

Measurement of mirror birefringence at the sub-ppm level: Proposed application to a test of QED

John L. Hall,* Jun Ye,* and Long-Sheng Ma†

JILA, National Institute of Standards and Technology and University of Colorado, Boulder, Colorado 80309-0440

(Received 6 January 2000; published 15 June 2000)

We present detailed considerations on the achievable sensitivity in the measurement of birefringence using a high finesse optical cavity, emphasizing techniques based on frequency metrology. Alternative approaches of laser locking and cavity measurement techniques are discussed and demonstrated. High-precision measurements of the cavity mirror birefringence have led to the interesting observations of photorefractive activities on mirror surfaces.

PACS number(s): 42.50.-p, 42.79.-e, 42.81.-i, 12.20.-m

In the last decade, progress in the preparation and understanding of mirrors of exceedingly low reflection losses has been spectacular, leading to a feasible cavity finesse approaching 10^6 . Measurements of optical phase anisotropy across the mirror surface can thus be enhanced by a similar factor. Shot-noise-limited measurement of the cavity resonance frequency with orthogonal polarizations can potentially resolve birefringence effects another factor of 10^6 smaller, limited by one's ability to split the cavity linewidth. Such measurement capability will open up interesting measurements regarding intrinsic mirror properties and their modification by light beams: the Cotton-Mouton effect in various gases [1], the influence of parity-nonconserving effects in chiral molecules, and an interesting test of QED based on magnetically induced birefringence of the vacuum, the so-called "light-by-light"—scattering Feynman diagram [2]. With the intracavity light beam of a high finesse cavity threaded through a string of strong dipole magnets, the shot-noise-limited measurement sensitivity of the cavity resonance would allow for the detection (and measurement) of the predicted Δn (birefringence) due to QED vacuum polarization (1.4×10^{-22}). Furthermore, such an experiment would allow for a search for light scalar and pseudoscalar particles (such as the axion), which can couple by a two-photon vertex. The limit for the axion–two-photon coupling constant measured by this technique should be comparable to the bound set from astrophysical arguments.

Several industrial organizations have independently invested adequate efforts to develop the art of "superpolishing" substrates to angstrom-level surface roughness, augmented by the art of depositing 40–50 alternating layers of high and low index dielectric materials, leading to the commercial availability of mirrors with losses guaranteed to be below five parts in 10^6 . Losses of ~ 1 ppm have been documented over selected submillimeter areas [3]. The resulting sharpness of the associated Fabry-Perot fringes is breathtaking to contemplate: one full fringe width is represented by a

distance below $10^{-2} \text{ \AA} = 10^{-10} \text{ cm}$. When we now feed this interferometer with a mW of technically quiet coherent light, in a one second averaging time—if all goes well and we have only shot noise as the limitation—these fringes can be effectively subdivided into about ten million parts. The resulting distance resolution is 10^{-17} cm . Sensitivity to these incredibly small distance changes has attracted wide attention because of their many potential applications, including the possibility of building interferometric antennas for gravitational wave radiation [4], and several groups worldwide are now designing and building such major facilities [5]. The three purposes of this paper are (i) to consider other attractive physical measurements as enabled by this mirror and laser-locking technology, namely, measurements based on birefringence interferometry, (ii) to document our present status of precisely measuring mirror birefringence, and (iii) to discuss the considerations that govern the achievable sensitivity in the measurement of birefringence.

We begin with a discussion of the operational interferometric experience using high grade mirrors: All such "gyro-quality" mirrors are observed to have a different reflection phase shift (or effective plane of reflection) depending upon the state of polarization of the incident light. Perhaps during polishing some microscopic stress fields were written into the substrate and not fully erased by the randomizing process of the slow spindle rotation. Perhaps during the coating process the flying molecules were incident at some inclined angle to the surface, leading to some small level of preferred-direction effects in the coating's otherwise glassy, isotropic, and nearly structureless coated layers. In any event, it is a fact of experience that all mirrors tested show some level of birefringent behavior that can be as small as a $0.1\text{-}\mu\text{rad}$ differential phase shift/bounce. Indeed, it is extremely difficult to robustly mount fused silica mirrors without inducing stress-related birefringence. A cavity formed with birefringent mirrors will accumulate the differential phase shift incurred per mirror bounce and magnify it by a factor of $(2 \times \text{finesse} / \pi)$ at the cavity transmission. This mirror-related cavity birefringence presents an adversity to many high-precision measurements using a high finesse cavity. For example, in the measurement of parity nonconservation in Cs atoms, a small systematic error is introduced when the two counterpropagating waves inside the cavity have slightly dif-

*Also at the Quantum Physics Division, National Institute of Standards and Technology, Boulder, CO.

†Permanent address: Laboratory for Quantum Optics, East China Normal University, Shanghai, China.

ferent polarizations, which result from the existence of cavity birefringence [6]. In the experiment involving optical cavity QED, cavity birefringence prevents the realization of the full potential of strong coupling between an atom and the cavity mode [7]. However, one could also take advantage of this intrinsic birefringence to boost the useful signal level. With an appropriate polarimetric setup, this dc bias can be gainfully employed to convert a quadratic signal of interest to a linearized one, with a scaling factor equal to the magnitude of cavity birefringence [8]. Accurately measuring the mirror birefringence could also lead to useful information on surface science. As for scale, a reasonable contemporary gyro-quality mirror may show about 10^{-6} waves of phase difference at normal incidence for light polarized in two perpendicular linear polarizations. Exceptional mirrors may be threefold or even tenfold less. The residual circular birefringence for almost all mirrors is at least an order smaller than the birefringence observed with linearly polarized light. Realizing that the high signal-to-noise ratio enabled by a cavity configuration (splitting of the linewidth) offers us access to another factor of $\sim 10^7$ in terms of sensitivity enhancement, it seems likely that some interesting things can be turned up.

How can we precisely measure these subtle optical phase shifts? Ideally, we will be able to measure them sufficiently accurately so that two measurements taken with some time separation can be found to agree at some attractive level of precision. Then we will be ready to jump into our applications. But first, operationally, how can we measure the effect accurately? What are the possible tools? It is instantly clear from the numbers quoted above that it is only with frequency-based metrology that we can hope to have the accuracy necessary to deal with an effect that can range from ~ 1 to $\sim 10^{-13}$ waves as the dynamic range. The possibility for the two polarizations to coexist between the cavity mirrors and then be separated externally with a polarizing prism is very attractive since it enables the measurement of the birefringent effects *differentially* between the two polarization modes. If we were to have a minor amount of residual mirror axial motion (very small relative to the laser wavelength), there would be a tremendous reduction of its consequences because of the “common-mode” nature of this excursion as viewed by the spatially coincident but separable orthogonally polarized beams. This is exactly the difference between this approach and the gravitational wave experiments such as the Laser Interferometer Gravitational Wave Observatory (LIGO) that must use the Michelson geometry: In that case the two distant mirrors exist at different spatial locations and the common-mode concept is not present. Only complex and expensive vibration isolation will work to produce the necessary inertial frame mirrors. By contrast, for our high finesse birefringence interferometer, the two beams will be spatially overlaid and sampling the same mirror surfaces. We need isolation only as good as the fringe width, while gravitational wave mirrors need isolation as good as the minimum detectable signal size, a vastly smaller quantity.

In many cases in precision measurement physics, it turns out that an equivalent signal-to-noise performance can be

realized in several different ways. For example, in the proposed QED birefringence experiment, two polarization states of light are resonantly interacting with mirrors with small polarization phase shifts. Theoretically, one finds that good performance and sensitivity can be obtained by measuring the differences in the apparent transmission when it is tuned near the high slope regions around the half height of the resonance line shape. A simpler method would be to illuminate the cavity with light polarized at $+45^\circ$ to the birefringence axes, and analyze the transmitted light with a crossed polarizer [9,10]. Another technique is to add some external ellipticity modulation to the beam in order to again linearize the signal [11]. Because of the cavity resonance effect, a very small birefringence will detune the two-cavity polarization modes by a significant part of a linewidth. A compensating phase plate in the exit beam can delay the faster component so that the dark fringe condition can be reestablished in the polarizer-transmitted light. The birefringence is then the measured phase, divided by the cavity finesse F . One sees an inconvenience for this method also: What happens if our birefringence phase could be larger than π/F ? Now we cannot excite both modes of the cavity simultaneously, so a lower finesse—broader cavity resonance—must be used.

However, it is clear we prefer to use a very high finesse—this directly increases the sensitivity. We want to use a frequency based scheme to provide the large dynamic range. For the locking, one also prefers to use some appropriate modulation method so that the desired antisymmetric resonance curve is produced by synchronous detection at the central tuning condition. A particularly attractive modulation/detection scheme was described in 1983 by Drever *et al.* [12] for producing this “discriminator” line shape using laser light reflected from the cavity. They showed that it was desirable to use a modulation frequency well above the cavity resonance response width so that the modulation sidebands would be nonresonant and would therefore be essentially reflected from the cavity input mirror. Upon being steered to the detector, these two optical frequencies form the “local oscillator” for heterodyne detection of the resonant electric field at the carrier frequency. For the birefringence measurement, there are actually two optical carriers of crossed polarizations, which are presented to the cavity. They are near or at resonance, leading to a strong resonant power buildup inside and a resulting phase shift of the reflected light that is related to the detuning. Some parts of these two internal fields leak out, returning to the detector along with the sidebands that were directly reflected as noted before. In this way one can generate the desired two polarization-separated discriminator signals. These detector voltages could be digitized and analyzed for the subtle difference in the detunings for the two polarization modes, as brought about by the mirror birefringence. For small signals, this would work well. But we need a seven decade dynamic range to cover the range between the shot-noise level and the full signal of one fringe width, and another six or so decades to deal with larger fractional signals where the equivalent birefringence is approaching one wavelength. To deal with such an extreme dynamic range, clearly a possible better choice is to lock two tunable coherent optical sources onto these cavity reso-

nances, one in each polarization, and heterodyne their outputs to recover our signal as a frequency to be counted by a contemporary frequency counter, capable of offering, for example, 12-digit resolution and accuracy in 1-s averaging time.

Conceptually, this is close to the right idea, but a clear problem will come up if the birefringence is very small: it could be that only 5% of a beat cycle would evolve during the counter gate time, making it very difficult to know the actual beat frequency accurately. Clearly we would prefer a higher frequency for the measurement. We would also prefer a much higher frequency for the physics reason that, because these two polarization modes are unlikely to be rigorously orthogonal, there will be a time-dependent thermal input via the absorption losses of the mirror coatings. If the birefringence-determined beat frequency were high enough, and the power and the losses low enough, this “zero-beat” method could perhaps be made to work via an optical frequency offset before detection. Another strategy would be to lock onto cavity modes of differing geometry, as they have different resonant frequencies. However, in the lab it is very easy to observe differential heating shifts in such a case, since our prized differential measurement principle is no longer fully valid: the heat inputs are in different spatial locations on the mirrors. We will return later to this multifrequency measurement approach.

Still, it is often useful to have a baseline for performance and performance limitations, so we first consider an alternative to the frequency-counter readout approach. This can be based on the readily available digitally synthesized frequency sources, which can offer a sinewave output of up to 30 MHz and can be programmed in steps as small as $1 \mu\text{Hz}$. So this frequency synthesis approach also offers the $>10^{12}$ dynamic range we estimate to be required. Of course we need to address the question of how to lock the laser onto the cavity for these two polarization modes. One sees two-cavity polarization modes and thus two required optical frequencies for resonance. A better strategy is to rediagonalize these two degrees of freedom into a common laser frequency variation to keep its wavelength accurately on the cavity resonance condition, with the second degree of freedom being the difference in the two polarization-mode frequencies supplied to the cavity. The average of the two locking signals will be information about the laser wavelength relative to the cavity length, while the difference between the two locking signals will be the desired birefringence. The information bandwidths can be vastly different; for example, with $B = 100 \text{ kHz}$ chosen as the locking bandwidth for the laser, a 1-Hz bandwidth for the birefringence signal would be a reasonable choice. With the lower bandwidth there is lower noise, so we see here already a factor of $\sqrt{10^5}$ in signal-to-noise ratio (S/N).

Consider the following reasonable laboratory case: $L = 1 \text{ m}$, $\lambda = 0.5 \mu\text{m}$, and finesse ($=F$) $= 10^5$. The fringe-order numbers are $\sim 2 \times 10^6$, and a change by unity results in the corresponding frequency change of $c/2L = 150 \text{ MHz}$. This “free spectral range” (fsr) is the optical frequency difference we would have if the birefringence were a 1/2 wavelength (one fringe). The fringe width is $\text{fsr}/F = 1.5 \text{ kHz}$. For

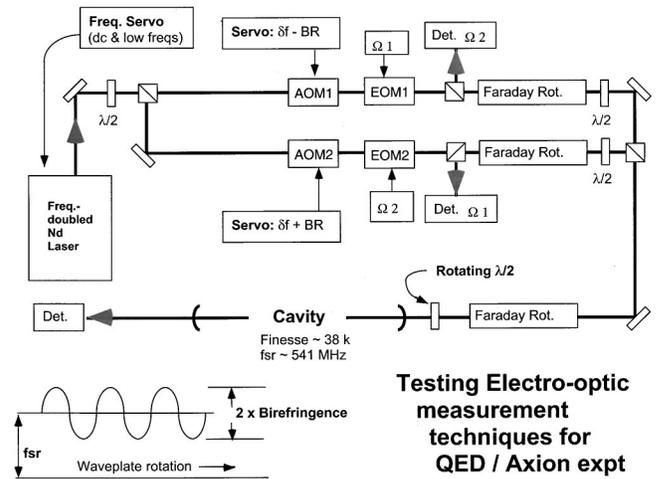


FIG. 1. A first model laser spectrometer to measure mirror or intracavity sample birefringence. The rotating wave plate exchanges the crossed linear polarized beams four times per rotation cycle. Here we use a computer-based servo to program the synthesizer with the presently estimated “correct” birefringence signal. This estimate is refined over time. The frequency input δf is associated with rapidly tracking the laser frequency to match the cavity resonance condition.

$\sim 1\text{-mW}$ detected power, the shot-noise-determined S/N is $\sim 6 \times 10^4$ in a typical 100-kHz control bandwidth, leading to frequency excursions of $\Delta = 25 \text{ mHz rms}$. The equivalent noise spectral density is $\sqrt{S_\nu} = \Delta/\sqrt{B} = 80 \mu\text{Hz}/\sqrt{\text{Hz}}$. Since the modulations at the fast Fourier frequencies accumulate little optical phase, one can expect that the Lorentzian linewidth of the stabilized laser under these ideal conditions would be $\delta\nu_L = \pi(\sqrt{S_\nu})^2 = \pi S_\nu = 0.02 \mu\text{Hz}$ [13]. In principle we can lock the laser this accurately to the cavity fringes, which appear superposed because the synthesizer(s) generate the correct offset between the two polarization modes. The shot-noise-limited equivalent birefringence for a 1-s measurement would be $\sqrt{2}\sqrt{S_\nu}/\sqrt{2\pi\tau} = 45 \mu\text{Hz}$. (The $\sqrt{2}$ comes from subtracting two somewhat noisy values to obtain the birefringence.) This 45- μHz sensitivity is to be compared to the 150 MHz that correspond to the fsr (i.e., a phase change of π), giving an ideal sensitivity of 3×10^{-13} of a fringe in 1 s, or equivalently 1×10^{-13} rad. The cavity resonance linewidth can be split with a fraction of $45 \mu\text{Hz}/1.5 \text{ kHz} = 3 \times 10^{-8}$. Of course, the ultimate resolution limit of the minimum detectable birefringence change can be as small as $\sqrt{2} \times 20 \text{ nHz} \approx 28 \text{ nHz}$, provided that the experimental integration time can be extended to where the discussion of the coherent laser linewidth becomes meaningful and can be gainfully employed for the measurement process. However, given the linewidth of 20 nHz, this integration process does not seem practical. In terms of a differential index of refraction sensitivity, if the entire cavity were filled with some gas with magnetically induced birefringence (the Cotton-Mouton effect), for example, this rather straightforward approach would bring us, in a 1-s averaging time, a birefringent index of refraction sensitivity of 5×10^{-23} . Here we have used the relation $\Delta n/n = \Delta f/f$. So it seems interesting to pursue this project further.

In Fig. 1, we show a feasible configuration of our first spectrometer, which would incorporate these concepts. Considering the extremely narrow resonance fringes, it is useful to separate the frequency servo into slow and fast subsystems. The slow frequency corrections are fed back to the neodymium-doped yttrium aluminum garnet (Nd:YAG) laser to keep it basically at the correct wavelength to match the cavity. Vibrations of the cavity mirrors and the fast frequency noise of the laser cause us to need a fast control loop, here implemented with two acousto-optic modulators, AOM1 and AOM2. Their input frequencies are synthesized from a common stem providing the variations indicated by δf in Fig. 1. Basically, this is a rapidly tunable voltage-controlled oscillator that is fed with the *sum* of the locking error signals from the two polarization-mode discriminators and that serves to make the laser frequency track (where it is incident on the cavity) very accurately any cavity noise. The *difference* of the two polarization-mode error signals is related to birefringence. As the wave plate turns, it exchanges the light from channels 1 and 2 between the two polarization modes of the cavity. Thus the signal in Fig. 1 labeled BR needs to know the wave plate's orientation to apportion the BR variation into the AOM's drive frequency. In addition, there is a common scale parameter, the birefringence, which is actually what we want to recover. Evidently, if we have set the rotating wave plate and the software requested frequency changes to be exactly tracking, and if the magnitude scale is set correctly in the software, then the error signals at the two polarization-mode detectors show no variation with the polarization plate rotation. If the BR scale is too large there are four lobes per wave plate rotation where the *H* mode has too high a frequency. If the BR scale is too small, the error sign is reversed. A nice background task for the computer then is to match the synthesizers' excursion to the physical birefringence, thus suppressing any variation synchronous with (four-times) the wave plate rotation. In analogous frequency-locking tasks using atomic references, we find a PI, I², D (P, proportional; I, integration; D, differentiation) servo works extremely well for this task. Other harmonics can be detected, related to errors in the wave plate, and taken into account. In one minute's averaging, the equivalent noise should be reduced by another attractive factor $\sim 1/\sqrt{\tau}$.

A further possible refinement of the system would be to modulate the polarization in pseudosquare wave fashion instead of sinusoidal modulation. This would improve the data gathering efficiency substantially since during sinusoidal modulation most of the time one is not looking at the maximum birefringence signal. However, some transient effects may intervene.

This system would work very nicely, but has a shortcoming when one attempts to push it toward higher sensitivity by increasing the power. In such a case, it can easily happen that the BR difference frequency is so low that the mirrors can thermally follow the time-dependent dissipation. If the absorption is not polarization sensitive, the differential character of the system strongly suppresses the associated variations. But it is clear that the approach of Fig. 1 will have a

fundamental sensitivity limit associated with the thermal response at the actual value of the birefringence-induced difference frequency.

The next idea clearly is to offset one polarized optical source by one exact free-spectral-range frequency interval. This frequency will lock almost identically to the previous one, and we will now have the desired high-frequency difference between the two intracavity waves, leading to no thermal response. As an example, in our previous effort of locking two different lasers on two adjacent modes of a single cavity, a linewidth of 70 mHz was demonstrated for the beat note of 243.735 600 MHz between the two lasers [14]. However, this approach will yield a beat frequency that also contains information about the full cavity length, as well as the birefringence. Immediately, some problems begin to appear. For one, the frequency equivalent of the shot-noise level of the detection is now some 12 or 13 decades below this beat frequency, producing serious stress in the choice of an adequately low-noise frequency reference for our counter. For another, the two optical wavelengths *are* actually a little bit different by the factor $(1+\lambda/L)$, so that the previous nearly perfect isolation from vibrations is now urgently degraded. Consider again the laboratory case: $L=1$ m, $\lambda=0.5$ μ m, finesse ($=F$) $=10^5$, and $S/N=10^6$. The fringe-order numbers are $\sim 2 \times 10^6$, and they differ by unity. The corresponding beat frequency is $c/2L=150$ MHz, while the fringe width is 1.5 kHz. The shot-noise equivalent frequency noise is still 45 μ Hz for a 1-s measurement. Since we would like to have the vibration noise as small as shot noise, we would need the vibration amplitude (measured in fringes) *a priori* to be the 3×10^{12} ratio between the fsr (the optical frequency shift per order) and the shot noise. This is mitigated by the $1:2 \times 10^6$ "nearly equal wavelengths" factor, leading us to need "only" a mirror stability of $2 \times 10^6 \times 1/(3 \times 10^{12}) = 0.7 \times 10^{-6}$ orders, corresponding to 0.07 fringe widths. We have just done this calculation for a 1-s excursion frequency, perhaps similar to the pendulum mode of the suspended mirror. This may be a possible level to achieve, but it will certainly require some effective vibration damping of the swing frequency. If we suppose that the frequency servo begins tracking the cavity length at ~ 100 kHz, using a 9-dB/octave servo filter design, we will have a gain of $10^{7.5}$ at 1 s. This means we could stand a swing amplitude as much as $10^{7.5} \times 0.7 \times 10^{-6}$ orders = 21 interference orders. Of course, at some higher frequency the gain is greatly reduced, and the tolerable vibration level would be greatly reduced. As a general conclusion, we have lost a lot by this choice of two unequal frequencies. A double-frequency modulation strategy will be discussed later.

It is instructive to put the aforementioned ideas under an experimental test. To postpone the issues about relative motion between the mirrors, the cavity actually employed is formed with a pair of superpolished "gyro" mirrors optically contacted to a hollow ULE cylinder. This robust interferometer is isolated passively from the environmental temperature changes inside an evacuated double-shelled chamber. The cavity is vibrationally isolated in the sensitive axial direction by hanging it on a thin ribbon inside the chamber. The molybdenum ribbon is centered on the cavity's

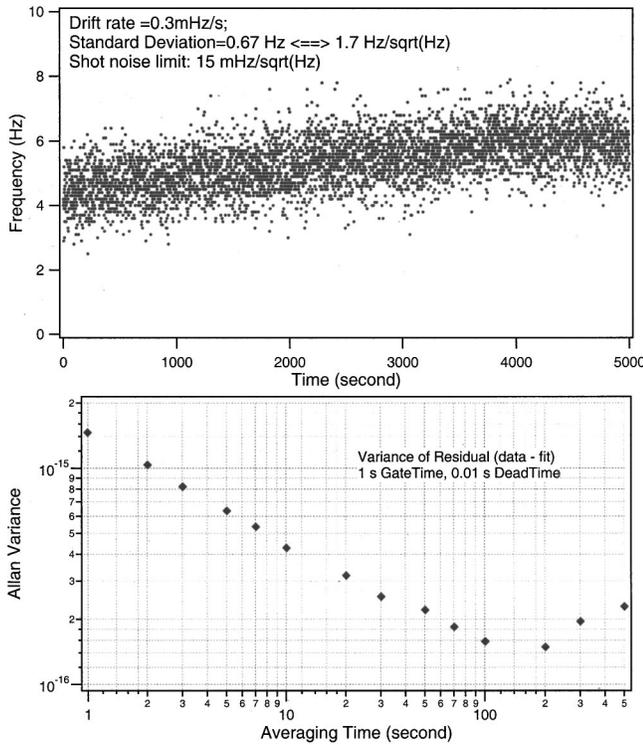


FIG. 2. Frequency readout of the VCXO when both beams are locked on adjacent cavity modes. The orthogonal linear polarizations match the cavity birefringence eigenaxis. Input powers are 10 and 2 μW . Counter gate time is 1 s.

length by a shallow groove cut into the ULE spacer. The cavity has a fsr of 541.596 MHz, a finesse of 40 720 at 633 nm, and a cavity transmission efficiency above 29%. The illustrative calculations we presented earlier can be directly applied to this cavity using the appropriate cavity parameters. A He-Ne laser beam is split into two by a polarizing cube. Each beam (a few to a few tens of μW) goes through respective optical isolators before recombining with a second polarizing cube in front of the cavity. The second beam goes through an additional frequency-offsetting AOM so that a frequency difference equal to the cavity fsr interval is introduced between the two beams. The combined beam goes through a half wave plate before entering the cavity. The half wave plate can be set to map the input polarization to two orthogonal eigendirections of the cavity birefringence. Alternatively, we can rotate the wave plate to find peak-to-peak variations of the cavity resonance due to birefringence. The cavity discrimination signals are arranged in the following way. Radio-frequency sidebands at 150 kHz are placed on the first beam that is detected in cavity reflection and used to lock the laser to the cavity. The resonance information of the second beam is recovered with a slow frequency dither on the frequency offset AOM and detected in cavity transmission. This information about the cavity fsr, which includes the contribution from the birefringence effect, is processed with a rather small bandwidth (on the order of a few hertz) and fed to a voltage-controlled crystal oscillator (VCXO) that drives the offset AOM. Figure 2 shows the time record of this VCXO frequency under the locked condition. Apart

from the apparent drift of 0.3 mHz/s, we obtain a frequency noise density of $1.7 \text{ Hz}/\sqrt{\text{Hz}}$, which is about 100 times above the shot-noise limit. The noise is in part due to problems of inadequate optical isolation and seismic isolation, and excessive acoustic noise in the laboratory. (Remember, in this experiment the seismic contribution is not common mode and remains an important noise source.) The associated Allan variance reaches 1.5×10^{-16} at a 100-s averaging time. It is interesting to note that this noise level is reduced when the input beams have circular polarization, leading to the conclusion that mirror birefringence also contributes to the measured noise level. We will elaborate more on this later when we attend to the issue of the dynamic noise of cavity birefringence.

In a subsequent experiment, instead of having the input polarization fixed, we continuously rotate the half wave plate in front of the cavity to exchange the polarization of the two input beams. The resultant variation of the cavity fsr is shown in Fig. 3(a), with a static birefringence amplitude of 304 Hz. A direct fast Fourier transform (FFT) of the original data [Fig. 3(b)] shows a noise floor 85 dB below the dc amplitude of 304 Hz, leading to a noise density of $2 \text{ Hz}/\sqrt{\text{Hz}}$, which is close to that of Fig. 1. However, since a full 360° rotation of the half wave plate results in four sinusoidal cycles in the variation of fsr, a more accurate approach to finding the noise floor is to average over the neighboring four cycles to smooth out any defects associated with the wave plate. We therefore perform a sinusoidal fit to each data cycle and extract the corresponding average amplitude. Figure 3(c) plots the time variation of the four-cycle averaged birefringence amplitude and in Fig. 3(d) we have the corresponding FFT spectrum. The lowest noise approaches 100 dB below the static 304-Hz level, corresponding to a $0.4 \text{ Hz}/\sqrt{\text{Hz}}$ noise level. The Fourier spectrum also indicates the optimum frequency ($\sim 2 \text{ mHz}$) that we should choose to modulate the magnetic field if we decide to search for a magnetic-field-induced physical signal. Equivalently, the birefringent phase difference per mirror bounce is $304 \text{ Hz}/541.6 \text{ MHz} \times \pi \sim 1.8 \mu\text{rad}$ and the associated birefringent phase noise is $2 \text{ nrad}/\sqrt{\text{Hz}}$. In terms of measuring the difference in the index of refraction between the two orthogonal polarizations, this mirror birefringence noise sets a lower limit of $\Delta n/n \sim 8 \times 10^{-16}/\sqrt{\text{Hz}}$ on the attainable sensitivity. Compared against the calculated result of the illustrative 1-m cavity, our experimental data show a degradation of the lower limit of the birefringence noise by a factor of 10^4 . This is partly due to the lowered finesse (2.5 times), smaller optical power (100 times), and shorter cavity length (3 times). However, the main limitation is due to the fact that we are still 100 times above the shot-noise limit. The signal level associated with the magnetically induced birefringence of the vacuum is predicted to be around 1.4×10^{-22} ; hence a serious task remains for apparatus improvement, and later for signal averaging. Two previously mentioned issues, namely, optical isolation and vibration isolation, will be far more demanding when we scale up the interferometer size to a few tens of meters and the input optical power to nearly 1 W.

The next interesting and important issue concerns the dynamic behavior of this birefringence noise associated with

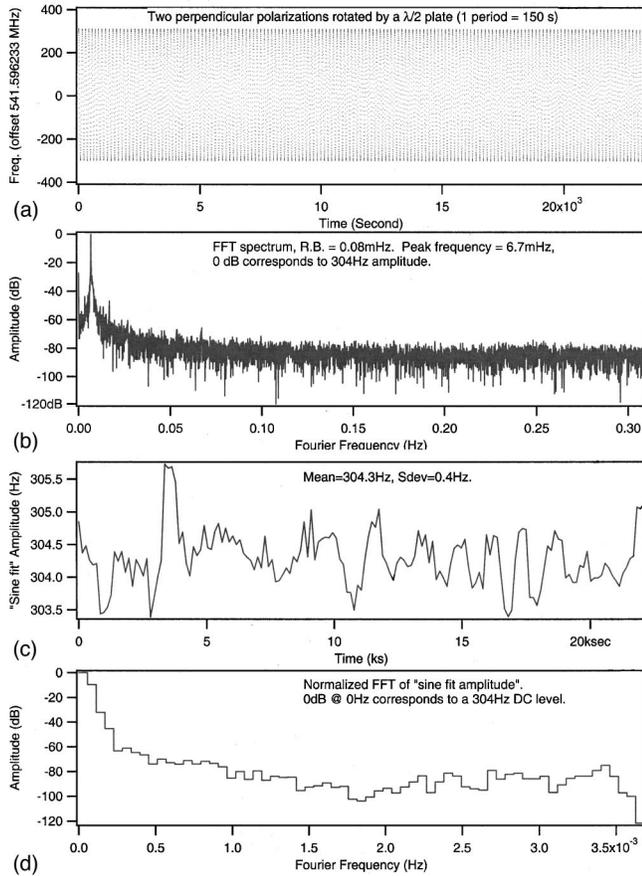


FIG. 3. (a) Time series showing 160 cycles of sinusoidal frequency change between two different cavity eigenfrequencies, dependent upon the rotation of the input linear polarization relative to the cavity birefringence axes. For the $15\text{-}\mu\text{W}$ utilized light, the noise density is $1.5\text{ Hz}/\sqrt{\text{Hz}}$, compared with the shot-noise limit of $15\text{ mHz}/\sqrt{\text{Hz}}$. (b) Fast Fourier transform of the data in (a). (c) Amplitude of sinusoid fit to data. The original data are fitted cycle by cycle and averaged over four cycles, corresponding to a full 360° rotation of the half wave plate. (d) Fourier distribution [transformed from the data in (c)] of the measured birefringence. A sensitivity of $\sim -100\text{ dB}$ relative to the mirror static birefringence can be achieved. This 4-mHz birefringence noise gives an index of refraction measurement sensitivity $\Delta n/n$ of 6.5×10^{-18} and can be readily reduced using a narrower linewidth cavity and higher laser power.

the mirror coating. Any time-dependent variation of the noise needs to be explored to find the optimum operating time scale for signal averaging. Besides, we need to know in advance the power-dependent noise fluctuations so that we can stabilize the input laser power to an acceptable level. Indeed, in our first experiment with power variation, we already have strong evidence that some dynamical photorefractive process in the mirror coatings is playing a role in the change of mirror birefringence.

To have sufficiently high power to explore the power-related birefringence noise, we next used the frequency doubled output of a Nd:YAG laser. The cavity finesse at 532 nm is similar to that at 633 nm. However, the static birefringence amplitude is increased to about 1.5 kHz for the green light. With input power now at the level of a few hundred

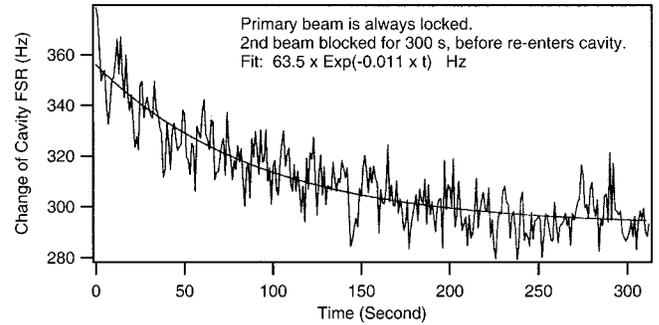


FIG. 4. Time-dependent cavity birefringence signal. The first 532-nm beam (0.25 mW) is left on all the time, locking the laser to the cavity. The second beam (0.5 mW) is switched on at 0 s and locked to the adjacent cavity mode. The VCXO frequency, representing the cavity fsr, shows a relaxation process with a step size of 64 Hz and a time constant of $\sim 90\text{ s}$.

μW to 1 mW, we observe very interesting characters of the cavity birefringence. First, as in Fig. 2, two beams of orthogonal polarization (each aligned to the cavity birefringence eigenaxis) are injected and locked to the cavity, and the cavity fsr information is read out through the VCXO. When the power level is similar to that used in the He-Ne case, the resulting frequency noise of the VCXO, representing the cavity fsr, is about 1.8 times worse than that in Fig. 2. However, it is clear that the frequency noise increases (by a factor of 2) as the input power levels are raised (by ten times). This noise becomes much worse when we rotate the two input polarization directions by 45° with respect to the eigen axes of cavity birefringence. In that case, the time record of the VCXO frequency actually shows a bistable behavior with two states separated by about 160 Hz. The second interesting experiment is based on beam switching. By leaving the first beam on all the time, the servo keeps the laser locked to the cavity. The second beam, which probes the next neighboring cavity resonance, can be switched on and off with the AOM in its path. With the second beam just switched on, we observe the cavity fsr frequency jumps by $\sim 70\text{ Hz}$ and then relaxes down to the equilibrium value with a time constant of 70 s. Figure 4 shows one representative trace of this switching transient. The power of the switched second beam is 0.5 mW, which is twice that of the first beam. If we choose to switch both beams, then the initial switching step size is reduced to $\sim 40\text{ Hz}$ while the relaxation time is extended to $\sim 500\text{ s}$. This is understandable since the effects from both switching beams will compensate each other to some degree. The most interesting observation is related to the use of circular polarization for both input beams. For circular polarization we simply replace the half wave plate positioned just in front of the cavity with a quarter wave plate. In this case, we observe absolutely no switching transients beyond the rms noise level. With both beams on at all times, the circular polarization also provides a more stable fsr reading, with the frequency noise cut down by a factor of 3 to 5 compared to the most stable case in which orthogonal linear polarizations are used.

The conclusion from these tests is that the mirrors develop memories when they are fluxed with linearly polarized

light, given that the light power is sufficiently high to initiate the dynamical process. This is the reason we believe that the observed phenomenon originate from photorefractive activities [15]. Circular polarization can erase that memory by sampling all polarization directions and averaging them with an infinitely faster rate than the time scale associated with the mirror photorefractivity. The lesson we draw from the experiment here is that one needs to be very careful when dealing with polarization issues in other precision measurements using a high finesse optical cavity. The associated polarization noise/transients can present a much more serious problem than the traditional concerns of thermal stability. Weak depolarized components of the input light beam can produce nonthermal time-dependent birefringence noise. Input light power as well as beam steering (which could effect the intracavity buildup power) need to be stabilized to acceptable levels. On the other hand, it is fascinating to think about the possibility of exploring photorefractivity on mirror surfaces with a spatial resolution of a few micrometers by simply displacing mirrors in corresponding steps transverse to the cavity optical axis.

Coming back to the issue of measuring the birefringent noise, another problem as seen in Fig. 2 is the drift between the two length/frequency standards in our experiment: the crystal oscillator resonator in the frequency counter and the light stabilized frequency/wavelength linked to the mirror separation. How will we know if the change over some time was associated with our changing the gas physics inside the cavity, or only some drift of the length? Or of the counter's reference crystal?

A clearly better strategy is to use *two* sidebands in one polarization—let us say *V* polarized—one above and one below the frequency of the reference beam in the other polarization, the *H* mode. Each *V* component is shifted in frequency by exactly one fsr and so is in resonance with the cavity on the adjacent interference order (including the small frequency shift due to birefringence that we wish to measure). One component frequency is below the reference and has a longer wavelength, and gathers a smaller optical phase shift when the mirror swings; the other component is symmetrically above the reference frequency and so gathers a little extra phase from its shorter wavelength. Evidently, if we use equal weights for these two detected signals, we will have an average wavelength. One may speak of a synthetic wavelength, ideally equal to that of the reference beam in the orthogonal polarization. In this case we again recover the exceedingly strong isolation from the vibrational noise.

One can notice that there is now an added degree of freedom in the problem, which can usefully be taken as the common separation of the fsr sidebands from the *V* carrier (which was their source). This fsr frequency, under servo lock conditions, will accurately represent information about the mirror spacing, while the other degree of freedom represents the small difference in the cavity resonance condition due to the birefringence we wish to measure. This can be read as one-half the difference from the *H*-polarized reference beam up to the *V*-polarized upper sideband, *minus* the frequency difference from the *H*-polarized reference down to the lower *V*-polarized sideband.

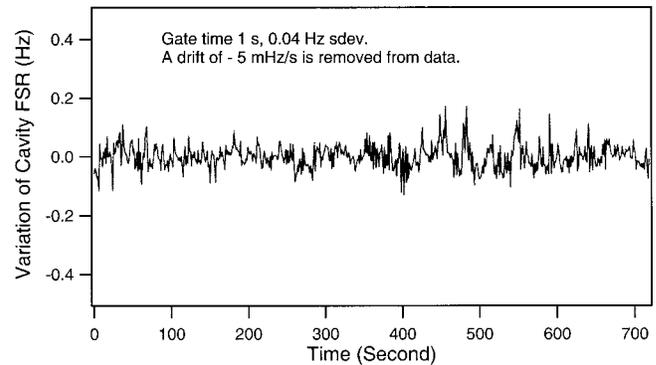


FIG. 5. Cavity fsr measurement with a FM modulation frequency matching the cavity fsr. Counter gate time is 1 s. The average value of the counted frequency (the fsr) was $319.694\,953\text{ MHz} \pm 3\text{ Hz}$.

Toward this approach we have already developed the technique of measuring the cavity fsr with rf sidebands matching the neighboring cavity modes. In fact, we have used this technique to measure the intracavity dispersion of weak molecular resonances [16]. As the laser itself is locked on the cavity, the information about the match between the sideband frequency and the fsr cannot be obtained by simply detecting and demodulating it at the sideband frequency in the cavity reflected light, which always yields a null signal. We choose to use an additional FM dither modulation on the sideband frequency itself, along with synchronous detection of this component in the cavity reflected or transmitted light. The rf sideband frequency is synthesized with two signal sources, a stable frequency synthesizer and a tunable frequency VCXO. Their sum frequency is bandpass filtered and amplified before it is fed to the corresponding electro-optic modulator. This configuration allows us to utilize the FM dither signal of the frequency synthesizer source while we separately servo control the frequency of the VCXO. In doing so we are able to obtain a better modulation quality and also avoid any change of modulation parameters when the VCXO is tuned. This additional dither frequency is low, typically in the range of 10–300 kHz. Phase-sensitive detection is at the dither frequency and allows locking for the maximum transmission of the sideband. This leads to the tracking between the sideband frequency and the cavity fsr. Figure 5 shows the counted (gate time 1 s) frequency record of the VCXO under the locked condition. After removal of a -5-mHz/s drift, we have a rms noise in the frequency read-out of about 40 mHz at the 1-s averaging time. This is 15 times more stable than the data shown in Fig. 2. Here the input beams, including the carrier and the two sidebands, are circularly polarized. This last described work was done at $1.064\ \mu\text{m}$ using a Nd:YAG laser, and the cavity had a fsr around 320 MHz and a finesse of 96 000. Again, we have not achieved the shot-noise limit in this case with the technical noise arising from the residual amplitude modulation (RAM) associated with the FM sideband generation. The RAM also degrades the long term stability of this VCXO-fsr tracking system. At the present stage, we have already made good progress toward reduction/control of the RAM [17]. It is

clear that a future birefringence measurement system can benefit a great deal from this initial demonstration.

In summary, we have made measurements on cavity mirror birefringence with the emphasis on the attainable sensitivity in future birefringence experiments involving optical cavities. Some important issues regarding how to proceed with these high-precision measurements are addressed. And we believe that a frequency metrology based approach offers a unique advantage in recovering the magnetic-field-induced

vacuum birefringence. Alternative cavity laser locking schemes are discussed and demonstrated.

We acknowledge useful discussions with colleagues in the P877 Fermilab QED magnetic birefringence collaboration, Dr. S. A. Lee, Dr. W. M. Fairbank, and Dr. F. A. Nezrick. This work was supported in part by the National Institute of Standards and Technology, the U.S. Air Force Office of Scientific Research, and the National Science Foundation.

-
- [1] A. Cotton and H. Mouton, C. R. Hebd. Seances Acad. Sci. **141**, 317 (1905); C. Rizzo and A. Rizzo (unpublished).
- [2] For theory and QED calculation, see G. Raffelt and L. Stodolsky, Phys. Rev. D **37**, 1237 (1988); S. L. Adler, Ann. Phys. (N.Y.) **67**, 599 (1971). For previous experimental efforts, see Y. Semertzidis *et al.*, Phys. Rev. Lett. **64**, 2988 (1990); R. Cameron *et al.*, Phys. Rev. D **47**, 3707 (1993). Our current experimental undertaking is motivated by collaboration with S. A. Lee *et al.* to measure the magnetic-field-induced birefringence using a high finesse optical resonator (S. A. Lee and F. A. Nezrick *et al.*, Proposal No. P877, Fermi National Accelerator Laboratory, 1998).
- [3] G. Rempe, R. J. Thompson, H. J. Kimble, and R. Lalezari, Opt. Lett. **17**, 363 (1992).
- [4] A. Abramovici *et al.*, Science **256**, 325 (1992).
- [5] Interested readers are referred to the internet web site of LIGO: <http://www.ligo.caltech.edu>, and the links contained therein.
- [6] C. S. Wood, S. C. Bennett, D. Cho, B. P. Masterson, J. L. Roberts, C. E. Tanner, and C. E. Wieman, Science **275**, 1759 (1997); C. Wood *et al.*, Opt. Photonics News **54** (1996).
- [7] H. Mabuchi, J. Ye, and H. J. Kimble, Appl. Phys. B: Lasers Opt. **B68**, 1095 (1999); C. J. Hood, J. Ye, and H. J. Kimble (unpublished).
- [8] D. Chauvat *et al.*, Appl. Phys. Lett. **73**, 1032 (1998).
- [9] S. Moriwaki, H. Sakaida, T. Yuzawa, and N. Mio, Appl. Phys. B: Lasers Opt. **B65**, 347 (1997).
- [10] F. Brandi *et al.*, Appl. Phys. B: Lasers Opt. **B65**, 351 (1997).
- [11] D. Bakalov *et al.*, Quantum Semiclass. Opt. **10**, 239 (1998); R. Cameron *et al.*, Phys. Rev. D **47**, 3707 (1993).
- [12] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward, Appl. Phys. B: Photophys. Laser Chem. **B31**, 97 (1983).
- [13] M. Zhu and J. L. Hall, J. Opt. Soc. Am. B **10**, 802 (1993).
- [14] Ch. Salomon, D. Hils, and J. L. Hall, J. Opt. Soc. Am. B **5**, 1576 (1988).
- [15] D. Anderson (private communication).
- [16] J. Ye, L.-S. Ma, and J. L. Hall, J. Opt. Soc. Am. B **15**, 6 (1998); L.-S. Ma, J. Ye, P. Dubé, and J. L. Hall, *ibid.* **16**, 2225 (1999).
- [17] J. L. Hall, J. Ye, L.-S. Ma, K. Vogel, and T. Dinneen, in *Laser Spectroscopy XIII*, edited by Z.-J. Wang, Z.-M. Zhang, and Y.-Z. Wang (World Scientific, Singapore, 1998), p. 75.