Heterodyne laser frequency stabilization for long baseline optical interferometry in space-based gravitational wave detectors

Johannes Eichholz,* David B. Tanner, and Guido Mueller

Department of Physics, University of Florida, PO Box 118440, Gainesville, Florida 32611, USA

(Received 31 March 2015; published 21 July 2015)

The European Space Agency (ESA) selected the gravitational universe as the science theme for L3, a large space mission with a planned launch in 2034. NASA expressed a strong interest in joining ESA as a junior partner. The goal of the mission is the detection of gravitational waves of frequencies between 0.1 mHz and 0.1 Hz, where many long-lived sources are expected to be steady emitters of gravitational waves. Most likely, the mission design will evolve out of the earlier Laser Interferometer Space Antenna (LISA) concept. The interferometric heterodyne phase readout in LISA is performed by phase meters developed specifically to handle the low light powers and Doppler-drift of laser frequencies that appear as complications in the mission baseline. LISA requires the frequency noise of its seed lasers to be below 300 $\mathrm{Hz}/\sqrt{\mathrm{Hz}}$ throughout the measurement band due to uncertainties in the absolute interferometer arm lengths. We have developed and successfully demonstrated Heterodyne Stabilization (HS), a novel cavity-laser frequency stabilization method that integrates well into the LISA mission baseline due to similar component demand. The cavities for the test setup were assembled with Clearceram-Z spacers, an ultralow thermal expansion coefficient material with potential applicability in interferometric space missions. Using HS, we were able to suppress the frequency noise of two lasers in a bench-top experiment to a level that meets the LISA requirement, suggesting both HS and Clearceram-Z can be considered in future mission concepts.

DOI: 10.1103/PhysRevD.92.022004 PACS numbers: 04.80.Nn, 07.87.+v, 95.55.Br, 95.75.Kk

I. INTRODUCTION

Gravitational waves (GWs) are a central feature of metric theories of gravity and an important test for general relativity. These propagating distortions of space-time are generated in observable magnitude only by violent astronomic events such as binary mergers of extremely massive, dense objects like neutron stars or black holes. GWs have great potential to complement conventional astronomy, since they are completely distinct from electromagnetic radiation.

A handful of GW observatories are being constructed or upgraded around the globe (LIGO, VIRGO, KAGRA, and GEO600), and coordinated network efforts are working toward a first direct observation of GWs [1]. However, seismic and gravity-gradient noise prevent Earth-bound detectors from sensing GWs in the sub-Hertz regime [2], where some important verification binaries—systems with known orbital periods that have been observed with radio/optical telescopes—are expected to be steady, monochromatic sources of gravitational radiation [3].

The Laser Interferometer Space Antenna (LISA) mission is a proposed space-borne GW detector that has been in development since the 1990s. It aims to observe GWs in the frequency band from 0.1 mHz to 0.1 Hz with three spacecraft that trail the Earth around the Sun [4]. Their orbits are optimized to maintain a nearly equilateral triangular constellation with as little deviation as possible, while interferometrically monitoring variations in the inter-spacecraft distances.

The several million kilometer long interferometer arms between the spacecraft are generally unequal in length, which means that common-mode laser frequency noise is not rejected, a major difference from Earth-bound GW detectors [5]. However, a near-equal arm length interferometer can be synthesized in postprocessing with time-delay interferometry (TDI) [6]. The impact of laser frequency noise in the TDI output scales with the residual uncertainty ΔL in the arm length mismatch. For a ΔL of 1 m, the LISA requirement for the frequency noise $\tilde{\nu}$ of a prestabilized laser demands that

$$\tilde{\nu} \leq 300 \times \sqrt{1 + \left(\frac{3\ \mathrm{mHz}}{f}\right)^4} \times \frac{(1\ \mathrm{m})}{(\Delta L)} \left[\frac{\mathrm{Hz}}{\sqrt{\mathrm{Hz}}}\right] \qquad (1)$$

in the range $0.1\ \mathrm{mHz} < f < 0.1\ \mathrm{Hz}$ so as not to bury potential GW signals [7]. An optical cavity made from an ultralow coefficient of thermal expansion (CTE) material such as ULE [8] or Zerodur [9] is commonly used as a long-term stable length reference for the wavelength of a laser. We investigate Clearceram-Z [10] as a potential alternative material.

We further propose Heterodyne Stabilization (HS) as a novel laser frequency stabilization method that can be

*eichholz@phys.ufl.edu

integrated into the LISA mission baseline with a minimal additional component count.

II. HETERODYNE STABILIZATION METHOD

We developed HS in an effort to embed the laser frequency control scheme as a subcomponent into the LISA interferometric readout system, as opposed to having a separate, independently operating laser stabilization unit. This could reduce the amount of auxiliary components in the spacecraft and simplify the laser phase modulation spectra, at the cost of a slightly increased complexity of the optical bench layout and the additional digital circuitry in the phase measurement hardware.

The principle of HS is similar to that of the Pound–Drever–Hall (PDH) technique [11]. The frequency of a laser is locked to a resonance of an optical cavity using the interaction phase shift, which the laser field experiences in reflection from the cavity, as a discriminator. The key difference from PDH is that, rather than phase modulating the to-be-stabilized carrier field, HS utilizes the fact that interferometric measurements in LISA are obtained by means of heterodyne interferometry.

A. Heterodyne interferometry

A superposition $E = E(x, y, z, t)$ of two laser fields, traveling in the z direction with amplitude profiles $E_i = E_i(x, y)$, angular frequencies ω_i, wave numbers $k_i = \omega_i/c$, and phases ϕ_i, respectively,

$$E = E_1 e^{-i(\omega_1 t - k_1 z + \phi_1)} + E_2 e^{-i(\omega_2 t - k_2 z + \phi_2)}, \tag{2}$$

produces a voltage $V(z,t) \propto \iint |E(x,y,z,t)|^2 dx dy$ as the output of a photo detector (PD) at location z that reads

$$V(z,t) = \alpha[P_1 + P_2 + 2\kappa\sqrt{P_1 P_2}\cos(\Delta\omega t + \Delta\phi(z))]. \tag{3}$$

Here P_i is the power in beam i, the constant α depends on the quantum efficiency and transimpedance gain of the detector, and $0 \le \kappa \le 1$ is a visibility factor that is determined by the spatial distributions of the individual fields and their overlap with each other. For the remainder of this discussion, we assume $\kappa = 1$. The time-dependent term, which oscillates at the difference angular frequency $\Delta\omega = (\omega_1 - \omega_2)$, is the beat note between the lasers, and the propagation phases $k_i z$ have been absorbed into the phase difference

$$\Delta\phi(z) = -(k_1 - k_2)z + \phi_1 - \phi_2. \tag{4}$$

In heterodyne interferometry, beat notes of the superposition (2) are sampled at probe points z_i in the interferometric setup, and further demodulation by a phase meter (PM) extracts the $\Delta\phi(z_i)$'s. Since ϕ_1, ϕ_2, k_1, and k_2 are independent of the location, differential variations between pairs of $\Delta\phi(z_i)$ reveal optical path length fluctuations in the interferometer.

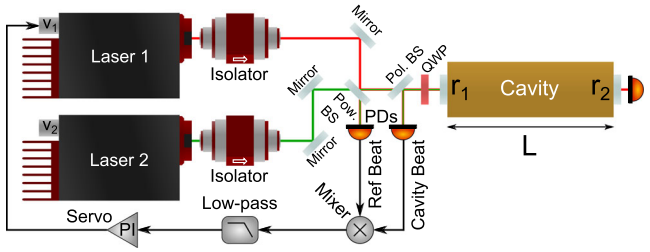

FIG. 1 (color online). Basic HS setup. Lasers 1 (resonant) and 2 (off resonant) are superimposed and reflect off an optical cavity. Their beat notes are formed with and without cavity interaction. The reflected beat carries the cavity interaction phase of the resonant laser, which is isolated via demodulation against the reference beat. The sequence of the mixer and low-pass filter generates the error signal, which is driven to zero by actuating laser 1's frequency to follow the cavity resonance.

In the LISA mission, baseline laser beat signals formed on the optical bench come overlaid with secondary beat notes for inter-spacecraft clock synchronization and broadband phase modulation of the carrier for ranging and communications [12]. Furthermore, a pilot tone may be added electronically to compensate for timing jitter in the analog-to-digital converters (ADCs) [13]. Stabilization with the PDH technique was used in earlier approaches [14,15], but it requires an additional set of sidebands to be added to the existing zoo of frequency components. In contrast, the HS method, which is illustrated in Fig. 1, uses only existing beat signals and does not require modulation of the carrier to isolate the cavity interaction phase.

B. Cavity interaction phase

Assume the superposition (2) reflects off a linear cavity of length L between two mirrors with transmissivities T_1 and T_2, located at $z_1 = 0$ and $z_2 = L$, respectively. Neglecting losses in the mirrors, their reflection coefficients are given by $r_1 = \sqrt{1 - T_1}$ and $r_2 = \sqrt{1 - T_2}$, such that in the reflected field at the front mirror

$$E_r(t) = F_r(\nu_1)E_1 e^{-i\omega_1 t} + F_r(\nu_2)E_2 e^{-i\omega_2 t} \tag{5}$$

each individual field is modified by the reflective transfer function

$$F_r(\nu) = \frac{r_1 - r_2 e^{2\pi i \frac{\nu}{\mathrm{FSR}}}}{1 - r_1 r_2 e^{2\pi i \frac{\nu}{\mathrm{FSR}}}}, \tag{6}$$

where $\nu = \omega/2\pi$ is the laser frequency and $\mathrm{FSR} = c/2L$ is the free spectral range of the cavity. $F_r(\nu)$ is periodic with resonance frequencies $\nu^n_{\mathrm{res}} = n \cdot \mathrm{FSR}$. Far off resonance, one finds that $F_r(\nu) \approx 1$, but in close proximity, $\nu - \nu^n_{\mathrm{res}} = \delta\nu \ll \Delta\nu$, which is the linewidth of the cavity, $F_r(\nu)$ can be approximated by

$$F_r(\delta\nu) \approx \frac{r_1 - r_2}{1 - r_1 r_2} - 2\pi i \frac{r_2(1 - r_1^2)}{(1 - r_1 r_2)^2} \frac{\delta\nu}{\mathrm{FSR}}. \quad (7)$$

Using the finesse $\mathcal{F} = \mathrm{FSR}/\Delta\nu$, which can be calculated from $\mathcal{F} = \frac{\pi\sqrt{r_1 r_2}}{1 - r_1 r_2}$ for high finesse cavities, the imaginary part of the response function can be rewritten as

$$\Im\{F_r(\delta\nu)\} \approx -\frac{2}{\pi}\left(\frac{1}{r_1} - r_1\right)\mathcal{F}\frac{\delta\nu}{\Delta\nu} = -\mathcal{G}\frac{\delta\nu}{\Delta\nu}, \quad (8)$$

where we have introduced the optical gain $\mathcal{G}$ as a shorthand. In overcoupled cavities, where $T_2 \ll T_1$, one can show that $\mathcal{G} \lesssim 4$. The expression in (8) is proportional to $\delta\nu$ and therefore responds linearly to differential changes between laser frequency and cavity resonance.

C. Error signal generation

In the HS scheme, one laser resonates in the cavity, and the second laser provides the reference phase for the demodulation of the interaction phase. If we assume only laser 1 to be near resonant in (5), with $\nu_1 - \nu_{\mathrm{res}} = \delta\nu$, but laser 2 to be off resonant, the reflected field becomes

$$E_r(t) = F_r(\delta\nu)E_1 e^{-i\omega_1 t} + E_2 e^{-i\omega_2 t}. \quad (9)$$

Following (2) and (3), the output $X(t)$ of a PD that samples $E_r(t)$ reads

$$X(t) = \alpha[|F_r(\delta\nu)|^2 P_1 + P_2 + 2\sqrt{P_1 P_2} \times [\Re\{F_r(\delta\nu)\}\cos(\Delta\omega t) + \Im\{F_r(\delta\nu)\}\sin(\Delta\omega t)]]. \quad (10)$$

Because $\sin^2(\varphi) = \frac{1}{2}[1 - \cos(2\varphi)]$, $\Im\{F_r(\delta\nu)\}$ can be isolated via multiplying $X(t)$ with the output $Y(t)$ of a reference PD of relative phase shift $\Delta\phi(z) = -\pi/2$ in (3),

$$Y(t) = 2\alpha\sqrt{P_1 P_2}\sin(\Delta\omega t). \quad (11)$$

The product contains terms that oscillate at $\Delta\omega$ and $2\Delta\omega$, which need to be removed via subsequent low-pass (LP) filtering, yielding

$$X(t) \times Y(t) \overset{\mathrm{LP}}{=} 2\ell\alpha^2 P_1 P_2 \Im\{F_r(\delta\nu)\}. \quad (12)$$

Here, ℓ factors in the conversion power loss of the mixer, which depends on the frequency and the amplitude of either beat, but usually lies between 4 and 6 dB, resulting in $\ell \approx 0.5\ \mathrm{V}^{-1}$. The error signal $e(\delta\nu)$ becomes

$$e(\delta\nu) = -2\mathcal{G}\ell\alpha^2 P_1 P_2 \frac{\delta\nu}{\Delta\nu}[\mathrm{V}]. \quad (13)$$

Driving $e(\delta\nu)$ to zero is equivalent to $\delta\nu = 0$, which establishes the lock between the resonant laser and cavity.

III. CONTROLLER IMPLEMENTATION IN PM HARDWARE

The LISA phase measurement system is a directly digitizing, dynamically tracking PM. It has been shown to cope well with low-light powers and Doppler drifts of heterodyne frequencies [13,16]. The noise budget requires the PM to extract the beat note phases with no readout noise in excess of 1 $\mu\mathrm{cycle}/\sqrt{\mathrm{Hz}}$ at 3 mHz or above, but relaxed quadratically toward lower frequencies identically to the frequency stability requirement in (1).

PM prototyping is based on field programmable gate arrays (FPGAs) [17], programmable microchips normally used for custom interfacing and streamlining calculations in real-time systems. The beat notes are digitized by ADCs at a sampling rate in the 50 to 80 MHz range, depending on the actual implementation, the data bus of which leads directly to an FPGA. The architecture of a PM channel inside the FPGA is outlined in Fig. 2. A digital phase-lock loop forces a number controlled oscillator (NCO), which consists of a phase accumulator and a lookup table (LUT) for the sinusoid generation, to run at a constant phase offset with respect to the sampled beat note. Within the bandwidth of the feedback controller, all phase information is copied onto the NCO. Phase reconstruction for the interferometric measurement is performed by reading out the frequency value of the NCO, which is stored in the phase increment register (PIR)—the signal phase can be reconstructed through integration of the frequency-time series.

Because LISA is a deep-space mission, all locks need to be acquired automatically and managed autonomously. Therefore, the primary frequency actuation of all lasers will

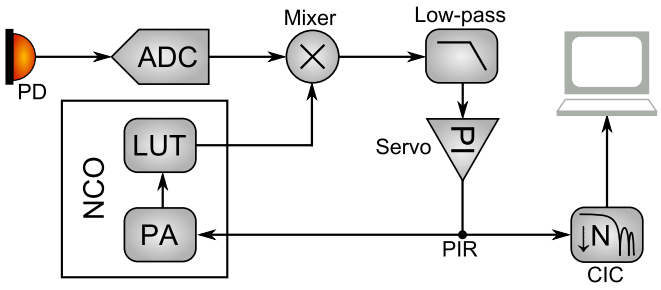

FIG. 2 (color online). Simplified PM channel layout. The digitized signal is mixed with the NCO, low-pass filtered, and refined by a feedback servo (PI). The PIR stores the value of the NCO frequency. By updating it, the PI controller forces the NCO to track the incoming signal. The PIR value is the primary readout point for the signal phase. Since the frequencies of interest at which variations in the signal occur lie far below the sampling frequency, the PIR information is heavily low-pass filtered and down sampled to frequencies of only several Hertz. Series of cascaded-integrator-comb (CIC) filters have been found to be most efficient for this task [18].

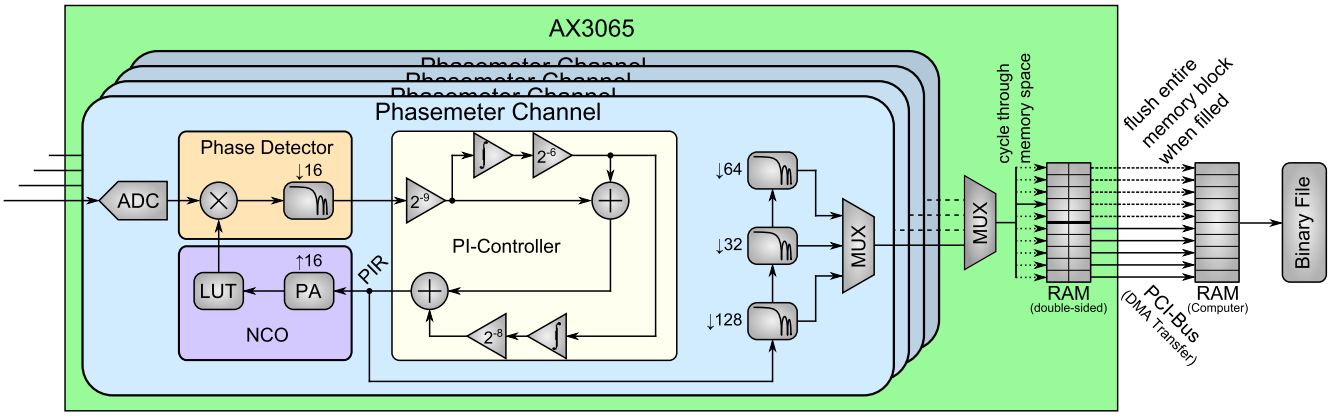

FIG. 3 (color online). Top-level view of the AX3065 PM architecture. The phase detector error signal is down sampled in a CIC filter by a factor of 16, such that the PI controller generates the frequency estimate (PIR) for the NCO at a rate of 4 MHz. The AX3065 can sample up to four beat notes simultaneously and extract their phase information. The acquired data are down sampled in three more steps of CIC filtering, first to 31.25 kHz, then 977 Hz, and finally 15.26 Hz. Data can be logged at any of these readout rates. We implemented the faster rates for diagnostic purposes and for possible application in ground-based interferometers. A multiplexer (MUX) cycles through the PM channels and writes their frequency estimates into the on-board RAM. The AX3065 features double-sided RAM, such that the PM channels can write their data uninterruptedly while the previously stored information is streamed to the host computer. The RAM address space is divided into two blocks. Once a block is completely filled, the AX3065 requests a direct-memory-access (DMA) transfer, takes control of the PCI bus, and flushes the entire memory block into a reserved memory region of the computer's RAM, from where it is logged to the disk in binary format. The flushed block is then cleared to be overwritten once needed.

be directly controlled by FPGA input. Because all beat notes are digitized and available to the PM, implementing HS to give the lasers the needed frequency stability only requires the addition of a cavity to the optical bench. No additional phase modulation is needed to extract the cavity interaction phase, and the feedback controls can be placed in the PM logic.

PM development at the University of Florida was initially focused on supplying a processing unit for the University of Florida LISA Interferometry Simulator (UFLIS) [19] with a PM that meets the mission μcycle requirement. UFLIS was used successfully to demonstrate experimentally the use of TDI [20,21], arm-locking [22,23], and GHz-sideband clock synchronization [24] on an optical test bench with real-time electronic phase delays of up to 32 sec, which is the round-trip time in a 5×10^6 km long interferometer arm.

We developed a second PM as a simplified version of the UFLIS PM specifically for the demonstration of HS and prototyping of advanced controller functionality. For the host system, we selected the Acromag module PMC-AX3065, which features a single Xilinx Virtex-II FPGA (XC2V3000), to which two dual channel 14-bit ADCs and a dual channel 16-bit DAC are connected. The AX3065 digitizes up to four beat notes at 64 MHz.[1] The propagation delay for signals from the analog-to-digital (A/D) conversion to D/A back conversion is 1.1 μs, which we verified with a timing measurement. This sets an upper limit for the bandwidth of the closed-loop operation of the HS lock of about 200 kHz, which is beyond the 30 kHz intrinsic bandwidth limit of the lasers we are using and therefore not a bottleneck for the design. Figure 3 shows a top-level schematic of the AX3065 PM architecture.

The PM controller suppresses the differential phase error between beat note and NCO by more than 300 dB in the measurement band, such that digitized signals with phase noise levels as high as 10^9 cycles/$\sqrt{\text{Hz}}$ can accurately be copied to the NCO to within the LISA μcycle requirement. However, the instrumental noise of the AX3065 does not meet this requirement, as can be seen in Fig. 4. Because of its scaling with the used signal frequencies across several test runs, we were able to trace the noise floor to timing jitter in the ADCs. This apparent limitation also appeared in the UFLIS PM and was remedied by electronically adding a common-source ultrastable pilot tone to all signals before their digitization. Additional PM instances can monitor this pilot tone and reveal discrepancies in the sample timing, allowing for small corrections to the signal phases.

Unfortunately, the lone FPGA of the AX3065 with its four PM channels is about 80% filled by the existing programming and almost at logic capacity, such that the addition of more channels to track a potential pilot tone is impossible. Nevertheless, the instrumental noise is about 8 orders of magnitude better than what is needed to observe laser frequency noise at the TDI requirement level. Therefore, we were still able to use the AX3065 for the demonstration of HS and the development of a digital control scheme and use it as a diagnostic tool to assess the

[1]In our setup, the ADCs and the FPGA are clocked by a Stanford Research Systems CG635 clock generator which is stabilized to a Stanford Research Systems FS725 rubidium frequency standard.

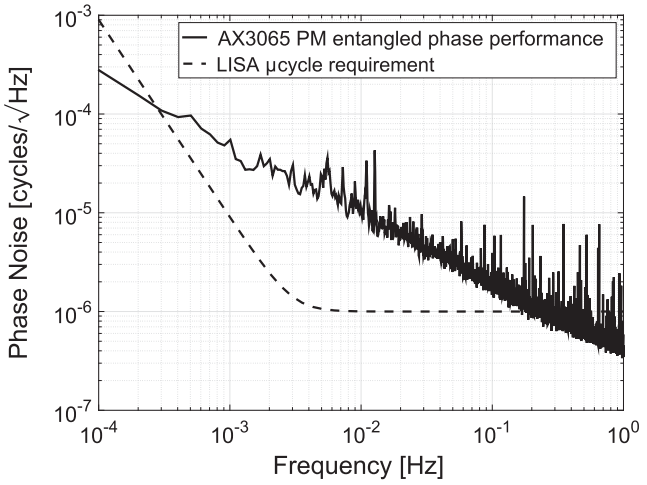

FIG. 4. AX3065 PM instrumental noise. The performance curve was assessed with an entangled phase measurement, in which a 19, a 14, and a 5 MHz signal with fixed phase relations were measured by three separate PM channels. Subtracting the phase information obtained from the 14 and the 5 MHz signals from that of the 19 MHz signal provides a null measurement that gives insight into the PM capabilities. The plotted solid curve is the linear spectral density of the null time series. The LISA μcycle requirement has been added as a dashed line.

stability of the lasers. The controller design is not specific to the AX3065 and can be copied to other PM implementations with ease.

IV. DIGITAL HETERODYNE STABILIZATION

In an all-digital locking scheme, the error signal (13) is generated in a similar fashion to the analog implementation. The cavity signal (10) and the reference beat (11) are both digitized to fractions of the ADC full-scale voltage $\pm V_{\mathrm{ref}}$ and then multiplied digitally. The conversion loss ℓ in (13) is replaced by the scaling of both signals with V_{ref}, and the now dimensionless error signal reads

$$e(\delta\nu) = -\frac{2\mathcal{G}\alpha^2 P_1 P_2}{V_{\mathrm{ref}}^2}\frac{\delta\nu}{\Delta\nu}. \tag{14}$$

The translation of analog low-pass filters and feedback controls to FPGA components is a straightforward process. The number generated by the proportional-integral (PI) controller is converted to a voltage by a DAC, which is connected to the laser frequency modulation port. To reduce the refresh rate of the controller output to a sampling frequency below the maximum DAC update rate of 900 kHz, the signal is down sampled from 64 MHz by a factor of 80 to 800 kHz.

In early demonstration steps, the phase of the reference beat note was adjusted via the placement of the reference PD. The correct demodulation phase $\Delta\phi_Y(z) - \Delta\phi_X(z) = (k_2 - k_1)\Delta z = -\pi/2$ can only be achieved by introducing an intentional distance offset $\Delta z \neq 0$ between the sampling

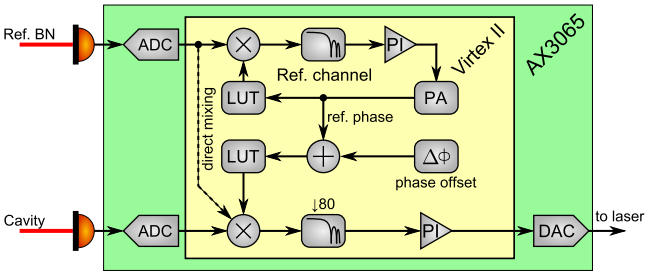

FIG. 5 (color online). Digital heterodyne controller. A high-bandwidth phase-meter channel locks to the reference beat note, and the cavity signal is mixed with an NCO copy of the reference beat. The phase offset $\Delta\phi$ can be adjusted to get the right demodulation phase. The down-converted signal is low-pass filtered, down sampled to 800 kHz, and passed to the DAC, which is connected to the laser frequency modulation port.

points of the cavity signal and the reference beat note. Unfortunately, since

$$k_2 - k_1 = \frac{2\pi}{c}(\nu_2 - \nu_1), \tag{15}$$

the demodulation phase is then subject to change if the beat frequency changes. To avoid this, the secondary laser was offset phase locked to the stabilized primary laser such that any frequency fluctuations are common mode to both lasers and do not change the demodulation phase.

With the high-bandwidth tracking AX3065 PM, a different solution to achieve the correct demodulation phase presents itself, which is illustrated in Fig. 5. The NCO of a PM channel is a high-precision copy of the beat note it is locked on, such that the HS demodulation can be performed against a secondary NCO instantiation of the reference beat, rather than the digitized reference beat itself. Since the NCO copy is created from addressing values that are stored in a LUT, its phase can be shifted by adding a phase offset before the lookup process for the sinusoid generation.

This makes $\Delta z = 0$ possible, in which case the demodulation phase becomes entirely independent of the beat frequency. For some residual static $\Delta z \neq 0$ on the optical bench, changes of $\delta\omega$ of the beat frequency can still be accounted for by adjusting the phase offset by $\delta\phi = \Delta z/c \times \delta\omega$.

Differential phase noise between the reference and the cavity beat that does not originate in the laser-cavity interaction will negatively affect the stabilization scheme. Additionally, since the reference phase for the demodulation of the error signal is that of the NCO, which is phase locked to the reference beat, residual phase noise in the NCO phase lock loop will similarly couple into the feedback loop. If we assume that all nonintended phase deviations are integrated into a phase variation $\varphi_{\mathrm{R}}(t)$ of the NCO signal $Y_{\mathrm{NCO}}(t)$ which is phase locked to the reference beat from (11),

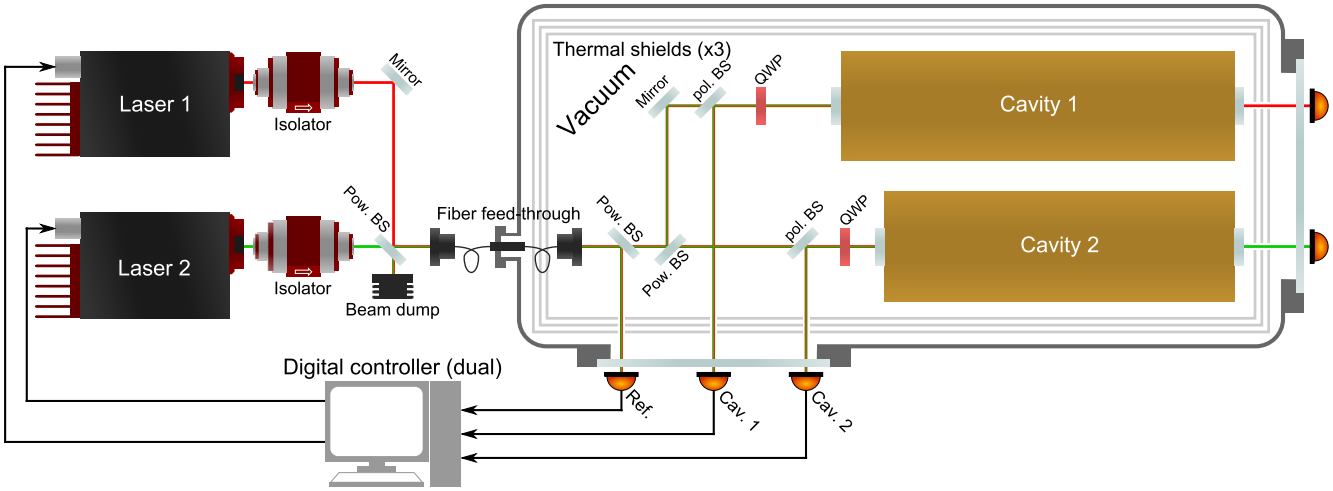

FIG. 6 (color online). Dual digital heterodyne stabilization setup. Two lasers are superimposed and sent into vacuum through a single fiber feed through. The superposition is split into a reference path and two paths that bounce off the optical cavities. All three instances of the superposition leave the vacuum through an optical window and generate beat notes. The dual digital HS controller samples them and actuates the laser frequencies, locking each laser to its respective cavity.

$$Y_{\mathrm{NCO}}(t) = \sin\left(\Delta\omega t + \varphi_{\mathrm{R}}(t)\right), \tag{16}$$

the demodulated signal in (14) receives an additional contribution proportional to $\Re\{F_r(\delta\nu)\} \times \varphi_{\mathrm{R}}(t)$. This can be seen as a time variable offset for the error signal, which changes the locking point of the feedback loop, since

$$\frac{2\alpha^2 P_1 P_2}{V_{\mathrm{ref}}^2}\left[-\mathcal{G}\frac{\delta\nu}{\Delta\nu} + \frac{r_1 - r_2}{1 - r_1 r_2}\varphi_{\mathrm{R}}(t)\right] \xrightarrow{\text{driven to}} 0. \tag{17}$$

This expression is equivalent with modulating the laser frequency to not only follow the cavity resonance but also track $\varphi_{\mathrm{R}}(t)$ as in

$$\nu_{\mathrm{Laser}} \xrightarrow{\text{driven to}} \nu_{\mathrm{Cavity}} + \frac{\Delta\nu}{\mathcal{G}}\frac{r_1 - r_2}{1 - r_1 r_2}\varphi_{\mathrm{R}}(t). \tag{18}$$

This residual modulation of the laser frequency can be reduced by increasing the optical gain $\mathcal{G}$ or impedance matching of the cavity, such that there is vanishing signal contribution in the quadrature $\Re\{F_r(\delta\nu)\}$. Figure 4 showed the maximum absolute phase noise we can expect to see in a single PM channel of the AX 3065. At 3 mHz, the phase noise is 30 μcycles/$\sqrt{\mathrm{Hz}}$, and therefore it is imperative to assert

$$\left|\frac{\Delta\nu}{\mathcal{G}}\frac{r_1 - r_2}{1 - r_1 r_2}\right| < 1.6\ \frac{\mathrm{Hz}}{\mu\mathrm{rad}} \tag{19}$$

when choosing cavity parameters to achieve the TDI frequency noise requirement of 300 Hz/$\sqrt{\mathrm{Hz}}$.

V. RESULTS

To demonstrate that the HS method is able to give the lasers in the LISA mission the frequency stability required by TDI, we stabilized two lasers with HS and determined their differential frequency fluctuations by measuring the frequency noise of their beat note. Additionally, we compared their stability against an independent PDH-stabilized cavity-laser pair that was integrated into UFLIS. The dual HS setup that we implemented is shown in a simplified manner in Fig. 6.

A superposition of two Coherent Mephisto S series Nd:YAG NPRO lasers is sent to the test bench, where it is split into three separate heterodyne fields. A reference path is picked off, and the two other fields reflect off their respective cavities. All three beams leave the vacuum through an optical window, behind which their beat notes are recorded.

The AX3065 digitizes the heterodyne signals and dedicates a PM channel to the reference beat. The two cavity signals are demodulated against phase-shifted copies of the NCO in the reference channel by a dual version of the controller presented in Fig. 5. PI controllers apply feedback gain and pass their output to the DACs, the voltages of which are applied directly to their respective laser's fast modulation input, closing both HS loops.

The two cavities have lengths 24 and 27 cm, with free spectral ranges of 625 and 555 MHz, respectively. The FSR mismatch of 70 MHz guarantees fundamental resonances in both cavities that are no farther apart than 35 MHz. In the experiment, we were always able to find pairs that were separated by less than 20 MHz.

The spacers are cylindrical with a diameter of 7.5 cm to keep spacer thermal noise low and have a 1 cm clearance

hole along their axis. They are made from ClearCeram-Z (CCZ) HS, a low-expansion ceramic produced by Ohara, which has been engineered to have a local maximum between two zero crossings of its CTE near room temperature, such that $|\alpha| < 0.2 \times 10^{-7}\ \mathrm{K}^{-1}$ is guaranteed between 0° and 50 °C. Besides Corning's ULE and Schott's Zerodur, ClearCeramZ is a low-CTE candidate material that we are investigating for its potential application in low-frequency optical interferometry.

The cavities were placed in vacuum within three layers of thermal shielding to mitigate the influence of ambient temperature fluctuations. Each of the three cascaded layers consisted of a cubical aluminum frame with aluminized Mylar spanned across the cube faces.

The mirrors—fused silica substrates with a reflective dielectric coating of alternating layers of Ta_5O_2 and SiO_2 polished and coated by Coastline Optics—are optically contacted to the end faces of the spacers. The cavities are half-symmetric with flat front mirrors and 1 m radius of curvature back mirrors. The nominal transmission losses are $T_1 = 280$ ppm for the front and $T_2 = 10$ ppm for the back mirrors as measured by Coastline. The expected losses due to scatter and absorption are in the 10 ppm range, yielding a finesse value for the cavities of 20,000.

The reason for overcoupling the cavities was our intention to simultaneously track the phase fluctuations of the cavity signals with dedicated phase-meter channels. This can only be realized with a beat signal of significant amplitude, which prohibits the use of an impedance matched cavity. We find that

$$\left|\frac{\Delta\nu_{1,2}}{\mathcal{G}}\frac{r_1 - r_2}{1 - r_1 r_2}\right| \approx 7.2\ \frac{\mathrm{mHz}}{\mu\mathrm{rad}} \tag{20}$$

for the given spacer and mirror specifications, which provides sufficient reference channel phase noise suppression according to (19) in either cavity.

An issue that arises is the masking of reference channel phase noise induced frequency fluctuations in a differential measurement. Since both lasers receive their error signal from the demodulation of their cavity signal against the same reference NCO, phase noise $\varphi_R(t)$ as introduced in (16) will drive the frequency of both laser 1 and laser 2 according to (18). This will affect their absolute stability, but in a measurement of their differential frequency noise, there is a strong common-mode rejection of this phase noise driven frequency modulation. The difference of the impact factors of the two cavities is on the order of

$$\left|\frac{1}{\mathcal{G}}\frac{r_1 - r_2}{1 - r_1 r_2}\right| \times |\Delta\nu_1 - \Delta\nu_2| \approx 0.8\ \frac{\mathrm{mHz}}{\mu\mathrm{rad}}, \tag{21}$$

and therefore the suppression ratio can be as high as a factor of 10. This effect would need to be considered if frequency noise levels below 1 $\mathrm{Hz}/\sqrt{\mathrm{Hz}}$ are observed in our setup.

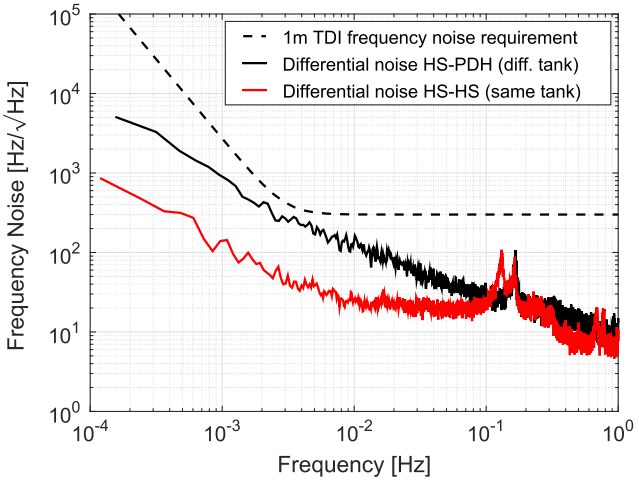

FIG. 7 (color online). Frequency stability of the HS stabilized lasers. Shown is the frequency noise of the beat between the two HS test systems, as well as the beat between the 27 cm cavity laser and an independently PDH-stabilized UFLIS reference laser. The LISA TDI requirement (1) was added as a dashed line. The peaklike features between 0.1 and 0.2 Hz emerge at the crossover between piezo- and laser crystal temperature feedback due to the limited resolution of the D/A conversion in the temperature feedback.

The two laser fields were sent into the vacuum chamber via a polarization maintaining single-mode fiber feed through made by Schäfter + Kirchhoff GmbH. A future plan is to suspend the entire test bench to isolate it against building vibrations and use it also to measure coating thermal noise for future ground-based GW detectors [25]. The fibers prevent any alignment noise of the lasers into the cavities as a result of differential motion of the test bench against the tank outside. Because of the mode-cleaning properties that fiber injection has for the individual laser beams, this method also forces the two lasers into the same spatial mode for use in the interferometer, maximizing the contrast on all photodetectors.

Figure 7 displays the differential frequency stability we achieved with the two HS stabilized lasers. Shown is the frequency noise linear spectral density of the beat note between the lasers in the reference channel. The time series from which this graph was generated was taken over the course of a 61 h period, during which the lasers stayed locked continuously. The locking scheme proved to be robust enough to not lose lock for weeks at a time.

The previously described common-mode rejection of residual frequency modulation is of no concern at the shown level of differential frequency noise. However, the two cavities are located in the same thermally insulated environment, such that they are subject to largely identical temperature variations. These couple to the absolute length stability of the cavities through the spacer material's CTE but show reduced impact in a differential measurement because of the similar dimensions of the cavities. Since their lengths differ by about 10% of their total length, one

can expect a suppression of roughly a factor of 10 in differential noise compared to absolute noise in regions where temperature noise is the dominating contribution.

To confirm the absolute stability of our system, we recorded the beat note between the 27 cm cavity HS laser and a third, PDH stabilized laser that originated on the UFLIS optical bench and was transferred to the HS setup with an optical fiber. The frequency noise, which was then measured between two cavities that were located in separate vacuum tanks within their own respective thermal shielding, is also shown in Fig. 7. There is no more thermal correlation in this measurement, and additionally it is no longer subject to the common-mode rejection of reference channel phase noise induced frequency modulation.

The frequency stability requirement (1) has been added to the graph, showing that we achieved an absolute noise level better than the TDI 1 m requirement. At frequencies below 10 mHz, we see indeed a separation of about a factor of 10 between the two measurements, which indicates that residual temperature fluctuations inside the thermal shielding are the limiting noise source. Above 10 mHz, the HS-HS noise begins to more closely resemble the HS-PDH curve. We determined that laser power fluctuation driven temperature noise due to absorption in the coating becomes the limiting factor in this region.

VI. IMPLEMENTATION IN LISA

To reach the LISA design sensitivity, all beat note phases are recorded with μcycle precision. TDI then applies linear combinations of appropriately time-shifted phase data streams to cancel out laser frequency noise. The uncertainty in the absolute arm length, which translates to an uncertainty in sample timing, determines the level of laser frequency noise that TDI is able to remove from the data. Which particular frequency stabilization method is implemented into the interferometric scheme to meet the frequency noise requirement (1) does not affect the effectiveness of TDI and how it unfolds its potential.

The LISA mission interferometric baseline design is to prestabilize a single master laser that feeds into the interferometer. Its frequency stability is transferred to other lasers through offset phase locking. The sufficiently accurate phase recovery of an optical carrier at laser powers of 100 pW, which is a lower bound of the expected power to be captured by the beam telescopes in LISA, has been demonstrated experimentally [26]. Therefore, even lasers on a different spacecraft can be stabilized using only the weak-field beat with the master laser. The lasers that do not directly interfere with the master laser can still inherit the inferred frequency stability from their beat with other secondary lasers.

The three interferometer arms in LISA are formed between pairs of drag-free inertial proof masses. Each spacecraft houses a pair of proof masses for its two

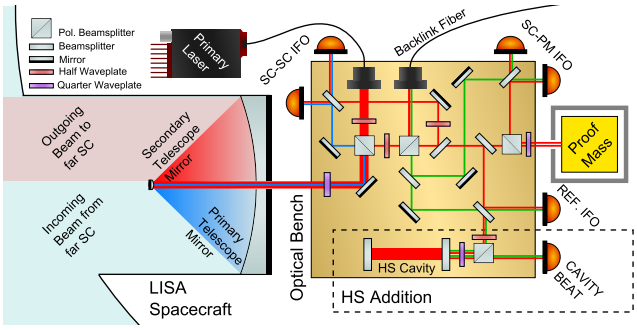

FIG. 8 (color online). Simplified, conceptually functional representation of the local interferometry on the OB of a LISA spacecraft (SC). Most of the primary laser power is sent out to the far SC, and a smaller fraction is sent to the other OB through the bidirectional backlink fiber, which also carries the secondary local laser to this OB. The received light from the far SC is beat with the primary laser in the SC-SC IFO to observe the inter-SC distance variation. The primary laser also reflects off the proof mass and is then interfered with the second laser for a beat with phase information about the proof mass position. The reference IFO provides the reference phase for the two heterodyne measurements. A possibility for the realization of HS is to use the superposition at its other port and add a cavity to the OB.

associated interferometer links, each with their own dedicated laser unit. The interferometric distance between pairs of proof masses on different spacecraft is split into three stages. The vast majority of the arm length is spanned between two monolithic, ultralow CTE optical benches (OBs), located as transponder units between the beam telescopes and the proof masses. An exemplary layout for the OB design in LISA is displayed in Fig. 8. The small fraction of the incoming light of the far laser that is captured by the primary telescope mirror is interfered with a strong local beam to recover a beat note of sufficient SNR for a phase measurement with μcycle precision. The remaining distance from the OB to the proof mass is monitored using the beat of the primary OB laser with the second laser coming from the other OB. Laser light is exchanged between same-spacecraft OBs via a bidirectional backlink fiber, which enables local heterodyne measurements. The fact that superpositions of two local lasers are readily available on every optical bench led to the development of HS.

The dashed box in Fig. 8 suggests a possible way to integrate HS into a given OB design. The beam splitter that creates the superposition of the two local lasers in the reference interferometer (IFO) has two output ports—only one of which is needed in the interferometric scheme. The other output is available for laser frequency stabilization using the HS method as is. A cavity along with some optical components for polarization control needs to be added to the OB, but there is no additional signal to be processed by the PM, as the beat would be recorded with or without the implementation of HS. The

reference beat for the cavity signal demodulation can be supplied by the reference IFO.

Care needs to be taken that the secondary laser does not resonate on a higher-order mode in the cavity, especially when tuning the phase-lock frequency to adjust for the interspacecraft Doppler shifts. This problem could be prevented if the master laser were reflected off the cavity alone, and only afterward interfered with the second laser. Since all optics are located on the monolithic OB, the differential length noise between reference and cavity path, which in this case scale as $k_{1,2} = 2\pi/\lambda_{1,2}$ rather than $(k_1 - k_2)$ as in (15), would be sufficiently low for the separate propagation of the two lasers. However, a more sophisticated modification of the OB design would be necessary to implement this version of HS.

VII. CONCLUSION

We have developed and demonstrated experimentally Heterodyne Stabilization, a novel laser frequency stabilization method. HS finds application in long-baseline interferometric space missions such as LISA, where the spread of Gaussian beam propagation, coupled with limited mirror size, requires heterodyne detection schemes.

HS uses a low-CTE optical cavity as a frequency reference and transfers its stability to the wavelength of a laser that is kept on resonance with it. It does not require dedicated phase modulation to extract the cavity interaction phase; instead, it uses the already existing beat note with a second laser as a phase reference.

In an all-digital locking scheme, which we implemented using LISA PM technology, we stabilized two laser systems to a frequency noise level below the requirement of the LISA mission to employ TDI. We confirmed the frequency stability in comparison with an independent reference and showed that HS is able to provide the necessary laser frequency stability on board the LISA spacecraft.

A secondary result of our investigation is that Clearceram-Z offered sufficient dimensional stability when subjected to the residual temperature fluctuations in our experiment. It may thus be a potential alternative to other low-expansion materials.

ACKNOWLEDGMENTS

This work was supported by NASA Contract No. 00078244 and by NSF Grants No. PHY-0969935, No. PHY-1205512, and No. PHY 1306594.

[1] B. F. Schutz, Classical Quantum Gravity **28**, 125023 (2011).
[2] D. A. Shaddock, Classical Quantum Gravity **25**, 114012 (2008).
[3] A. Stroeer and A. Vecchio, Classical Quantum Gravity **23**, S809 (2006).
[4] O. Jennrich, LISA Mission Science Office, *LISA: Assessment Study Report (Yellow Book)* (European Space Agency, Noordwijk, 2011).
[5] M. Tinto and S. V. Dhurandhar, Living Rev. Relativity **17**, 6 (2014).
[6] D. A. Shaddock, M. Tinto, F. B. Estabrook, and J. W. Armstrong, Phys. Rev. D **68**, 061303 (2003).
[7] The LISA Frequency Control Team, Report No. ESA-LISA-JPL-TN-823, 2009.
[8] Corning, ULE Corning Code 7972 Ultra Low Expansion Glass (2015), http://www.corning.com/docs/specialtymaterials/pisheets/ulebro91106.pdf; Accessed: 14-January-2015.
[9] Schott, ZERODUR Zero Expansion Glass Ceramic (2015), http://www.us.schott.com/advanced_optics/english/syn/advanced_optics/products/optical-materials/zerodur-extremely-low-expansion-glass-ceramic/zerodur/index.html; Accessed: 14-January-2015.
[10] Ohara, ULTRA-LOW EXPANSION GLASS-CERAMICS CLEARCERAM-Z (2015), http://www.oharacorp.com/pdf/ccz-2013-nav.pdf; Accessed: 14-January-2015.
[11] R. Drever, J. Hall, F. Kowalski, J. Hough, G. Ford, A. Munley, and H. Ward, Appl. Phys. B **31**, 97 (1983).
[12] G. Heinzel, J. J. Esteban, S. Barke, M. Otto, Y. Wang, A. F. Garcia, and K. Danzmann, Classical Quantum Gravity **28**, 094008 (2011).
[13] D. Shaddock, B. Ware, P. G. Halverson, R. E. Spero, and B. Klipstein, AIP Conf. Proc. **873**, 654 (2006).
[14] P. W. McNamara, H. Ward, J. Hough, and D. Robertson, Classical Quantum Gravity **14**, 1543 (1997).
[15] R. Thompson, W. Folkner, G. de Vine, W. Klipstein, K. McKenzie, R. Spero, N. Yu, M. Stephens, J. Leitch, R. Pierce, T.-Y. Lam, and D. Shaddock, *Frequency Control and the European Frequency and Time Forum (FCS), 2011 Joint Conference of the IEEE International* (IEEE, Piscataway, 2011) p. 1.
[16] O. Jennrich, Classical Quantum Gravity **26**, 153001 (2009).
[17] I. Bykov, J. J. E. Delgado, A. F. G. Marn, G. Heinzel, and K. Danzmann, J. Phys. Conf. Ser. **154**, 012017 (2009).
[18] O. Gerberding, B. Sheard, I. Bykov, J. Kullmann, J. J. E. Delgado, K. Danzmann, and G. Heinzel, Classical Quantum Gravity **30**, 235029 (2013).
[19] R. J. Cruz, J. I. Thorpe, A. Preston, R. Delgadillo, M. Hartman, S. Mitryk, A. Worley, G. Boothe, S. R. Guntaka, S. Klimenko, D. B. Tanner, and G. Mueller, Classical Quantum Gravity **23**, S751 (2006).
[20] S. J. Mitryk, V. Wand, and G. Mueller, Classical Quantum Gravity **27**, 084012 (2010).

[21] S. J. Mitryk, G. Mueller, and J. Sanjuan, Phys. Rev. D **86**, 122006 (2012).
[22] Y. Yu, S. J. Mitryk, and G. Mueller, Classical Quantum Gravity **28**, 094009 (2011).
[23] Y. Yu, S. Mitryk, and G. Mueller, Phys. Rev. D **90**, 062005 (2014).
[24] D. Sweeney and G. Mueller, Opt. Express **20**, 25603 (2012).
[25] M. Hartman, J. Eichholz, P. Fulda, G. Ciani, D. B. Tanner, and G. Mueller, Measurement of Thermal Noise in Optical Coatings for Gravitational-Wave Detectors (2014), APS April Meeting 2014 Savannah, GA.
[26] J. J. Esteban, A. F. García, S. Barke, A. M. Peinado, F. G. Cervantes, I. Bykov, G. Heinzel, and K. Danzmann, Opt. Express **19**, 15937 (2011).