Nevertheless, when misalignments of the optical system cause changes in the phase of the open shutter data, the open shutter calibration will compensate for this and therefore these effects were not considered when calculating the uncertainty in the open shutter calibration. A more detailed discussion can be found in Section IID.

The calculation of the error bars on the open-to-closed shutter ratio also does not include effect from the wedge angles and path length changes of the half-wave plates on the COB as they are specified to have a parallelism better than 0.5 μ rad, with a thickness of less than 3 mm. Also, differential changes in free space path length of the arms of the MZ are not believed to be significant with respect to the other effects mentioned in this section and are therefore also not considered in the calculation of the error.

B. System Locked Flags

For the closed or open shutter data to be considered valid, the magnets must be operating at full current and all control loops must be engaged and stable at the correct frequencies. The loops include the PDH, PLL1, PLL2, and PLL3 control loops. The status of the system was recorded in the form of a number of ‘flags’. In addition to being stably locked, i.e. the control system engaged with a stable error signal, it was also important for the frequencies of all PLLs to have the correct sign such that the HPL would be resonant in the RC and the beatnote between the HPL and LO would be at the proper frequency. The sign of the frequency difference between the two fields generating the beatnote was checked before each time the system was locked by varying the temperature of one of the laser crystals with the PLL disengaged and examining the change in the beatnote frequency. Because the sign of the temperature to frequency calibration of the laser is known, it was possible to tell which laser had the higher frequency based on whether or not the beatnote frequency increased or decreased.

The flag for PLL1 was based on monitoring the frequency of the PLL1 beatnote. If it deviated from its nominal value the flag would mark the data as invalid. We should note though that for S_{\perp} this system was not yet implemented and the PLL1 flags are only used for S_{\parallel} . Based on the flag data from S_{\parallel} , it does not appear that there were any sections of closed shutter data during S_{\perp} in which PLL1 was unlocked, but that the data during that stretch was otherwise flagged as valid. The reason for this is that when the PLL1 lock is broken, this almost always breaks the PLL2 lock as well. Furthermore, examining the stray-light signal over the duration of the run indicates that the system flags for S_{\perp} are representative of the system status, as if the system is unlocked the stray-light signal disappears and the stray-light signal is ever present in the closed shutter data of S_{\perp} .

It was possible to distinguish if the PLL2 and PLL3 loops were unlocked by observing the control signals sent to the Piezoelectric and temperature actuators mounted to the laser crystal. The flags for these loops were therefore based on their control signals. Furthermore, the control signals were also inspected for times when they may have unlocked and then relocked on their

own and the data from the time periods after the initial unlock were disregarded.

The flag for the magnet current was set to be valid if the magnet current was larger than 5.6 kA. The magnet current, however did not vary over the duration of the run and was always within 0.0003 kA of 5.7118 kA during the closed shutter periods.

The open and closed shutter status flags were based on the recorded position of the shutter wheel such that it had to be oriented within 3 degrees of the center position of the open hole for the open status, or one of the closed holes for the closed status.

There was also a flag for the case where the operator was performing ‘maintenance’ on the system. Maintenance involved taking measurements used for monitoring purposes as well as realigning the optics. In this case the flag was manually set to be invalid if maintenance was being performed as all data acquired during this period are disregarded.

The Boolean product of each set of flags was used to form the complete set of flags for the shutter open and shutter closed periods. Out of caution, the complete flags were cut by 60 s such the data from the first and last minute of every open and closed shutter section were disregarded.

Over the course of S_{\perp} more than 580,000 s of closed shutter data with valid flags was collected, while for S_{\parallel} the total data from the closed shutter periods consisted of over 1,060,000 s of valid data.

C. Normalization by HPL Power

Recalling Equations 9 and 24, the HPL power exiting the RC in NR during the closed (P_{ϕ}) and open (P_{open}) shutter periods is

$$P_{\phi} = \eta^2 \beta_{\text{RC}} P_i \mathcal{P}_{\gamma \leftrightarrow \phi}^2 \quad (26)$$

and

$$P_{\text{open}} = \eta^2 \beta_{\text{RC}} T_{\text{COB}} T_{\text{RC2}} P_i. \quad (27)$$

As stated earlier, T_{COB} and T_{RC2} are static parameters that did not need to be measured constantly during the run, while the electromagnetic-BSM field conversion rate $\mathcal{P}_{\gamma \leftrightarrow \phi}$ is the measurable for the experiment. The total field overlap η and the resonant

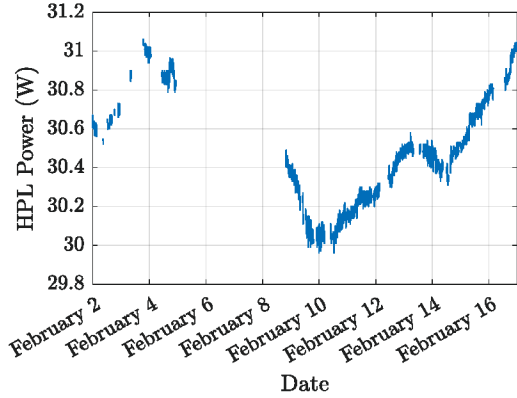

(a) HPL power (S_{\perp})(b) HPL power (S_{\parallel})

FIG. 7: The HPL power measured exiting the production magnet string in CH during the open and closed shutter measurement periods of S_{\perp} and S_{\parallel} .

enhancement β_{RC} are not static parameters and depend on attributes of the optical system, such as the alignment of the HPL and storage time of the RC, that could change throughout the run.

P_i is also a dynamic variable, but in this case it was constantly measured during the entirety of the runs. By normalizing the H -function by the time series of the HPL power, the effect of the changes in the power of HPL on the closed/open shutter ratio can be removed, especially over long closed shutter periods. For clarity, consider the following example. Suppose the HPL power was measured during a maintenance period that was immediately followed by

an open shutter period that is then used to calibrate a long closed shutter period. During the closed shutter period the HPL power fluctuates while all other variables remain the same. If the closed shutter data are not normalized to the HPL power, the ratio between any point in the closed shutter time series and the open shutter period measured earlier will be affected by these fluctuations. Likewise, these fluctuations would also cause the apparent conversion rate, $\mathcal{P}_{\gamma \leftrightarrow \phi}$, to fluctuate, when in reality it should remain constant. If the entire data set including the open and closed shutter data is normalized to the measured HPL power, the changes in the power will not affect the ratio or apparent coupling of the BSM fields. The signal to noise ratio of the experiment will fluctuate due to changes in the HPL power, however that would be the case with or without this normalization.

As the amplitude of the heterodyne beatnote scales with $\sqrt{P_i}$; dividing the H -function by $\sqrt{P_i}$ removes the effect of fluctuations in the injected power,

$$H_p[n] = \frac{H[n]}{\sqrt{P_i[n]}}. \quad (28)$$

Here, $H_p[n]$ refers to the normalized H -function, while $P_i[n]$ is the HPL power measured at sample ‘ n ’ by the out-of-vacuum power meter in CH. The closed and open shutter powers will both be proportional to the product of the absolute square of the expectation value of H_p and the injected power, as shown in the expressions

$$P_{\phi} \propto |\langle H_{p,\text{closed}}[n] \rangle|^2 P_i \quad (29)$$

and

$$P_{\text{open}} \propto |\langle H_{p,\text{open}}[n] \rangle|^2 P_i, \quad (30)$$

with the only difference between $H_{p,\text{closed}}$ and $H_{p,\text{open}}$ being the status of the shutter. By this definition H_p is a complex number with units of V/\sqrt{W} for either the open or closed shutter case.

The measured power injected into the production magnet string is shown in Figures 7a and 7b for S_{\perp} and S_{\parallel} , respectively. These measurements were used to calculate H_p . The power ranged from 29.9 W to 31.1 W for S_{\perp} with a mean value of 30.4 W during the closed shutter periods. For S_{\parallel} , the range was 19.4 W to 25.9 W with a mean value of 23.1 W during the closed



(a) Open shutter magnitude squared (S_{\perp})(b) Open shutter phase (S_{\perp})

FIG. 8: Magnitude squared (a) and phase (b) of $C_{\text{open},j}$ measured during S_{\perp} for the science detector (blue) and veto detector (orange).

shutter periods. The fluctuations in power during the runs are related to changes in the alignment of the HPL to the mode cleaning cavity on the NL table either due to environmental noise or a manual realignment of the system.

There was also a drop in power between S_{\perp} , which saw a maximum power of 31 W, and S_{\parallel} , with a maximum power of 26 W. This was related to the lowering of the injection current sent to pump diodes for S_{\parallel} because the HPL began to overheat in the period between S_{\perp} and S_{\parallel} due to a heatsink malfunctioning.

D. Closed to Open Shutter Ratio

In Section 1C, the integration of the time series of the H -function to produce the ‘ Z -function’ was discussed. In order to remove systematic drifts and uncertainty from the coherent sum, the H -function must first be expressed in terms of the closed/open shutter ratio before the closed shutter data series is integrated, by calculating the Z -function for each open shutter section of data. Then each closed shutter section of data is divided by its nearest valid point in the set of the open shutter data denoted by $C_{\text{open},j}$, which can be expressed as

$$C_{\text{open},j} = \langle H_{\text{p,open}} \rangle_j \quad (31)$$

$$= \frac{1}{m[j] - k[j]} \sum_{n=k[j]}^{m[j]} H_{\text{p,open}}[n]. \quad (32)$$

Here, the index j is assigned to each specific period of continuous open shutter section of data such that $j = 1$ refers to the first period of open shutter data. The series $k[j]$ gives the sample number corresponding to the start of each open shutter section according to the flags discussed in the previous section, while $m[j]$ is the end of each open shutter period of data. The sum is normalized by $1/(m[j] - k[j])$ such that $C_{\text{open},j}$ is proportional to the average value of H_{p} in that section of data. Like $H_{\text{p,open}}$, $C_{\text{open},j}$ is a complex number with units of $\text{V}/\sqrt{\text{W}}$ as $m[j]$ and $k[j]$ are unitless.

The phase of $C_{\text{open},j}$ represents the phase of the mean value of each of the open shutter periods while Equation 20 (with P_{s} referring to P_{open}) can be used to show that the magnitude squared of $C_{\text{open},j}$ is proportional to the ratio

$$|C_{\text{open},j}|^2 \propto \frac{P_{\text{open}}}{P_1} = \eta^2 \beta_{\text{RC}} T_{\text{COB}} T_{\text{RC2}}. \quad (33)$$

This demonstrates why periodically measuring $C_{\text{open},j}$ is important as it represents a direct measurement of the dynamic variables η , the field overlap, and β_{RC} , the resonant enhancement. The transmissivity of the COB T_{COB} and the flat RC mirror T_{RC2} are, on the other hand, static variables that were measured independently from the open shutter periods (See Section 1D and Ref. [20] for more details on these measurements).

The magnitude squared and the phase of the open shutter measurements of both science and

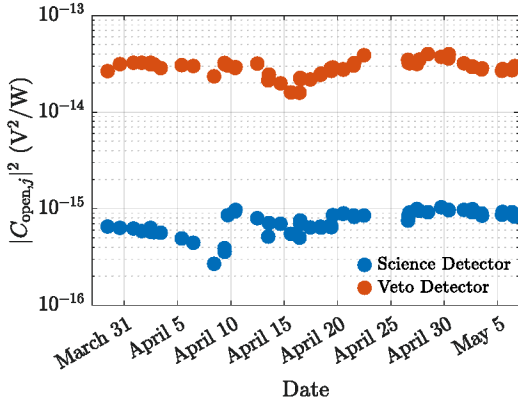
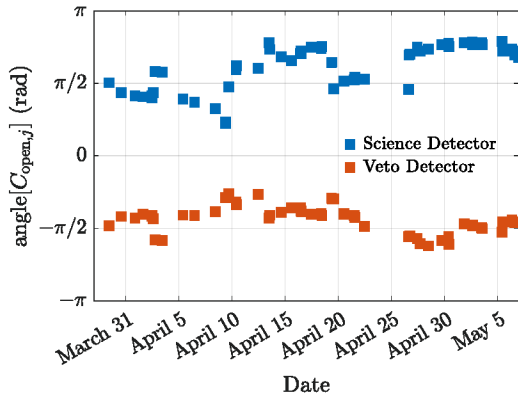
(a) Open shutter magnitude squared (S_{\parallel})(b) Open shutter phase (S_{\parallel})

FIG. 9: Magnitude squared (a) and phase (b) of $C_{\text{open},j}$ measured during S_{\parallel} for the science detector (blue) and veto detector (orange).

veto detector are shown in Figures 8a, 8b, 9a, and 9b, for S_{\perp} and S_{\parallel} . Each data point represents the 19 and 57 distinct open shutter periods in S_{\perp} and S_{\parallel} respectively. These results show that $C_{\text{open},j}$ was extremely stable during S_{\perp} , especially the science detector, where the changes in $|C_{\text{open},j}|^2$ were less than 30% through the duration of the run and phase changes of roughly a tenth of a cycle. $C_{\text{open},j}$ appears to be less stable during S_{\parallel} .¹⁵ Here the main driver behind the changes in the magnitude and phase of $C_{\text{open},j}$ during S_{\parallel} were the length changes of the RC de-

¹⁵ This is believed to be due to changes in the weather over the duration of the run.

tuning the HPL from the cavity resonance. For more details please refer to Section IV B 1.

The time stamp associated with the midpoint of each open shutter period $C_{\text{open},j}$ is used to assign each point in the closed shutter data to its nearest point in $C_{\text{open},j}$, provide no realignment of the system was performed between the closed shutter data point and $C_{\text{open},j}$. The calibration series $C[n]$, is the series composed of these assigned values of $C_{\text{open},j}$ for each value of $H_{\text{p,closed}}[n]$. For the specific case in which no valid open shutter period was measured between two different times in which the system was aligned, the closed shutter data between these times is disregarded.

The summation of the closed shutter H -function is then performed:¹⁶

$$Z_{\text{closed}} = \frac{1}{N} \sum_{n=1}^N \frac{H_{\text{p,closed}}[n]}{C[n]}. \quad (34)$$

Here, the series corresponds to the product of the full normalized H -function $H_{\text{p,closed}}[n]$ and one over the calibration series, or $1/C[n]$, while N is the total number of samples in the series. It should be noted here that the Z_{closed} is a unitless complex number. The lack of units of Z_{closed} is a consequence of the fact that the calibration procedure not only compensates for the parameters η and β_{RC} , but also for attributes of the heterodyne detection system, such as the conversion gain of the photodetector or the optical losses between RC1 and PD_{sci}. This is because they are common to both the open and closed shutter signals and is an important benefit of the calibration procedure.

Working under the assumption that there are no uncontrolled effects that can change the relative amplitude and phase of the regenerated field with respect to the signal during the open shutter periods, this method provides an automatic correction for the coupling and phase of the HPL with respect to the RC and LO. If, for example, the spatial coupling of the LO to the RC drops, this will affect both $C_{\text{open},j}$ and

¹⁶ In this equation, the H -function has been named $H_{\text{p,closed}}[n]$ as it is assumed that the H -function has been truncated to remove the data which are invalid either because of the status of the system flags or the lack of an open shutter period between alignments.

$H_p[n]$, provided it does not change significantly between sample n and sample $(m[j] - k[j])/2$.¹⁷

Likewise, because $H_p[n]$ and $C_{\text{open},j}$ are left as complex numbers, this effect also compensates for any slow phase drifts that are common mode to both the open and closed shutter periods. An example of this compensation occurs if there is an offset between the frequency difference of the HPL and LO and the frequencies used in the double demodulation scheme. This is only possible if such detunings lead to fractions of a cycle in phase drift or less between the open shutter periods. Therefore, this method can only compensate for this effect up to the order of tens of μHz . (It should be clear that the unexplained phase effects seen in the open shutter periods in Figures 8b and 9b were much less than a full cycle and frequencies less than μHz .)

Another dynamic effect that this method compensates for is the slow changes in the cavity length that can detune the HPL from the resonance of the RC, this could lead to the cavity inducing a phase offset into the science signal. These changes in phase would be common mode to the phase changes introduced to the open shutter signal as it transmits through the cavity. Therefore, as long as the cavity length changes only induce phase changes on the order of fractions of a cycle between the open shutter periods, this method can compensate for these drifts.

In either case, the optical system was actively monitored and re-aligned during maintenance periods during the science campaign to mitigate the phase drifts and changes in the coupling as much as possible in order to reduce the systematic uncertainty in the measurement. An analysis of the results of this effort is discussed in Section IV.

The power ratio is equivalent to the squared magnitude of Z_{closed} :

$$\frac{P_\phi}{P_{\text{open}}} \equiv |Z_{\text{closed}}|^2. \quad (35)$$

Together with Equation 25, $\mathcal{P}_{\gamma\leftrightarrow\phi}$ can be expressed as the product of $|Z_{\text{closed}}|$ with the

square root of T_{COB} and T_{RC2} :

$$\mathcal{P}_{\gamma\leftrightarrow\phi} = \sqrt{T_{\text{COB}} T_{\text{RC2}}} |Z_{\text{closed}}|. \quad (36)$$

This equation represents the primary achievement of the calibration procedure as it provides an elegant method to directly translate the results measured by the heterodyne detection system in terms of the conversion rate between electromagnetic and BSM fields $\mathcal{P}_{\gamma\leftrightarrow\phi}$. This requires the use of only one other dynamic variable, the open shutter measurements, as well as the measured combined transmissivities of the optics between the two magnet strings along the optical axis of the system. The results presented in Section III D, in terms of $\mathcal{P}_{\gamma\leftrightarrow\phi}$, are then interpreted in Ref. [1] as limits on $g_{\phi\gamma\gamma}$.

E. Veto Detection System

A veto detection system was also operated during the first science campaign to check the validity of signals measured at the science detection system. The veto photodetector (PD_{veto}) is positioned in the CH cleanroom such that it would detect the light reconverged from the BSM signal that transmits RC2. As RC2 is nearly a factor of 20 less transmissive than RC1, for a given electromagnetic to BSM field conversion rate, the reconverged signal power on PD_{veto} would be significantly lower than the signal power present on the science photodetector. This factor is seen in the expression for the close shutter signal power transmitted by RC2 in the presence of a BSM signal:

$$P_{\phi,\text{veto}} = \eta^2 \beta_{\text{RC}} \frac{T_{\text{RC2}}}{T_{\text{RC1}}} P_i \mathcal{P}_{\gamma\leftrightarrow\phi}^2. \quad (37)$$

Although it is still capable of measuring the reconverged field, the true purpose of the veto detector is to observe the stray light, due to its proximity to the shutter system and COB, while also providing independent means of monitoring the performance of the optical system. Stray light incident on the veto detector does not need to pass through the beam tube or RC like the stray light present at the science detector. Furthermore, stray light on the optical axis of the system that directly couples to the RC should have a power roughly a factor of $R_{\text{RC,CH}}/T_{\text{RC}} \sim 30$ higher at the veto detector. Here $R_{\text{RC,CH}} \sim 0.95$ is the reflectivity of

¹⁷ Later in Section II F and Appendix A 2 it is shown that on time scales of roughly two days or less the system is stable and these changes are small compared to other sources of uncertainty in the measurement.

the RC for a laser that is perfectly spatially mode matched to the cavity eigenmode (see Section IV A 3 and Equation 56 for more details). $T_{\text{RC}} \sim 3\%$ is the transmissivity of the RC for a laser sharing the spatial eigenmode of the cavity.

The analysis for the data from the veto detection system is nearly the same as for the science detector. The signal from the veto photodetector is amplified and then digitized by an ADC and the digital signal is demodulated at the veto heterodyne frequency $f_{\text{veto}} = 33f_0$, where f_0 is the FSR of the RC. As the diagram of the laser frequencies in Figure 3 shows, the veto heterodyne frequency is different from the science heterodyne frequency because the interference beatnote measured at the veto detector is between the signal field and the AL light in transmission of the RC. The demodulated data are used to form the open and closed shutter arrays H_p^{veto} and the veto systems version of the calibration series $C_{\text{veto}}[n]$ is constructed with the open shutter veto series $C_{\text{open},j}^{\text{veto}}$. Just like the calibration procedure for the science detector data, the results at the veto detector are formulated in terms of the normalized closed shutter power $|Z_{\text{closed}}^{\text{veto}}|^2$.

When the shutter is open, the HPL light reflected by RC2 that is in the spatial mode of the AL¹⁸ can be expressed by

$$P_{\text{open}}^{\text{veto}} = \frac{\eta^2}{\eta_z^2} R_{\text{RC,CH}} T_{\text{COB}} P_1. \quad (38)$$

The spectral component of the field overlap is canceled from the total overlap as the term η^2 is divided by η_z^2 . Again, $R_{\text{RC,CH}}$ (Equation 56 in Section IV A 3) is the reflectivity of the RC for a laser occupying the spatial mode of the cavity. Here, the longitudinal overlap η_z^2 will play a small role in $R_{\text{RC,CH}}$. From this equation it is apparent that $|Z_{\text{closed}}^{\text{veto}}|^2$ has a very different meaning than its counterpart calculated from the science detector as the dynamic variables β_{RC} and η_{xy} do not perfectly cancel here.

F. Calibration Uncertainty

The optical system is not static and parameters such as the cavity length and HPL coupling to the RC did change over time, producing changes in the expected signal for a given BSM coupling during the run. These changes are encoded on the open shutter signal which is why open shutter data was taken as frequently as possible. The uncertainty in the projection of this data to the calibration series $C[n]$ for nearby closed-shutter periods is one of the leading contributors to the systematic uncertainty in the calibration of the regenerated power in terms of $\mathcal{P}_{\gamma \leftrightarrow \phi}$.

To set error bars on the open shutter calibration, the changes in the open shutter results were examined for points between which no realignment of the system was performed. This was done by compiling sets of data expressing the changes between the magnitude squared and phase of different open shutter sections as a function of the time between them.

The normalized absolute value of the difference in the magnitude squared for two points in the open shutter data $C_{\text{open},j}$ and $C_{\text{open},m}$, separated in time by $t(m) - t(j)$ (here $t(m)$ and $t(j)$ are the time at the midpoint of their corresponding open shutter periods), is

$$\Delta |C_{\text{open},j \rightarrow m}|^2 = \frac{||C_{\text{open},j}|^2 - |C_{\text{open},m}|^2|}{|C_{\text{open},j}|^2}. \quad (39)$$

A first order polynomial was then fit to $\Delta |C_{\text{open},j \rightarrow m}|^2$ as a function of $t(m) - t(j)$ to find the average normalized changes as a function of the time between the open shutter periods. This process gave a linear trend of 1.5%/day for S_{\perp} in the science detector with a static offset of 7.3%. That is, open shutter periods measured only minutes apart would vary on average by 7.3% in terms of $|C_{\text{open},j}|^2$, but if these measurements were separated by a day, they would vary on average by 8.8%. From these data it can be seen directly that it would take nearly 5 days for the uncertainty due to the long term drifts in the system to have a larger impact than the uncertainty measured over shorter times between the open shutter periods. The linear trend was 3.1%/day at the science detector for S_{\parallel} , with a static offset of 6.0%. The positive slope in all of the fits indicates that the changes between open shutter periods tend to be larger for larger pe-

¹⁸ The HPL light must be in the spatial mode of the AL to generate a beatnote measurable by the heterodyne detection system.

riods of time as expected. Figures 22a and 22c in Appendix A 2 show scatter plots for S_{\perp} and S_{\parallel} of the relative absolute change in $|C_{\text{open},j}|^2$ for different open shutter periods as a function of the time between them.

These trends were then used to estimate the uncertainty in $|C[n]|^2$. Here, each point in $C[n]$ is assigned error bars based on how far in time it is from the nearest open shutter period with no realignment of the system performed in between. The average of the error bars of $|C(n)|^2$ for all closed shutter points found with this process was then used to calculate contribution of the changes in the open shutter data to the relative error on $|Z_{\text{closed}}|^2$. This process gave an average relative uncertainty of 7.8% on $|Z_{\text{closed}}|^2$ with a range from 7.3% to 8.8% for S_{\perp} . For S_{\parallel} the average uncertainty due to the changes in the open shutter periods was 7.0% with a range from 6.0% to 11.5%.

The change in the phase of the open shutter periods was also calculated. The first-order polynomial fit to the data gave $0.123 \text{ rad/day} + 0.020 \text{ rad}$ for the science detector and $0.041 \text{ rad/day} + 0.033 \text{ rad}$ for the veto detector during S_{\perp} . In the case of S_{\parallel} , the first-order polynomial fit gave $-0.010 \text{ rad/day} + 0.152 \text{ rad}$ at the science detector and $0.035 \text{ rad/day} + 0.059 \text{ rad}$ for the veto detector. These trends shows that longer times between the open shutter periods corresponded to larger changes in the phase of the open shutter measurements for S_{\perp} for both the science and veto detectors, as well as S_{\parallel} at the veto detector. The phase changes measured at different open shutter periods of S_{\parallel} appear actually to decrease at the science detector as a function of the time between the points. This is believed to be due to several outliers biasing the data at earlier times. The average uncertainty on the angle of $C[n]$ was 0.062 rad for S_{\perp} with a range of 0.026 rad to 0.1394 rad . For S_{\parallel} the average uncertainty on the angle of $C[n]$ was 0.148 rad for with a range of 0.136 rad to 0.151 rad . This would correspond to an average uncertainty on the magnitude squared of $C[n]$ of 0.45% for S_{\perp} and 2.2% for S_{\parallel} . As these errors were insignificant compared to other sources of error in $C[n]$, the contribution of the changes in the phase between open shutter periods was not considered when calculating the error of $C[n]$ and $|Z_{\text{closed}}|^2$. Scatter plots of the absolute change in angle of the open shutter periods as a function of the time interval between them

can be found in Appendix A 2 with Figure 22b showing the results from S_{\perp} and Figure 22d S_{\parallel} .

It must be also noted that interpreting the changes in the open shutter data should be approached with a certain degree of caution. As the previous paragraph mentions, it is possible that the data have some bias related to the stability of the system. For example, when the system was unstable, it may have become unlocked more frequently during the open shutter periods. This would lead to more frequent open shutter periods during these times when the system was less stable because the operators were instructed to relock the system and restart the open shutter period if a certain integration time was not reached. Likewise, when the system was more stable, it would be more likely that no relock was required and therefore there may not even be any data at all for shorter durations between open shutter periods during the more stable periods. This may also explain why the shorter time periods appear to be less stable in some cases, than the longer time periods in the plots shown in Appendix A 2.

Later, in Section III E 2, the time evolution of the closed shutter data is examined on the complex plane. The effects in the open shutter data discussed in the previous paragraphs can also be examined on the complex plane. Figures 16b and 16d show plots of the real and imaginary components of $C_{\text{open},j}/C_{\text{open},m}$ during S_{\perp} , for the science and veto detectors, respectively, in comparison to a moving average of the closed shutter data (for more information see Section III E 2). Here, only the changes between open shutter periods in which the system was free-running and no alignment took place between them were considered. This means that these are not necessarily all consecutive measurements of the open shutter periods, as often realignments of the system occur in between the open shutter periods. In these plots, if no changes occurred in the open shutter results, all points would be at $1+0i$. In the science detector, the mean value of the data shown in these plots was $0.98 - 0.08i$ with a standard deviation of 0.23 . The mean of the absolute value of the real components of these data minus one was 0.04 , while for the imaginary components the mean absolute value was 0.18 . These changes can be used to calculate a rough estimation of the relative error in $|Z_{\text{closed}}|^2$ of 8% . The estimates of the errors from both methods are therefore

equal with respect to their significant digits.

The reason that this result is slightly different from the estimate of the error compiled using the trends earlier in this section is that this is just an average of all of the changes in the complex open shutter data with no attempt to project these changes on to the closed shutter periods or account for the time between the open shutter periods. In either case, t

Data for S_{\parallel} are shown in Figures 17b and 17d for the science and veto detector, respectively. In these plots, the mean value of the data at the science detector was $0.95 + 0.08i$ with a standard deviation of 0.26. For these data the mean of the absolute value of the real component of the data minus one was 0.06, while the mean of the absolute value for the imaginary component was 0.14 leading to a rough estimate of the relative error of 12% on $|Z_{\text{closed}}|^2$. This shows relatively good agreement with the estimate of the error using the linear trend of the magnitude square of $C_{\text{open},j}$ of 7.0%. The discrepancy between the two methods is believed to be due to the fact that this method simply finds an average of all of the changes in the complex open shutter data with no attempt to project these changes on to the closed shutter periods or account for the time between the open shutter periods as the method using the linear trends in the data does.

The technical noise of the system (discussed in more detail in Section III A 2) also contributed to the uncertainty calculated here. As the technical noise is symmetric about the actual signal, over long integration times it will average to zero. This can be seen explicitly in Figures 18a and 18c (discussed later in Section III E 3) as the $1-\sigma$ statistical error on the data converges proportionally to a trend of $1/\sqrt{t'}$ with t' being the integration time. Some of the open shutter periods were short enough that the impact of the technical noise was significant. The technical noise contribution to the uncertainty was calculated by performing a Fourier transform on the data from each of the open shutter periods and calculating the mean noise in the bins excluding the frequency bin where the actual signal resides.¹⁹ For S_{\perp} , the mean value of the relative error in $|C_{\text{open},j}|^2$ due to technical noise during

the open shutter periods was only 0.15% for the science detector however the shortest section of open shutter data at 493s gave an uncertainty of 3.99%. For S_{\parallel} the average uncertainty due to technical noise was only 0.30% for the science detector, but the shortest open shutter section lasted only 59s and had an uncertainty of 12.69% due to the technical noise.

The primary reason for the lower relative noise in the S_{\perp} open shutter periods was that the average duration was their longer median duration of 2070s for S_{\perp} compared to 960s for S_{\parallel} . In either case, technical noise appears to impact the uncertainty of the calibration only marginally. These effects are already accounted for in the analysis from this section as this will also play a role in the changes observed between different periods of open shutter data.²⁰

G. Uncertainty in the Results

Three sources are considered in the estimate of the uncertainty in $|Z_{\text{closed}}|^2$. Two of these, the wedge angles of the COB optics changing the spatial overlap of a BSM signal with respect to the open shutter signal (introducing a relative uncertainty of 10% on $|Z_{\text{closed}}|^2$) and the loss of phase coherence of potential BSM signals with respect to the open shutter data due to temperature changes on the COB (contributing a relative uncertainty of 4.8% to $|Z_{\text{closed}}|^2$ for S_{\perp} and 2.4% for S_{\parallel}), were discussed in Section II A. The third source of error in $|Z_{\text{closed}}|^2$ was the uncertainty of $C[n]$ discussed in the previous section.

²⁰ It is not entirely clear what the threshold is for when the drifts in the setup take over from the technical noise as the dominant source of noise in the open shutter data. From Figures 18a and 18c this appears to be longer than 1000s for sections of open shutter data where the system was continuously locked based on the fact that the error bars never stop converging. In Appendix A 2, Figures 22a-22d, the data compiled from periods in which the system was not continuously locked show that the drifts in the system are the dominant source of noise after only a few hours as the uncertainty on the shortest time scales shown in these plots is larger than the uncertainty shown on the longest time scales of Figures 22a-22d. Nevertheless, as the text states, the results over short time scales in Figures 22a-22d should be treated with caution, as there may be some bias in the data toward times when the system was less stable.

¹⁹ Only frequency bins within the bandwidth of the filters applied to the data were considered.

	Observed result: $ Z(f_k = 0)_{\text{closed}} ^2$	Estimated background: $\mathcal{M}_{\text{TH}}(f_k = 0)$	Estimated background: $\mathcal{M}_{\text{SG}}(f_k = 0)$
Science (S_{\perp})	$(1.95 \pm 0.26) \times 10^{-5}$	$(2.54 \pm 0.34) \times 10^{-5}$	$(2.28 \pm 0.31) \times 10^{-5}$
Veto (S_{\perp})	$(4.58 \pm 0.53) \times 10^{-5}$	$(2.59 \pm 0.30) \times 10^{-5}$	$(2.46 \pm 0.29) \times 10^{-5}$
Science (S_{\parallel})	$(2.87 \pm 0.36) \times 10^{-5}$	$(0.85 \pm 0.11) \times 10^{-5}$	$(1.07 \pm 0.13) \times 10^{-5}$
Veto (S_{\parallel})	$(4.90 \pm 0.57) \times 10^{-5}$	$(3.87 \pm 0.45) \times 10^{-5}$	$(4.09 \pm 0.52) \times 10^{-5}$

TABLE II: Summary of the heterodyne data in the signal bin for both detectors and both science runs are shown in the column labeled $|Z(f_k = 0)_{\text{closed}}|^2$. The errors are derived from the uncertainty in the calibration time series $C[n]$. The columns labeled $\mathcal{M}_{\text{TH}}(f_k = 0)$ and $\mathcal{M}_{\text{SG}}(f_k = 0)$ give the background estimates discussed in Section III B calculated using the top-hat and Savitzky-Golay filters respectively.

In this case the relative uncertainty introduced to $|Z_{\text{closed}}|^2$ by $C[n]$ was estimated to be 7.8% for S_{\perp} and 7.0% for S_{\parallel} . Summing these quadratically gives a total relative error on $|Z_{\text{closed}}|^2$ of 13.3% for S_{\perp} and 12.4% for S_{\parallel} .

The uncertainty of $\mathcal{P}_{\gamma \leftrightarrow \phi}$ will be affected by all of the sources discussed in the previous paragraph as well as the uncertainties on the transmissivities of the COB optics along the optical axis of the system (discussed later in Section III F) and the transmissivity of RC2 (discussed in Ref. [20]). With a relative uncertainty in T_{COB} of 7.1% and a 9.8% in T_{RC2} the relative error bars on $\mathcal{P}_{\gamma \leftrightarrow \phi}$ come out to 9.1% for S_{\perp} and 8.7% for S_{\parallel} .²¹

III. RESULTS AND BACKGROUNDS

The results acquired by the heterodyne detection system using the science and veto detectors are shown for both runs in terms of $|Z_{\text{closed}}|^2$ in Table II. The measured value given by the science detector data was $(1.95 \pm 0.26) \times 10^{-5}$ for S_{\perp} and $(2.87 \pm 0.36) \times 10^{-5}$ for S_{\parallel} . For the veto detector $|Z_{\text{closed}}|^2$ was $(4.58 \pm 0.53) \times 10^{-5}$ for S_{\perp} and $(4.90 \pm 0.57) \times 10^{-5}$ for S_{\parallel} . The error bars shown in Table II represent the estimated systematic uncertainty that was calculated with the methods discussed in the previous section.

The resulting values of $|Z_{\text{closed}}|^2$ in Table II lie in what we will refer to as the signal frequency bin, because the data was demodulated at the heterodyne frequencies given in Section I C. As Section III A explains, the Z -function was also evaluated at alternative frequencies where no BSM contribution was expected. This allowed the backgrounds to be estimated without requiring background measurements with the magnets off, under the assumption that the neighboring bins are indicative of the background in the signal bin. The estimated background in the signal bin is shown as the second and third columns of Table II and this process is summarized and compared with the measured data in Section III B.

Section III C explains how the statistical distribution of the background was calculated by comparing the estimated background to the measured values at frequencies outside the signal bin. This distribution was then used to set detection thresholds that determined whether or not a signal was present at a given confidence level.

The results in the signal bin given in Table II are consistent with the estimated background levels during both runs. Thus neither run showed an excess statistically significant enough to claim the presence of a signal generated by new physics. Section III C therefore also examines how exclusion limits on the conversion rate between electromagnetic and BSM fields were set based on the background distribution. The resulting exclusion limits are summarized in terms of $\mathcal{P}_{\gamma \leftrightarrow \phi}$ in Section III D.

Section III E discusses how the behavior and data measured at the science and veto detection systems indicate that the backgrounds could be due to stray light. This is done by examining the

²¹ It must be considered that because $\mathcal{P}_{\gamma \leftrightarrow \phi}$ is the square root of the product of the transmissivities of these optics and the $|Z_{\text{closed}}|^2$, its relative error is roughly half the quadratic sum of the errors of these parameters.


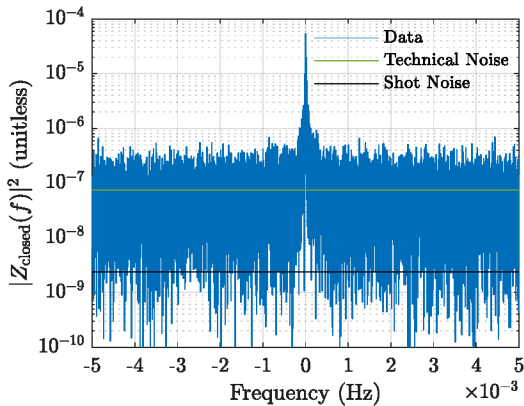
(a) Science detector (S_{\perp})(b) Veto detector (S_{\perp})

FIG. 10: $|Z(f)_{\text{closed}}|^2$ measured by the science (a) and veto (b) detectors during S_{\perp} are shown in blue over the frequency range from -5 mHz to $+5$ mHz around the signal frequency bin. The mean technical noise is shown in green, while the shot noise limit is shown in black.

time evolution of $|Z_{\text{closed}}|^2$ as well as the ratio of the measured signals and backgrounds at the science and veto detectors.

Measurements of the transmissivities of the COB optics are then described in Section III F, as these parameters are critical to calculating $\mathcal{P}_{\gamma \leftrightarrow \phi}$ from the results of $|Z_{\text{closed}}|^2$.

A. Alternative Frequency Analysis

As mentioned earlier, the signal strengths at frequencies other than the BSM heterodyne frequency were also examined. The Z -function for

a frequency f away from the signal bin (in this equation the signal bin is at $f = 0$) is

$$Z(f)_{\text{closed}} = \frac{1}{N} \sum_{n=1}^N \frac{H_p[n] e^{i2\pi \frac{f}{f_n} n}}{C[n]}. \quad (40)$$

Here f_n is the sampling frequency (1 Hz) of the down-sampled data series. To be clear, H_p is still the same H -function that has undergone double demodulation, down sampling, and normalization by the HPL injected power. Also, it should be noted that for calculating the signal strength at other frequencies, there is no mathematical difference between this method and changing one of the heterodyne demodulation frequencies given in Section I C.

The data were compiled with a bin size in the frequency series of one over the total time elapsed between the first and last samples of valid closed shutter data. With 1.2 million seconds elapsed between the first and last samples of closed shutter data for S_{\perp} , the frequency resolution for S_{\perp} was $0.85 \mu\text{Hz}$. S_{\parallel} saw 3.3 million seconds elapse between the first and last valid closed shutter samples and therefore the frequency series is sampled every $0.30 \mu\text{Hz}$ for this run. Using these resolutions in the frequencies does mean that the Fourier series does not form a linearly independent basis set, as the gaps in the data during periods in which the system was unlocked cause the Fourier components to lose their linear independence over the data run. The consequences of this are discussed in more detail in Appendix A 3.

Because H_p/C is a complex number, $Z(f)_{\text{closed}}$ is not necessarily symmetric about $f = 0$ and it is important to examine both the positive and negative frequencies. The results of S_{\perp} are shown from -5 mHz to $+5$ mHz in Figures 10a and 10b for the science and veto detectors, respectively. Likewise, the results for S_{\parallel} over the same frequency range are shown in Figures 11a and 11b. All of these plots are centered on the heterodyne frequency such that $f = 0$ is the signal frequency bin.

In all four plots, the measured data, shown in blue, is dominated by technical noise, whose mean is shown in green, for frequencies more than $100 \mu\text{Hz}$ away from the heterodyne fre-

quency.²² In the case of S_{\perp} the mean technical noise in terms of $|Z_{\text{closed}}|^2$ was 4.2×10^{-7} , while for S_{\parallel} it was 3.3×10^{-7} .

The derivation of the shot noise limits, shown in black, and a comparison to the technical noise are discussed further in Section III A 1. A peak above the technical noise floor is apparent in the data at frequencies less than 100 μHz away from the signal bin. The fact that these signals are spread over many frequency bins and not constrained to the signal bin indicates that this peak is the result of a background in the detection system. One possible source of background is stray light that scatters out of the HPL and takes a random uncontrolled path through the optical system to the detection system. Because the length of this path will fluctuate in time due to environmental noise, the stray light can be Doppler shifted by μHz into the frequency bins neighboring the heterodyne frequency. This hypothesis is corroborated by the fact that similar peak amplitudes in terms of the closed to open shutter ratio are seen at the veto and science detectors. This point is discussed in more detail in Section III E. A statistical analysis of the background and plots showing the data from frequencies within 100 μHz of the signal bin is discussed in Section III B.


1. Quantum Noise

The shot noise limit for the single sided power spectral density (PSD) of the photocurrent produced by a photodiode can be expressed using measured parameters by the equation

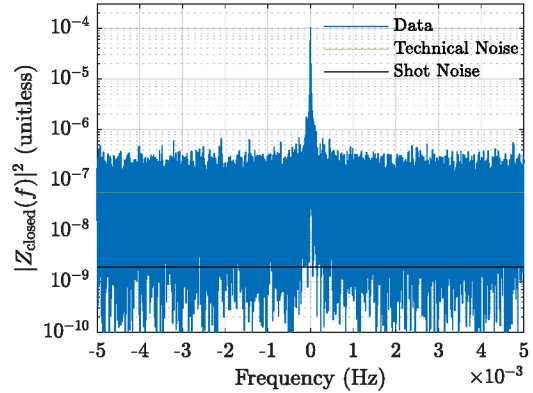
$$S_v(f) = \frac{2\bar{P}G^2h\nu}{\epsilon}, \quad (41)$$

in terms of V^2/Hz . Here \bar{P} is the average laser power on the science PD, G is the gain in V/W of the detection system, h is Planck's constant, ν is the laser frequency, and ϵ is the quantum efficiency of the photodetector.

²² The mean technical noise is calculated for frequencies within the bandwidth of the filters applied to the data and more than 1 mHz away from the heterodyne frequency to avoid the peak near the heterodyne frequency having an effect on the calculation.



(a) Science detector (S_{\parallel})



(b) Veto detector (S_{\parallel})

FIG. 11: $|Z(f)_{\text{closed}}|^2$ measured by the science (a) and veto (b) detectors during S_{\parallel} are shown in blue over the frequency range from -5 mHz to +5 mHz around the signal frequency bin. The mean technical noise is shown in green, while the shot noise limit is shown in black.

Assuming an average power of 270 μW for the S_{\parallel} , a conversion gain at the heterodyne frequency of 3.00 V/mW , a frequency of 282 THz, and a quantum efficiency of 0.7, the voltage shot noise limited PSD would be $1.30 \times 10^{-15} \text{V}^2/\text{Hz}$. After 1.06 million seconds of integration time this would produce a single sided flat power spectrum with an expectation value of $1.22 \times 10^{-21} \text{V}^2$. Normalizing by the average HPL power and open shutter calibration amplitude, and then also dividing by two as our measurement is equivalent to a double sided spectrum, gives the shot noise limit on $|Z_{\text{closed}}|^2$ of

5.2×10^{-8} for the science detector during S_{\parallel} . Calculating this noise limit using the parameters for the science detector during S_{\perp} gives 6.3×10^{-8} . The mean technical noise, derived in Section III A 2, in terms of $|Z_{\text{closed}}|^2$ was 3.3×10^{-7} for S_{\parallel} , a factor of 6 above the shot noise limit, and 4.2×10^{-7} for S_{\perp} , a factor of 7 above shot noise.

For the veto detector the expected single sided shot noise PSD can be found to be 2.5×10^{-15} V²/Hz for S_{\parallel} and 2.2×10^{-15} V²/Hz for S_{\perp} . This leads to a shot noise limit of 1.8×10^{-9} on $|Z_{\text{closed}}^{\text{veto}}|^2$ for S_{\parallel} and 2.2×10^{-9} on $|Z_{\text{closed}}^{\text{veto}}|^2$ for S_{\perp} . In terms of $|Z_{\text{closed}}|^2$, the shot noise limits for the veto detector are lower than at the science detector simply because the open shutter signal does not need to transmit the RC. If both detectors reach the shot noise limit the veto detector will be more sensitive with respect to $|Z_{\text{closed}}|^2$ by a factor of $1/T_{\text{RC}}$ or roughly 30. This does not mean the veto detector is more sensitive to the coupling of BSM fields. On the contrary, it is significantly less sensitive as the coupling of the BSM signal to the detector in terms of the closed to open shutter power ratio is roughly a factor of 600 lower than at the science detector. This is explained in more detail in Section III E 1, however it is due to the fact that reflectivity of the RC, and hence the open-shutter power of the HPL in the RC's spatial eigenmode that is sent to the veto detector, is a factor of 30 larger than the transmissivity of the RC, or the in-mode HPL open-shutter power sent to the science detector. Likewise, the power reconverted from an actual BSM signal would be roughly 20 times lower in transmission of RC2 (the cavity port directed to the veto detector) compared to RC1 (the cavity port directed to the science detector) due to the transmissivities of the mirrors.

2. Technical Noise

Measurements of conventional sources of technical noise such as the noise of the ADCs used in the data acquisition system and the noise equivalent power of the science PD gave power spectral densities at the heterodyne frequency lower than the technical noise measured in the experiment or even the projected shot noise limit at that frequency. Therefore, a more thorough in-

vestigation was required to identify the source of this noise. This investigation included tests that revealed that the noise was related to the interaction of the LO and the AL when the AL was frequency stabilized to a cavity resonance. One set of measurements was performed in reflection of the RC with the circulating field of the cavity blocked, but the AL reduced in power to the same power level as if it were on resonance. With the LO at its normal power, the resulting PSD of the voltage noise output by the photodetector was within 20% of the predicted shot noise level at the heterodyne frequency. However, measurements performed with the system in its normal configuration (cavity unblocked, AL resonant with the RC) the PSD of the voltage noise output by the science PD was a factor of 6 above the predicted shot noise level, identical to the results achieved during S_{\parallel} . This was despite the fact that power incident on the science PD was the same in both measurements. A third measurement with the AL stabilized to the RC, but the LO blocked also produced a PSD in the voltage noise output by the science PD at the heterodyne frequency that was consistent with the predicted shot noise limit. It was not possible to measure the power noise at the science PD with the AL off resonance from the RC and both it and the LO at nominal power as this saturated the photodetector. While this points to the interaction between the LO, AL, and the RC being the source of the technical noise, these effects are not entirely understood and therefore are still under investigation.

While it was not possible to increase the LO power due to saturation effects in the science PD, future upgrades plan for a new design of the science PD that can accommodate higher powers to ensure that shot noise from the detection of the LO is the dominant source of noise at the detector. This would lead to an apparent gain in sensitivity of a factor of roughly 2.5 with respect to $\mathcal{P}_{\gamma \leftrightarrow \phi}$ and 1.6 with respect to $g_{\phi \gamma \gamma}$ under the assumption the current science runs were limited by technical noise. However, both S_{\perp} and S_{\parallel} appear to be limited by an additional background discussed in the next section.

B. Background Estimation

An estimation of the expected measured noise in the absence of a BSM signal is performed by



(a) Science detector (S_{\perp})(b) Veto detector (S_{\perp})

FIG. 12: $|Z(f)_{\text{closed}}|^2$ measured by the science (a) and veto (b) detectors during S_{\perp} are shown in blue over the frequency range from $-100 \mu\text{Hz}$ to $+100 \mu\text{Hz}$ away from the signal frequency bin. The data point in the signal bin is not considered when calculating the filtered data to avoid the possibility that an actual signal would influence the estimation of the background (the estimated background at the signal frequency is still calculated).

applying a filter to the data to calculate the expected value in the signal bin based on results in the frequency bins neighboring it. Two different filters were used in this process. One was a top-hat filter that produces a moving average of the eleven frequency bins centered on the bin in which the noise is being estimated for while the other was a Savitzky-Golay filter [26] that fits a 6-th order polynomial to the 101 points centered

on that same bin.²³ In all cases, the signal bin was not included in the data considered by the filters to ensure that an actual signal would not influence the estimated background.²⁴

The results of applying the filters to $|Z(f)_{\text{closed}}|^2$ at frequencies $100 \mu\text{Hz}$ above and below the signal bin are shown in Figures 12a and 13a. Figures 12b and 13b show the same process applied to the data from the veto detector. In all of these plots the data are shown in blue, while the results from the top-hat filter are represented by the purple lines with the orange lines giving the results of using the Savitzky-Golay filter. In the case of the veto detector results, a narrower Savitzky-Golay filter that considered the 41 bins and a top-hat filter that averaged 5 frequency bins were used to estimate the background as the data showed a narrower peak.

The resulting expectation values of the background in the signal bin for S_{\perp} and S_{\parallel} estimated by the filters are shown in Table II in the two columns to the right of the results in the signal bin. In this table $\mathcal{M}_{\text{TH}}(f_k = 0)$ refers to the background estimated by the top-hat filter and $\mathcal{M}_{\text{SG}}(f_k = 0)$ is the background estimated by the Savitzky-Golay filter. The error bars here come from the uncertainty in the calibration array discussed earlier. From these results, it is apparent that S_{\perp} did not show an excess in the signal bin with respect to the expected background, while the S_{\parallel} data showed a slight excess. Significant background signals were also seen in the veto detectors. The ratio of backgrounds measured at the science and veto detectors along with the same ratio of the results measured in the signal bin is discussed more in Section III E 1, but is interpreted as evidence that these signals and background are the result of stray light scattering from the HPL to the detection system and not the result of the coupling of the laser to BSM fields.

The number of points considered in the filtering was selected based on Monte-Carlo simulations of randomly generated data with a probability distribution modeled after the general

²³ The effects of spectral leakage discussed in Appendix A 3 are not considered.

²⁴ While the signal bin is excluded from consideration in each of the points of the filtered data, this should not be mistaken for the background in the signal bin not being estimated.

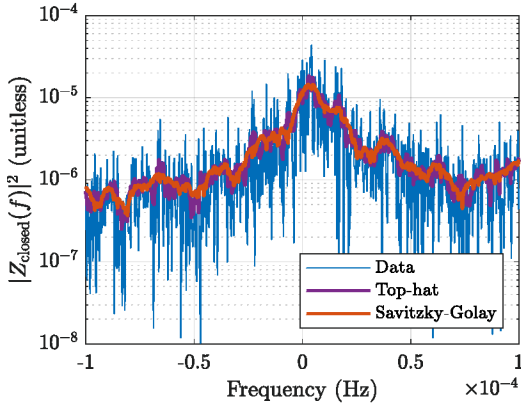

(a) Science detector ($S_{||}$)(b) Veto detector ($S_{||}$)

FIG. 13: $|Z(f)_{\text{closed}}|^2$ measured by the science (a) and veto (b) detectors during $S_{||}$ are shown in blue over the frequency range from $-100 \mu\text{Hz}$ to $+100 \mu\text{Hz}$ away from the signal frequency bin. The data point in the signal bin is not considered when calculating the filtered data to avoid the possibility that an actual signal would influence the estimation of the background (the estimated background at the signal frequency is still calculated).

shape seen in the frequency spectrum of the measured data within $100 \mu\text{Hz}$ of the signal bin. For each iteration of the simulations, a Gaussian shape was used to model the background peak observed in the data. The width of the Gaus-


sian, its frequency shift from 0 Hz ,²⁵ peak amplitude, and baseline amplitude were set to be similar to what was measured during the run, but with random noise giving the values a standard deviation of 20% from their nominal values from iteration to iteration. This was done to mimic the changes seen in the shape of the background spectra between the runs.²⁶

While the measured data showed the broad features that were modeled in the Monte Carlo simulations by the process outlined in the previous paragraph, it also showed random behavior from bin to bin that adhered to a non-central χ^2 distribution with two degrees of freedom. Therefore, each of the modeled spectra was then multiplied by a random series of numbers drawn from a χ^2 distribution. After this, the filter being tested was applied to the modeled spectra and the deviation of the filtered value from the nominal value of the spectra before the noise was included was calculated. With this technique, it was possible to assess whether the filters produced an accurate representation of the underlying spectral features, and the filter parameters could be tuned for optimal performance based on the results.


The results of the Monte-Carlo simulation can be seen in Figures 14a for the top-hat filter and 14b for the Savitzky-Golay filter. Here, the number of frequency bins being filtered is represented by the horizontal axis. \mathcal{M}_{TH} and \mathcal{M}_{SG} are meant to represent the filtered values of the noisy spectra in central bin, while $\mathcal{M}_{\text{true}}$ is the nominal ‘noiseless’ value. The normalized deviation is therefore given by $(\mathcal{M}_{\text{TH/SG}} -$

²⁵ In Figures 12a and 13a it is apparent that the measured distributions are not centered on the signal bin.

²⁶ Because the background is believed to be due to stray light, this behavior could originate from μHz Doppler shifts in the stray light as its optical path length either expands or contracts. Over long enough integration times the Doppler shifts are expected to be symmetric about zero as the optical path length of the stray light to the detector should be just as likely to contract as it should be the expand. However, the measured background spectra only represent a snapshot of the background during the period in which the data runs were taking place and biases could exist due to the times of day when the system was in operation. The offsets seen in these measurements are therefore interpreted to be representative of the expectation value for magnitude of the shifts present in the center of background spectra over these time scales.



(a) Top-hat Monte Carlo results



(b) Savitzky-Golay Monte Carlo results

FIG. 14: Results of Monte-Carlo simulations of applying a top-hat (a) and Savitzky-Golay (b) filters to data with a noise distribution modeled after data measured in S_{\perp} and S_{\parallel} . The Gaussian background was randomly varied in its width, central frequency, peak amplitude and baseline amplitude with a standard deviation of 20% of their observed values in the data, and noise from a non-central χ^2 distribution was applied to the resulting spectra of each iteration.

$\mathcal{M}_{\text{true}})/\mathcal{M}_{\text{true}}$ and is shown as the solid orange and purple lines for the top-hat and Savitzky-Golay filters, respectively. The dashed lines show the average standard deviations of the filtered data between iterations of the Monte Carlo simulations. The results of filtering the modeled background spectra are compared with the results of applying the filter to a flat background frequency spectrum with the same noise characteristics shown as the blue lines. Ideally, the

filter parameters should be tuned such that data with the modeled spectrum match the data from the flat spectrum. If that is not the case, it can indicate that the filter is somehow biasing the data. The black dashed lines show the expected standard deviations for the filtered flat distribution as a function of the number frequency bins filtered. Both plots show that these traces agree well with the Monte Carlo results.

These plots show that the top-hat filter has a tendency to underestimate the value of the distribution in the signal bin after the number of points considered goes beyond 20, while the Savitzky-Golay filter was accurate out to filter sizes of 150 data points. It is also apparent that as the filter size increases, the variance in the results decreases for both filters. Therefore, it is advantageous to choose the largest filter size possible before the filter begins to skew the average of the data. With this in mind, the top-hat filter width was chosen to be ten points, while the 6th-order Savitzky-Golay filter was chosen to be 100 points wide to allow for sufficient margin from the filter parameters where the biases begin to appear in the simulations. It is also important to note that using the Savitzky-Golay filter is considered to be the more robust method due to the lower standard deviation of the results from the Monte-Carlo simulations at the selected filter size. For this reason, using the Savitzky-Golay with a filter width of 101 points (the central data point as well as the 100 neighboring frequency bins) was adopted as the formal method in the analysis. The top-hat filter is only included to allow in the following to show a comparison of the results with an alternative method.

It should be emphasized again that the individual data value in the signal bin is not considered. This is to avoid the possibility that an actual signal would influence the estimated background.

C. Statistics of the Background

Statistics on the data measured at frequencies outside the signal bin were also compiled. With these statistics, detection thresholds and exclusion limits on the signal rate were calculated using the probability distribution of this data. The calculation of the probability distribution of the data is discussed in Section III C 1

while the process of deriving the exclusion limits is discussed in Section III C 2.

1. Detection Thresholds

Because the filtered spectrum of the closed to open power ratio gives an estimate of the expectation value of the background spectrum in terms of $|Z_{\text{closed}}|^2$, it can be used to compile statistics on the noise in the data relative to that expectation value. In turn, these statistics can be used to assess the probability of measuring a given result in the signal bin and calculate a threshold on the signal strength for claiming a detection at a given level of confidence. To do this, the ratio of measured data to the filtered data was calculated over the frequency span of $-100 \mu\text{Hz}$ to $+100 \mu\text{Hz}$. This ratio $W(f)$, is

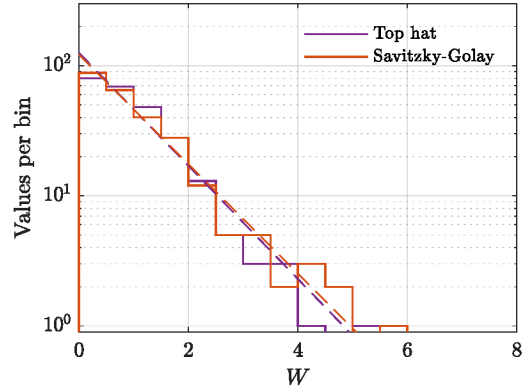
$$W(f) = \frac{|Z(f)_{\text{closed}}|^2}{\mathcal{M}(f)}. \quad (42)$$

The further the value of W is from one the less likely the result.

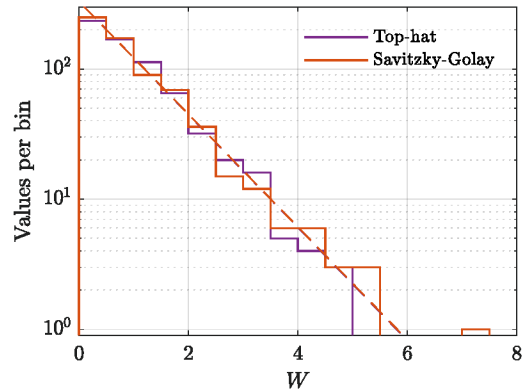
This helps reduce the impact of the shape of the frequency spectrum of the background near the signal bin. If only the average value of the data were used for normalization, the peak in the frequency spectrum would have a major influence on the resulting probability distribution. By filtering, the effects from the broad features of the spectrum can be removed, leaving only the random variance of the data from point to point. The resulting histograms of W from $-100 \mu\text{Hz}$ to $+100 \mu\text{Hz}$ are shown in Figure 15a for S_{\perp} and Figure 15b for S_{\parallel} . In these plots, the purple line is the histogram of W using the top-hat filter and the red-orange line is the histogram when using the Savitzky-Golay filter. At low values of W , the histograms show a similar shape that follows an exponential distribution.

A non-central χ^2 distribution was fit to the histograms produced by the different filters because $|Z_{\text{closed}}|^2$ is in terms of a squared magnitude and the noise on the signal is assumed to have the properties of a normal distribution on the complex plane [27]. The results of the fitting procedure is shown in Figures 15a and 15b as the dashed lines.

The non-central χ^2 distribution is a generalization of the χ^2 distribution. Here we use a scaled version in which Gaussian noise on the complex plane with a variance in the real and



(a) Histogram of the W data (S_{\perp})



(b) Histogram of the W data (S_{\parallel})

FIG. 15: Histograms of the W data from S_{\perp} (a) and S_{\parallel} (b) using the top-hat filter (purple) and the Savitzky-Golay filter (red-orange). The dashed lines show the resulting non-central χ^2 distribution of two degrees of freedom by performing an unbinned maximum likelihood estimation of the non-centrality parameter for the W data using the two filtering methods.

imaginary components of 0.5 each, is offset from the origin by some value ζ giving rise to the non-centrality parameter $|\zeta|^2$. For two degrees of freedom, the probability density function (PDF) of the non-central χ^2 distribution is given by the equation

$$f_{\text{nc}\chi^2}(x|\zeta) = e^{-\left(x + \frac{|\zeta|^2}{2}\right)} I_0\left(\sqrt{2|\zeta|^2 x}\right). \quad (43)$$

In this equation I_0 is the modified Bessel function of the first kind. In this formulation the

	S _⊥ Run	
	W _{TH}	W _{SG}
Observed result: $W(f_k = 0)$	0.76	0.85
Non-centrality parameter: $ \zeta ^2$	$0.000_{-0.000}^{+0.126}$	$0.067_{-0.067}^{+0.114}$
Std. dev. of noise (W)	1.00	1.03
5σ det. Thresh. (W)	14.37	14.81
Significance of observed result	0.7σ	0.8σ
5σ Det. Thresh. ($ Z_{\text{closed}} ^2$)	$(3.7 \pm 0.5) \times 10^{-4}$	$(3.4 \pm 0.5) \times 10^{-4}$
95% C.L. 5σ Ex. Lim (W)	23.97	24.54
95% C.L. 5σ Ex. Lim. ($ Z_{\text{closed}} ^2$)	$(6.1 \pm 0.8) \times 10^{-4}$	$(5.6 \pm 0.8) \times 10^{-4}$

	S _∥ Run	
	W _{TH}	W _{SG}
Observed result: $W(f_k = 0)$	3.35	2.67
Non-centrality parameter: $ \zeta ^2$	$0.000_{-0.000}^{+0.080}$	$0.005_{-0.005}^{+0.086}$
Std. dev. of noise (W)	1.00	1.00
5σ det. Thresh. (W)	14.37	14.41
Significance of observed result	2.1σ	1.8σ
5σ det. Thresh. ($ Z_{\text{closed}} ^2$)	$(1.2 \pm 0.2) \times 10^{-4}$	$(1.5 \pm 0.2) \times 10^{-4}$
95% C.L. 5σ ex. lim (W)	23.97	24.01
95% C.L. 5σ ex. lim. ($ Z_{\text{closed}} ^2$)	$(2.0 \pm 0.3) \times 10^{-4}$	$(2.6 \pm 0.3) \times 10^{-4}$

TABLE III: Summary of the results at the science detector of W for the different analysis methods, including the unbinned maximum likelihood estimate of the non-centrality parameter, the 5σ detection thresholds in terms of W and $|Z_{\text{closed}}|^2$ that are calculated from the resulting probability distributions, as the 95% confidence level 5σ exclusion limits.

non-central χ^2 distribution has a mean value of $1 + |\zeta|^2/2$ and a variance of $1 + |\zeta|^2$, where $|\zeta|^2$ is the non-centrality parameter, .

Table III gives a summary of the statistical analysis of the background data. In this table, the first row gives the ratio of the result in the signal bin versus the expectation value for the top-hat filter W_{TH} and the Savitzky-Golay filter W_{SG} . The second row gives the non-centrality parameter resulting from an unbinned maximum likelihood estimation using the W values from each of the filtering methods within 100 μHz of the signal bin with 95% confidence intervals from this process. As the table shows, all of the fitted non-centrality parameters are extremely close to zero. The standard deviation of the noise in W is calculated from the non-centrality parameter. This gives a measure of the statistical uncertainty of the results. Here,

the sample standard deviation of the W data is not shown because it can be biased due to the correlations between bins that is introduced by dividing the alternative frequency data by the filtered data. Nevertheless, the Monte Carlo simulations also showed that these effects do not significantly affect our ability to calculate the probability distribution of the noise.

The 5σ detection threshold in terms of W are displayed in the fourth row. This corresponds to the value of W that a non-central χ^2 distribution with the non-centrality parameter shown above in the table would produce with a probability of less than 5.73×10^{-7} .

The values of W observed in the signal bin, shown in the first rows of each of the tables, are all far below than the 5σ detection thresholds derived in the sixth row. Therefore, there were no signals significant enough at the expected fre-

quency to claim a discovery for either of the science runs. The significance of the result for S_{\perp} was 0.7σ for the top-hat filter and 0.8σ for the Savitzky-Golay filter, while for S_{\parallel} it was 2.1σ for the top-hat filter and 1.8σ for the Savitzky-Golay filter. Here, it should be emphasized that, despite the presence of a minor excess above the average background, the results of S_{\parallel} are not interpreted as a hint of an interaction between electromagnetic and BSM fields and is believed to be simply a statistical anomaly.

The detection thresholds are also shown in terms of $|Z_{\text{closed}}|^2$ in the sixth row of the table. This is calculated by multiplying the 5σ detection threshold in terms of W by the expectation value of $|Z_{\text{closed}}|^2$ calculated by applying the filters.

2. Setting Exclusion Limits

The non-central χ^2 distribution can also be used to set exclusion limits $|Z_{\text{closed}}|^2$ based on the 5σ detection thresholds described in the previous paragraph. These are found for a 95% confidence level by calculating the non-centrality parameter, or signal strength, which would have a 95% probability of producing a result above the 5σ detection threshold, given the measured probability distribution of the background. The resulting 95% confidence level 5σ exclusion limits are shown in terms of the ratio of signal to estimated background W in the seventh row of Table III, and in terms of the closed to open shutter power ratio $|Z_{\text{closed}}|^2$ in the eighth row. For S_{\perp} , a ratio of the signal to estimated background of 23.97 for the top-filter and 24.54 for the Savitzky-Golay filter can be excluded with 95% confidence, while its possible to be excluded at the same confidence level closed to open shutter power ratios of $(6.1 \pm 0.8) \times 10^{-4}$ for the top-hat filter and $(5.6 \pm 0.8) \times 10^{-4}$ for the Savitzky-Golay filter. For S_{\parallel} , the 95% confidence exclusion limits on the signal to background were 23.97 for the top-filter and 24.01 for the Savitzky-Golay filter, while the same exclusion limits interpreted a closed to open shutter power ratio of $(2.0 \pm 0.3) \times 10^{-4}$ can be excluded for the top-hat filter and $(2.6 \pm 0.3) \times 10^{-4}$ can be excluded for the Savitzky-Golay filter.

D. Summary of Results

Table IV gives a summary of the results S_{\perp} and S_{\parallel} in terms of $\mathcal{P}_{\gamma\leftrightarrow\phi}$. All results in this table are obtain using the Savitzky-Golay filter as this method is considered to be more robust than when the top-hat filter is used. Here the signal bin showed an electromagnetic power to BSM power conversion rate of $(4.4 \pm 0.4) \times 10^{-14}$ for S_{\perp} and $(5.3 \pm 0.5) \times 10^{-14}$ for S_{\parallel} . The average estimated background at the detection system for S_{\perp} was $(4.7 \pm 0.4) \times 10^{-14}$ and for S_{\parallel} was $(3.2 \pm 0.3) \times 10^{-14}$. The estimated backgrounds were then used to calculate a 5σ detection threshold of $(18 \pm 2) \times 10^{-14}$ for S_{\perp} and $(12 \pm 1) \times 10^{-14}$ for S_{\parallel} . In S_{\perp} no excess was seen above the estimated background. Although a slight excess with a significance of 1.8σ was observed in S_{\parallel} (the significance of the result for S_{\perp} was 0.8σ), this is well below the 5σ detection threshold. Therefore, the discovery of new physics cannot be claimed for either measurement run. Signals with a $\mathcal{P}_{\gamma\leftrightarrow\phi}$ above 25×10^{-14} for S_{\perp} and 17×10^{-14} for S_{\parallel} can be excluded with a confidence level of 95%. These limits were calculated using the upper limit of the systematic uncertainty interval such that for S_{\perp} $(22 \pm 3) \times 10^{-14} \rightarrow 25 \times 10^{-14}$ and for S_{\parallel} $(15 \pm 2) \times 10^{-14} \rightarrow 17 \times 10^{-14}$.

E. Verifying the Result

For both S_{\perp} and S_{\parallel} , background signals well above the technical noise were observed at $100\mu\text{Hz}$ above and below the heterodyne frequency. This is not believed to be evidence of BSM field, but rather due to stray light scattering out of the HPL beam and eventually coupling to the detection system. One indication of this is that the background spectrum is spread over many frequencies rather than being concentrated within just the signal bin. The open shutter measurements over the course of both runs showed no indication for this to be the case for a true science signal. The reason why stray light may not remain at the signal frequency is that it takes a random path to the detection system and is Doppler shifted by μHz , due to uncontrolled environmental noise driving mechanical motion of reflecting surfaces and changing the path length to the detection system.

	$\mathcal{P}_{\gamma\leftrightarrow\phi} (S_{\perp})$	$\mathcal{P}_{\gamma\leftrightarrow\phi} (S_{\parallel})$
Observed Result	$(4.4 \pm 0.4) \times 10^{-14}$	$(5.3 \pm 0.5) \times 10^{-14}$
Estimated Background	$(4.7 \pm 0.4) \times 10^{-14}$	$(3.2 \pm 0.3) \times 10^{-14}$
5σ Det. Thresh.	$(18 \pm 2) \times 10^{-14}$	$(12 \pm 1) \times 10^{-14}$
95% CL 5σ Ex. Lim.	25×10^{-14}	17×10^{-14}

TABLE IV: Summary of the results in terms of $\mathcal{P}_{\gamma\leftrightarrow\phi}$.

Other evidence that the signals measured during the science run are not the result of the conversion and reconversion of electromagnetic fields to BSM fields are discussed in this section. This includes the ratio of the results at the veto detector system versus the science detector as well as the average time evolution of the variance of the closed shutter signal. A discussion of how alternative calibration series $C[n]$ are used to check the results can be found in Appendix A 4.

1. Science-to-Veto Ratio

As Table II shows, the results in the signal bin and the expected background are very similar for the science and veto detectors in terms of $|Z_{\text{closed}}|^2$. This is not the expected result if the signals were due to the conversion of BSM fields to electromagnetic fields inside the RC. Based on Equations 37 and 38, a true science signal would produce a value of $|Z_{\text{closed}}|^2$ over 600 times larger at the science detector. This is due to the fact that the cavity mirror at the science port RC1 is roughly 20 times more transmissive than the cavity mirror at the veto port RC2, and also the open shutter signal at the veto detector is roughly a factor of 30 larger due to the reflectivity of the RC itself.

This discrepancy between the actual results and the expected results given the presence of BSM fields could be explained by stray light coupling to both detectors. Scattered light from the HPL still must pass through the optics that the open shutter beam passes in order to couple to the detectors. It is therefore reasonable to assume that the average background seen at the veto and science detectors would be similar if the cause is stray light.

If the background seen at the veto and science detectors is indeed because of stray light, the shape of the background spectrum, as well

as the moving average of the closed shutter data, indicate that the stray light is not scattering into the RC eigenmode before the flat cavity mirror on the COB. Otherwise, there should be correlations present between the spectra measured at the science and veto detectors, as well as the time evolution of their data series demodulated at the heterodyne frequency. As the next section shows, such correlations do not appear to exist.

2. Moving Averaged Data

The time evolution of the data demodulated at the signal frequency was also examined by performing a moving average of the data. This time series can be expressed as

$$Z_{\text{move}}^{10 \text{ ks}}(t) = \frac{1}{N_{10 \text{ ks}}} \sum_{n=n_t-5 \text{ ks}}^{n_t+5 \text{ ks}} \frac{H_p[n]}{C[n]}, \quad (44)$$

with $N_{10 \text{ ks}}$ representing the number of point in a 10 ks period of data. With a moving average time constant of 10 ks the summation on the heterodyne data takes place for the 5,000 seconds before and the 5,000 seconds after time t . The resulting moving average signal is shown on the complex plane for the science and veto detectors in Figures 16a and 16c, respectively for S_{\perp} , and for S_{\parallel} in Figures 17a and 17c.

In these plots the colorbar to the right gives the dates associated with Z_{move} . Breaks in the data indicate periods in time when the system was not locked. Here, only segments of data in which the system was locked for the entire duration of the 10,000 s of averaging appear on the plot. Here, 10,000 s is chosen because it is a long enough time that the stray light is well above the technical noise of the system, but short enough that the stray-light signal remains relatively coherent (this is explained in more detail in the

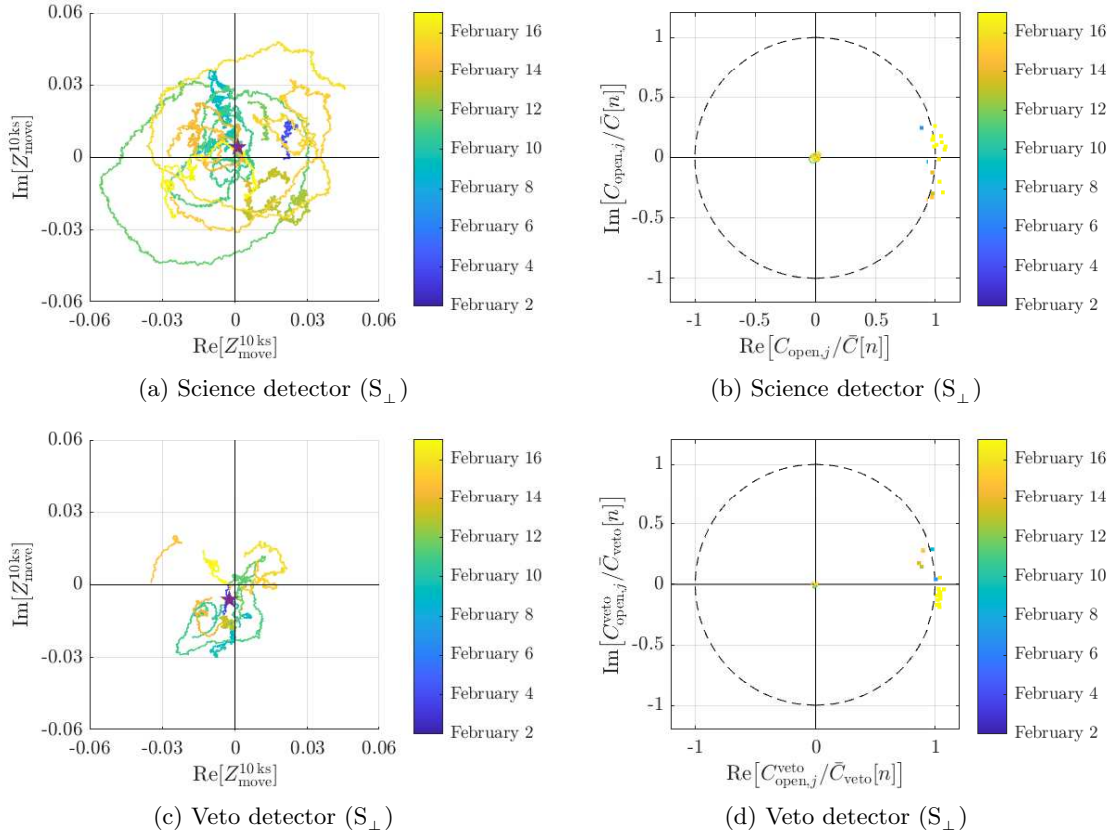


FIG. 16: Moving average of the calibrated closed to open shutter ratio on the complex plane measured by the science detector during S_{\perp} . The color bar indicates the time since the beginning of data taking and the moving average filters over 10,000 seconds. The purple star indicates the mean value of the data.

next section). The axes on the plots are unitless, as the magnitude of the data represents the square root of the closed to open shutter ratio. The purple star gives the result of integrating Z_{closed} for each run, which is analogous to the mean value of the complex time series shown in the plots.

From these data, it is also interesting to consider that for time periods of over 10,000 s, the signals measured in the heterodyne data appear to be more stable in phase at the veto detectors in both runs. This can be seen in Figures 16a, 16c, 17a, and 17c as Z_{move} is much more dynamic on the complex plane for the science detectors in both runs. If one were to assume that stray light is responsible for the signals seen in these plots, this effect could be explained by the shorter path to the veto detector providing

fewer opportunities for the phase of the stray-light signal to be modulated. Because of this, when $|Z_{\text{closed}}|^2$ is calculated the constructive interference produces a larger value at the veto PD relative to the science PD.²⁷ This can be seen directly in the ratio of the average of $|Z_{\text{move}}|$ measured at the science and veto versus the same ratio for $|Z_{\text{closed}}|^2$ and \mathcal{M} in the signal bin. As Table II shows $|Z_{\text{closed}}|^2$ is roughly two times larger at the veto detector compared to the sci-

²⁷ The fluctuations in the phase of the signal measured at the heterodyne frequencies are believed to be the dominant source of this effect in comparison to fluctuations in the amplitude as the phase of the moving average time series appears to travel through several cycles over the duration of the run and in general more so than the time series moves in amplitude.

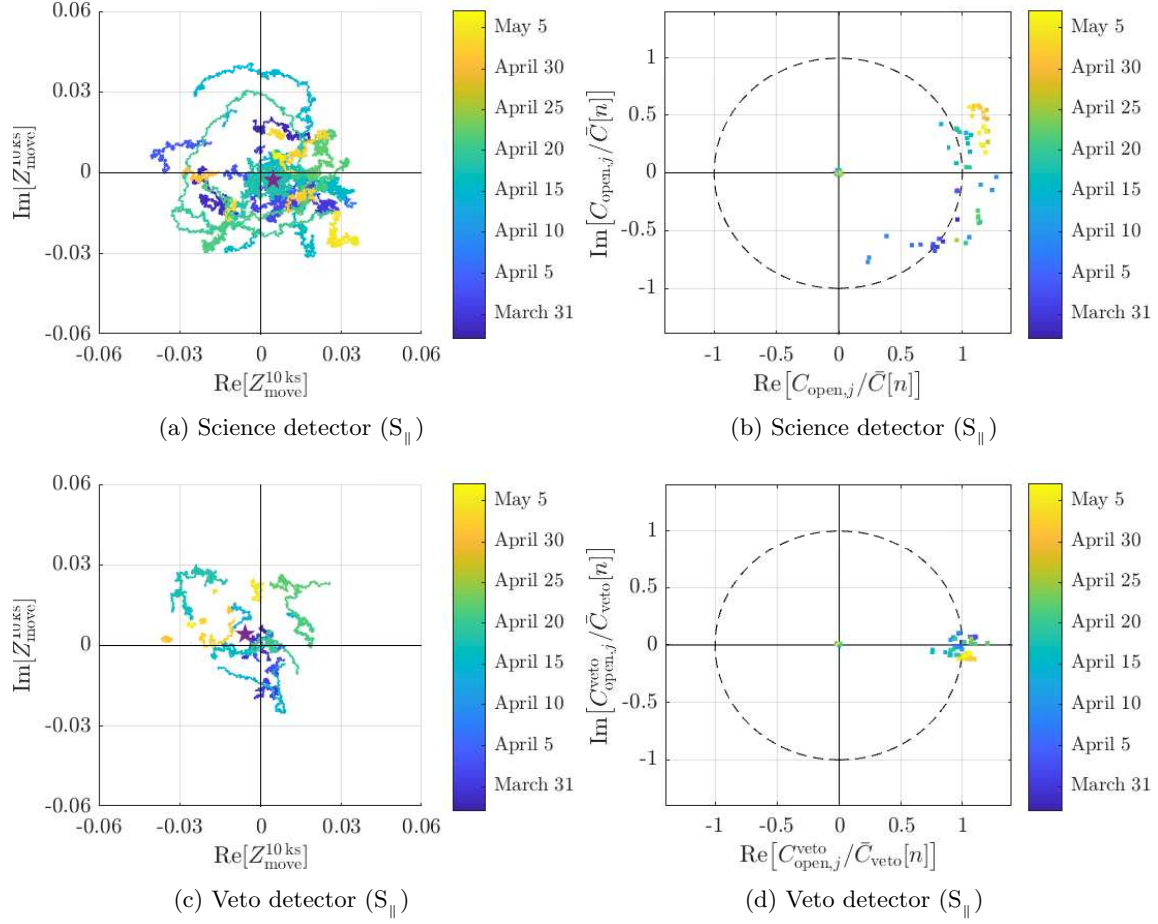


FIG. 17: Moving average of the calibrated closed to open shutter ratio on the complex plane measured by the veto detector during S_{\parallel} . The color bar indicates the time since the beginning of data taking and the moving average filters over 10,000 seconds. The purple star indicates the mean value of the data.

ence detector for both runs, corresponding to a roughly 40% larger magnitude. For S_{\perp} , \mathcal{M} is very similar for both filters although the results are slightly larger at the veto detector. For S_{\parallel} , \mathcal{M} is nearly four times larger for both filters at the veto detector leading to the magnitude of Z_{closed} being twice as large at the veto detector. In contrast, the results of Z_{move} discussed in the previous paragraph show a mean magnitude that was 30% higher for the science detector during S_{\perp} compared to the veto detector and a roughly equal average magnitude between the two detectors for S_{\parallel} . This indicates that there may be more instantaneous stray-light power relative to the open shutter power present at the

science detector compared to the veto detector, but the coherent sum shows a lower closed-to-open-shutter power ratio over the duration of the run at the science detector due to the larger phase fluctuations in the stray-light there.

In either case, the dynamic nature of the signals on short time scales (10,000s) and their relatively large size compared to the expectations of the technical noise are an indication that these signals are due to the spurious coupling of HPL power to the detection systems when the shutter is closed and not the reconversion of power from the BSM fields.

Figures 16b, 16d, 17b, 17d show a comparison of the moving average of the closed shutter

data versus complex value of $C_{\text{open},j}$ normalized by the average value of $C[n]$ for each search. This is meant to convey what the phase evolution of an actual signal would look like compared to the data measured during the closed shutter periods. Here it is apparent that the values of $C_{\text{open},j}$ are well offset from origin clustered, for the most part, near $1 + 0i$. The difference between the $C_{\text{open},j}$ data measured at the veto detector, which in general has much smaller deviations from $1 + 0i$ than the data measured at the science detector, shows the effect that the detuning of the HPL from the resonance of the RC has on the complex phase. This effect is discussed in more detail in Section IV B 1. Again, the open shutter calibration will compensate for these effects.

3. Time Evolution of Backgrounds

One potential risk in the experiment is that unknown systematic effects could alter the phase of an actual science signal and prevent it from coherently summing when $|Z_{\text{closed}}|^2$ is integrated. This can be checked by examining stretches of data with shorter time scales where it is clear that the optical system is maintaining the coherence between the HPL and the oscillators of the detection system, and keeping it on resonance with the RC. Potential signals that may not integrate coherently on longer time scales would have a better chance to be measured in this case. The open shutter periods can show what these signals may look like. The blue traces in Figures 18a and 18c show, for both runs, the average of every segment of open shutter data with a length given by the x -axis for S_{\perp} and S_{\parallel} respectively. In these plots, the open shutter periods were normalized by the calibration time series. There are also lines for 1σ uncertainty (between the 16th and 84th percentile) and 2σ uncertainty interval (between percentiles 2.5 and 97.5). It is important to note that these uncertainty intervals converge toward the average of the data as the time increases, indicating the presence of a stable signal in the data.

Figures 18b and 18d show the same analysis, but instead performed on the closed shutter periods of both science runs. The range of integration times is much wider on these plots as there were more long stretches of data during the

closed shutter periods. One notable difference between these plots and the plots of the open shutter periods is the total lack of convergence of the uncertainty interval in the closed shutter data. This could be interpreted as an indicator that the blue traces are not showing a coherent signal emerging from the noise in the system, but rather a random signal with mean magnitude squared that is equal to the blue line. This is strong evidence that the results in the signal bin are due to the presence of the background noise and not the conversion and reconversion of BSM fields to an electromagnetic field.

It is also interesting to compare the mean backgrounds between S_{\perp} and S_{\parallel} . For integration times of 3000 s, the mean background is nearly a factor of two lower for S_{\parallel} , but by 10,000 s integration times they are only on the order of 40% lower. These results appear to roughly agree with the heterodyne moving average analysis in Section III E 2 and the estimate of the backgrounds from the stray-light models in Section III B, that both appear to indicate a reduction of the stray light between S_{\perp} and S_{\parallel} . Again, this is notable as the shutter material (see Section I D) between the runs was changed from the silver mirror to the plate coated with absorptive black foil. Additionally, all but one of the open slots in the shutter wheel were closed for S_{\parallel} , while for S_{\perp} , only the slot with the actual shutter was occupied.

F. Central Optical Bench Transmissivity

Precise measurements of combined transmissivity of the COB mirrors mounted on the optical axis of the system (see Figure 6),

$$T_{\text{COB}} = T_{\text{LT1}} T_{\text{MZ3}} T_{\text{MZ4}} T_{\text{LT2}}, \quad (45)$$

are essential to calibrate the ratio of the closed to open shutter power in terms of the conversion rate between electromagnetic and BSM fields $\mathcal{P}_{\gamma \leftrightarrow \phi}$. This can be seen in Equation 25 with the critical parameter being $T_{\text{COB}} T_{\text{RC2}}$. T_{COB} is also an important to calculating the field overlap η as Section IV B will explain.

Ex-situ measurements of these optics were performed after the completion of the first science campaign with the COB removed from the vacuum system. As the transmissivities of these optics were strongly dependent on the angle of incidence (~ 1.5 ppm per degree for the mirrors




FIG. 18: The mean open shutter signal (a) and closed shutter background (b) for S_{\perp} and open shutter signal (c) and closed shutter background (d) for S_{\parallel} is shown in blue as a function of the integration time. The black lines show the 1σ uncertainty (between the 16% and 84%) and the gray lines show the 2σ uncertainty interval (between 2.5% and 97.5%).

sealing the light-tight housings and ~ 150 ppm per degree for the MZ mirrors) the flat cavity mirror of the RC was used as an alignment reference to ensure that the laser used in the setup passed through the optics at the same angle as the HPL light during the open shutter periods of the science runs. The measurement procedure involved shining a laser coupled to an optical fiber on the mirror being tested. The light reflected from the optic was then incident normal to the RC flat mirror and reflected back through the fiber. This ensured that the incident beam was aligned to the angle of the RC eigenmode (assuming no changes in the alignment of the mirror after the science campaign)

to better than 1 mrad. The results are shown in Table V. Similar results were expected for the pair of light tight housing mirrors and the pair of Mach-Zehnder mirrors, and the measurements of LT1 and LT2 both resulted in a transmissivity of 12.4 ± 0.5 ppm. MZ3 was measured to have a transmissivity of 452 ± 10 ppm while the transmissivity of MZ4 was 275 ± 10 ppm. While MZ3 and MZ4 were expected to have the same transmissivity as their dielectric coatings were applied during the same coating run, the discrepancy between them is believed to be related to differences in their angles of incidence with respect to the optical axis of the system. The uncertainty in these measurements is due to a

Ex-situ		In-situ	
Optic	Transmissivity	Optic	Transmissivity
LT1	(12.4 ± 0.5) ppm	LT1	(11 ± 1) ppm
LT2	(12.4 ± 0.5) ppm	LT2	(11 ± 1) ppm
MZ3	(452 ± 10) ppm	MZ3 MZ4	(0.137 ± 0.007) ppm
MZ4	(275 ± 10) ppm		
T_{COB}	$(1.9 \pm 0.1) \times 10^{-17}$	T_{COB}	$(1.7 \pm 0.2) \times 10^{-17}$
$T_{\text{COB}} T_{\text{RC2}}$	$(9.7 \pm 1.2) \times 10^{-23}$	$T_{\text{COB}} T_{\text{RC2}}$	$(8.5 \pm 1.4) \times 10^{-23}$

TABLE V: The transmissivity of the on-axis COB optics for p -polarized light. Here the combined transmissivity of the COB on-axis optics, T_{COB} , is calculated from the product of the transmissivities of the individual components.

combination of the uncertainty in the calibration of the power meters used in the setup, the presence of stray light at the power meters measuring the weak transmitted signals, and the uncertainty of the angle of incidence of the laser on the mirrors.

In-situ measurements of the on-axis COB optics were also performed in between S_{\perp} and S_{\parallel} . Here, the transmissivity of LT1 was observed via a measurement of the ratio of HPL power directly at the input window to the vacuum system in NL and then at the High-Power port of the CH vacuum tank (see Figure 6). This gave a value of 11 ± 1 ppm. The transmissivity of LT2 was measured with a similar method, but instead measuring the RL power just before the Injection Port at the CH tank, also giving a value of 11 ± 1 ppm. The combined transmissivity of MZ3 and MZ4 was measured using a power meter to measure the HPL power at the PLL2 port, and then at the Mach-Zehnder port by measuring the amplitude of the RL-HPL beatnote. The spatial overlap of the RL-HPL beatnote at optimal alignment could be directly measured at the PLL2 port because it was possible to measure the static power of both lasers there. This value could then be used to compensate for these effects in the measurement of the beatnote amplitude at the Mach-Zehnder port.²⁸ The combined transmissivity

of the MZ mirrors found in this measurement was 0.137 ± 0.007 ppm.

The result of the combined transmissivities of the MZ mirror from the ex-situ measurements gave 0.124 ± 0.005 ppm. Thus, it is clear that the individual measurements of each of the on-axis COB optics agree between the two methods. Even more relevant is the total transmissivity T_{COB} , in which a value of $(1.9 \pm 0.1) \times 10^{-17}$ was found for the ex-situ measurement with $(1.7 \pm 0.2) \times 10^{-17}$ measured for the in-situ measurements. In short, these measurements agree within their uncertainty.

Finally, the most relevant parameter for the calibration of the results in terms of $\mathcal{P}_{\gamma \leftrightarrow \phi}$ is the product of the transmissivity of COB optics and the cavity mirror RC2. For this, the ex-situ measurements gave $T_{\text{COB}} T_{\text{RC2}} = (9.7 \pm 1.2) \times 10^{-23}$ while the in-situ results were $T_{\text{COB}} T_{\text{RC2}} = (8.5 \pm 1.4) \times 10^{-23}$. Again, the measurements agree within their uncertainty. This gives us confidence in the accuracy of the overall calibration of the results. We consider the method for the ex-situ measurements to be more robust as we were able to directly access the optics. Therefore, the ex-situ measurements were used for the actual calibration of the results.

IV. PERFORMANCE OF THE OPTICAL SYSTEM

Using the open shutter periods to calibrate the science data means that measuring the parameters of the optical system, such as the coupling of the HPL to the RC, is not strictly necessary. However, these measurements are impor-

²⁸ It should be noted that to find the beatnote amplitude at each port for the ‘optimal’ alignment, RL was realigned to maximize the the beatnote amplitude in both cases. This was necessary to compensate for the alignment drift of the MZ on the COB.

tant for verifying the performance of the optical system and also help identify the cause of changes in the calibration series. Calculating the important metrics to assess the performance of the optical system, such as the resonant enhancement of the RC and the total power overlap between the HPL and the RC, required measuring a number of other parameters. Some of these were dynamic parameters that were measured during maintenance periods during the run, while other static parameters were measured before and/or after the run. This section discusses how these measurements were made and gives a summary of the performance of the optical system during the science runs.

A. Regeneration Cavity Parameters

The RC is a critical component of the optical system. Its resonant enhancement factor and the power overlap between its eigenmode and the HPL directly affect the sensitivity of the experiment. A more complete characterization of the RC can be found in Ref. [20]. This section will focus on how the RC performance was characterized during the first science campaign and the results of that effort. Explanations of the resonant enhancement factor, transmissivity, and reflectivity of the RC are given in Sections IV A 1, IV A 2, and IV A 3 respectively, along with a discussion of the measurements of these parameters during the science runs.

1. Resonant Enhancement

The resonant enhancement factor of the RC, β_{RC} , can be defined as the ratio of the power in the signal field that exits the cavity at the mirror coupled to the detection system, to the signal power at the same location if the cavity were not there at all, for a signal field that is on resonance and in the spatial mode of the RC. β_{RC} therefore gives a measure of the amplification of the science signal. In this case, the BSM field reconverts to an electromagnetic field inside the cavity and is stored there while being reflected back and forth between the mirrors. During this time the field experiences the cavity gain as it constructively interferes with the electromagnetic field that is newly reconverted from the BSM field. In this way, the field is



(a) S_{\perp} (b) S_{\parallel}

FIG. 19: The estimated resonant enhancement of the RC is shown for S_{\perp} (a) and S_{\parallel} (b) with error bars based on the uncertainty of the measured storage time. The black vertical lines in the plot of the S_{\parallel} data give the times at which the RC mirrors were realigned to optimize its resonant enhancement. No vertical black lines are shown in the plot of the S_{\perp} data as the RC mirrors were never manually realigned during the course of this run.

injected via the reversion of the BSM field to an electromagnetic field, rather than transmitting through an input coupling mirror. The electromagnetic field that has built up in the cavity is then transmitted by RC1 and directed to the science detector. The resonant enhancement, β_{RC} , can be expressed as [13]

$$\beta_{\text{RC}} = \frac{T_{\text{RC1}}}{(1 - \sqrt{1 - A})^2} \simeq \frac{4T_{\text{RC1}}}{A^2}. \quad (46)$$

In this equation, T_{RC1} is the power transmissivity of RC1 while A is the total round-trip attenuation of the power circulating in the cavity. This is due to the combination of mirror transmissivities and other optical losses including scattering from the mirror surface roughness, absorption in the mirror coatings, and clipping losses on the combined free apertures of the magnet string and reflective surfaces of the mirrors. The approximation assumes $A \ll 1$ and therefore $\sqrt{1-A} \simeq 1 - A/2$, leading to an error in the resonant enhancement on the order of 0.01%.

Assessing the total attenuation A is critical to knowing not only the resonant enhancement factor during the run, but a number of other parameters as well. Fortunately, A can be measured via a combination the cavity FSR f_0 and storage time τ . The relationship between these parameters is

$$A = 1 - e^{-\frac{2}{f_0\tau}} \simeq \frac{2}{f_0\tau}, \quad (47)$$

resulting in

$$\beta_{\text{RC}} \simeq T_{\text{RC1}} f_0^2 \tau^2, \quad (48)$$

allowing it to be determined by also measuring the FSR and storage time of the RC. In this case, using the approximate equation leads to a relative error in the measured resonant enhancement factor on the order of 0.01%. An estimation of the uncertainty in these measurements is discussed later in this section, but using this approximation leads to errors well below the other sources of error in the measurement.

The relationship between the finesse of the RC and β_{RC} can be approximated by

$$\mathcal{F} = \frac{f_0}{2\nu_0} = \frac{\pi\sqrt{1-\sqrt{A}}}{1-\sqrt{1-A}} \quad (49)$$

$$\simeq \pi \frac{2}{A} \simeq \pi \sqrt{\frac{\beta_{\text{RC}}}{T_{\text{RC1}}}}, \quad (50)$$

while the finesse can also be defined in terms of the storage time τ and FSR f_0 of the cavity [28, 29] as

$$\mathcal{F} = \pi f_0 \tau. \quad (51)$$

Details on the technique for measuring the storage time and the values for T_{RC1} and T_{RC2} can be found in the Refs. [20] and [30]. To measure the

cavity FSR f_0 during the run, the laser AL was PDH locked to the cavity. Then the PLL3 frequency was tuned near 45 FSRs, until the power transmitted by the cavity was at a maximum. The value of the PLL3 frequency at that point was then divided by 45 to give the FSR of the RC. The uncertainty in this measurement was 1 Hz, leading to an uncertainty on the FSR of 0.02 Hz (see Table I).

The measurements of β_{RC} during the maintenance periods are shown in Figures 19a and 19b for S_{\perp} and S_{\parallel} respectively. For S_{\parallel} , the black vertical lines in the plot indicate times at which the RC eigenmode position on the mirrors was aligned to improve the storage time and in turn the resonant enhancement factor (no such realignments were required for S_{\perp}). For S_{\perp} the mean value of β_{RC} associated with the closed shutter periods was 7150 with a range from 6810 to 7120 and RMS fluctuations from its mean value of 90. The uncertainty on the measurements of β_{RC} was 80, mostly due to the uncertainty on the transmissivity of RC1 and the uncertainty on the measurements of the storage time. For the closed shutter periods of S_{\parallel} the mean value of β_{RC} was estimated to be 6870 with a range from 6640 to 6910 with an RMS of 70. These measurements had the same uncertainty as those performed during S_{\perp} . While a value below 5900 was measured during one of the periods of S_{\parallel} , and a value above 7200 was measured during one of the periods of S_{\perp} , there were no closed shutter periods during time periods associated with these measurements. For this reason this data point is not considered in the range of data mentioned above.

2. Transmissivity

The transmissivity of the cavity for a laser occupying its spatial eigenmode (the spatial coupling will be examined later) can be expressed as [28]

$$T_{\text{RC}} = \frac{T_{\text{RC1}} T_{\text{RC2}}}{\left| 1 - \sqrt{1-A} e^{i2\pi \frac{\Delta\nu}{f_0}} \right|^2} \quad (52)$$

$$\simeq \frac{4T_{\text{RC1}} T_{\text{RC2}}}{A^2} \frac{1}{1 + \left(\frac{4\pi \Delta\nu}{A f_0} \right)^2}, \quad (53)$$

with $\Delta\nu$ being the offset of the laser frequency from its nearest resonance. Assuming that the laser is very nearly on resonance with the cavity ($\Delta\nu \ll Af_0$), the maximum transmissivity of the cavity is

$$T_{\text{RC,max}} = \frac{T_{\text{RC1}} T_{\text{RC2}}}{(1 - \sqrt{1 - A})^2} = T_{\text{RC2}} \beta_{\text{RC}}. \quad (54)$$

In this equation it is clear that due to the common factors, $T_{\text{RC,max}}$ can be expressed as the product of the resonant enhancement and the transmissivity of the flat cavity mirror on the COB (again ignoring the spatial or longitudinal field overlap). It should also be noted that $T_{\text{RC,max}}$ is independent of which side the input field is incident on. The maximum transmissivity of the RC on resonance can therefore be expressed as the product

$$T_{\text{RC,max}} \simeq T_{\text{RC1}} T_{\text{RC2}} f_0^2 \tau^2. \quad (55)$$

For S_{\perp} , during the closed shutter periods $T_{\text{RC,max}}$ was estimated to have an average value of 3.6% with a range from 3.5% to 3.6% and RMS deviations from the mean of 0.05%. The uncertainty of T_{RC2} was the primary contributor to the absolute uncertainty of $\pm 0.4\%$ on $T_{\text{RC,max}}$. During the closed shutter periods of S_{\parallel} the average of $T_{\text{RC,max}}$ was estimated to be 3.5%, with a range from 3.4% to 3.5%, an RMS of 0.04%, and an absolute uncertainty of $\pm 0.3\%$.

As Equation 53 shows, detuning the incident laser frequency from the cavity resonance will re-

duce its transmissivity. Measurements of the detuning $\Delta\nu$ between the HPL frequency and resonance of the RC are discussed in Section IV B 1 and were used to infer the transmissivity of the RC during the first science campaign. During the closed-shutter periods of S_{\perp} , T_{RC} had a mean value of 3.3%, with RMS fluctuations about the mean of 0.2%, a maximum value of 3.6%, and a minimum value of 2.6%. The uncertainty of these measurements was 0.3%. For S_{\parallel} , the closed-shutter periods gave a mean value of 3.0% for T_{RC} , with RMS fluctuations of 0.5%, a maximum of 3.5%, and a minimum of 1.4%. The uncertainty was the same as the measurements during S_{\perp} at 0.3%. In both runs the uncertainty was almost entirely driven by the uncertainty in the measurement of the transmissivity of the flat RC mirror mounted to the COB, RC2.

3. Reflectivity

The reflectivity of the RC is used in the calibration of several critical parameters used to evaluate the performance of the optical system. Unlike the transmissivity, the reflectivity differs depending on the side of the cavity on which the laser is incident. The reflectivity experienced by the HPL when incident on the flat RC mirror in the CH cleanroom (RC2) will be referred to as $R_{\text{RC,CH}}$ while the reflectivity for lasers incident on the curved cavity mirror in the NR cleanroom (RC1) will be referred to as $R_{\text{RC,NR}}$. These quantities can be expressed as [28]

$$R_{\text{RC,CH}} = \left| \frac{1}{R_{\text{RC2}}} \left(R_{\text{RC2}} - \frac{T_{\text{RC2}}}{1 - \sqrt{1 - A} e^{i2\pi \frac{\Delta\nu}{f_0}}} \right) \right|^2 \simeq \left| \left(1 - \frac{2T_{\text{RC2}}}{A} \frac{1}{1 - i \frac{\Delta\nu}{\nu_0}} \right) \right|^2, \quad \text{and} \quad (56)$$

$$R_{\text{RC,NR}} = \left| \frac{1}{R_{\text{RC1}}} \left(R_{\text{RC1}} - \frac{T_{\text{RC1}}}{1 - \sqrt{1 - A} e^{i2\pi \frac{\Delta\nu}{f_0}}} \right) \right|^2 \simeq \left| \left(1 - \frac{2T_{\text{RC1}}}{A} \frac{1}{1 - i \frac{\Delta\nu}{\nu_0}} \right) \right|^2. \quad (57)$$

In both cases, the approximations assume the reflectivity of RC1 and RC2 are one. This can be expressed in terms of measured parameters

as

$$R_{\text{RC,CH}} \simeq \left| 1 - \frac{T_{\text{RC2}} f_0 \tau}{1 - i2\pi\tau\Delta\nu} \right|^2, \quad \text{and} \quad (58)$$

$$R_{\text{RC,NR}} \simeq \left| 1 - \frac{T_{\text{RC1}} f_0 \tau}{1 - i2\pi\tau\Delta\nu} \right|^2. \quad (59)$$

In Sections IV B 4 and IV B 5 $R_{\text{RC,CH}}$ is used to calibrate the spatial overlap and field overlap measured at the veto detection system respectively. During the closed shutter sections of S_{\perp} average value of $R_{\text{RC,CH}}$ was 92%, with a range from 92% to 94%, and RMS fluctuations from the mean of 0.6%. For S_{\parallel} the RC reflectivity gave an average value of 93%, with values ranging between 92% to 97%, and an RMS of 1%. The uncertainty of these measurements was 1%, largely driven by the uncertainty in measurements of the transmissivity of RC2.

The reflectivity of the RC for the AL is used in Section to calculate its spatial overlap with the cavity eigenmode. Since AL is injected to the RC via RC1 and is frequency stabilized directly to a cavity resonance, Equation 59 for $R_{\text{RC,NR}}$ can be used with $\Delta\nu = 0$. In this case the reflectivity of the cavity had an average value of 2.6% during the closed shutter periods of S_{\perp} with a range from 2.4% to 3.1% and an RMS of 0.2%. During S_{\parallel} the average reflectivity was 3.0% while the range was from 2.8% to 3.4%. The RMS of the reflectivity during these measurements was 0.2%. For both runs the absolute uncertainty in the reflectivity was 1.3% and was due to the uncertainty in the measurements of the cavity storage time and the transmissivity of RC1.

B. Field overlap

The field overlap between the HPL and the RC, η , is a critical metric for evaluating the performance of the optical system. As explained in Section I A, it can be expressed as the product of spectral (η_z), polarization (η_{pol}), and spatial (η_{xy}) components.

The spectral component, also called the longitudinal overlap, is discussed in Section IV B 1 and can be calculated exclusively from measurements of the RC storage time and the frequency detuning of HPL from its resonance. The polarization overlap, discussed in Section IV B 2, is calculated from the orientation of polarization state of the HPL with respect to the magnetic field.

The spatial overlap is measured by calibrating the results of the open shutter periods in terms of the HPL power leaving the cavity P_{open} , and dividing by the expected HPL power at this point for a laser occupying the spatial eigenmode

of the RC. At the science detector, this can be expressed as

$$\eta_{xy}^2 = \frac{P_{\text{open}}}{T_{\text{COB}} T_{\text{RC}} P_i}. \quad (60)$$

The expected power is therefore the product of the combined transmissivities of the COB optics T_{COB} , the transmissivity of the RC, T_{RC} , and the HPL power traveling through the production magnet string P_i .

The veto detector can also be used to measure the field overlap between the HPL and the RC. There are even some advantages to using the veto detector as the HPL does not need to pass through the cavity to reach it resulting in a HPL power is higher at the veto PD. Therefore, the following sections include a discussion of the field overlap measured at both the science and veto detectors. The process of measuring the ratio P_{open}/P_i is explained in Section IV B 3. Projections of the spatial overlap during the first science campaign are discussed in Section IV B 4, while the total field overlap η is discussed in Section IV B 5. The methods used to optimize the spatial overlap are described in Section IV B 6.

1. Longitudinal Overlap

The longitudinal field overlap of the RC η_z , expresses the relative reduction in the resonant enhancement factor experienced by a light field which is detuned from the cavity resonance by some frequency $\Delta\nu$, the difference between the frequency of the HPL and the nearest resonance of the RC. It can be expressed as

$$\eta_z = \frac{1 - \sqrt{1 - A}}{\left| 1 - \sqrt{1 - A} e^{i2\pi \frac{\Delta\nu}{f_0}} \right|} \simeq \frac{1}{\sqrt{1 + \left(\frac{4\pi\Delta\nu}{Af_0} \right)^2}}. \quad (61)$$

This can also be expressed in terms of the measured parameters τ and $\Delta\nu$ as

$$\eta_z \simeq \frac{1}{\sqrt{1 + (2\pi\tau\Delta\nu)^2}}. \quad (62)$$

A laser that is detuned from the cavity resonance will also accumulate some additional phase offset during its transmission through the cavity. This phase offset θ_z , can be found using

the equation

$$\theta_z = \arg \left[\frac{1 - \sqrt{1 - A}}{1 - \sqrt{1 - A} e^{i2\pi \frac{\Delta\nu}{f_0}}} \right] \quad (63)$$

$$\simeq \arctan(2\pi\tau\Delta\nu). \quad (64)$$

This effect is important as it means a detuning of the HPL from the resonance of the cavity can distort the open shutter phase. This is one of the reasons why the calibration array is expressed in terms of a complex variable as it will be able to compensate for these phase changes when the closed shutter data are normalized.

The initial FSR of the RC was used to establish the frequencies of the PLLs that control the frequency of the HPL relative to the frequency of the AL light transmitted by the RC. At 33 FSRs, this static frequency difference meant that the HPL frequency was tuned to be on resonance with the RC when the run was initially set up. However, as the length of the cavity drifted and the FSR changed, the HPL frequency was detuned from resonance and the longitudinal overlap η_z , decreased. These effects also led to changes in θ_z , which was exhibited as a change in the phase of the open shutter measurements C_{open} .

As Section IV A 1 described, it was possible to measure the FSR with a precision a precision of the 0.02 Hz. The frequency detuning $\Delta\nu$ of the HPL with respect to the RC could therefore be measured with a precision of 0.7 Hz because the frequency spacing between AL and the HPL was 33 FSRs.

Figures 20a and 20b show $\Delta\nu$ for S_{\perp} and S_{\parallel} , as a result of the cavity FSR measurements. For S_{\perp} , the average absolute detuning of the HPL from the RC resonance (the mean value of $|\Delta\nu|$) estimated during the closed shutter periods was 5.4 Hz with a minimum detuning of 0 Hz, a maximum detuning of 14.5 Hz, and RMS deviations from the mean of 4.0 Hz. S_{\parallel} showed an average absolute detuning of the HPL frequency of 8.9 Hz with a minimum detuning of 0.3 Hz, a maximum detuning of 28.4 Hz, an RMS of 6.4 Hz. The uncertainty in these measurements was 0.7 Hz as explained in the previous paragraph.

Figures 20c and 20d show the projected HPL-RC longitudinal power overlap η_z^2 , from measurements of the cavity FSR and storage time. The black vertical lines in Figure 20d show periods when the cavity length was adjusted to try

to bring η_z^2 back to one during S_{\parallel} . No such readjustments were performed on the length of the cavity during S_{\perp} . In these plots S_{\perp} showed an estimated average value of η_z^2 during the closed shutter periods of 92.8% with a minimum of 71.9%, a maximum value of 100%, and RMS fluctuations from the mean of 7.1%. For S_{\parallel} the average value of η_z^2 was 85.5%, while its minimum value was 40.6%, its maximum value was 100%, and its RMS was 14.4%. The uncertainty on these measurements were also primarily limited by the precision of the measurement of the RC FSR and was between 0.2% and 0.3%.

The projections of θ_z based on measurements of the cavity FSR and storage time are shown in Figures 20e and 20f. Again, the black lines show periods where the cavity length was readjusted. For S_{\perp} , the mean absolute deviation of the longitudinal phase offset from zero estimated during the closed shutter periods was 0.22 rad and had a minimum value of 0 rad with a maximum value of 0.56 rad. The RMS fluctuations from the mean in this case were 0.16 rad. For S_{\parallel} the average absolute deviation of θ_z was 0.33 rad with a minimum of 0.01 rad, a maximum of 0.88 rad and an RMS of 0.22 rad. Like the previous calculations, the uncertainty in these measurements was mainly due to the uncertainty of the FSR measurement and was between 0.01 and 0.03 rad.

It is interesting to compare the data shown in Figures 20c, 20d, 20e, and 20f, to their corresponding plots of the open shutter data C_{open} measured at the science detectors in Figures 8a, 8b, and 9a and 9b. Even though these two sets of plots were compiled from different sources of data, the relationship between C_{open} , and η_z and θ_z is obvious. This demonstrates how the detuning of the HPL frequency from the RC was perhaps the most significant source of drifts in C_{open} .

2. Polarization Overlap

The polarization overlap η_{pol} , quantifies the mismatch between the ideal polarization state of the HPL for a given run (\hat{x} for S_{\perp} and \hat{y} for S_{\parallel}) and the orientation of the magnetic field (defined to be in the \hat{y} direction). The error in the orientation of the optical tables with respect to the magnetic fields is on the order of 100 μrad . The polarization state of the HPL for S_{\perp} was

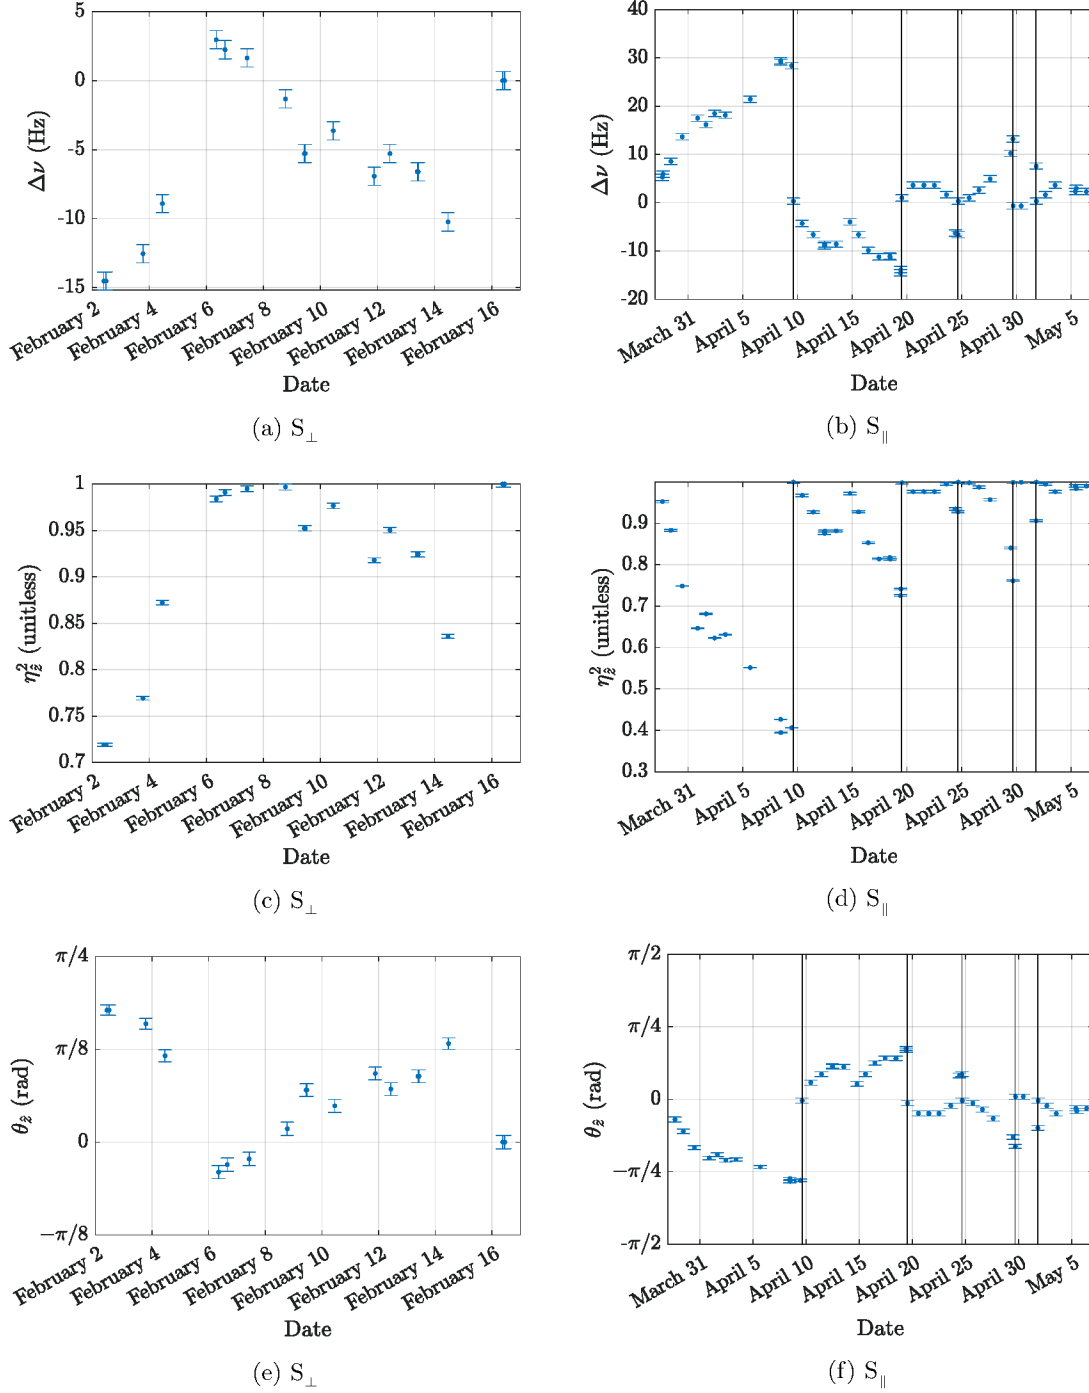


FIG. 20: The measured detuning $\Delta\nu$ of the HPL with respect to the RC resonance is shown for S_{\perp} (a) and S_{\parallel} (b) with error bars based on the uncertainty of the measured FSR. The corresponding longitudinal overlaps in terms of power η_z^2 are shown in (c) and (d), while the phase shift induced by the detuning θ_z is shown for each run in (e) and (f). The vertical black lines in (b), (d), and (f) show times during S_{\parallel} in which the FSR was manually reset to eliminate the detuning by adjusting the physical length of the cavity. No such adjustments were performed during S_{\perp} .

tuned for maximum transmission through LT1 (which was rigidly mounted to the COB). For S_{\parallel} an additional $\lambda/2$ wave-plate, whose orientation was precisely aligned, was also located on the COB in the path to LT1. The polarization state of the AL and LO are set by a rotatable half wave-plate in NR. The exact angle of the ‘roll’ of the COB with respect to the optical axis was not measured.

A measurement of the polarization state of the HPL before being injected to the magnet string via the vacuum window in NL showed an azimuthal angle in the polarization state of $7.9 \pm 0.1^\circ$ immediately following S_{\perp} . This would correspond to a power loss in the desired polarization state of less than 2%. As this is well below the uncertainty from other sources, the polarization overlap was not considered in the analysis. For S_{\parallel} the polarization overlap was also not considered as the azimuthal polarization angle was $89.6 \pm 0.1^\circ$ corresponding to a power loss in the desired polarization state of roughly 50 ppm.

3. Open Shutter Power Ratio

The ratio of HPL power leaving the cavity through RC1 during the open shutter periods to the power injected to the production magnet string P_{open}/P_i , must be calculated to calibrate the field overlap measured at the science detector. While the open shutter calibration array from Section IID $|C_{\text{open}}|^2$, gives a measure of the ratio between the signal power and the injected power, it is still expressed in terms of the ratio of the squared voltage measured by the heterodyne detection system to HPL power injected to the production magnet string (or units of V^2/W). Nevertheless, $|C_{\text{open}}|^2$ can be used with the expression

$$\frac{P_{\text{open}}}{P_i} = \frac{|C_{\text{open}}|^2}{P_{\text{LO}} G_{\text{AC}}^2 T_{\text{sci}} \eta_{xy,\text{LO}}^2}, \quad (65)$$

to calculate P_{open}/P_i for a given open shutter period. Here P_{LO} is the LO (the laser used as the heterodyne local oscillator), G_{AC} is the conversion gain of the science detection system including the full electronic chain at the first heterodyne demodulation frequency (~ 14 MHz), T_{sci} is the power transmissivity of the optical path from the RC to the science PD, and $\eta_{xy,\text{LO}}$ is

the spatial overlap between the LO and the RC eigenmode.

The spatial overlap between the LO and the RC is analogous to the spatial overlap between the LO and the HPL in transmission of the cavity. $\eta_{xy,\text{LO}}$ was measured by tuning the frequency of PLL3 such that both the LO and AL were at a cavity resonance. Then the drop in the amplitude of the PLL3 beatnote at the science PD in reflection of the cavity could be measured and compared to the amplitude of the PLL3 beatnote in reflection of the cavity when the intracavity field was blocked (while also considering the reflectivity of the RC). We assume that the LO and AL share the same spatial eigenmode because they are initially interfered using a fiber beamsplitter (fused fiber coupler). The static reflected power from a cavity should adhere to the following equation.

$$P_{\text{ref}} = P_0 R_{\text{RC,NR}} + P_1, \quad (66)$$

Here, P_{ref} is the power reflected from the cavity, P_0 is the power in the spatial mode of the cavity, $R_{\text{RC,NR}}$ is the cavity reflectivity in terms of power of a laser incident on the mirror in NR (Equation 59), and P_1 is the power outside the spatial mode of the cavity. With this, the spatial overlap can be defined using the expression

$$\frac{m_b}{m_a} = \frac{P_0 R_{\text{RC,NR}} + P_1}{P_0 + P_1} \quad (67)$$

$$= \eta_{xy,\text{LO}}^2 R_{\text{RC,NR}} + 1 - \eta_{xy,\text{LO}}^2, \quad (68)$$

where m_a represents the power measured when the cavity is blocked while m_b is the value of the reflected power when the laser is on resonance. (This approximation assumes the reflectivity of RC1 is one for the measurements of m_a .) Knowledge of the PD gain is unnecessary because the equation is expressed in terms of the ratio m_b/m_a , provided the same PD is used to sense m_a and m_b . Thus, the spatial overlap in terms of power $\eta_{xy,\text{LO}}^2$, can be calculated with the equation

$$\eta_{xy,\text{LO}}^2 = \frac{\frac{m_b}{m_a} - 1}{R_{\text{RC,NR}} - 1}. \quad (69)$$

The mean value of $\eta_{xy,\text{LO}}$ estimated during the closed shutter periods of S_{\perp} was 94.5%, with a range from 92.2% to 95.7% and RMS from the mean of 0.9%. During S_{\parallel} the mean was

93.8%, with a minimum value of 87.2%, a maximum value of 96.4% and an RMS of 2.1%. In both runs the uncertainty on the measurements of $\eta_{e,y,LO}$ was between 0.2% and 0.3%, and was driven primarily by the uncertainty in the measurements of the beatnote amplitude.

P_{LO} is also an important parameter as it directly affects the amplitude of the heterodyne signal at the science detector. It was measured by calibrating the static voltage output by the science detector by the conversion gain and subtracting the power of the AL that was also present at this detector. The change in the power during the runs is believed to be related to changes in the alignment of each laser to its corresponding optical fiber, as well as changes in T_{sci} . For S_{\perp} the mean value of P_{LO} during the closed shutter runs was 310 μ W with a range of 296 μ W to 322 μ W and RMS fluctuations from the mean of 10 μ W. During S_{\parallel} closed shutter periods P_{LO} is estimated to have a mean value of 385 μ W, with a minimum of 354 μ W, a maximum of 408 μ W, and an RMS of 15 μ W. The uncertainty in the measurements of P_{LO} during both runs was between 10 μ W and 12 μ W. Although both the uncertainty in the measurements of the DC voltage from the science detector, as well as the uncertainty in the measurement of the reflectivity of the RC, played a role in the overall uncertainty on P_{LO} , the dominant source of uncertainty in these measurements was the uncertainty in the measurements of the DC gain of the science detector electronic chain. This is discussed in more detail in the following section.

For the science detection system, both the DC conversion gain and the gain of the system at the heterodyne frequency, G_{AC} , had to be measured to calculate P_{open} . The DC conversion gain of the science PD was measured by varying the power of a single laser incident on the science PD, measuring the power with a power meter and recording the voltage level output by the detector. The slope of a linear fit to the voltage as a function of the incident power gave an average value of 3.9 ± 0.1 V/mW over several measurements made at different times. The uncertainty is believed to be related to the variance in the measured gain as a function of the position of the beam on the active area of the PD, as well as the uncertainty in the power meter calibration.

The conversion gain of the science detector

electronic chain at the heterodyne frequency G_{AC} , was also measured under conditions meant to resemble the actual conditions of the science campaign. The LO and AL were both incident on the curved RC mirror in NR with PLL3 locked at its operating frequency. To make the measurements independent of the parameters of the RC the circulating field of the cavity was blocked between NR and CH and the AL power was reduced to 25.9 μ W, a value near the mean of what was seen during the science campaign when the AL was stabilized to a resonance of the RC. A signal was sent to a fiber electro-optic modulator (EOM) to generate phase modulation sidebands on AL at the heterodyne frequency with roughly 0.01% of its carrier power. Therefore, like S_{\perp} and S_{\parallel} , the light on the science PD had a large interference beatnote at the PLL3 frequency and a low amplitude signal at the heterodyne frequency. With this setup, the conversion gain of the full detection chain at 14.603 MHz, G_{AC} , could be measured. This was done using the ratio of the sideband to carrier beatnote amplitude measured at the PLL3 PD (see Figure 2) to estimate the relative power in the sideband. A previous measurement of the frequency-dependent gain of the PLL3 PD at these two frequencies was used to correct for the difference in the gain of the PD at 14.603 MHz and 54.95 MHz. The power of the LO was then measured at the science PD along with the amplitude of the signal at the heterodyne frequency. Thus the sideband power (with the sideband at the heterodyne frequency) could be calculated along with the conversion gain of the detection system, independent of the measurement of the DC conversion gain of the photodetector. These measurements resulted in a conversion gain of 2.9 ± 0.1 V/mW. This is slightly lower than the DC conversion gain as a result of the electronic transfer function of the photodetector. Like the measurements of the DC conversion gain, the primary uncertainty in these measurements was related to the dependence of the gain on the position of the beam on the active area of the PD, as well as the calibration of the power meter.

To check for saturation effects, the LO power at the science PD was varied in power between 172 μ W and 1 mW and measured, holding the AL power at 25.9 μ W, while measuring the beatnote amplitudes at 54.95 MHz and 14.603 MHz at both the PLL3 PD and the science PD. Be-

tween 172 μW and 680 μW no evidence of saturation was seen, however the saturation effects became clear once the LO power reached 1 mW. However, the power of LO was at most 400 μW for both runs and therefore the science detection chain is not believed to have ever experienced saturation during the closed shutter periods.

T_{sci} is the total percentage of power exiting the RC that reaches the science photodetector. It was measured intermittently over the course of the science campaign by measuring the power at various points in the optical system in the NR cleanroom. These points included before and after the Faraday isolator, at RC1, directly after the Faraday isolator at the detection port with the cavity circulating field blocked, and at the science PD.

To find T_{sci} , the ratio of power incident on the RC cavity mirror versus the power incident on the science detector was calculated. This was done several times before, during, and after the first science campaign and gave values between 80% and 90%. The value used to calculate the field overlap was $85\% \pm 5\%$. The losses were mostly due to optical pickoffs used to sample the light reflected from the cavity for the PDH frequency stabilization system as well as other diagnostic purposes. Changes in the total reflectivity of the input optics that appear to be humidity dependent were also observed, along with clipping losses on the free aperture of the input optics after the beam expanding telescope in NR. Both of these effects played a role in the uncertainty of the results of T_{sci} .

The open shutter calibration series for the veto detection system $C_{\text{open}}^{\text{veto}}$, is defined in the same way as for the science detector and is also in units of V^2/W . The power ratio $P_{\text{open}}^{\text{veto}}/P_i$ for the veto detector can be expressed as

$$\frac{P_{\text{open}}^{\text{veto}}}{P_i} = \frac{|C_{\text{open}}^{\text{veto}}|^2}{P_{\text{AL,veto}} G_{\text{AC,veto}}^2 T_{\text{veto}}}. \quad (70)$$

Here $P_{\text{AL,veto}}$ is the AL (the laser used as the heterodyne local oscillator for the veto detector) power measured at the veto detector in units of W, $G_{\text{AC,veto}}$ is the conversion gain of the full veto detection system electronic chain at the veto heterodyne demodulation frequency ($\simeq 40.35$ MHz) in V/W , T_{veto} is the transmissivity of the optical path from mirror RC2 to the veto PD. There is no need for a spatial overlap term between AL and the RC in this case be-

cause AL is transmitted by the cavity and shares its spatial eigenmode.

The AC gain of the veto detector chain was measured by setting the PLL3 frequency to the veto heterodyne frequency of $f_{\text{veto}} = 40.34920$ MHz such that both AL and LO were resonant with the cavity. Then the power of LO was varied and the peak to peak amplitude of the signal at f_{veto} output by the veto detection chain was measured. At each point the power at the veto detector was measured with a power meter with and without the LO resonant (AL was always resonant in these measurements). With this the beatnote amplitude in power could be projected because in transmission of the RC, AL and LO are both entirely in the cavity eigenmode. Using this method a gain of $G_{\text{AC,veto}} = 25 \pm 1$ V/mW was measured. In the measurement the absolute accuracy of the power meter was the limiting factor in the uncertainty.

The power measured at the veto photodetector $P_{\text{AL,veto}}$ is an important parameter as it plays a direct role in the calibration of $C_{\text{open}}^{\text{veto}}$ in terms of optical power. It can also be used to estimate the optical losses in the path from the RC output mirror to the veto PD discussed in the following section. $P_{\text{AL,veto}}$ can be calculated by dividing the DC voltage output by the veto PD $V_{\text{DC,veto}}$, by its DC conversion gain $G_{\text{DC,veto}}$:

$$P_{\text{AL,veto}} = \frac{V_{\text{DC,veto}}}{G_{\text{DC,veto}}}. \quad (71)$$

The DC conversion gain $G_{\text{DC,veto}}$ of the veto detection system was measured by varying the power incident on the veto detector, and measuring the voltage at the monitoring point used during the maintenance periods. This gave a value of $G_{\text{DC,veto}} = 86 \pm 1$ V/mW. The error in this measurement is due to the variance in the results over multiple measurements. The primary reason for the difference between the AC and DC gain, was the presence of a three-way splitter in the AC path that allowed the veto PD to also be used as the sensor for PLL1.

One complication in the measurement of $P_{\text{open}}^{\text{veto}}/P_i$ arises from the fact that the transmissivity of the optics on the path from the output mirror of the RC to the veto PD, T_{veto} , can not be directly measured, as there are a significant number of optics in the portion of this path inside the central vacuum chamber. T_{veto} can be calculated using the ratio of the power of the AL measured at the veto PD to the expected

power of AL given its input power and coupling efficiency to the RC, as well as the transmissivity of the RC. This is measured by dividing the drop in the DC voltage output at the science PD $\Delta V_{AL,sci}$, when AL is frequency stabilized to the cavity, by the DC gain of the science PD detection chain ($G_{DC,sci}$) and by T_{sci} as expressed by

$$\Delta P_{AL,NR} = \frac{\Delta V_{AL,sci}}{T_{sci} G_{DC,sci}}. \quad (72)$$

As the following shows, this definition of $\Delta P_{AL,NR}$ is useful as it already takes the spatial overlap into account. Assuming minimal absorption and scattering losses in the substrate, the light incident on the cavity can either be reflected by the cavity, scattered out of the eigenmode or absorbed while circulating in the cavity, or transmitted by the cavity. Equation 72 represents the power in the latter two categories, light leaving the cavity eigenmode due to excess optical losses inside the cavity, or light transmitted by the cavity.

Because the total round trip attenuation A is measured via the cavity storage time τ during the maintenance periods and the transmissivity of the cavity mirrors were measured before the science runs took place, the power directly in transmission of the RC ($P_{AL,CH}$) can be estimated from

$$P_{AL,CH} = \frac{T_{RC2}}{A - T_{RC1}} \Delta P_{AL,NR} \quad (73)$$

$$= \frac{T_{RC2}}{\frac{2}{f_0\tau} - T_{RC1}} \frac{\Delta V_{AL}}{T_{sci} G_{DC,sci}}. \quad (74)$$

It should be noted that this equation can also be derived using the expressions for the spatial overlap between AL the RC along with the RC transmissivity and that Equation 74 is the simplest expression for the AL power in transmission of the RC.

T_{veto} can then be derived by calculating the ratio between the power measured at the veto PD and the power transmitted by the cavity in CH that is predicted by the Equation 74. This gives

$$T_{veto} = \frac{P_{AL,veto}}{P_{AL,CH}} \quad (75)$$

$$= T_{sci} \frac{V_{DC,veto}}{\Delta V_{AL}} \frac{G_{DC,sci}}{G_{DC,veto}} \frac{\left(\frac{2}{f_0\tau} - T_{RC1}\right)}{T_{RC2}}. \quad (76)$$

This equation shows that T_{veto} can be expressed as the product of T_{sci} and several ratios. The ratio of the DC voltage measured at the veto PD when AL is on resonance to the drop in DC voltage measured at the science PD, the ratio of the science PD DC conversion gain to the veto PD conversion gain, and finally the ratio of the losses minus the transmissivity of the RC input mirror to the transmissivity of the RC output mirror.

For the closed shutter periods of S_{\perp} , the average value of T_{veto} was estimated to be 68% with a minimum value of 65%, a maximum of 71%, RMS fluctuations from the mean of 2%, and an uncertainty of 9%. During S_{\parallel} , the average value of T_{veto} estimated for the closed shutter sections was 65% with a range from 55% to 70%, an RMS of 4%, and an uncertainty of 9%. In both runs, the primary source of the uncertainty was the uncertainty in the measurement of T_{RC2} .

After the conclusion of the first science campaign the power of the AL in transmission of the RC was measured immediately after the vacuum flange in the veto detector path to be $9.0 \mu\text{W}$, while the power at the veto PD was measured to be $7.4 \mu\text{W}$, indicating that 82% of the power leaving the COB at the veto port reaches the PD. These losses are mostly due to a diagnostic pick-off that sends some light to a camera to examine the position and shape of the RC eigenmode. The Brewster window at this port is also oriented such that it has an angle with respect to the vertical direction such that it has a transmissivity estimated to be 88% for horizontally polarized light. In addition to this, the mirror before the quadrant PD on the COB was measured to have a reflectivity of 95%. Therefore the path that the RC transmitted light takes to the veto PD is expected to have a transmissivity of 68%. While this agrees well with the measured values from science runs, in principle it should not be possible to measure values T_{veto} higher than the measured transmissivity of the optics. This could indicate a source of systematic error in the measurements of the transmissivities of the optics or T_{veto} . Nevertheless, the transmissivity measured after the first science campaign is well within the error bars of the T_{veto} values measured during the science run.

4. Spatial Overlap Results

The spatial coupling of the HPL to the RC measured during the open shutter periods should be representative of the coupling of the BSM field as long as there are no significant effects that can alter the amplitude or alignment of the HPL as it travels between the production and regeneration areas. The spatial component of the field overlap between the HPL and the RC measured by the science detector was expressed earlier as Equation 5 in Section I A. The square of the spatial field overlap, η_{xy}^2 , represents the spatial overlap in terms of power and can be expressed as a function of measured parameters as:

$$\eta_{xy}^2 = \frac{|C_{\text{open}}|^2}{T_{\text{COB}} T_{\text{RC}} P_{\text{LO}} G_{\text{AC}}^2 T_{\text{sci}} \eta_{xy,\text{LO}}^2}. \quad (77)$$

Here, T_{RC} is used instead of $T_{\text{RC,max}}$ because η_{xy}^2 is defined as the spatial overlap rather than the total field overlap.

The results of measuring η_{xy}^2 are shown in Figures 21a and 21b respectively for S_{\perp} and S_{\parallel} . For S_{\perp} , the average spatial overlap between the HPL and the RC during the closed shutter periods was estimated to be 47% with a range of 41% to 54% and RMS fluctuations from the average of 4%. The uncertainty in these measurements varied from 7% to 10%, with an average value

of 9%, and was due mostly to the uncertainties in the measured transmissivities of the on-axis COB mirrors and the flat cavity mirror mounted to the COB (RC2), with some additional contribution due to the uncertainty on the calibration series. During the closed shutter periods of S_{\parallel} the mean value of η_{xy}^2 was estimated to be 48% with a minimum of 42%, a maximum of 60%, and an RMS of 6%. The uncertainty of these measurements was again driven by the uncertainties on the mirror transmissivities like in S_{\perp} , along with some additional uncertainty due to the calibration series. For S_{\parallel} the uncertainty on η_{xy}^2 varied between 6% and 11% with an average value of 9%.

To calculate the HPL-RC spatial power overlap measured by the veto detection system $\eta_{xy,\text{veto}}^2$, the open shutter power defined in Section IV B 3 is divided by the expected power of the HPL passing through the production magnet string and the COB optics, along with the reflectivity of the cavity from the direction of CH:

$$\eta_{xy,\text{veto}}^2 = \frac{P_{\text{open}}^{\text{veto}}}{R_{\text{RC,CH}} T_{\text{COB}} P_1}. \quad (78)$$

The spatial overlap measured by the veto detection system can be expressed in terms of measured parameters as:

$$\eta_{xy,\text{veto}}^2 \simeq \frac{|C_{\text{open}}^{\text{veto}}|^2 \Delta V_{\text{AL}} G_{\text{DC,veto}}^2 T_{\text{RC2}} (1 + (2\pi\tau\Delta\nu)^2)}{G_{\text{AC,veto}}^2 T_{\text{sci}} ((1 - T_{\text{RC2}} f_0\tau)^2 + (2\pi\tau\Delta\nu)^2) T_{\text{COB}} V_{\text{DC,veto}}^2 G_{\text{DC,sci}} \left(\frac{2}{f_0\tau} - T_{\text{RC1}} \right)}. \quad (79)$$

Equations 58, 62, 70, 71, and 76 are used to derive these expressions. Figures 21a and 21b show the HPL-RC spatial overlap measured at the veto detector during the maintenance periods of the science runs as the orange data points and the results of $\eta_{xy,\text{veto}}^2$ agree very well with the data acquire at the science detector. During the closed shutter periods of S_{\perp} , the estimated spatial overlap had an average value of 53% with a range of 40% to 63%, RMS fluctuations from the mean of 6%, and an uncertainty of between 6% and 10% with an average uncertainty of 8%. For S_{\parallel} the average value was 49% with a minimum of 34%, a maximum value of 66%, an RMS

of 8%, and an uncertainty between 5% and 10% with an average value of 7%. For both runs the uncertainty was primarily limited by the uncertainty on the measured transmissivities of the COB optics and mirror RC2. While RC2 does not play a direct role in the HPL power measured at the veto PD, its uncertainty does effect the measurement of T_{veto} .

5. Field Overlap Results

The total power overlap between the HPL and RC is the product of the spatial overlap and the

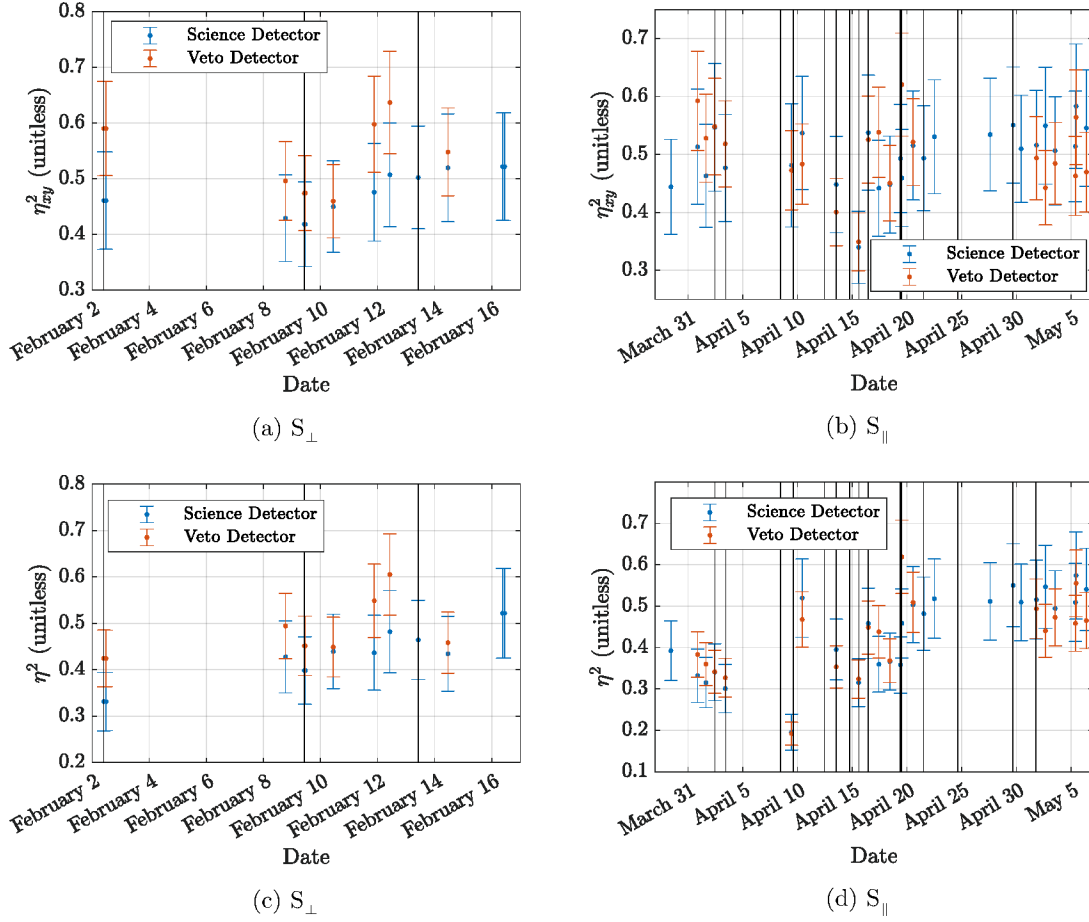


FIG. 21: The field overlap between the HPL and RC measured at the science detector and veto detectors respectively, from data taken during the maintenance periods of S_{\perp} and S_{\parallel} . The squared spatial overlap η_{xy}^2 is show in (a) for S_{\perp} and (b) for S_{\parallel} . The squared total field overlap η^2 is show in (c) for S_{\perp} and (d) for S_{\parallel} . In all four plots the black lines indicate any times at which the cavity alignment was actively changed, either to optimize the spatial field overlap or the resonant enhancement, while in (d) they also indicate times in which the length of the cavity was change to optimize the spectral component of the field overlap. In the case of (c) the length of the cavity was never changed during S_{\perp} .

longitudinal overlap, or

$$\eta^2 = \eta_{xy}^2 \eta_z^2 \quad (80)$$

$$= \frac{|C_{\text{open}}|^2}{T_{\text{COB}} T_{\text{RC,max}} P_{\text{LO}} G_{\text{AC}}^2 T_{\text{sci}} \eta_{xy,\text{LO}}^2}. \quad (81)$$

In this case $T_{\text{RC,max}}$ is used because the effects of the longitudinal overlap must be included in this calculation.

Measurements at the science detector of the total overlap η^2 are shown for S_{\perp} and S_{\parallel} in Fig-

ures 21c and 21d. During the closed shutter periods of S_{\perp} , the mean value of the total overlap was estimated to be 43% with a range from 33% to 51%, RMS fluctuations from the mean of 4%, with an uncertainty that varied between 6% and 10% and had an average value of 8%. These measurements were limited by the same sources of uncertainty as η_{xy}^2 . For S_{\parallel} , the average estimated value of η^2 at the science detector during the closed shutter periods was 41% with a range from 21% to 58%, an RMS of 9%, and an un-

certainty between 4% and 10% with an average value of 8%. The uncertainty on these measurements was also due to the same effects as were just mentioned for S_{\perp} .

Using Equation 79 to express $\eta_{xy,\text{veto}}^2$, the total field overlap between the HPL and the RC can be expressed in terms of measured parameters as

$$\eta_{\text{veto}}^2 \simeq \frac{\eta_{xy,\text{veto}}^2}{\eta_{\tilde{z}}^2} \simeq \frac{\eta_{xy,\text{veto}}^2}{1 + (2\pi\tau\Delta\nu)^2}. \quad (82)$$

The total overlap between the HPL and RC measured at the veto detector is shown in Figures 21c and 21d for S_{\perp} and S_{\parallel} respectively. During the closed shutter periods of S_{\perp} , the average value estimated for η_{veto}^2 was 48%. The minimum value over this time was 39%, while the maximum value was 60%, the RMS fluctuations from the mean were 5%, and the uncertainty varied between 6% and 10% with an average value of 8%. For S_{\parallel} the average value was 41%, and the data varied over a range from 20% to 66% with an RMS of 8% and an uncertainty between 3% and 10% with an average value of 6%. In general, the spatial overlap and total overlap measured at the veto and science detectors showed good agreement and the differences between the results were on the order of 10% or less with the exception of a few outliers.

Finally, it should be noted that in both runs, like $\eta_{xy,\text{veto}}^2$, η_{veto}^2 agree very well with the results of the spatial overlap and total overlap measured at the science detector and discussed in Section IV B 4.

6. Optimizing the Spatial Overlap

The spatial coupling of the HPL to the RC can be affected by the relative position, alignment, and size of the HPL and RC eigenmodes. The initial mode-matching of the HPL was performed by first building a 246 m cavity with the curved mirrors of the PC and RC located on the optical tables in the NL and NR cleanrooms respectively (see Figure 2 for a diagram of the cleanrooms). Then the AL was coupled to the cavity from NR and the light in transmission of the cavity in NL served as a reference for the eigenmode of the cavity. This light was then mode matched, using the HPL input optics, to a triangular cavity which served as an input mode cleaner for the HPL [31]. The spatial mode of

the 246 m cavity should be representative of the RC eigenmode because it had a nearly confocal geometry with twice the length of the RC.

A spatial overlap of roughly 90% was achieved. The input mode cleaner was then left in place to serve as a reference for the mode of the input laser. The higher order mode content related to the waist size and position of the HPL with respect to the RC should have therefore remained optimized, although drifts in the alignment of the mirrors after the input mode cleaner could still cause a decoupling of the HPL to the RC. Therefore during the maintenance periods the HPL alignment was corrected using the following procedure. After the eigenmode position of the RC was optimized using the storage time, the RL was manually realigned to maximize the PLL1 beatnote amplitude. After this the HPL input optics were tuned to optimize the beatnote amplitude between HPL and RL at the MZ port. This method was limited in its accuracy as the maximum power spatial overlap achieved between RL and the AL the light transmitted by the RC was on the order of 50% due to the astigmatism of the RL beam generated by a pair of curved mirrors in its mode matching telescope. Mirrors were used instead of lenses to avoid creating a direct path for HPL stray light to couple to the RC via backreflections from these lenses.

V. CONCLUSION

Optical light-shining-through-a-wall experiments are a powerful tool to search for pseudo-Goldstone bosons that could reveal a new energy scale between the electroweak and Planck scales. With optical cavities before and after the wall these experiments offer an unprecedented photon flux along with a long interaction length and boosted signal strength. The ALPSII first science campaign took place from February to May of 2024 with runs using the light-shining-through-a-wall technique to search for beyond the standard model bosonic fields. These searches were performed with a regeneration cavity resonantly enhancing the electromagnetic field potentially regenerated behind the wall, the first time this has been done with optical cavities. This campaign consisted of two science runs, one with the polarization states of the lasers perpendicular to the direction of the magnetic field (S_{\perp}) and one with the polariza-

tion states aligned parallel to the magnetic field (S_{\parallel}). The signals measured by the heterodyne detection system over the duration of the runs were coherently summed, allowing the experiment to have a frequency resolution on the order of μHz . Therefore, background signals that are not at the exact expected frequency can be distinguished from signals that are the result of the coupling between electromagnetic and beyond the standard model fields.

The calibration of the data was performed using measurements of the transmission of the high-power input laser through the system when a shutter in the wall was open. This allowed the laser light to directly couple to the regeneration cavity after passing through the optics on the central optical bench. The signal rates measured when the shutter was closed were then normalized by the measured open shutter transmission. The use of this ratio cancels most experimental parameters and systematic uncertainties except the transmissivity of the central optical bench optics. The relevant optics showed a combined transmissivity of $(9.7 \pm 1.2) \times 10^{-23}$, resulting in less than ten attowatts reaching the detection system when the shutter is open. Such low powers highlight the challenges faced when using this technique and make it clear that the operation and calibration of such a system is in itself an accomplishment. Via this procedure, it is possible to express the results in terms of $\mathcal{P}_{\gamma\leftrightarrow\phi}$, the rate at which energy converts between electromagnetic and beyond the standard model fields after a single pass through one of the magnet strings.

The results at the expected signal frequency did show an excess relative to the technical noise of the system, however, when this ratio was evaluated at alternative frequencies, a broad peak in the spectrum was evident. This indicated the presence of a background believed to be caused by stray light coupling from the high-power laser to the detection system. This was corroborated by an analysis of the time evolution of the closed shutter signals.

For each run, a statistical analysis was performed on the background signals within $100 \mu\text{Hz}$ of the signal frequency giving an expectation value for the background and its probability distribution. Based on this, the values measured at the signal frequency were well below 5σ detection thresholds. We therefore set limits that exclude conversion rates above

a significance of 5σ with a confidence level of 95%. The resulting exclusion limits are $\mathcal{P}_{\gamma\leftrightarrow\phi} < 2.5 \times 10^{-13}$ for S_{\perp} and $\mathcal{P}_{\gamma\leftrightarrow\phi} < 1.7 \times 10^{-13}$ for S_{\parallel} .

These results compare favorably to previous generations of light-shining-through-a-wall experiments as the limits on $\mathcal{P}_{\gamma\leftrightarrow\phi}$ for scalar and pseudoscalar bosons surpass OSQAR [10] by a factor of 18 and 31 respectively, despite the fact that much stricter exclusion criteria were applied in this analysis. We would also like to emphasize that the sensitivity in terms of $\mathcal{P}_{\gamma\leftrightarrow\phi}$ is independent of the properties of the magnet string. It is therefore a useful metric to compare the optical systems of different light-shining-through-a-wall experiments. The exclusion limits on $\mathcal{P}_{\gamma\leftrightarrow\phi}$ reported here are interpreted in terms of the coupling between electromagnetic and pseudoscalar, scalar, vector, and tensor bosons in Ref. [1]. Those results show that ALPS II improved the limits on the electromagnetic coupling of scalar bosons by a factor of 18 and pseudoscalar bosons by a factor of 23 with respect to the OSQAR results. In both limits, the improvement in sensitivity due to the longer magnet strings used by ALPS II and its more sophisticated optical system is roughly the same.

The optical system was constantly monitored during the first science campaign. Although parameters such as the resonant enhancement of the regeneration cavity were not directly used in the calibration of the science data, they provided useful information on the status of the optical system. With a world record storage time of 7.17 ms [20], the regeneration cavity sustained a resonant enhancement factor of roughly 7000 during the science campaign.

In the future, we plan to implement an optical cavity before the wall, further reduce backgrounds due to stray light and technical noise, boost the resonant enhancement factor of the regeneration cavity and increase its overlap with the beyond the standard model field. With these upgrades, the signal to noise ratio with respect to the power measured at the detection system will increase by seven orders of magnitude allowing ALPS II to probe $\mathcal{P}_{\gamma\leftrightarrow\phi}$ down to rates of 3×10^{-17} , an improvement of nearly four orders of magnitude.

ACKNOWLEDGMENTS

We would like to thank the ALPS II collaboration, especially David Reuther, Sandy Croatto, and Sven Karstensen, the FTX group, especially, Karsten Gadow and Maik Dinter, the MVS group, especially Antonio de Zubiaurre-Wagner, MKS, FEB, MPC, MPS, ZM1 and ZM5, and the MEA2 and MEA5 groups at DESY as well as the FH mechanical workshop, for their support during the commissioning and operation of the experiment. We are also grateful to Manuel Meyer for the many valuable discussions and his feedback regarding the manuscript.

We acknowledge the support of the National Science Foundation (Grant No. 1802006), of the Heising-Simons Foundation (Grant No. 2015-154 and 2020-1841), of the Deutsche Forschungsgemeinschaft through project grant WI 1643/2-1 and EXC 2121 “Quantum Universe” – 390833306, of the Science and Technology Facilities Council (UK), grants ST/T006331/1 and ST/Y004515/1, as well as support by the German Volkswagen Stiftung and the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program Grant agreement No. 948689.

VI. BIBLIOGRAPHY

- [1] D. C. Brotherton, S. Croatto, J. Egge, A. Ejlli, H. Frädrieh, J. Gleason, H. Grote, A. Hallal, M. T. Hartman, H. Hollis, K.-S. Isleif, F. Januschek, K. Karan, S. Karstensen, T. Kozlowski, A. Lindner, M. Meyer, G. Müller, G. Othman, J. H. Pöld, D. Reuther, A. Ringwald, E. Rivasto, J. A. R. Gimeno, J. Schaffran, U. Schneekloth, C. Schwemmbauer, A. D. Spector, D. B. Tanner, D. Trines, L.-W. Wei, B. Willke, and R. Wolf (ALPS Collaboration), Any light particle searches with ALPS II: first science results (2025), arXiv:2512.14110 [hep-ex].
- [2] L. B. Okun *et al.*, *Limits of electrodynamics: paraxial photons*, Tech. Rep. (Institute of Theoretical and Experimental Physics, Moscow, 1982).
- [3] A. A. Anselm, Arion \leftrightarrow Photon Oscillations in a Steady Magnetic Field. (In Russian), *Yad. Fiz.* **42**, 1480 (1985).
- [4] K. Van Bibber, N. Dagdeviren, S. Koonin, A. Kerman, and H. Nelson, Proposed experiment to produce and detect light pseudoscalars, *Phys. Rev. Lett.* **59**, 759 – 762 (1987).
- [5] F. Hoogeveen and T. Ziegenhagen, Production and detection of light bosons using optical resonators, *Nucl. Phys. B* **358**, 3 – 26 (1991).
- [6] P. Sikivie, D. B. Tanner, and K. Van Bibber, Resonantly enhanced axion-photon regeneration, *Phys. Rev. Lett.* **98**, 172002 (2007).
- [7] G. Mueller, P. Sikivie, D. B. Tanner, and K. van Bibber, Detailed design of a resonantly enhanced axion-photon regeneration experiment, *Phys. Rev. D* **80**, 072004 (2009).
- [8] R. Bähre, B. Döbrich, J. Dreyling-Eschweiler, S. Ghazaryan, R. Hodajerdi, D. Horns, F. Januschek, E.-A. Knabbe, A. Lindner, D. Notz, *et al.*, Any light particle search II—technical design report, *J. Instrum.* **8** (09), T09001.
- [9] K. Ehret, M. Frede, S. Ghazaryan, M. Hildebrandt, E.-A. Knabbe, D. Kracht, A. Lindner, J. List, T. Meier, N. Meyer, *et al.*, New ALPS results on hidden-sector lightweights, *Phys. Lett. B* **689**, 149 (2010).
- [10] R. Ballou, G. Deferne, M. Finger, M. Finger, L. Flekova, J. Hosek, S. Kunc, K. Macuchova, K. A. Meissner, P. Pugnât, M. Schott, A. Siemko, M. Slunicka, M. Sulc, C. Weinsheimer, and J. Zicha (OSQAR Collaboration), New exclusion limits on scalar and pseudoscalar axionlike particles from light shining through a wall, *Phys. Rev. D* **92**, 092002 (2015).
- [11] P. Arias, J. Jaeckel, J. Redondo, and A. Ringwald, Optimizing light-shining-through-a-wall experiments for axion and other weakly interacting slim particle searches, *Phys. Rev. D* **82**, 115018 (2010).
- [12] C. García-Cely and A. Ringwald, Stellar bounds on light spin-2 particles in bimetric theories (2025), arXiv:2511.03707 [hep-ph].
- [13] M. D. Ortiz, J. Gleason, H. Grote, A. Hallal, M. Hartman, H. Hollis, K.-S. Isleif, A. James, K. Karan, T. Kozlowski, *et al.*, Design of the ALPS II optical system, *Phys. Dark Universe* **35**, 100968 (2022).
- [14] M. Frede, B. Schulz, R. Wilhelm, P. Kwee, F. Seifert, B. Willke, and D. Kracht, Fundamental mode, single-frequency laser amplifier for gravitational wave detectors, *Opt. Express* **15**, 459 (2007).
- [15] R. Meinke, Superconducting magnet system for hera, *IEEE transactions on magnetics* **27**, 1728 (1991).
- [16] C. Albrecht, S. Barbanotti, H. Hintz, K. Jensch, R. Klos, W. Maschmann, O. Sawlanski, M. Stolper, and D. Trines, Straightening of superconducting HERA dipoles for the any-light-particle-search experiment ALPS II, *Eur. Phys. J. Tech. Instrum.* **8**, 5 (2021).
- [17] R. Paschotta, Mode matching, *RP Photonics Encyclopedia* (2007).

- [18] Z. R. Bush, S. Barke, H. Hollis, A. D. Spector, A. Hallal, G. Messineo, D. B. Tanner, and G. Mueller, Coherent detection of ultraweak electromagnetic fields, *Phys. Rev. D* **99**, 022001 (2019).
- [19] A. Hallal, G. Messineo, M. D. Ortiz, J. Gleason, H. Hollis, D. B. Tanner, G. Mueller, and A. Spector, The heterodyne sensing system for the ALPS II search for sub-ev weakly interacting particles, *Phys. Dark Universe* **35**, 100914 (2022).
- [20] T. Kozłowski, L.-W. Wei, A. D. Spector, A. Hallal, H. Frädrieh, D. C. Brotherton, I. Oceano, A. Ejlli, H. Grote, H. Hollis, *et al.*, Design and performance of the ALPS II regeneration cavity, *Opt. Express* **33**, 11153 (2025).
- [21] E. D. Black, An introduction to pound–drever–hall laser frequency stabilization, *Am. J. Phys.* **69**, 79 (2001).
- [22] R. V. Pound, Electronic frequency stabilization of microwave oscillators, *Rev. Sci. Instrum.* **17**, 490 (1946).
- [23] R. Drever, J. L. Hall, F. Kowalski, J. Hough, G. Ford, A. Munley, and H. Ward, Laser phase and frequency stabilization using an optical resonator, *Appl. Phys. B* **31**, 97 (1983).
- [24] D. B. Leviton and B. J. Frey, Temperature-dependent absolute refractive index measurements of synthetic fused silica, in *Optomechanical Technologies for Astronomy*, Vol. 6273 (SpIE, 2006) pp. 800–810.
- [25] T. Hahn and R. Kirby, Thermal expansion of fused silica from 80 to 1000 k-standard reference material 739, in *AIP Conference Proceedings*, Vol. 3 (American Institute of Physics, 1972) pp. 13–24.
- [26] A. Savitzky and M. J. Golay, Smoothing and differentiation of data by simplified least squares procedures., *Anal. Chem.* **36**, 1627 (1964).
- [27] P. B. Patnaik, The non-central χ^2 - and f -distribution and their applications, *Biometrika* **36**, 202 (1949).
- [28] A. E. Siegman, *Lasers* (Oxford University Press, University Science Books, 1987).
- [29] T. Isogai, J. Miller, P. Kwee, L. Barsotti, and M. Evans, Loss in long-storage-time optical cavities, *Opt. Express* **21**, 30114 (2013).
- [30] A. D. Spector and T. Kozłowski, Optical cavity characterization with a mode-matched heterodyne sensing scheme, *Opt. Express* **32**, 27112 (2024).
- [31] B. Willke, N. Uehara, E. K. Gustafson, R. L. Byer, P. J. King, S. U. Seel, and R. L. Savage, Spatial and temporal filtering of a 10-w nd:yag laser with a fabry–perot ring-cavity premode cleaner, *Opt. Lett.* **23**, 1704 (1998).

Appendix A: Appendix

1. Digital Frequency Representation

Table VI summarizes the frequencies as they are defined in the user interface of the devices generating them in the column labeled ‘Set frequency’, while the offset from those frequencies due to their representation in terms of the clock frequencies is shown in the column ‘Digitization offset’. The row label ‘Total’ for the science and veto detectors shows the result of the linear combination of frequencies that produces the signal that is demodulated at the second stage in post-processing. In this case the input frequencies of the PLLs are tuned to produce a 2.4 Hz signal frequency after the first demodulation. The digital offset on that signal is a result of the same linear combination of the digital offsets. The second demodulation should therefore be performed at the total set frequency plus its digital offset, so in all cases 2.4 Hz + 198.58 nHz. This is the digital representation of a 2.4 Hz signal using a clock with a rate of 1 GHz and a precision of 48 bits. In post-processing, the frequency of the second demodulation can be represented with double-precision (64-bit) floating-point because this process is no longer limited by the precision of the oscillators driving the PLLs. While the numbers in this table are only shown with a precision of tens of pHz, the actual second demodulation frequency was defined to tens of fHz.²⁹

2. Changes in Open Shutter Data

Scatter plots showing the changes between open shutter periods are shown in Figure 22. From the plots in 22a and 22c it is apparent that the absolute value of the changes between the magnitude squared of two points in the open shutter calibration series $||C_{\text{open},j}|^2 - |C_{\text{open},j-1}|^2|$ increased as the time between the points $j - 1$ and j increased. The data in these plots are also normalized to the value of $|C_{\text{open},j-1}|^2$ to allow the relative error in the calibration series due to these effects to be

²⁹ This precision is only defined with respect to the clock used for the system and should not be mistaken in terms of absolute frequency.

System	S_{\perp}	
	Set frequency	Digitization offset
PLL1	12.2 MHz	-11.37 nHz
PLL2	52.5468728 MHz	699.04 nHz
PLL3	54.95 MHz	1608.67 nHz
DM1 (Science)	14.6031296 MHz	1096.83 nHz
DM1 (Veto)	40.3468752 MHz	908.99 nHz
Total (Science)	2.4 Hz	198.58 nHz
Total (Veto)	2.4 Hz	198.58 nHz

System	S_{\parallel}	
	Set frequency	Digitization offset
PLL1	12.2 MHz	-11.37 nHz
PLL2	52.54686689 MHz	1098.23 nHz
PLL3	54.95 MHz	1608.67 nHz
DM1 (Science)	14.60313551 MHz	697.65 nHz
DM1 (Veto)	40.34686929 MHz	1308.17 nHz
Total (Science)	2.4 Hz	198.58 nHz
Total (Veto)	2.4 Hz	198.58 nHz

TABLE VI: Summary of the frequencies and their digital offsets for S_{\perp} and S_{\parallel} .

calculated (discussed in Section IIF). The data for the science detector are shown as the blue points, while the veto detector data are shown in orange. The dashed lines give the fit of a first order polynomial to the data. The result of this fit is then used to calculate the relative uncertainty for each point in the calibration series $C[n]$, based on its length in time from the nearest open shutter measurement.

The changes in the angle of $C_{\text{open},j}$ between open shutter points are also shown in Figures 22b and 22d, again with blue represent the data measured at the science detector, orange the data measured and the veto detector, and the dashed lines giving the linear fit to the data. Interestingly, the fit of the data from S_{\parallel} shows a negative slope, indicating that changes in the angle decrease as the time between the points increases. We do not believe that it is actually the case that the uncertainty in the complex angle of the open shutter measurements decrease over time, but rather this is due to a bias in the data. This bias comes from the fact that the operation procedure required an open shutter period of at least 15 minutes before starting a closed shutter

measurement. As the shorter periods are still used to assess the uncertainty, times in which the system was unstable would have produced more open shutter data. Partially because of this reason, the linear fit of the angular changes are not used in the calculation of the uncertainty on $C[n]$, and displayed here rather to illustrate the remarkable stability of the system.

3. Signal Leakage

Irregular gaps in the closed shutter data exist due to the need for open shutter and maintenance periods, as well as the system coming unlocked during the closed shutter periods. These gaps lead to signals at a given frequency leaking into neighboring frequency bins as the integration of the product a static constant with these neighboring frequencies is no longer performed over a finite number of cycles and these components will not sum to zero. Figures 23a and 23b show plots of this leakage into the different Fourier components. Here the spectrum of $Z(f)_{\text{closed}}$ is calculated with $H_p[n]$ set to a se-

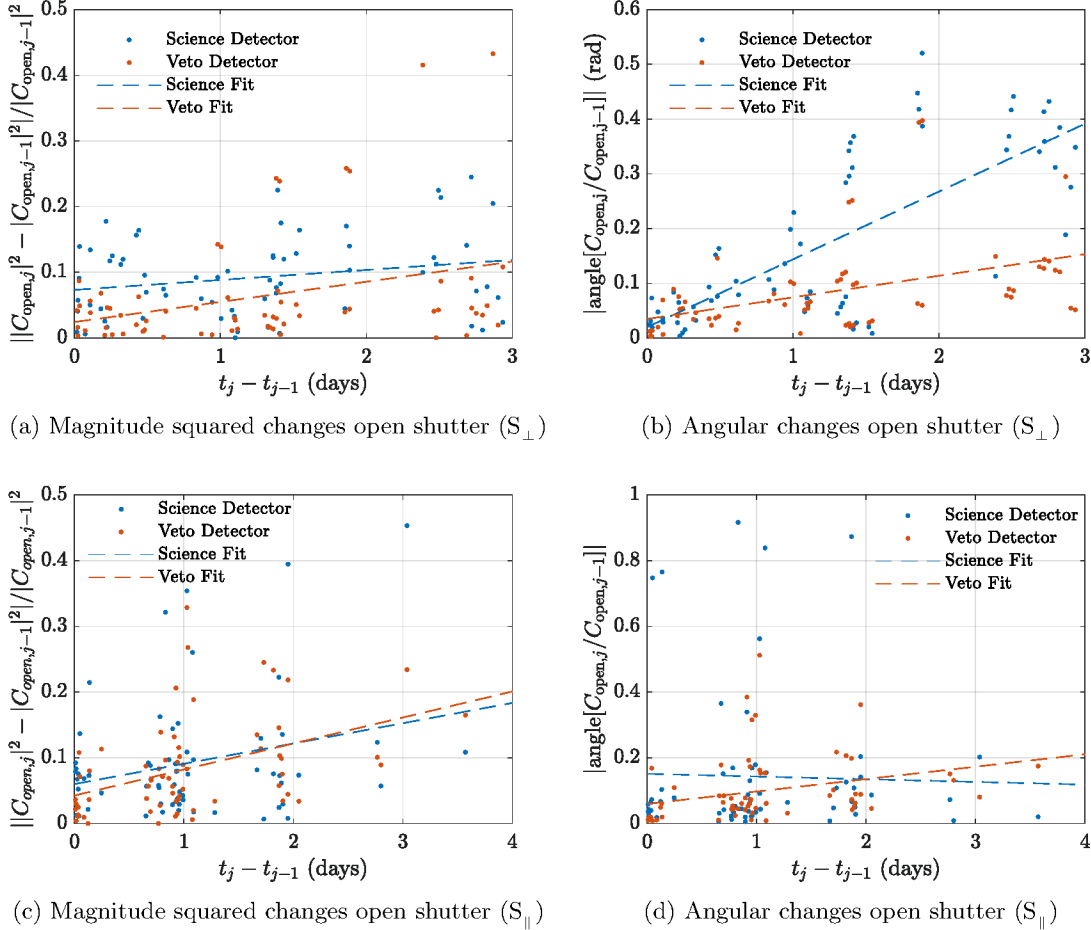


FIG. 22: Changes in the magnitude and angle of $C_{\text{open},j}$ plotted as a function of the time between the points for S_{\perp} . Science detector data are shown in blue and the veto detector data are shown in orange. The dashed lines show the linear fits to the data. The units of the y -axis in this plot are rad.

ries of ones for both S_{\perp} (blue) and S_{\parallel} (orange). Note that the data for S_{\parallel} are offset downward by a decade to help show both sets of data on the same plots.

In these plots it can be seen that the neighboring frequency bins at $f = \pm 0.85 \mu\text{Hz}$ contain a value of $|Z|^2$ that is 0.093 times the value in the signal bin for S_{\perp} and 0.214 at $f = \pm 0.60 \mu\text{Hz}$ for S_{\parallel} . While in both cases none of these Fourier components are present in a continuous series of ones, they are present in a series of ones with gaps similar to our data sets. S_{\parallel} has more of this cross coupling between the frequency bin at $f = \pm 0.60 \mu\text{Hz}$ due to the fact that over 97% of the valid closed shutter data came in three

distinct periods each with nearly a week in between, from March 31 to April 6, from April 12 to April 22, and from April 29 to May 6. The closed shutter data of S_{\perp} , on the other hand, was acquired more regularly.

Gaps in the data are not expected to lead to excess coupling of white Gaussian noise into the signal bin or across other frequency bins with respect to the expected noise level if there were no gaps in the data (effects from the integration time notwithstanding). It is therefore believed that any stochastic noise in the data should produce a spectrum in $Z(f)_{\text{closed}}$ in accordance with its power spectral density. Deterministic signals, on the other hand, may couple into alternate fre-




FIG. 23: Leakage from the signal bin into neighboring frequency bins for S_{\perp} (blue) and S_{\parallel} (orange) from -5 mHz to $+5$ mHz (a) and -100 μ Hz to $+100$ μ Hz (b). Note that the S_{\parallel} data are offset a decade lower to show both data sets on the same plot.

quency bins. Based on Figures 23a and 23b this is not expected to reduce the power measured at the signal frequency as here the value at $f = 0$ is 1 for both runs.

4. Double Checking the Calibration

Three alternative versions of $C[n]$ were applied to the data to investigate different effects in the calibration array. For the first alternative series, the mean value of the calibration series, or $\langle C[n] \rangle$ was calculated over the full run, the coherent sum was performed independently, and $\langle C[n] \rangle$ was used to calibrate the result of the summation. The second version of the calibration array simply used the absolute value of the calibration array and performed no correction to the phase of the heterodyne data and is referred to as $|C[n]|$. The phase of the veto calibration series is used in the last alternative calibration series. Here the $C_v[n]$ refers to using the magnitude of the normal calibration series, but the phase of the veto calibration series.

The results of the alternative frequency analysis are shown in Figures 24a and 24b, while Table VII shows a summary of the results in the signal bin using the alternative calibration arrays. The data for the actual calibration array is shown as the blue trace in these plots. The results of $\langle C[n] \rangle$ and $|C[n]|$ demonstrate that the calibration method makes only a marginal difference in the results. This makes sense as the changes in the amplitude and phase of the open shutter data throughout the runs was not so large as to substantially effect the results. The results of using $C[n]_v$ were also very similar to the results obtained with the normal calibration array.

For both runs it is apparent that the background levels do not show significant differences when one of these alternative calibration arrays was used.

	Scalar	Pseudoscalar
	$ Z(f_j = 0)_{\text{closed}} ^2$	$ Z(f_j = 0)_{\text{closed}} ^2$
$C[n]$	1.94×10^{-5}	2.87×10^{-5}
$\langle C[n] \rangle$	2.61×10^{-5}	2.52×10^{-5}
$ C[n] $	2.66×10^{-5}	1.95×10^{-5}
$C[n]_v$	2.24×10^{-5}	1.80×10^{-5}

TABLE VII: Summary of the heterodyne data in the signal bin for both detectors and both science runs.



(a) S_{\perp} (b) S_{\parallel}

FIG. 24: The stray-light signal from S_{\perp} (a) and S_{\parallel} (b) using different calibration functions.