

Coherent detection of ultraweak electromagnetic fields

Zachary R. Bush,¹ Simon Barke,¹ Harold Hollis,¹ Aaron D. Spector,² Ayman Hallal,¹
Giuseppe Messineo,¹ D. B. Tanner,¹ and Guido Mueller¹

¹*Department of Physics, University of Florida, P.O. Box 118440, Gainesville, Florida 32611, USA*

²*Deutsches Elektronen-Synchrotron (DESY), Notkestrae 85, D-22607 Hamburg, Germany*



(Received 3 April 2018; published 2 January 2019)

We explore the application of heterodyne interferometry for a weak-field coherent detection scheme. The methods detailed here will be used in ALPS II, an experiment designed to search for weakly interacting, sub-eV particles. For ALPS II to reach its design sensitivity this detection system must be capable of accurately measuring fields with equivalent amplitudes on the order of 10^{-5} photons/s or greater. We present initial results of an equivalent dark count rate on the order of 10^{-5} photons/s as well as successful generation and detection of a signal with a field strength equivalent to 10^{-2} photons/s.

DOI: [10.1103/PhysRevD.99.022001](https://doi.org/10.1103/PhysRevD.99.022001)

I. INTRODUCTION

A. Axions and axionlike particles

The Standard Model (SM) incorporates our current knowledge of subatomic particles as well as their interactions via three of the four fundamental forces of nature. The SM is not complete, however, as it does not contain gravity and does not explain certain observations. One notable unresolved issue is that of charge-conjugation parity symmetry (CP -symmetry) violation. The QCD Lagrangian includes terms capable of breaking CP symmetry for the strong force. In contrast, experiments found that the strong forces respect CP symmetry to a very high precision [1].

The most prominent proposed solution, introduced by Peccei and Quinn [2], involves spontaneously breaking a global $U(1)$ symmetry leading to a new particle, named the axion [3,4]. Interactions with the QCD vacuum cause the axion to have a nonzero mass, m_a [2]. While axions may interact with SM particles, the interactions can be weak. Most notably for experimental purposes, axion mixing with neutral pions leads to a characteristic two-photon coupling, $g_{a\gamma\gamma}$ [5]. This, in turn, constrains the product of the axion mass and coupling such that these two parameters are dependent. Experimental and observational factors place the axion mass between 1 and $1000 \mu\text{eV}$. The corresponding range for $g_{a\gamma\gamma}$ is 10^{-16} to $10^{-13} \text{ GeV}^{-1}$.

While the QCD axion is confined to a specific band in the parameter space, it might just be a member of a larger class of axionlike particles, some with a stronger two-photon coupling [6,7]. The interactions between these axions/axionlike particles and photons may possibly explain unanswered astronomical questions including TeV photon transparency in the Universe [8] and anomalous white dwarf cooling [9]. The intrinsic properties of

axions and axionlike particles also make them prime candidates for cold dark matter. This theoretical motivation has led to the formulation of various experiments designed to detect axions and axionlike particles by utilizing their coupling to photons.

Although axions can naturally decay into two observable photons, the rate at which this occurs is extremely low, making detection by observing this decay impossible. Axion search experiments therefore also rely on the inverse Primakoff or Sikivie effect in which a strong static magnetic field acts as a high density of virtual photons. This field stimulates the axion/axionlike particle to convert into a photon carrying the total energy of the axion/axionlike particle [10,11]. A number of strategies have been employed by these experiments to search for axions from several sources. Haloscope experiments, such as ADMX, use resonant microwave cavities and strong superconducting magnets to search for axions comprising the Milky Way's cold dark-matter halo [12]. Helioscope experiments, such as CAST, look for relativistic axions originating from the Sun that convert into detectable x-rays as they pass through a supplied magnetic field [13]. Differing from these types of axion searches that rely on astronomical sources, "light shining through walls" (LSW) experiments attempt to generate and detect axions in the laboratory and therefore have the advantage of independence from models of the galactic halo and models of stellar evolution [14–19].

B. ALPS II

LSW experiments use the axion-photon coupling first to convert photons into axions under the presence of a strong magnetic field. These axions then pass through a light-tight barrier and enter another strong magnetic field where some

are converted back into detectable photons. Energy is conserved in the process so that the regenerated photons behind the wall have the same energy as those incident in front of it. The Any Light Particle Search (ALPS) is one such LSW experiment. The first generation of this experiment, ALPS I, set the most sensitive laboratory experimental limits of its time on the coupling of axions to two photons, $g_{a\gamma\gamma}$, for a wide range of axion masses [18]. ALPS I used a single optical cavity placed before a light-tight barrier to increase the circulating power on the axion production side of the magnet. The second iteration of the experiment, ALPS II, will improve the sensitivity further with the addition of an optical cavity after the barrier. The presence of this cavity will resonantly enhance the probability that axions/axionlike particles will reconvert to photons [20–24]. ALPS II is currently being developed in two stages. The first stage, ALPS IIa, will use two 10 m long resonant cavities without magnets [25]. The second stage, ALPS IIc, will use two 100 m long cavities with 5.3 T superconducting HERA dipole magnets. Longer cavity lengths increase the interaction time between the photons and the magnetic field.

Figure 1 shows a simplified layout of the ALPS IIc experiment. Infrared laser light is injected into an optical cavity whose eigenmode is immersed in a 5.3 T magnetic field. The polarization of the injected light is set to be linear. The direction of the polarization is oriented either parallel or orthogonal to the direction of the external magnetic field in order to search separately for pseudoscalar or scalar axionlike particles. Power buildup from this cavity causes a high circulating laser power, increasing the flux of axionlike particles through the wall. After these particles traverse the light-tight barrier they enter a second cavity, called the regeneration cavity, also subject to a 5.3 T magnetic field. The particles then have the same probability to reconvert back into photons having an energy identical to those in the production cavity.

The exclusion limits (95% confidence level) measured by ALPS I for a 31 h data run in vacuum is shown by the green region of Fig. 2. Improvements in the optical design from ALPS I to ALPS IIc lead to a projected 2000-fold increase in sensitivity to the coupling parameter, $g_{a\gamma\gamma}$.

ALPS IIc will inject a 30 W laser field into the 100 m long production cavity (PC), which is immersed in a 5.3 T magnetic field. The circulating power inside the PC is




FIG. 1. Simplified model of the ALPS IIc experiment. Axions generated in the left-hand side cavity (the production cavity) traverse the wall and some turn back into detectable photons in the right-hand side cavity (the photon regeneration cavity).




FIG. 2. Parameter space of the axion mass (m_a) and coupling ($g_{a\gamma\gamma}$) showing projected improvements in sensitivity from ALPS I (in vacuum) to ALPS IIc [24].

expected to reach 150 kW. The 100 m long regeneration cavity (RC) on the other side of the wall will have a finesse of 120,000. The RC is also placed inside a similar 5.3 T magnetic field. Assuming a coupling strength of $g_{a\gamma\gamma} \sim 2 \times 10^{-11} \text{ GeV}^{-1}$, the employed photon detector has to be able to measure fields with a photon rate as low as 2×10^{-5} photons/s [23]. For 1064 nm light, this is equivalent to an average power on the order of 10^{-24} W. ALPS II is exploring two technologies for detecting such weak signals. The first of these uses a transition edge sensor [26]. This technology utilizes a superconducting thin film operating near its critical temperature to absorb the regenerated photons, thereby changing its temperature and thus its resistance. An alternative approach, heterodyne interferometry, is the subject of this paper [23].

C. Heterodyne detection principles

The principle of heterodyne interferometry requires interfering two laser fields at a nonzero difference frequency. Let one laser, L_1 , have frequency f , phase ϕ_1 , and average power \bar{P}_1 and a second laser, L_2 , have frequency $f + f_0$, phase ϕ_2 , and average power \bar{P}_2 . Optically mixing these lasers at a photodetector yields the following expression:

$$\begin{aligned} & \left| \sqrt{\bar{P}_1} e^{i(2\pi f t + \phi_1)} + \sqrt{\bar{P}_2} e^{i[2\pi(f+f_0)t + \phi_2]} \right|^2 \\ &= \bar{P}_1 + \bar{P}_2 + 2\sqrt{\bar{P}_1 \bar{P}_2} \cos(2\pi f_0 t + \Delta\phi). \end{aligned} \quad (1)$$

Here, we have written the laser field amplitudes as proportional to the square root of the average power and have set $\Delta\phi = \phi_2 - \phi_1$. While the first two terms on the right side of the equation are the individual dc powers, the third term is a time varying signal, called a beat note, at the difference frequency, f_0 .

In our implementation of the heterodyne readout, the detector photocurrent, represented by Eq. (1), is digitized satisfying the Nyquist criterion for sampling signals at f_0 . The band-limited signal is then mixed to an intermediate frequency and written to file using a field programmable gate array (FPGA) card. Then, a second mixing stage in postprocessing shifts the signal to dc, splitting it into two quadratures. Each resultant quadrature is continuously integrated over the measurement time. In order for the signal to accumulate, phase coherence between the two laser fields must be maintained during this entire process. The two quadratures are then combined to give a single quantity proportional to the product of the photon rate of each laser.

Implementation of a heterodyne detection scheme in ALPS II will involve injecting a second laser, phase coherent with the signal field and resonant in the regeneration cavity at a known offset frequency. The overlapped beams are transmitted out of the cavity and are incident onto the heterodyne detector.

In this report we present results from a test setup which validates the approach and will guide its implementation in ALPS IIc. To begin, in Sec. II we mathematically demonstrate how a coherent signal is extracted from the input. In Sec. III, we then discuss the optical design created to test this technique and the digital design which forms the core of heterodyne detection. Finally, in Sec. IV we present results on device sensitivity and coherent signal measurements.

II. MATHEMATICAL EXPECTATIONS

A. Signal behavior

In our stand-alone experiment, two lasers are interfered and incident onto a photodetector. Laser 1 acts as our local oscillator (LO) with average power \bar{P}_{LO} , while laser 2 provides the signal field we wish to measure with average power \bar{P}_{signal} . The difference frequency is set such that the generated beat note has frequency f_{sig} . Once the combined beam is incident onto a photodetector with gain G in V/W, it is digitized into discrete samples, $x[n]$, where n is the individual sample number, at sampling frequency f_s . Sampling is done using an analog-to-digital converter (ADC) with a 1 V reference voltage. In the absence of noise, the ac component becomes

$$x_{\text{sig}}[n] = 2G\sqrt{\bar{P}_{\text{LO}}\bar{P}_{\text{signal}}}\cos\left(2\pi\frac{f_{\text{sig}}}{f_s}n + \phi\right), \quad (2)$$

where ϕ is an unknown but constant phase.

In order to recover amplitude information, the digitized beat note signal is separately mixed with a cosine/sine wave at frequency $f_d = f_{\text{sig}}$ in a process known as in-phase and quadrature (I/Q) demodulation:

$$\begin{aligned} I[n] &= x_{\text{sig}}[n] \times \cos\left(2\pi\frac{f_d}{f_s}n\right), \\ Q[n] &= x_{\text{sig}}[n] \times \sin\left(2\pi\frac{f_d}{f_s}n\right). \end{aligned} \quad (3)$$

Each quadrature is individually summed from $n = 1$ to N samples. The squared sums are added together and normalization is done through division by N^2 . This entire quantity is given by the following expression:

$$Z(N) = \frac{(\sum_{n=1}^N I[n])^2 + (\sum_{n=1}^N Q[n])^2}{N^2}. \quad (4)$$

The numerator is in fact the square of the magnitude of the discrete Fourier transform (DFT) of the digitized input¹ evaluated at f_d/f_s :

$$Z(N) = \frac{|X[\frac{f_d}{f_s}]|^2}{N^2}, \quad (5)$$

where

$$X\left[\frac{f_d}{f_s}\right] = \sum_{n=1}^N x[n]e^{-i2\pi\frac{f_d}{f_s}(n-1)}. \quad (6)$$

Setting $f_d = f_{\text{sig}}$ and solving for $Z(N)$ with an input given by Eq. (2) yields

$$Z_{\text{sig}}(N) = G^2\bar{P}_{\text{LO}}\bar{P}_{\text{signal}}. \quad (7)$$

Demodulating at the beat note signal frequency causes $Z(N)$ to be proportional to \bar{P}_{signal} and constant with integration time, τ . The power in the local oscillator amplifies the beat note amplitude and will be set to overcome all technical noise sources.

B. Noise behavior

We wish to set the signal field to compare with the projected sensitivity of ALPS IIc on the order of a few photons per week. Therefore, we must consider the influence of important noise sources such as laser relative intensity noise and optical shot noise. In order to understand the influence of such noise, let us determine $Z(N)$ in the absence of an rf signal ($\bar{P}_{\text{signal}} = 0$) but in the presence of noise.

Consider the input $x[n]$ to be a random stationary process. The quantity $Z_{\text{noise}}(N)$ can be written in terms of the single-sided analog power spectral density (PSD)

¹It must be noted that this is only exactly true in the case that $\frac{f_d}{f_s} = \frac{k}{N}$ for some integer k . If this requirement is not met, then the windowing process results in spectral leakage and $Z(N)$ becomes an estimate of the DFT. However, in the large N limit this leakage becomes negligible.

evaluated at the demodulation frequency, f_d . To do so, we first convert the analog PSD in $\text{V}^2/(\text{cycles per second})$ to the digitized power spectral density (DPSD) in $\text{V}^2/(\text{cycles per sample})$ using the sampling frequency f_s [27]:

$$\text{DPSD}\left(\frac{f_d}{f_s}\right) = f_s \text{PSD}(f_d). \quad (8)$$

The DPSD is related to the expectation, \mathcal{E} , of the DFT of $x[n]$ [27]:

$$\text{DPSD}\left(\frac{f_d}{f_s}\right) = \lim_{N \rightarrow \infty} \mathcal{E} \left[\frac{|X[\frac{f_d}{f_s}]|^2}{N} \right]. \quad (9)$$

Using Eq. (5) we can solve for $Z(N)$,

$$\lim_{N \rightarrow \infty} \mathcal{E}[Z_{\text{noise}}(N)] = \frac{\text{PSD}(f_d)}{\tau}, \quad (10)$$

where we use the substitution $N = \tau f_s$. It is important to note that this only depends upon the PSD evaluated at f_d and not across the entire spectrum.

Although Eq. (10) exactly relates the expectation value of $Z_{\text{noise}}(N)$ to the analog PSD, we are interested in the result of a single trial. For such an individual trial, $Z_{\text{noise}}(N)$ provides only an estimate of the analog PSD. Because the noise is assumed to be stationary, the PSD is by definition constant with time. The behavior of $Z_{\text{noise}}(N)$ for a single trial therefore tends to fall off as $1/\tau$. However, for a set integration time the outcome of multiple trials of $Z_{\text{noise}}(N)$ will have some nonzero variance [27,28],

$$\lim_{N \rightarrow \infty} \sigma_Z^2 = \left(\frac{\text{PSD}(f_d)}{\tau} \right)^2. \quad (11)$$

A confidence threshold for a single run must therefore be determined in order to distinguish between coherent detection of a signal and the random nature of this noise. From this point forward we assume N to be sufficiently large such that Eq. (10) and its derivatives provide good approximations to real world applications.

C. Detection threshold

To simplify this calculation let us assume that the input is appropriately bandpass filtered around f_d and down-sampled such that the resulting frequency spectrum is locally flat. It has been shown that in the large N limit $X[f_d/f_s]$ is a Gaussian random variable, independent of the other $X[f/f_s]$ due to the central limit theorem [28]. $Z_{\text{noise}}(N)$ therefore follows an exponential distribution. Using the cumulative distribution function [29], the probability, \mathcal{P} , of measuring a final value of $Z_{\text{noise}}(N)$ between 0 and an upper limit u for a given τ is

$$\mathcal{P}(u) = 1 - e^{-u/\sigma_Z}. \quad (12)$$

From the inverse of Eq. (12), we can define a probability range for individual outcomes of $Z_{\text{noise}}(N)$ to fall between 0 and an upper limit for any given probability \mathcal{P} . For the 5-sigma limit ($\mathcal{P}_{5\sigma} = 0.9999994$) this is

$$u(\mathcal{P}_{5\sigma})[Z_{\text{noise}}(N)] = -\ln(6 \times 10^{-7}) \frac{\text{PSD}(f_d)}{\tau}. \quad (13)$$

Consequently, when $Z(N)$ has a value above this limit for a predefined number of samples N , we can claim with 99.99994% confidence that a coherent signal is present.

The expected behaviors of $Z(N)$ and the 5-sigma limit are plotted vs integration time τ in Fig. 3. When a beat note signal is present at frequency $f_{\text{sig}} = f_d$ [Eq. (7)], the expectation value, shown in red, is constant with integration time and scales linearly with the power of the signal field, \bar{P}_{signal} . This power can be expressed in terms of photons per second, our quantity of interest.

Following Eq. (10), the expectation value of $Z_{\text{noise}}(N)$ (signal absent), shown as the solid green line, goes as $1/\tau$. Similarly the 5-sigma limit falls off as $1/\tau$ according to Eq. (13).

D. Fundamental limits

From now on, we will scale $Z_{\text{signal}}(N)$ to photons per second in the signal field, $\bar{P}_{\text{signal}}/h\nu$. A scaling factor of $1/(G^2 h\nu \bar{P}_{\text{LO}})$ is applied to Eq. (7) such that




FIG. 3. Expected behavior of noise, signal, and the 5-sigma limit when plotting $Z(N)$ vs integration time τ . Noise and the 5-sigma limit both go as $1/\tau$, whereas the signal stays flat with time. Because $Z(N)$ is proportional to the power in the signal field we can scale the y axis accordingly using the gain factors within our system in order to obtain a meaningful photon rate of the weak field. Noise-level-dependent integration times τ_x (where the signal crosses the expected value of noise) and $\tau_{5\sigma}$ (where a detection can be claimed with 5-sigma confidence) are highlighted.

$$\frac{Z_{\text{signal}}(N)}{G^2 h\nu \bar{P}_{\text{LO}}} = \frac{\bar{P}_{\text{signal}}}{h\nu}, \quad (14)$$

where h is the Planck constant and ν is the laser frequency, so that $h\nu$ is the photon energy. Scaling the noise [Eq. (10)] and 5-sigma limit [Eq. (13)] by the same factor yields

$$\frac{\mathcal{E}[Z_{\text{noise}}(N)]}{G^2 h\nu \bar{P}_{\text{LO}}} = \frac{\text{PSD}(f_d)}{G^2 h\nu \bar{P}_{\text{LO}} \times \tau} \quad (15)$$

and

$$\frac{u(P_{5s})}{G^2 h\nu \bar{P}_{\text{LO}}} = \frac{-\ln(6 \times 10^{-7}) \text{PSD}(f_d)}{G^2 h\nu \bar{P}_{\text{LO}} \times \tau}. \quad (16)$$

The fundamental noise source in our optical heterodyne detection setup as well as in ALPS IIc is shot noise (sn). This type of noise is well characterized and follows Poisson statistics [30]. Experimentally, we ensure that our system is shot-noise limited at the demodulation frequency. We may then use the known PSD for shot noise in A^2/Hz [31],

$$\text{PSD}_{\text{sn}} = 2qI_{\text{dc}}, \quad (17)$$

where q is the electron charge. The dc photocurrent, I_{dc} , is related to the total input average optical power. With $\bar{P}_{\text{LO}} \gg \bar{P}_{\text{signal}}$ the photocurrent becomes

$$I_{\text{dc}} = \eta \frac{q}{h\nu} \bar{P}_{\text{LO}}, \quad (18)$$

where η is the quantum efficiency of the photodetector. Finally, we use the photodetector gain G in order to convert this to V^2/Hz ,

$$\text{PSD}_{\text{sn}} = 2G^2 h\nu \bar{P}_{\text{LO}} \frac{1}{\eta}. \quad (19)$$

Substituting this quantity into Eq. (15) yields the expected behavior when shot noise is the dominant source at the demodulation frequency,

$$\frac{\mathcal{E}[Z_{\text{sn}}(N)]}{G^2 h\nu \bar{P}_{\text{LO}}} = \frac{2}{\eta\tau}. \quad (20)$$

Because the left-hand side of this equation is equal to the photon rate of the signal field if a signal is present, using Eq. (14) we can predict the time at which a signal will cross the expected value of this fundamental noise limit,

$$\tau_{\text{x,sn}} = 2 \frac{h\nu}{\eta \bar{P}_{\text{signal}}}. \quad (21)$$

Similarly from Eq. (16), we find that the time required for the signal to cross the 5-sigma detection threshold is

$$\tau_{5s,\text{sn}} = -2 \ln(6 \times 10^{-7}) \frac{h\nu}{\eta \bar{P}_{\text{signal}}} \approx 29 \frac{h\nu}{\eta \bar{P}_{\text{signal}}}, \quad (22)$$

in the case of a shot-noise-limited input.

In conclusion, for a weak field with a power equivalent to 1 photon/s it takes 2 s for the expected value of shot noise to decrease to the signal level with $\eta = 1$. However, it takes ~ 29 s in order to claim a detection of a signal with 5-sigma confidence. For arbitrary noise input, both integration times, as depicted in Fig. 3, can be generalized to

$$\tau_{\text{x}} = \frac{\text{PSD}(f_d)}{G^2} \times \frac{1}{\bar{P}_{\text{LO}} \bar{P}_{\text{signal}}} \quad (23)$$

and

$$\tau_{5s} = \frac{\text{PSD}(f_d)}{G^2} \times \frac{-\ln(6 \times 10^{-7})}{\bar{P}_{\text{LO}} \bar{P}_{\text{signal}}} \quad (24)$$

if the noise is locally flat around f_d . The factor between τ_{5s} and τ_{x} ,

$$\frac{\tau_{5s}}{\tau_{\text{x}}} = -\ln(6 \times 10^{-7}) \approx 14, \quad (25)$$

does not depend on the PSD, the average power of either laser, or the sampling frequency f_s .

Additionally, Eq. (23) shows the importance of a higher power \bar{P}_{LO} when the system is not dominated by shot noise. The larger the LO power, the less time required for the signal to cross the expected noise limit, thus improving the SNR. However, once \bar{P}_{LO} is large enough such that the system is shot-noise limited, τ_{x} and, consequently, the SNR no longer depend on the LO power.

III. EXPERIMENTAL SETUP

A. Optical design

To demonstrate this concept experimentally, we assembled the optical setup shown in Fig. 4. This apparatus allows us to measure the resultant beat note generated from interfering a weak signal field with our LO. There are two 1064 nm lasers used, L_1 and L_2 . L_1 is our LO and L_2 provides the field used for weak signal generation. A half-wave plate and polarizing beam splitter (PolBS) pair is placed at the start of each beam path for power control purposes. This combination also causes the outgoing light to be linearly polarized. Laser 2 passes through an electro-optic modulator (EOM) which generates sidebands to be used as the weak signal. This will be discussed in more detail later in this section. Laser 2 then passes through two neutral density (ND) filters with a combined attenuation factor of $\sim 2 \times 10^5$ in order to further reduce the weak-field signal to an appropriate level.




FIG. 4. Optical layout of the heterodyne interferometer used for single photon detection. $\lambda/2$ denotes a half-wave plate, PolBS refers to a polarizing beam splitter, BS denotes a 50/50 power beam splitter, ND refers to the neutral density filters, EOM is the electro-optic modulator, PM fiber is the polarization-maintaining optical fiber, and PD is a photodetector.

The two fields are interfered at a 50/50 power beam splitter (BS) and the combined beam is sent into a single-mode polarization-maintaining fiber. By sending both beams into the same single-mode fiber we ensure complete overlap of the spatial eigenmodes at the output coupler. After the fiber, the combined beam passes through another 50/50 power BS. Each path is then focused individually onto separate photodetectors. PD1 is used to lock the two lasers to the constant difference frequency. This is done via feedback to the laser controller for L_2 using a phase lock loop (PLL) setup. PD2 is a homemade photodetector used for our signal measurements. For a large enough local oscillator power the shot-noise level exceeds the noise equivalent power (NEP) of the photodetector and PD2 produces shot-noise-limited signals. We set the average local oscillator power to 5 mW and observe a shot noise-to-NEP ratio of 6 at the measurement frequency.

Overlapping the two lasers generates a beat note between L_1 and L_2 , called the carrier-carrier (CC) beat note at frequency f_{CC} . The error signal for the PLL feedback comes from mixing the carrier-carrier beat note with a numerically controlled oscillator (NCO), also at frequency f_{CC} , synchronized to a master clock. This feedback is controlled by the FPGA card and keeps the CC beat note stable at frequency f_{CC} .

B. Digital design

The electrical signal from PD2 is digitized via an ADC on board a FPGA card at a rate of $f_s = 64$ MHz. A simplified digital layout following the path of the photodetector signal is detailed in Fig. 5.

The signal at frequency f_{sig} is mixed down to an intermediate frequency, f_δ , on the order of a few hertz. This is done via multiplication with a sinusoid from a NCO at frequency $f_1 = f_{sig} - f_\delta$ generated with a look-up table on the FPGA card.

While it is possible to demodulate directly down to dc during the first demodulation stage simply by setting the




FIG. 5. Digital layout of detection scheme describing the digital processing techniques involved. The photodetector signal is digitized via an analog-to-digital converter at a rate of 64 MHz, after which it is mixed with a sine wave, produced by a numerically controlled oscillator, at frequency f_1 . A cascaded integrated comb filter is used to remove the higher frequency components due to mixing and downsample the data to 20 Hz, where it is written to file. f'_s and n' are used to reference the lower sampling rate. I/Q demodulation is done onboard a desktop computer, and the quadratures are individually summed and $Z(N)$ is computed.

NCO frequency to $f_1 = f_{sig}$, we observed spurious dc signals generated within the FPGA card when tested with this configuration. The strength of these signals are orders of magnitude larger than the beat notes of interest, thus preventing any useful measurements. This issue is solved by mixing the beat note signal down to the intermediate frequency, writing the data to file, and performing a second demodulation stage on a desktop personal computer. This double demodulation shifts the unwanted spurious signal to a nonzero frequency where it integrates away. With this configuration, the beat note can be accurately measured.

A cascaded integrated comb (CIC) filter [32] removes the higher frequency components resulting from the mixing process. The CIC filter also downsamples the data to a rate of $f'_s \approx 20$ Hz at which they are written to file.

The signal at f_δ is decomposed into its I and Q components via separate mixing with a cosine and sine NCO at $f_2 = f_\delta$, respectively. Considering the same signal input from Eq. (2), this process referenced to the higher sampling rate, f_s , is equivalent to

$$\begin{aligned} I_2[n] &= x_{sig}[n] \sin\left(2\pi \frac{f_1}{f_s} n\right) \times \cos\left(2\pi \frac{f_2}{f_s} n\right), \\ Q_2[n] &= x_{sig}[n] \sin\left(2\pi \frac{f_1}{f_s} n\right) \times \sin\left(2\pi \frac{f_2}{f_s} n\right). \end{aligned} \quad (26)$$

The DFT of the recorded data at the lower sampling rate, $|X[f_2/f'_s]|^2$, is then computed. The expectation values of $Z(N)$ from Sec. II must be rewritten to include this second demodulation stage. We denote these equations with a subscript “2.” The total number of samples written to file is $N' = \tau f'_s$.

With a signal present at the demodulation frequency in the absence of noise we find

$$Z_{2,\text{sig}}(N') = \frac{G^2}{4} \bar{P}_{\text{LO}} \bar{P}_{\text{signal}}. \quad (27)$$

The photon rate in the signal field is

$$\frac{4Z_{2,\text{sig}}(N')}{G^2 h\nu \bar{P}_{\text{LO}}} = \frac{\bar{P}_{\text{signal}}}{h\nu}. \quad (28)$$

Using this new scaling factor of $4/(G^2 h\nu \bar{P}_{\text{LO}})$, we obtain a quantity equal to the photon rate of the signal field after two demodulation stages.

In the case where there is only noise at the input, the PSD when the data are recorded (DPSD') must be related to the original DPSD right after the ADC. Multiplication by a sine wave reduces the DPSD by a factor of 2. The decimation stage reduces the level of the DPSD by a factor of f'_s/f_s . For $|f_2| \leq f'_s/2$,

$$\text{DPSD}'\left(\frac{f_2}{f'_s}\right) = \frac{1}{2} \frac{f'_s}{f_s} \text{DPSD}\left(\frac{f_d}{f_s}\right) = \frac{f'_s}{2} \text{PSD}(f_d). \quad (29)$$

This quantity is related to the DFT by

$$\text{DPSD}'\left(\frac{f_2}{f'_s}\right) = \mathcal{E}\left\{\frac{|X[\frac{f_2}{f'_s}]|^2}{N'}\right\} = \mathcal{E}\{Z_2(N') \times N'\}. \quad (30)$$

Solving for $\mathcal{E}[Z_2(N')]$ in terms of the analog PSD evaluated at $f_d = f_1 + f_2$ gives

$$\mathcal{E}[Z_{2,\text{noise}}(N')] = \frac{\text{PSD}(f_d)}{2\tau}, \quad (31)$$

where we use the substitution $N' = \tau f'_s$. In order to compare the expectation value of noise to that of the signal, we must apply the new scaling factor of $4/(G^2 h\nu \bar{P}_{\text{LO}})$. Doing so, we arrive at

$$\frac{4\mathcal{E}[Z_{2,\text{noise}}(N')]}{G^2 h\nu \bar{P}_{\text{LO}}} = \frac{2\text{PSD}(f_d)}{G^2 h\nu \bar{P}_{\text{LO}} \times \tau}. \quad (32)$$

For the shot-noise-limited case with PSD_{sn} given by Eq. (19), this calculation yields

$$\frac{4\mathcal{E}[Z_{2,\text{sn}}(N')]}{G^2 h\nu \bar{P}_{\text{LO}}} = \frac{4}{\eta\tau}. \quad (33)$$

Comparing it to Eq. (20), the introduction of a second demodulation stage causes the sensitivity to decrease by a factor of 2. This decrease, in turn, also causes the 5-sigma limit to increase by a factor of 2. Therefore, using two demodulation stages requires twice as long an integration

time (when compared to a single stage setup) to detect confidently a signal.²

Signal and noise add linearly in the PSD:

$$\frac{4\mathcal{E}[Z_{2,\text{total}}(N')]}{G^2 h\nu \bar{P}_{\text{LO}}} = \frac{\bar{P}_{\text{signal}}}{h\nu} + \frac{4}{\eta\tau}. \quad (34)$$

For short integration times and a low photon rate, $4/\eta\tau$ is the dominating term. After long enough integration, $\bar{P}_{\text{signal}}/h\nu$ becomes dominant, causing the curve to remain constant with time.

These equations now reflect the result of adding a second demodulation stage; however, one final experimental consideration must be taken into account. Simply lowering the power of laser 2 to subphoton per second levels reduces the CC beat note below the point at which the PLL becomes unstable. Experimentally, a stable lock can be maintained with $\bar{P}_{\text{LO}} = 5$ mW and $\bar{P}_{\text{L2}} = 60$ pW measured at PD1. This leads to a minimum CC beat note amplitude on the order of $1 \mu\text{W}$. Increasing \bar{P}_{LO} any further pushes the photodetector past the level at which it begins to saturate. Therefore, the minimum photon rate of laser 2 at PD2, such that the PLL remains stable, is 3×10^8 photons/s. In order to generate signals with field strengths below this value while still maintaining a stable PLL between the two lasers, we make use of phase modulation from an EOM.

C. EOM sideband generation

As mentioned earlier, the EOM shown in Fig. 4 was used to generate sidebands on laser 2. The EOM is driven at frequency f_{EOM} using a sine wave from a function generator that is synchronized to a maser clock. This voltage modulates the phase of the beam as it passes through the EOM. Phase modulation generates sidebands both above and below the laser frequency. These sidebands occur at k integer multiples of the drive frequency, f_{EOM} . The amount of light power in the k th order sideband is [34]

$$\bar{P}_{\text{SB},k} = J_k(m)^2 \bar{P}_{\text{L2}}, \quad (35)$$

where $J_k(m)$ is the k th order Bessel function and m is the modulation depth, dependent on the drive amplitude of the modulation. All of these sidebands beat with the LO to produce ac signals with peak amplitudes given by the following,

²In principle, it is possible to regain the earlier sensitivity while still using two demodulation stages. This is done by taking both I and Q out of the FPGA. Then a second I/Q demodulation is performed on each output channel. This results in four terms II', IQ', QI', and QQ', where the prime indicates the second demodulation stage. Using a specific combination of these terms yields the same set of equations described in Sec. 2 [33]. This concept is currently being tested and has not yet been implemented.

$$A_k = 2\sqrt{\bar{P}_{\text{SB},k}\bar{P}_{\text{LO}}}. \quad (36)$$

The two ND filters directly after the EOM attenuate the power of laser 2 and all of the subsequent sidebands by a factor of $\sim 2 \times 10^5$. The addition of these ND filters is necessary to reduce the sideband power of interest to the desired level.

The power in the k th sideband, $\bar{P}_{\text{SB},k}$, can be fine-tuned by adjusting the drive amplitude to the EOM, thus changing the modulation depth, m . To set the modulation depth to a specific value, the two ND filters are removed such that both the CC and sideband beat notes are visible on a spectrum analyzer. The ratio between the two beat note amplitudes is adjusted in order to obtain the desired modulation depth. The ND filters are then placed back into the beam path.

Using this configuration, the average power of laser 2 is set to maintain a stable PLL. Higher order sidebands fall off in power to levels comparable to the projected sensitivity of ALPS IIc. The interference between these sidebands and the LO form beat notes at known, fixed frequencies.

IV. RESULTS

A. Noise behavior and device sensitivity

We first performed a measurement with no signal field present to study the behavior of the noise in our system. Only the LO beam with power $\bar{P}_{\text{LO}} = 5$ mW is incident onto PD2. The photodetector is shot-noise limited at this level of incident light power. The photodetector has gain $G = 1.44 \times 10^3$ V/W and quantum efficiency $\eta = 0.7$. After both demodulation stages, $Z_2(N')$ is computed and the result is scaled to an equivalent photon rate using the factor stated in Eq. (33). The result of this 19-day measurement, plotted against integration time τ , is shown in Fig. 6.

$Z_2(N')$ was computed 50 additional times using different second demodulation frequencies near 2.5 Hz. The results are then scaled to the photon rate and averaged. This average is identical to the curve representing the expectation values for different integration times. Both have essentially the same amplitude and fall off as $1/\tau$ as expected. The data stream demodulated at $f_2 = 2.5$ Hz shows one representation of a shot-noise dominated signal over the integration time. In addition, the 5-sigma threshold is plotted.

The solid orange line shows the expected fundamental shot-noise limit for the given optical power if only one demodulation stage was used. As our measurement requires a second demodulation stage, the amount of shot noise returned by the measurement, scaled to photons per second, increases by a factor of 2 [Eq. (33)], shown as the dashed red line. Because the expectation value of our data lies on top of the theoretical shot-noise limit after the second




FIG. 6. Shot-noise-limited measurement with no signal field present. After the second demodulation at $f_2 = 2.5$ Hz, $Z_2(N')$ is computed and the resultant is scaled to an equivalent photon rate, shown in light blue. $Z_2(N')$ is also computed for 50 separate demodulation frequencies near 2.5 Hz. These data are then averaged to produce the dark blue curve. This average follows the expected value line, shown in solid green, based on the PSD of the noise. The dashed green line shows the 5-sigma limit that the measurement curve would cross if a signal were present. The fundamental shot-noise limit (if only one demodulation stage were required) is drawn as the solid orange line for comparison. The second demodulation stage increases the shot-noise limit by a factor of 2 (dashed red). Because the expected value of the measurement sits on top of this theoretical limit we show that shot noise is the dominant noise source in our setup.

demodulation stage, shot noise is in fact the dominant noise source in our setup.

This measurement verifies that our system is shot-noise limited and behaves as expected. Because the measurement does not cross the 5-sigma threshold, this also shows that no spurious signals are picked up over the entire 19-day integration time when laser 2 was turned off.

B. Weak signal generation and detection

In order to demonstrate that a signal is observable using heterodyne detection, we generate a beat note between the LO and an ultraweak sideband of the second laser. We choose a sideband power equivalent to $\sim 10^{-2}$ photons/s. Reducing the signal further was not possible in our current setup, as we started to pick up spurious signals electronically. While this has been observed we want to stress that the spurious signal vanishes when the EOM phase modulation is turned off. Thus, it is not an artifact of the second laser field but rather a result of the modulation itself. We assume the issue to be cross talk between the function generator driving the EOM and the FPGA data acquisition and signal processing card. Further work on generating ultraweak laser fields without electrical interference is required.

In order to generate a sideband with the specified power, we first remove the ND filters and set the local oscillator to $\bar{P}_{LO} = 5$ mW and L_2 to $\bar{P}_{L2} = 6$ μ W. Both of these measurements are taken at the photodetector input. The modulation depth is set to $m = 0.0109$ by adjusting the drive amplitude to the EOM. Using Eq. (35), the power in the second order sideband ($k = +2$) is calculated to be on the order of 10^{-15} W. The ND filters are placed back into the beam path attenuating the sideband by a factor of $\sim 2 \times 10^5$, yielding $\bar{P}_{\text{signal}} = \bar{P}_{\text{SB},2} = 6.33 \times 10^{-21}$ W. For 1064 nm light, this is equivalent to 3.39×10^{-2} photons/s in the sideband we wish to measure.

The CC beat note between L_1 and L_2 is set to 30 MHz. Phase modulation is done by driving the EOM with a sine wave at 23 MHz + 1.2 Hz. This choice of frequency sets the beat note between the second order sideband and L_1 to be at $f_{\text{sig}} = 16$ MHz + 2.4 Hz. With the first demodulation frequency set to $f_1 = 16$ MHz, the beat note of interest is therefore at 2.4 Hz when the data are written to file. These data are then imported into MATLAB, where the second demodulation is performed. Finally, we compute $Z_2(N')$ and scale the result to photons per second.

The results of this measurement are shown in Fig. 7. Demodulating at a frequency not equal to any beat note signal frequency demonstrates the expected behavior of noise. This is shown by the light blue curve for which a




FIG. 7. Shot-noise-limited signal measurement scaled to photons per second. Two demodulation stages are used with a signal present at 2.4 Hz when the data are written to file. When second demodulation is at a frequency $f_2 \neq 2.4$ Hz, the result yields the behavior of noise, shown in light blue. The trend of the expectation value for this level of noise is shown in solid green. The 5-sigma confidence line is shown in dashed green. The result when demodulating at the beat note signal frequency, $f_2 = 2.4$ Hz, is shown as the dark blue curve. This curve crosses the 5-sigma line, demonstrating a confident detection. The level that this curve flattens out to yields a rate in the sideband of interest of 3.33×10^{-2} photons/s, shown as the red line. Demodulating at a frequency 300 μ Hz away from the beat note signal, shown in yellow, highlights the energy resolution of this detection method.

demodulation 0.1 Hz away from the 2.4 Hz beat note signal was used. In this case, no coherent signal can accumulate and the only influence at the demodulation frequency is noise. This matches the trend of the expectation value of the noise, shown in solid green.

Demodulating at the beat note signal frequency of $f_\delta = 2.4$ Hz, shown as the dark blue curve, initially behaves as noise. The noise dominance continues until the signal begins to take over, causing the curve to flatten out and subsequently cross the 5-sigma threshold. The level at which this curve flattens out yields a rate for the sideband of 3.33×10^{-2} photons/s. The measured photon rate differs from expectation by 2%, a difference that we find acceptable. This error arises from both laser power measurements and modulation depth measurements. The constant level crosses the solid green expected noise curve at ~ 170 s, in agreement with the expected τ_x . A 5-sigma confidence detection is made after ~ 2500 s of integration time, in agreement with the expected $\tau_{5\sigma}$. We therefore demonstrate that our experimental setup is viable for both generating and detecting subphoton per second level signals using optical heterodyne interferometry.

Demodulation 300 μ Hz away from the beat note signal demonstrates the importance of maintaining phase coherence throughout the entire measurement. In this case, shown in yellow, the demodulation waveform drifts in and out of phase with the beat note signal. When this happens, the integrated I and Q values begin to oscillate with the difference frequency, $|f_\delta - f_2|$. This causes $Z_2(N')$ to fall off as a sinc function, preventing it from crossing the 5-sigma threshold.

V. CONCLUSION

These measurements demonstrate that heterodyne interferometry can be applied as a single photon detector. It, however, requires that the demodulation waveform maintains phase coherence with the signal during the entire integration time. Measurements at the shot-noise limit with laser 2 off did not reveal any spurious signals that would degrade the sensitivity of our setup after 19 days of integration. Therefore, we can detect coherent signals with field strengths equivalent to about 10^{-5} photons/s (1-sigma limit). In order to claim 5-sigma confident detection for such signals we require an integration time of approximately 47 days.

We also demonstrate successful generation and detection of a signal with a field strength on the order of 10^{-2} photons/s. Longer integration times and improvements in the generation of ultraweak laser fields are required to achieve lower power levels which are comparable to the projected sensitivity of ALPS IIC. Work on the generation, implementation, and detection of weaker signal fields is currently ongoing.

Our results also highlight the importance of maintaining phase coherence and stability throughout the measurement. These limitations to heterodyne detection must be taken

into account during implementation into ALPS II. For example, while our measurements are performed using free field propagating beams, ALPS II will make use of two Fabry Perot cavities. The phase noise induced by these cavities must be kept at a low enough level to prevent any notable degradation of the signal.

While this detection method emerged from the need of a single photon detector for the ALPS II experiment, heterodyne interferometric detection of weak fields can be modified for a variety of applications. Although this technique is demonstrated here using 1064 nm laser light, it can be extended to any wavelength provided that noise and the coherent signal can be decoupled. The versatility of

heterodyne detection makes it applicable to a broad range of fields including astronomy, classical communications, and biomedical imaging [35], as long as the signal field is coherent and its frequency is known.

ACKNOWLEDGMENTS

The authors would like to thank Johannes Eichholz for the phase meter design used for real time signal demodulation. This material is based upon work supported by the National Science Foundation under Grants No. PHY-1505743 and No. PHY-1802006 and by the Heising-Simons Foundation under Grant No. 2015-154.

-
- [1] C. A. Baker, D. D. Doyle, P. Geltenbort, K. Green, M. G. D. van der Grinten, P. G. Harris, P. Iaydjiev, S. N. Ivanov, D. J. R. May, J. M. Pendlebury, J. D. Richardson, D. Shiers, and K. F. Smith, *Phys. Rev. Lett.* **97**, 131801 (2006).
- [2] R. D. Peccei and H. R. Quinn, *Phys. Rev. Lett.* **38**, 1440 (1977).
- [3] S. Weinberg, *Phys. Rev. Lett.* **40**, 223 (1978).
- [4] F. Wilczek, *Phys. Rev. Lett.* **40**, 279 (1978).
- [5] K. Olive (Particle Data Group), *Chin. Phys. C* **38**, 090001 (2014).
- [6] P. Svrcek and E. Witten, *J. High Energy Phys.* **06** (2006) 051.
- [7] P. W. Graham, I. G. Irastorza, S. K. Lamoreaux, A. Lindner, and K. A. van Bibber, *Annu. Rev. Nucl. Part. Sci.* **65**, 485 (2015).
- [8] M. Meyer, D. Horns, and M. Raue, *Phys. Rev. D* **87**, 035027 (2013).
- [9] J. I. S. Catalán, E. García-Berro, and S. Torres, *Astrophys. J. Lett.* **682**, L109 (2008).
- [10] P. Sikivie, *Phys. Rev. Lett.* **51**, 1415 (1983).
- [11] H. Primakoff, *Phys. Rev.* **81**, 899 (1951).
- [12] N. Du *et al.* (ADMX Collaboration), *Phys. Rev. Lett.* **120**, 151301 (2018).
- [13] K. Zioutas *et al.* (CAST Collaboration), *Phys. Rev. Lett.* **94**, 121301 (2005).
- [14] K. Van Bibber, N. R. Dagdeviren, S. E. Koonin, A. K. Kerman, and H. N. Nelson, *Phys. Rev. Lett.* **59**, 759 (1987).
- [15] C. Robilliard, R. Battesti, M. Fouché, J. Mauchain, A.-M. Sautivet, F. Amiranoff, and C. Rizzo, *Phys. Rev. Lett.* **99**, 190403 (2007).
- [16] A. S. Chou, W. Wester, A. Baumbaugh, H. R. Gustafson, Y. Irizarry-Valle, P. O. Mazur, J. H. Steffen, R. Tomlin, X. Yang, and J. Yoo, *Phys. Rev. Lett.* **100**, 080402 (2008).
- [17] A. Afanasev, O. K. Baker, K. B. Beard, G. Biallas, J. Boyce, M. Minarni, R. Ramdon, M. Shinn, and P. Slocum, *Phys. Rev. Lett.* **101**, 120401 (2008).
- [18] K. Ehret, M. Frede, S. Ghazaryan, M. Hildebrandt, E.-A. Knabbe, D. Kracht, A. Lindner, J. List, T. Meier, N. Meyer, D. Notz, J. Redondo, A. Ringwald, G. Wiedemann, and B. Willke, *Phys. Lett. B* **689**, 149 (2010).
- [19] P. Pugat *et al.*, *Eur. Phys. J. C* **74**, 3027 (2014).
- [20] F. Hoogeveen and T. Ziegenhagen, *Nucl. Phys.* **B358**, 3 (1991).
- [21] Y. Fukuda, T. Kohmoto, S. Nakajima, and M. Kunitomo, *Prog. Cryst. Growth Charact. Mater.* **33**, 363 (1996).
- [22] P. Sikivie, D. B. Tanner, and K. van Bibber, *Phys. Rev. Lett.* **98**, 172002 (2007).
- [23] G. Mueller, P. Sikivie, D. B. Tanner, and K. van Bibber, *Phys. Rev. D* **80**, 072004 (2009).
- [24] R. Bähre, B. Döbrich, J. Dreyling-Eschweiler, S. Ghazaryan, R. Hodajerdi, D. Horns, F. Januschek, E. A. Knabbe, A. Lindner, D. Notz, A. Ringwald, J. E. von Seggern, R. Stromhagen, D. Trines, and B. Willke, *J. Instrum.* **8**, T09001 (2013).
- [25] A. D. Spector, J. H. Pöld, R. Bähre, A. Lindner, and B. Willke, *Opt. Express* **24**, 29237 (2016).
- [26] J. Dreyling-Eschweiler, N. Bastidon, B. Döbrich, D. Horns, F. Januschek, and A. Lindner, *J. Mod. Opt.* **62**, 1132 (2015).
- [27] P. Stoica and R. L. Moses, *Spectral Analysis of Signals* (Prentice Hall, Englewood Cliffs, NJ, 2005).
- [28] M. Peligrad and W. B. Wu, *Ann. Probab.* **38**, 2009 (2010).
- [29] A. Papoulis and S. Pillai, *Probability, Random Variables, and Stochastic Processes*, McGraw-Hill Electrical and Electronic Engineering Series (McGraw-Hill, New York, 2002).
- [30] M. Fox, *Quantum Optics: An Introduction*, Oxford Master Series in Physics Vol. 15 (Oxford University Press, New York, 2006).
- [31] P. M. Mayer, F. Rana, and R. J. Ram, *Appl. Phys. Lett.* **82**, 689 (2003).
- [32] E. Hogenauer, *IEEE Trans. Acoust., Speech, Signal Process.* **29**, 155 (1981).
- [33] E. Ragaini and R. F. Woodman, *Radio Sci.* **32**, 783 (1997).
- [34] B. Saleh and M. Teich, *Fundamentals of Photonics*, Wiley Series in Pure and Applied Optics (Wiley, New York, 2013).
- [35] R. H. Hadfield, *Nat. Photonics* **3**, 696 (2009).