FERMILAB-THESIS-2016-20

THE UNIVERSITY OF CHICAGO

EXPERIMENTAL CONSTRAINTS ON THE EXOTIC SHEARING OF SPACE-TIME

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF ASTRONOMY & ASTROPHYSICS

BY
JONATHAN WILLIAM RICHARDSON

CHICAGO, ILLINOIS
AUGUST 2016

Operated by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the United States Department of Energy.

Copyright © 2016 by Jonathan William Richardson

All Rights Reserved

To Zrria, the light of my life.

"We shall not cease from exploration

And the end of all our exploring

Will be to arrive where we started

And know the place for the first time."

T. S. Eliot

Table of Contents

LIST OF FIGURES . . . vii

ACKNOWLEDGMENTS . . . ix

ABSTRACT . . . xi

1 INTRODUCTION . . . 1

2 EXOTIC QUANTUM POSITION NOISE . . . 4
- 2.1 Transverse Position Uncertainty . . . 4
- 2.2 Gauss-Markov Representation . . . 5
  - 2.2.1 Stochastic Spatial Fluctuations . . . 6
  - 2.2.2 Generalizing to Continuous Coordinates . . . 7
  - 2.2.3 Spatial Fluctuation Rate . . . 9
- 2.3 Interferometric Detection . . . 9
  - 2.3.1 The Observable . . . 10
  - 2.3.2 Effect on the Observable . . . 12
  - 2.3.3 Spatial Entanglement . . . 17
  - 2.3.4 Single-Interferometer Statistics . . . 21
  - 2.3.5 Dual-Interferometer Statistics . . . 23
- 2.4 Instrument Response . . . 24
  - 2.4.1 Optical Geometry . . . 25
  - 2.4.2 Coupling to Shear Effect . . . 27
  - 2.4.3 Coupling to Rotational Effect . . . 27
  - 2.4.4 Auto-Response Functions . . . 28
  - 2.4.5 Cross-Response Functions . . . 30
- 2.5 Previous Experimental Limits . . . 32

3 THE HOLOMETER DETECTOR . . . 36
- 3.1 Physical Environment . . . 36
- 3.2 Optical System . . . 37
  - 3.2.1 Power-Recycling Cavity . . . 39
  - 3.2.2 Antisymmetric-Port Response . . . 40
  - 3.2.3 Optical Elements . . . 41
- 3.3 Vacuum System . . . 48
- 3.4 Interferometer Control System . . . 52
  - 3.4.1 Piezoelectric Drive System . . . 54
  - 3.4.2 Low-Speed Digitization Electronics . . . 55
  - 3.4.3 Controls Software . . . 56
  - 3.4.4 Feedback Control Loops . . . 57
- 3.5 Data Acquisition System . . . 62
  - 3.5.1 High-Speed, High-Power Photodetectors . . . 63

3.5.2 High-Speed Digitization Electronics 64
3.5.3 Data Pipeline Software 66

4 MEASUREMENT CALIBRATION 70
4.1 Overview 70
4.2 Indirect Calibration 71
4.3 Direct Low-Frequency Calibration 73
4.3.1 Calibration Line Injection and Monitoring 73
4.3.2 Reference to Optical Wavelength 74
4.4 Signal-Readout Transfer Functions 76
4.4.1 Measurement Below 900 kHz 78
4.4.2 Measurement Above 900 kHz 81
4.5 Systematics & Data Quality Control 84
4.5.1 Calibration Drift in the PZT Drives 84
4.5.2 Photodetector Dependence on Power 85
4.5.3 ADC Timing Instability 94

5 ANALYSIS & RESULTS 101
5.1 Sensitivity Enhancement by Averaging 101
5.2 Analysis of Optics Thermal Noise 103
5.2.1 Spectral Model of Thermal Excitations 105
5.2.2 Calibration Using the Optics Thermal Noise 107
5.3 Interferometer Signal Recombination 108
5.4 Spectral Leakage Estimation 112
5.5 Constraints on Exotic Spatial Shear 116
5.6 Physical Implications 120

A HOLOMETER PHASE II DESIGN 122
A.1 Optical Geometry 122
A.2 Response to Rotational Effect 125

REFERENCES 127

List of Figures

2.1 Illustration of shear versus rotational spatial fluctuations. 16
2.2 Illustration of the entanglement of rotational spatial fluctuations measured relative to multiple light origins. 18
2.3 Illustration of the entanglement of shear spatial fluctuations measured relative to multiple light origins. 20
2.4 Geometrical schematic of the optical paths through each Holometer interferometer. 26
2.5 Time- and frequency-domain response functions of each Holometer interferometer to exotic spatial shear noise. 29
2.6 Correlated time- and frequency-domain response functions of the Holometer interferometers to exotic spatial shear noise. 31
2.7 GEO600 upper limit on flat unexplained strain noise. 35

3.1 The Holometer site at Fermi National Accelerator Laboratory. 37
3.2 Climate-controlled enclosure housing the end stations of the above-ground interferometer arms. 38
3.3 Transfer function of the power-recycled interferometers, shown at several OPD operating offsets. 42
3.4 Laser launch table for preparation of the injection beam. 43
3.5 Sectioned view of the end mirror vibration isolation mount and PZT actuator. . 46
3.6 Illustration of the dual vacuum systems housing the high-power interferometer optics. 49
3.7 Installation of fiberglass pipe insulation on the interferometer arms. 50
3.8 Central vacuum vessel housing the beamsplitter and PRM. 51
3.9 End station vacuum assembly with top and side flanges removed. 53
3.10 PZT-actuated end mirror drive. 54
3.11 Transfer function of the PZT drive of one end mirror. 55
3.12 Common-mode loop error signals as a function of detuning. 58
3.13 Suppression of the seismic noise background by the interferometer control system. 61
3.14 Schematic of the modified New Focus 1811 photodetector circuit for high-power (100 mW) operation. 65
3.15 Schematic of the high-speed data acquisition system. 67

4.1 Time history of the continuous 1-kHz direct calibrations. 75
4.2 PZT-drive calibrations for the L interferometer over the six-month data collection period. 76
4.3 PZT-drive calibrations for the T interferometer over the six-month data collection period. 77
4.4 Transfer functions of the signal-readout chains measured via lock-in drive detection at low frequency. 80
4.5 Gains of the AC-to-AC transfer functions measured via optical shot noise. . . . 82
4.6 Phase differences between the AC-to-AC transfer functions measured via lock-in drive detection. 84

4.7 DC-channel response of a modified New Focus 1811 photodetector at different levels of DC illumination. 87
4.8 AC-channel response of a modified New Focus 1811 photodetector at different levels of DC illumination. 88
4.9 Transimpedance gain of the AC-coupled channel of a modified New Focus 1811 photodetector at different levels of DC illumination. 89
4.10 Time history of the optical power incident on the modified New Focus 1811 photodetectors. 92
4.11 Time history of the photodetector transimpedance ratios, a proxy for response stability. 95
4.12 Sampling time drift between two GPS-disciplined ADC units measured by coherent broadband noise injection. 97
4.13 Decoherence of time-averaged CSD measurements due to clocking noise between the ADC units. 98
4.14 Time history of the sampling synchronicity between the interferometer readout channels. 100

5.1 CSD matrix measured between the AS-port detector readout channels. 104
5.2 Thermal noise of the optics measured via the single-interferometer CSDs. 106
5.3 CSD matrix measured between the AS-port interferometer signals. 113
5.4 Coherence of the cross-interferometer CSD measurement. 114
5.5 Cross-interferometer CSD measurements rebinned to 250 kHz resolution. 116
5.6 Log-likelihood ratio of transverse position commutator normalizations relative to the zero-normalization model. 119

A.1 Rendering of the Phase II reconfiguration of the Holometer interferometers. . . . 123
A.2 The time- and frequency-domain response functions of each Holometer interferometer to exotic spatial rotational noise under the Phase II reconfiguration. . . 126

ACKNOWLEDGMENTS

I am indebted to countless individuals whose support, mentoring, and encouragement over the years has made this dissertation possible. It has been a true privilege to work alongside so many talented physicists in our exploration of the natural world. Your commitment to disciplined, evidence-based inquiry has fundamentally altered my way of thinking, and from you I draw continual inspiration.

I first thank my parents, William and Sandra Richardson, for their unfailing love and support from my earliest memories to the present. At every step of my educational career, you have provided me the means and encouragement to continue on to the next level. For that I will always be grateful, and I will one day seek to instill the same love of learning in my own children.

Back at Yale, I thank my first physics professor, Paul Tipton, for recognizing scientific potential in me, long before I recognized it in myself, and for introducing me to my undergraduate research advisor, Daisuke Nagai. I thank Daisuke for his three years of dedicated mentoring and for advocating fiercely on my behalf to place me at my dream graduate school. I also thank Suchetana Chatterjee and Zheng Zheng for their patient mentoring during my first research projects.

At the University of Chicago, I am grateful to the many people who helped me successfully transition from undergraduate to graduate life. In particular, I thank Don York and Rocky Kolb for funding my research expenses and conference travel during my first year, before I had found a graduate advisor. I also thank Angela Olinto for making available department funding to subsidize my frequent, long commutes to the experimental site at Fermilab.

To my Holometer collaborators, it has been a pleasure to work with you all. I am especially grateful to Steve Meyer, my experimental advisor, and Sharon Salveter for all of their support, both personally and professionally, during my graduate years. I am equally grateful to my theory advisor, Craig Hogan, for all of his insight and for originally recruiting

me into the project. Both of you have profoundly shaped how I think and see the world, and you are models of the scientific career I aspire to one day achieve. My own work stands on the shoulders of Bobby Lanza and Lee McCuller, the previous Holometer graduate students, from each of whom I have learned much.

Finally, I thank my soon-to-be wife, Zrria, for her incredible support—and patience with many late nights—as I prepared this work. You are truly the light at the center of my life, a constant, loving source of stability throughout my often tumultuous graduate career. To Zrria, I dedicate this dissertation.

ABSTRACT

The Holometer program is a search for first experimental evidence that space-time has quantum structure. The detector consists of a pair of co-located 40-m power-recycled interferometers whose outputs are read out synchronously at 50 MHz, achieving sensitivity to spatially-correlated fluctuations in differential position on time scales shorter than the light-crossing time of the instruments. Unlike gravitational wave interferometers, which time-resolve transient geometrical disturbances in the spatial background, the Holometer is searching for a universal, stationary quantization noise of the background itself.

This dissertation presents the final results of the Holometer Phase I search, an experiment configured for sensitivity to exotic coherent shearing fluctuations of space-time. Measurements of high-frequency cross-spectra of the interferometer signals obtain sensitivity to spatially-correlated effects far exceeding any previous measurement, in a broad frequency band extending to 7.6 MHz, twice the inverse light-crossing time of the apparatus. This measurement is the statistical aggregation of 2.1 petabytes of 2-byte differential position measurements obtained over a month-long exposure time. At 3σ significance, it places an upper limit on the coherence scale of spatial shear two orders of magnitude below the Planck length.

The result demonstrates the viability of this novel spatially-correlated interferometric detection technique to reach unprecedented sensitivity to coherent deviations of space-time from classicality, opening the door for direct experimental tests of theories of relational quantum gravity.

CHAPTER 1

INTRODUCTION

Quantum principles are widely accepted to govern all physical systems. Like matter and energy, space-time is expected to exhibit quantum behavior. However, despite many decades of theoretical progress, there is no generally accepted theory of quantum geometry, the dynamical quantum system whose classical limit is general relativity[21, 23, 1]. Moreover, no experimental technique has emerged to study the exotic quantum nature of space-time itself, as distinguished from the quantum behavior of fields that inhabit a classical space-time.

The content of any quantum theory can be expressed in terms of correlations between observable quantities. In a space-time of quantum origin, positional relationships are expected to exhibit departures from classicality in the form of non-local correlations. The natural scale for these correlations is the Planck length, $l_p \equiv \sqrt{\hbar G/c^3} = 1.616 \times 10^{-35}$ m, where the gravitational self-energy of a single quantum approaches its total energy. This quantum behavior is widely thought to be undetectable, since the standard theory of quantum fields propagating in a classical background space-time[24, 25, 14], as well as its UV completions, such as string theory[19, 22, 5, 11], predicts that exotic correlations from Planck-length quantum gravity are negligible on scales accessible to experiment[12].

Another school of thought[21, 23, 1] holds that space-time is "relational" in nature. In this view, there is no fixed absolute background. Rather, all locations in space and time are defined from within the geometrical quantum system, relative to other locations in the system. A quantum system composed of many Planck-scale elements could display a large-scale behavior that closely resembles classical space-time, but cannot be described by quantum field states on a classical background. In principle, a relational quantum system can introduce much larger exotic spatial correlations on laboratory scales than standard quantum field theory, even with a Planck coherence length. A relational quantum geometry approaches the classical limit more gradually on large scales, like a Planck-scale random

walk, with covariances constrained by the symmetries of the system.

Although there is an extensive literature on such relational theories, there is no generally accepted framework for how the Planck-scale elements relate to each other on large scales, as this depends on the detailed spatial structure of their quantum entanglement. This entanglement determines the statistical covariance between observable operators. Experimental results already significantly constrain some forms of exotic large-scale correlations. Observations of distant astronomical objects (e.g., quasars, AGN, gamma-ray bursts) now constrain exotic correlations with longitudinal, or line-of-sight, symmetry to a coherence scale below the Planck length [13, 15, 20, 4, 3, 17, 18]. Under such an effect, light emitted by a distant source accumulates random phase deviations across the wavefront while propagating between source and observer. The optical phase decoherence would be detectable as a distance-dependent degradation of the intensity and/or resolution of the source image.

More recently, exotic correlations with symmetry *transverse* to the line-of-sight separation have been proposed[8, 9, 10, 12]. In principle, such an effect is detectable via laser interferometry, a technique which compares the optical phase simultaneously accumulated along two orthogonal light paths. Exotic correlations of this kind have not been previously tested to Planck-scale sensitivity, as existing gravitational wave interferometers are optimized for lower-frequency detection (corresponding to time scales much longer than the light-crossing time of the apparatus). High-fidelity constraints on time and position correlations in an extended volume of space-time require measurements sampled faster than a light-crossing time in more than one spatial direction. The Fermilab Holometer is the first experiment specifically designed to test this class of relational quantum gravity models. It consists of two co-located but independent interferometers employing a dual-detection strategy to isolate spatially-correlated effects.

This dissertation presents the final results of the Holometer Phase I experiment, an exhaustive search for exotic spatial shear correlations. The measurement will be shown to

constrain transverse shear correlations to a coherence scale two orders of magnitude below the Planck length. A novel path-integral-based formalism will also be introduced for modeling the effect of arbitrary spatial fluctuations on optical phase measurements. This treatment uses a bottom-up construction integrating the statistical effect of Planck-scale quantum elements, under an assumed large-scale covariance. The technique predicts the full functional form of the cross-correlation of coplanar interferometers with only one free parameter, the fundamental coherence scale of the exotic correlation. It thus presents a more complete, precisely formulated, and unified phenomenological framework for a comprehensive experimental program to explore relational quantum gravity. This analysis will be useful to design experiments and to interpret experimental results in terms of Planck-scale symmetries.

CHAPTER 2

EXOTIC QUANTUM POSITION NOISE

2.1 Transverse Position Uncertainty

Given some configuration of massive bodies, the mirrors in an interferometer, for example, a fundamental theory of quantum geometry should predict an exact form of the spatial correlations between the bodies in the system. Even in the absence of a fundamental theory, it has been argued by Hogan [8, 9, 10, 12] that the possible forms of spatial correlation are constrained by the known macroscopic structure of the system, without knowing the underlying degrees of freedom of the quantum space-time. Relying only on general principles of quantum mechanics and causality, Hogan proposes a simple, covariant model of position states of a massive body in flat space on large scales, excluding all standard quantum and gravitational degrees of freedom. The model is based on a quantum commutator of position analogous to the angular-momentum algebra,

$$[\hat{x}_1,\ \hat{x}_2] \quad = \quad i l_P R \ , \tag{2.1}$$

which relates the number of position states within a distance R to the transverse variance of a spatial wave function in physical length units. A non-zero value of the commutator normalization, l_P, imposes a limit on the density of position states, introducing deviations from classical space-time.

These deviations are shown to manifest as a new quantum geometrical uncertainty of the transverse position of two world lines separated by a distance R,

$$\langle X_\perp^2 \rangle_R = l_P R \ , \tag{2.2}$$

which dominates over the standard quantum position uncertainty for bodies much larger than

the Planck mass, $m_p = 2.176 \times 10^{-5}$ g. Based on heuristic arguments matching the number of spatial degrees of freedom to the limit implied by entropic gravity, Hogan estimates a normalization of $l_P \sim l_p$, where $l_p = 1.616 \times 10^{-35}$ m is the Planck length. Throughout, this work will regard the normalization as a free parameter to be constrained by experiment, with a nominal value of $l_P = l_p$ adopted for all numerical calculations.

For a configuration of multiple bodies, the commutator specifies the total number of position states available to the system, but not the degree of their entanglement. However, Hogan demonstrates that, in order to approximate locality on macroscopic scales, the position states of neighboring bodies separated by distances $\lesssim R$ must be highly entangled. This interpretation predicts spatially-coherent fluctuations in the transverse position of macroscopically-separated bodies on a timescale of R/c, with variance $\langle X_\perp^2 \rangle_R$. As will be shown, this interpretation leads to precisely predictable correlations in the position measurements made by co-located inteferometers. The detection of such a correlation would provide compelling first evidence of the quantum structure of space-time, despite the absence of a fundamental theory.

2.2 Gauss-Markov Representation

This section will demonstrate that the transverse position uncertainty relation proposed by Hogan, in fact, describes the integrated effect of a Gauss-Markov process at the Planck scale. Under this interpretation, the position uncertainty arises from a random-walk accumulation of spatial fluctuations over the macroscopic separation distance, R, where the coherence length, or "step size" of the walk, is the commutator normalization, l_P. A general statistical framework is developed for propagating such spatial fluctuations into the optical phase measurement of an interferometer.

2.2.1 Stochastic Spatial Fluctuations

Under a Gauss-Markov, or random walk, interpretation, the spatial separation of two world lines, R, is discretized into R/l_P individual steps. For a null-propagating wave crossing the spatial separation, each step corresponds to a different interval of time. Every coherence time, $t_P = l_P/c$, the classical spatial background is modeled to undergo a coherent fluctuation due to quantization noise. The transverse position of any step, j, is the transverse position of the previous step, $j-1$, plus a random, normally-distributed fluctuation

$$\delta X_\perp^{(j)} \;\sim\; \mathcal{N}\left(0, l_P^2\right) \,. \tag{2.3}$$

Since consecutive steps are separated by a full coherence length, the covariance between fluctuations is

$$\mathrm{cov}\left(\delta X_\perp^{(j)}, \delta X_\perp^{(k)}\right) \;=\; \delta_{jk}\, l_P^2 \,, \tag{2.4}$$

where δ_{jk} is the Kronecker delta. Over the first n Planckian steps, the relative transverse position of the two world lines is

$$\begin{aligned}
X_\perp^{(1)} &= X_\perp^{(0)} + \delta X_\perp^{(1)} &&= \delta X_\perp^{(1)} \\
X_\perp^{(2)} &= X_\perp^{(1)} + \delta X_\perp^{(2)} &&= \delta X_\perp^{(1)} + \delta X_\perp^{(2)} \\
\vdots & \qquad\quad \vdots && \quad\;\; \vdots \\
X_\perp^{(n)} &= X_\perp^{(n-1)} + \delta X_\perp^{(n)} &&= \sum_{j=1}^{n} \delta X_\perp^{(j)} \,.
\end{aligned} \tag{2.5}$$

It is clear that the transverse position, after any number of steps, is simply the sum of random deviations from all preceding steps.

As a sum of independent normally-distributed random variables, the transverse position

after n steps is itself a normally-distributed random variable with enhanced variance

$$\begin{aligned}
\mathrm{var}\left(X_\perp^{(n)}\right) &= \left\langle \left(X_\perp^{(n)}\right)^2 \right\rangle \\
&= \left\langle \sum_{j=1}^{n} \delta X_\perp^{(j)} \sum_{k=1}^{n} \delta X_\perp^{(k)} \right\rangle \\
&= \sum_{j=1}^{n}\sum_{k=1}^{n} \left\langle \delta X_\perp^{(j)} \delta X_\perp^{(k)} \right\rangle \\
&= \sum_{j=1}^{n}\sum_{k=1}^{n} \delta_{jk} l_P^2 \\
&= \sum_{j=1}^{n} l_P^2 \\
&= n\, l_P^2 \,. \qquad (2.6)
\end{aligned}$$

Over the full separation distance, reached after $n = R/l_P$ Planckian steps, the variance of transverse position is then

$$\mathrm{var}\left(X_\perp^{(R/l_P)}\right) = l_P\, R \,, \qquad (2.7)$$

thus recovering the variance of the transverse position commutator (eq. 2.2). This interpretation now explains the dependence of the variance on the system size, R. It enters simply from the length of the random walk, which sets the total "exposure time" to the stochastic process. The spatial fluctuations themselves can be viewed as a universal, stationary property of emergent space-time, with no dependence on the system size.

2.2.2 Generalizing to Continuous Coordinates

It will prove convenient to equivalently formulate the random walk in terms of the continuous coordinates of classical space-time. The transverse position fluctuations formerly indexed by discrete space-time location j will be promoted to functions of proper time t and spatial

position $\mathbf{x}$ as

$$\delta X_{\perp}^{(j)} \;\rightarrow\; \delta X_{\perp}(t,\,\mathbf{x}) \sim \mathcal{N}\left(0, l_P^2\right) \tag{2.8}$$

with covariance

$$\mathrm{cov}\left(\delta X_{\perp}(t,\,\mathbf{x}),\,\delta X_{\perp}(t+\tau,\,\mathbf{x}')\right) \;=\; \begin{cases} l_P^2\,, & |\tau| \leq \dfrac{t_P}{2} \\ 0\,, & |\tau| > \dfrac{t_P}{2} \end{cases} . \tag{2.9}$$

Embedded in this covariance structure is the core aspect of the model: non-local spatial correlations arising from entanglement of the quantum geometrical degrees of freedom. At a common time, the correlation coefficient of the fluctuations at any two spatial positions, $\mathbf{x}$ and $\mathbf{x}'$, is

$$\begin{aligned} \mathrm{corr}\left(\delta X_{\perp}(t,\,\mathbf{x}),\,\delta X_{\perp}(t,\,\mathbf{x}')\right) &= \frac{\mathrm{cov}\left(\delta X_{\perp}(t,\,\mathbf{x}),\,\delta X_{\perp}(t,\,\mathbf{x}')\right)}{\sqrt{\mathrm{var}\left(\delta X_{\perp}(t,\,\mathbf{x})\right)}\sqrt{\mathrm{var}\left(\delta X_{\perp}(t,\,\mathbf{x}')\right)}} \\ &= 1\,. \end{aligned} \tag{2.10}$$

This spatial coherence will be seen to make the effect measurable as time correlations in optical phase measurements on scales smaller than the light-crossing time of the apparatus. Since, at a given time, every spatial position undergoes a correlated fluctuation, measurements whose durations partially overlap measure a common set of spatial fluctuations during the period of overlap.

2.2.3 Spatial Fluctuation Rate

It will prove further convenient to define the *rate* of transverse position fluctuation, or the effective transverse velocity, as

$$\dot{X}_\perp(t) \equiv \frac{\delta X_\perp(t, \mathbf{x})}{t_P} , \tag{2.11}$$

which is explicitly independent of $\mathbf{x}$ due to the spatial coherence of the fluctuations. This variable is distributed as

$$\dot{X}_\perp(t) \sim \mathcal{N}\left(0, \left(\frac{l_P}{t_P}\right)^2\right) \tag{2.12}$$

with the covariance structure

$$\mathrm{cov}\left(\dot{X}_\perp(t), \dot{X}_\perp(t+\tau)\right) = \begin{cases} \left(\frac{l_P}{t_P}\right)^2 , & |\tau| \leq \frac{t_P}{2} \\ 0 , & |\tau| > \frac{t_P}{2} \end{cases} . \tag{2.13}$$

The characteristic rate of transverse position drift is thus one coherence length per coherence time.

2.3 Interferometric Detection

This section will now demonstrate that transverse spatial fluctuations can, in principle, impart measurable phase shifts on beams of freely-propagating light. These phase shifts are measured by an optical device known as a Michelson interferometer. A Michelson interferometer consists of a beamsplitter and two or more reflecting mirrors. Light from a coherent source is split into two beams which propagate through an arrangement of mirrors and back to the beamsplitter, where they recombine. Any difference in the distance propagated by the two beams introduces a corresponding difference in the accumulated optical phase. Interfer-

ence effects from this phase difference modulate the power of the recombined beam, which is read out via a photodetector at the output, or antisymmetric (AS), port of the instrument. In this section, the general response of a Michelson-class interferometer to transverse spatial fluctuations will first be derived. From this, the instrument response of the Holometer detector will be calculated. In the context of these calculations, the motivation of the Holometer design will be made apparent.

2.3.1 The Observable

Under the superposition principle of electromagnetism, the power at the AS port follows from the superposed electric field amplitude of the two individual beams, integrated over the image plane of the detector,

$$P_{\mathrm{AS}} = \frac{c\epsilon_0}{2} \iint \left|E_1 + E_2\right|^2 dS \; . \tag{2.14}$$

Assuming, for simplification, that the beamsplitter evenly splits the incident power, and neglecting losses, the electric field amplitudes of the beams at the AS port are

$$E_1 = \frac{E_0}{2} \exp\left(i\frac{2\pi}{\lambda} S_1\right) \tag{2.15}$$

and

$$\begin{aligned} E_2 &= \frac{E_0}{2} \exp\left(i\frac{2\pi}{\lambda}\left(S_2 + \frac{\lambda}{2} + S_0\right)\right) \\ &= -\frac{E_0}{2} \exp\left(i\frac{2\pi}{\lambda}\left(S_0 + S_2\right)\right) \; , \end{aligned} \tag{2.16}$$

where E_0 is the amplitude of the electric field incident on the beamsplitter, S_j is the round-trip optical path length (OPL), and λ is the wavelength of the light. By convention, a half-wavelength optical distance is added to one of the paths so that destructive interference

occurs when $S_1 = S_0 + S_2$. The second optical distance offset, $|S_0| << \lambda$, is added for control feedback purposes. Its role will be discussed in §3.4.4.

Expanding eq. 2.14, the power at the AS port is

$$\begin{aligned} P_{\mathrm{AS}} &= \frac{c\epsilon_0}{2}\left(\iint |E_1|^2\, dS + \iint |E_2|^2\, dS + 2\iint \mathrm{Re}\left[E_1^* E_2\right]\, dS\right) \\ &= \frac{c\epsilon_0}{4}\left[1 - \cos\left(\frac{2\pi}{\lambda}\left(S_0 + S_2 - S_1\right)\right)\right]\iint |E_0|^2\, dS \\ &= P_0 \sin^2\left(\frac{\pi}{\lambda}\left(S_0 + S_2 - S_1\right)\right) , \end{aligned} \tag{2.17}$$

where

$$P_0 \equiv \frac{c\epsilon_0}{2}\iint |E_0|^2\, dS \tag{2.18}$$

is the optical power incident on the beamsplitter. This establishes that a Michelson interferometer is sensitive to the length difference between two optical paths, with a measurement duration equal to the round-trip light-crossing time.

Each OPL will be promoted to a function of measurement time t as $S_j \to S_j(t)$. The measurement-time-dependent power at the AS port is then

$$P_{\mathrm{AS}}(t) = P_0 \sin^2\left(\frac{\pi}{\lambda}\left(S_0 + S(t)\right)\right) , \tag{2.19}$$

where

$$S(t) \equiv S_2(t) - S_1(t) \tag{2.20}$$

is the optical path difference (OPD). For small OPD, defined as $|S(t)| << \lambda$, eq. 2.19

simplifies to the linearized response

$$P_{\mathrm{AS}}(t) \approx \frac{\pi^2 P_0 S_0^2}{\lambda^2} + \frac{2\pi^2 P_0 S_0}{\lambda^2} S(t) \,. \tag{2.21}$$

In §2.4.4, the effect of transverse spatial fluctuations will be demonstrated to be firmly in the linear-response regime.

2.3.2 Effect on the Observable

As this section will demonstrate, transverse spatial fluctuations manifest as random deviations from the classical OPL. Taking each inteferometeter arm to have a total length of L, the two classical round-trip optical paths will be parameterized by propagation time $\tau \in [0,\, 2T]$, where $T = L/c$ is the light-crossing time. The proper time of a measurement maps to propagation time inside the instrument as $\tau = t - t_{\mathrm{inject}}$, where t_{inject} is the proper time at which the light first entered the interferometer. At propagation time τ, the classical position of a tracer photon traversing path i will be denoted by $\mathbf{x}_i(\tau)$. The tracer photons do not represent actual quanta of localized energy, but rather the causal bounds on any null-propagating wave through the instrument. At every point along optical path $\mathbf{x}_i(\tau)$, there exists a unit vector tangent to the path, $\hat{\mathbf{T}}_i(\tau) \equiv \dot{\mathbf{x}}_i(\tau)/c$, where $\dot{\mathbf{x}}_i(\tau)$ denotes the time derivative. This unit vector represents the instantaneous direction of light travel through interferometer.

The light origin, $\mathbf{x}_1(0) = \mathbf{x}_2(0)$, provides a common reference for the optical phase measurement made by either path. It can thus be regarded, under a relational theory, as a fixed reference point against which all other points in classical space appear to fluctuate. For each path, the optical distance propagated by the tracer photon over a classical round-trip light-crossing time is given by the path integral

$$S_i\left(t + 2T\right) = \int_t^{t+2T} \left[\dot{\mathbf{x}}_i\left(t' - t\right) + \mathbf{F}\left(t'\right)\right] \cdot \hat{\mathbf{T}}_i\left(t' - t\right) dt' \,, \tag{2.22}$$

where $\mathbf{F}(t)$ is a stochastic vector field representing the instantaneous position displacement relative to the light origin. The vector sum of this field and the classical velocity corresponds to the "effective" velocity of the tracer photon relative to the classical spatial coordinates. The additive field does not represent any physical motion of the test particle, but rather fluctuations of the spatial background itself due to quantization noise. Under this virtual process, no momentum exchange occurs. The instantaneous effect on the OPL is obtained by taking the component of this effective velocity in the direction of light travel[1].

The transverse position uncertainty relation constrains the general form of the stochastic vector field. Across all space, the commutator and requirements of spatial coherence (see §2.2.2) imply that

$$\mathbf{F}(t) = \dot{X}_{\perp}(t)\,\hat{\mathbf{d}}\,, \tag{2.23}$$

where $\dot{X}_{\perp}(t)$ is the instantaneous transverse spatial fluctuation (introduced in §2.2.3) and $\hat{\mathbf{d}}$ is a unit vector. Along a straight optical path, the commutator requires that $\hat{\mathbf{d}}$ correspond to the direction transverse to the path. However, the commutator does not uniquely specify the direction of spatial fluctuations at locations away from this path. The following sections will consider two candidate forms of $\mathbf{F}(t)$, each obeying the requirements of §2.1 under different interpretations.

Under the variable substitution $\tau \equiv t' - t$, eq. 2.22 becomes

$$\begin{aligned} S_i(t+2T) &= \int_0^{2T} \left[\dot{\mathbf{x}}_i(\tau) + \mathbf{F}(\tau+t)\right] \cdot \hat{\mathbf{T}}_i(\tau)\,d\tau \\ &= \int_0^{2T} \left[c\hat{\mathbf{T}}_i(\tau) + \mathbf{F}(\tau+t)\right] \cdot \hat{\mathbf{T}}_i(\tau)\,d\tau \\ &= 2L + \int_0^{2T} \mathbf{F}(\tau+t) \cdot \hat{\mathbf{T}}_i(\tau)\,d\tau\,. \end{aligned} \tag{2.24}$$

1. An angular effect analogously arises from the effective velocity component in the direction *transverse* to light travel. However, this effect is highly subdominant to the effect on OPL and will be neglected.

In this form, the spatial fluctuations accumulated over a measurement duration can be clearly seen to manifest as a deviation from the classical OPL. The difference of optical distances is then a pure measure of the exotic effect,

$$\begin{aligned} S(t+2T) &= \int_0^{2T} \mathbf{F}(\tau+t)\cdot\hat{\mathbf{T}}_2(\tau)\,d\tau - \int_0^{2T} \mathbf{F}(\tau+t)\cdot\hat{\mathbf{T}}_1(\tau)\,d\tau \\ &= \int_0^{2T} \mathbf{F}(\tau+t)\cdot\left[\hat{\mathbf{T}}_2(\tau)-\hat{\mathbf{T}}_1(\tau)\right]d\tau \\ &= \int_0^{2T} \dot{X}_\perp(\tau+t)\,\hat{\mathbf{d}}\cdot\left[\hat{\mathbf{T}}_2(\tau)-\hat{\mathbf{T}}_1(\tau)\right]d\tau \\ &= \int_0^{2T} \dot{X}_\perp(\tau+t)\,\Theta(\tau)\,d\tau\ , \end{aligned} \tag{2.25}$$

where

$$\Theta(\tau) \equiv \hat{\mathbf{d}}\cdot\left[\hat{\mathbf{T}}_2(\tau)-\hat{\mathbf{T}}_1(\tau)\right] \tag{2.26}$$

is the geometrical coupling of the apparatus to the stochastic field. As will be seen, only the calculation of this coupling function is required to solve for the statistical response of an instrument of arbitrary geometry.

Shear-Field Effect

Under a shear-field interpretation, one optical path experiences longitudinal fluctuations while the second path functions as a light clock, providing an independent reference against which the phase of first path is compared. The left panel of Fig. 2.1 depicts this scenario for a straight-arm Michelson interferometer. With no loss of generality, let the beamsplitter be centered on the origin of coordinates and oriented such that optical paths 1 and 2 emerge along the $\hat{\imath}$ and $\hat{\jmath}$ axes, respectively[2]. Choosing, arbitrarily, optical path 1 as the phase

2. Effects calculated under this assumption will be seen to be invariant under an arbitrary change of spatial coordinates.

reference constrains the spatial fluctuations to the direction transverse to the path, the $\hat{\mathbf{j}}$-direction. This implies the shear vector field

$$\mathbf{F}(t) = \dot{X}_{\perp}(t)\,\hat{\mathbf{j}}\,. \tag{2.27}$$

It is immediately clear that this field has no effect on the optical distance of path 1, which is oriented purely in the $\hat{\mathbf{i}}$-direction. However, the field does fully couple to the optical distance of orthogonally-oriented path 2, resulting in a phase shift relative to the independent reference provided by path 1.

This interpretation asserts a duality between the phase measurements of the two optical paths. From the perspective of optical path 1, as above, path 1 experiences no longitudinal displacement itself, but path 2 does. However, from the perspective of optical path 2, now chosen instead as the phase reference, path 2 experiences no longitudinal displacement itself, but path 1 does. Through this duality, a violation of the transverse position commutator, which forbids displacements along the line-of-sight separation from the light origin, is only apparent, as it is impossible to determine which path truly fluctuated.

It is important to recognize that this interpretation derives its notion of duality from the geometrical symmetry of the two optical paths. In order to provide an *independent* phase reference, one path must be oriented purely transverse to the stochastic shear field. And in order to be dual, the second path must be an identical, 90° rotation of the first. As such, the shear-field interpretation cannot be generalized to make predictions for arbitrary optical geometries, which do not exhibit the above symmetries.

Rotational-Field Effect

Relaxing the notion of optical-path duality invoked by the shear-field effect requires both optical paths to satisfy the transverse position commutator simultaneously. Again adopting

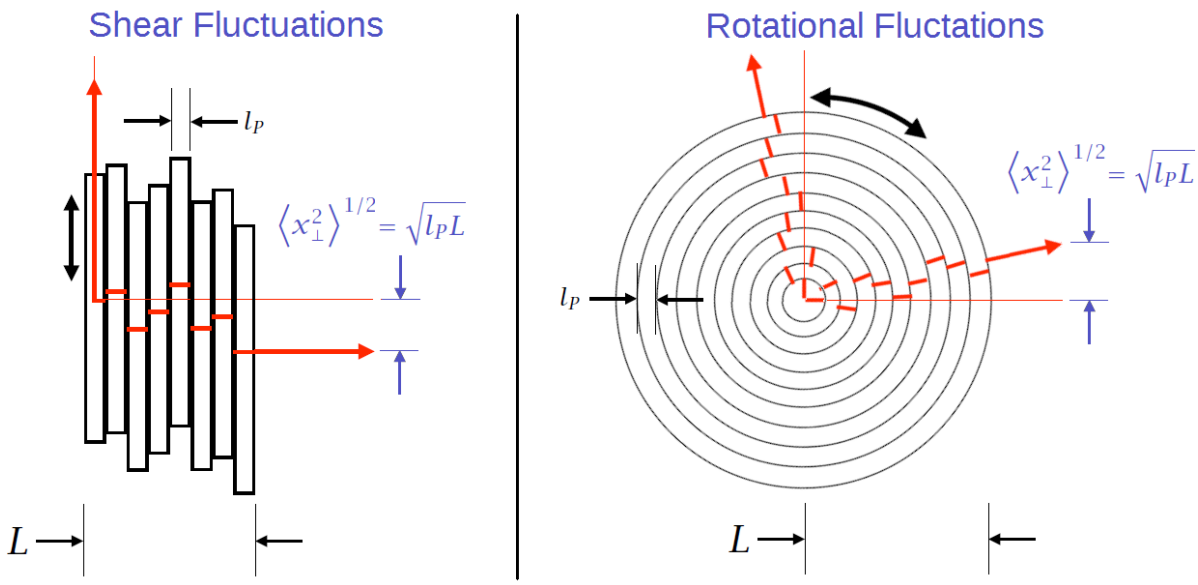

Figure 2.1: Illustration of shear versus rotational spatial fluctuations, each arising from Planck-width elements of a different spatial symmetry. Along the horizontal optical path, indicated by the narrow red line, both interpretations produce the same integrated transverse position uncertainty, $\langle x_\perp^2 \rangle^{1/2}$, over the path length L. However, the two interpretations dramatically differ in their predicted effect along the second, vertically-oriented optical path. The shear symmetry produces an apparent longitudinal uncertainty in the second path, while the rotational symmetry produces a purely transverse uncertainty in both paths.

the beamsplitter-centered coordinate system introduced in §2.3.2, any position is constrained to fluctuate in the direction transverse to its radial separation from the light origin, which is in the angular direction $\hat{\boldsymbol{\theta}}$. This implies the rotational vector field

$$\mathbf{F}(t) \;=\; \dot{X}_\perp(t)\,\hat{\boldsymbol{\theta}}\,, \tag{2.28}$$

as illustrated in the right panel of Fig. 2.1. Unlike the shear field, the rotational field naturally allows for calculations over optical paths of arbitrary geometry. The purely angular character of the field satisfies the commutator relation between the light origin and every spatial position simultaneously, requiring no additional symmetry from the optical paths.

2.3.3 Spatial Entanglement

Exotic spatial fluctuations measured by co-located interferometers are expected to exhibit a high degree of correlation due to the entanglement of neighboring quantum position states. The degree of entanglement between spatial positions be can expressed in terms of the covariance of their fluctuations. This section will demonstrate that the structure of this covariance is uniquely specified by the Gauss-Markov formalism.

Consider first a rotational-field effect as measured by the configuration of interferometers illustrated in Fig. 2.2. The thick blue line represents an interferometer with its beamsplitter located at the coordinate origin. At each point along its optical paths, this interferometer measures spatial fluctuations in the angular direction relative to its beamsplitter. Suppose, at a given time, that the direction of fluctuation is denoted by the blue circle. An orange arrow indicates the optical-axis displacement measured at a particular location due to this fluctuation.

Now consider a second interferometer, represented by the thick magenta line, shifted a distance x_0 along the $\hat{\mathbf{i}}$-axis. This interferometer, too, measures spatial fluctuations in the angular direction relative to its beamsplitter, as denoted by the magenta circle. However, there is a segment along the $\hat{\mathbf{i}}$-axis where the optical paths of both interferometers overlap. At a given point on this segment, denoted in Fig. 2.2 by two overlapping red dots, both interferometers measure the same observable, the local fluctuation of a common spatial position, simultaneously. Since these two measurements are, in fact, identical, consistency requires that the same optical-axis displacement must be measured relative to either beamsplitter. This consistency condition is only satisfied if the fluctuations around the two origins are perfectly correlated in their magnitude and direction of rotation (clockwise/counterclockwise). It can be seen that this correlation occurs at a measurement time offset of x_0/c, the difference in emission times for the two beams to reach the same point on the overlapping path segment.

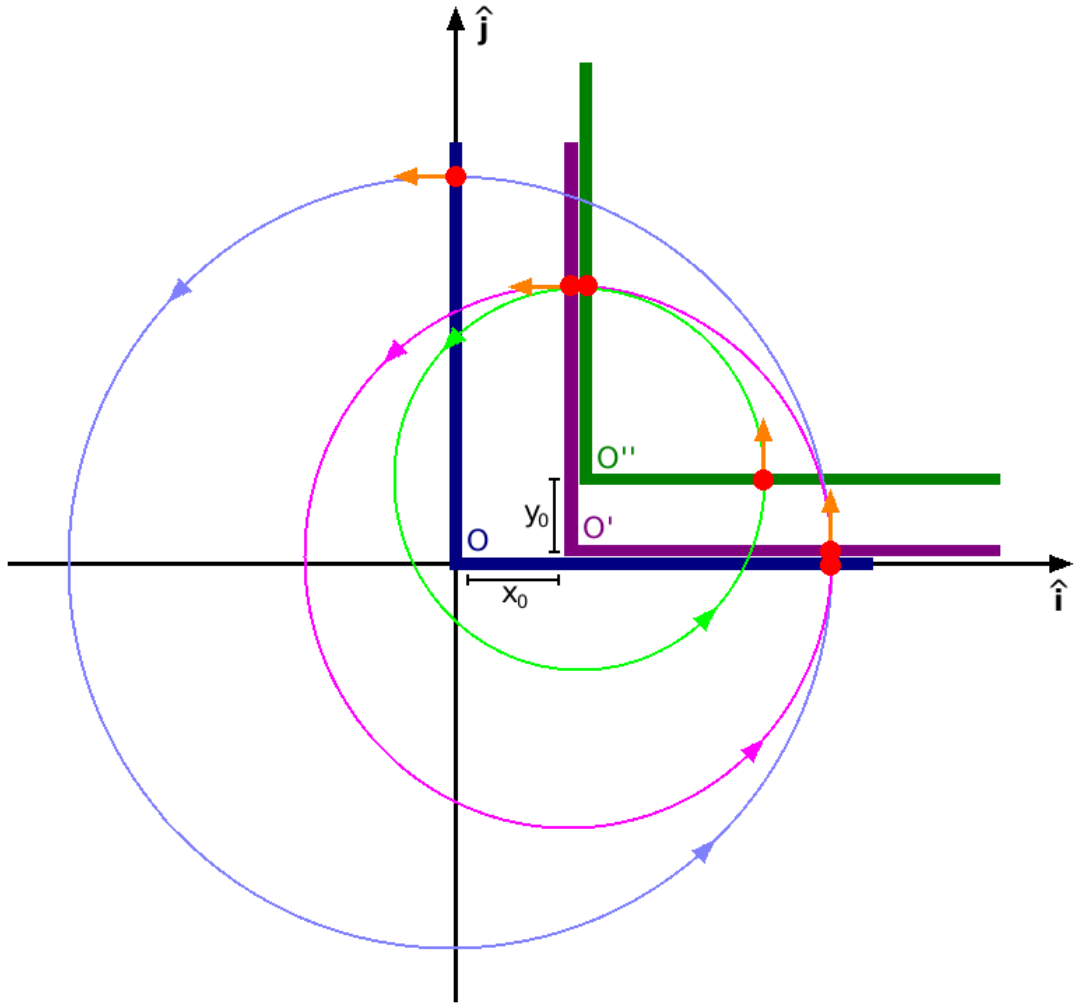

Figure 2.2: Illustration of the entanglement of rotational spatial fluctuations measured relative to multiple light origins. Three interferometers, represented by the thick colored lines, are shown spatially displaced from one another so that their light paths partially overlap. Along the overlapping segments, the interferometers make identical measurements of the local spatial fluctuation, denoted by an overlapping pair of red dots, which must be consistent. At this location, the local space must fluctuate identically relative to both light origins, as indicated by the co-directional field lines, on a correlation time scale equal to the light-time separation of the two origins.

Next, consider a third interferometer, represented by the thick green line, shifted a distance y_0 in the $\hat{\mathbf{j}}$-direction relative to the second interferometer. The green circle denotes the direction of spatial fluctuations relative to its beamsplitter. Since the optical paths of the second and third interferometers partially overlap, the above logic implies that fluctuations around the second and third beamsplitters must correlate at a measurement time offset of y_0/c. However, the fluctuations around the second beamsplitter were already established to correlate with those around the first beamsplitter. By transitivity, the fluctuations around the first and third beamsplitters must then also correlate, *even though no portion of their corresponding light paths overlap.* An identical conclusion is trivially reached when considering the same configuration of interferometers in the presence of a shear-field effect, as shown in Fig. 2.3.

This demonstrates that the spatial fluctuations measured by neighboring interferometers correlate purely due to the spatial proximity of the instruments, irrespective of whether their optical paths coincide, on a time scale equal to the light-crossing time of their separation. Formally, this statement is expressed by the covariance

$$\operatorname{cov}\left(\dot{X}_{\perp}(t),\, \dot{X}'_{\perp}(t+\tau)\right) = \begin{cases} \left(\dfrac{l_P}{t_P}\right)^2, & \left||\tau| - \dfrac{|\mathbf{x}'|}{c}\right| \leq \dfrac{t_P}{2} \\ 0\,, & \left||\tau| - \dfrac{|\mathbf{x}'|}{c}\right| > \dfrac{t_P}{2} \end{cases}, \tag{2.29}$$

where $\dot{X}'_{\perp}(t)$ is the spatial fluctuation rate relative to an origin translated by a distance $|\mathbf{x}'|$. The correlation time scale $|\mathbf{x}'|/c$ can be seen to impose causal bounds on the spatial entanglement of the two origins, as it specifies the time scale on which neighboring spatial positions entangle. When $|\mathbf{x}'|/c > 2L/c$, the time to entangle exceeds the measurement duration and this structure forces the correlation between fluctuations in the two interferometers to vanish.

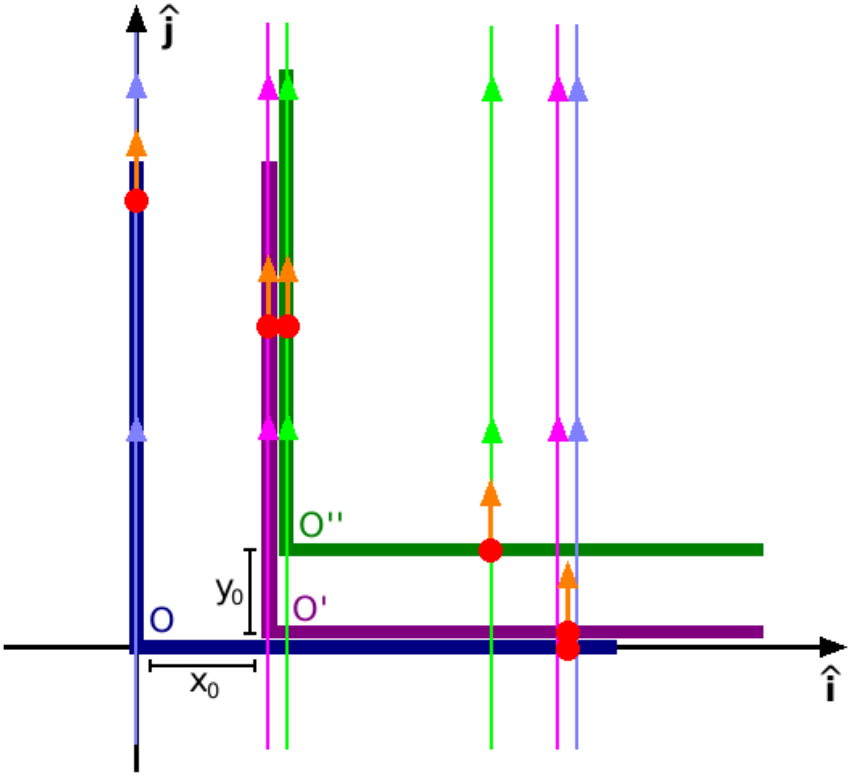

Figure 2.3: Illustration of the entanglement of shear spatial fluctuations measured relative to multiple light origins. Three interferometers, represented by the thick colored lines, are shown spatially displaced from one another so that their light paths partially overlap. Along the overlapping segments, the interferometers make identical measurements of the local spatial fluctuation, denoted by an overlapping pair of red dots, which must be consistent. At this location, the local space must fluctuate identically relative to both light origins, as indicated by the co-directional field lines, on a correlation time scale equal to the light-time separation of the two origins.

2.3.4 Single-Interferometer Statistics

Eq. 2.25 represents OPD measurements made over an interval of time as a set of random variables indexed by measurement time t. It is a straightforward exercise to calculate the general statistical moments of this measurement set in the presence of exotic spatial fluctuations, as the following sections now do.

Mean

The expected value of an OPD measurement is

$$\begin{aligned} \left\langle S(t) \right\rangle &= \left\langle \int_0^{2T} \dot{X}_\perp \left(\tau + (t - 2T)\right) \Theta\left(\tau\right) \, d\tau \right\rangle \\ &= \int_0^{2T} \left\langle \dot{X}_\perp \left(\tau\right) \right\rangle \Theta\left(\tau\right) \, d\tau \\ &= 0 \, , \end{aligned} \tag{2.30}$$

using the additive separability of expectation values. The zero mean of the measurement is seen to be a direct consequence of the zero-mean process generating the spatial fluctuations, independent of the geometry of the apparatus.

Autocovariance

The autocovariance of two OPD measurements separated in time by $|\tau| < 2T$ is

$$
\begin{aligned}
C_{SS}(\tau \,|\, l_P) &\equiv \big\langle S(t)\, S(t+\tau) \big\rangle - \big\langle S(t) \big\rangle \big\langle S(t+\tau) \big\rangle \\
&= \Bigg\langle \int_0^{2T} \dot{X}_\perp \left(\tau' + (t-2T)\right) \Theta\left(\tau'\right) \, d\tau' \\
&\quad \times \int_\tau^{2T+\tau} \dot{X}_\perp \left(\tau'' + (t-2T)\right) \Theta\left(\tau'' - \tau\right) \, d\tau'' \Bigg\rangle \\
&= \int_0^{2T} \int_\tau^{2T+\tau} \Big\langle \dot{X}_\perp \left(\tau'\right) \, \dot{X}_\perp \left(\tau''\right) \Big\rangle \Theta\left(\tau'\right) \, \Theta\left(\tau'' - \tau\right) \, d\tau'' \, d\tau' \\
&= c\, l_P \int_{|\tau|}^{2T} \Theta\left(\tau'\right) \, \Theta\left(\tau' - |\tau|\right) \, d\tau' \qquad (2.31)
\end{aligned}
$$

and identically zero for larger separations. This relation identically vanishes for $|\tau| > 2T$, the duration of an individual measurement.

Power Spectral Density

Under the Wiener-Khinchin theorem, an equivalent, frequency-space representation of the autocovariance is the power spectral density (PSD). The PSD is defined as the Fourier transform of the autocovariance,

$$
\begin{aligned}
\widetilde{C_{SS}}(f \,|\, l_P) &\equiv 2 \int_{-\infty}^{\infty} C_{SS}(\tau \,|\, l_P)\, e^{-i2\pi f \tau} \, d\tau \\
&= 2 \left[\int_{-\infty}^{\infty} C_{SS}(\tau \,|\, l_P) \, \cos\left(2\pi f \tau\right) \, d\tau - i \int_{-\infty}^{\infty} C_{SS}(\tau \,|\, l_P) \, \sin\left(2\pi f \tau\right) \, d\tau \right] \\
&= 2 \Bigg\{ \left[\int_{-\infty}^{0} C_{SS}(\tau \,|\, l_P) \, \cos\left(2\pi f \tau\right) \, d\tau + \int_{0}^{\infty} C_{SS}(\tau \,|\, l_P) \, \cos\left(2\pi f \tau\right) \, d\tau \right] \\
&\quad - i \left[\int_{-\infty}^{0} C_{SS}(\tau \,|\, l_P) \, \sin\left(2\pi f \tau\right) \, d\tau + \int_{0}^{\infty} C_{SS}(\tau \,|\, l_P) \, \sin\left(2\pi f \tau\right) \, d\tau \right] \Bigg\} \\
&= 4 \int_{0}^{\infty} C_{SS}(\tau \,|\, l_P) \, \cos\left(2\pi f \tau\right) \, d\tau \,, \qquad (2.32)
\end{aligned}
$$

where the final equality follows from the fact that $C_{SS}(\tau \,|\, l_P) = C_{SS}(-\tau \,|\, l_P)$. This PSD, written in the so-called engineering convention, is defined only for positive frequencies, into which the power contained in the redundant negative frequencies is folded via the multiplicative prefactor of two.

2.3.5 Dual-Interferometer Statistics

In §2.3.2, $S(t)$ was defined as the OPD measurement made by an interferometer whose beamsplitter is centered on the coordinate origin. Similarly, $S'(t)$ can be defined as the OPD measurement made by an interferometer translated to the coordinate position $\mathbf{x}'$. Analogously to the single-interferometer case, the cross-interferometer statistics of the two measurement sets can now be calculated via eq. 2.29.

Cross-Covariance

The cross-covariance of two spatially-separated OPD measurements at time separation $|\tau| < 2T - |\mathbf{x}'|/c$ is

$$\begin{aligned}
C_{SS'}(\tau \,|\, l_P) &\equiv \left\langle S(t)\, S'(t+\tau) \right\rangle - \left\langle S(t) \right\rangle \left\langle S'(t+\tau) \right\rangle \\
&= \left\langle \int_0^{2T} \dot{X}_\perp \left(\tau' + (t-2T)\right) \Theta\left(\tau'\right) d\tau' \right. \\
&\quad \left. \times \int_\tau^{2T+\tau} \dot{X}'_\perp \left(\tau'' + (t-2T)\right) \Theta\left(\tau'' - \tau\right) d\tau'' \right\rangle \\
&= \int_0^{2T} \int_\tau^{2T+\tau} \left\langle \dot{X}_\perp\left(\tau'\right) \dot{X}'_\perp\left(\tau''\right) \right\rangle \Theta\left(\tau'\right) \Theta\left(\tau'' - \tau\right) d\tau'' \, d\tau' \\
&= c\, l_P \int_{|\tau| + \frac{|\mathbf{x}'|}{c}}^{2T} \Theta\left(\tau'\right) \Theta\left(\tau' - \left(|\tau| + \frac{|\mathbf{x}'|}{c}\right)\right) d\tau'
\end{aligned} \tag{2.33}$$

and identically zero for larger separations. This "reduced" response function exhibits an equivalency between temporal and spatial separation, as the cross-covariance at zero delay

equals the auto-covariance of either interferometer at a delay of $|\mathbf{x}'|/c$.

Cross-Spectral Density

The equivalent frequency-space representation of the cross-covariance is the cross-spectral density (CSD). Analogously to the PSD, the CSD is defined as the Fourier transform of the cross-covariance,

$$\begin{aligned}
\widetilde{C_{SS'}}(f\,|\,l_P) &\equiv 2\int_{-\infty}^{\infty} C_{SS'}(\tau\,|\,l_P)\,e^{-i2\pi f\tau}\,d\tau \\
&= 2\left[\int_{-\infty}^{\infty} C_{SS'}(\tau\,|\,l_P)\,\cos(2\pi f\tau)\,d\tau - i\int_{-\infty}^{\infty} C_{SS'}(\tau\,|\,l_P)\,\sin(2\pi f\tau)\,d\tau\right] \\
&= 2\left\{\left[\int_{-\infty}^{0} C_{SS'}(\tau\,|\,l_P)\,\cos(2\pi f\tau)\,d\tau + \int_{0}^{\infty} C_{SS'}(\tau\,|\,l_P)\,\cos(2\pi f\tau)\,d\tau\right]\right. \\
&\quad \left. - i\left[\int_{-\infty}^{0} C_{SS'}(\tau\,|\,l_P)\,\sin(2\pi f\tau)\,d\tau + \int_{0}^{\infty} C_{SS'}(\tau\,|\,l_P)\,\sin(2\pi f\tau)\,d\tau\right]\right\} \\
&= 4\int_{0}^{\infty} C_{SS'}(\tau\,|\,l_P)\,\cos(2\pi f\tau)\,d\tau\,, \qquad (2.34)
\end{aligned}$$

whose final form again follows from the fact that $C_{SS'}(\tau\,|\,l_P) = C_{SS'}(-\tau\,|\,l_P)$.

2.4 Instrument Response

Phase I of the Holometer experiment, whose results are presented in this dissertation, is a search for a stochastic shear-field effect (see §2.3.2). The following sections describe the geometry of the interferometer light paths and calculate the statistical response of the instruments under the Phase I design.

2.4.1 Optical Geometry

Each of the Holometer interferometers consists of two straight 40-m arms oriented in orthogonal directions, as shown in Fig. 2.4. Adopting the beamsplitter-centered coordinate system introduced in §2.3.2, optical path 1 extends purely in the $\hat{\mathbf{i}}$-direction and optical path 2 extends purely in the $\hat{\mathbf{j}}$-direction. The classical round-trip light paths are

$$\mathbf{x}_i(\tau) = \begin{cases} \hat{\mathbf{i}}\,, & i = 1 \\ \hat{\mathbf{j}}\,, & i = 2 \end{cases} \times \begin{cases} c\tau\,, & \tau < T \\ 2L - c\tau\,, & \tau \geq T \end{cases}, \tag{2.35}$$

where the arm length $L = 40$ m. The time derivatives of these paths imply the unit tangent vector

$$\begin{aligned} \hat{\mathbf{T}}_i(\tau) &= \begin{cases} \hat{\mathbf{i}}\,, & i = 1 \\ \hat{\mathbf{j}}\,, & i = 2 \end{cases} \times \begin{cases} 1\,, & \tau < T \\ -1\,, & \tau \geq T \end{cases} \\ &= \begin{cases} \hat{\mathbf{r}}\,, & \tau < T \\ -\hat{\mathbf{r}}\,, & \tau \geq T \end{cases}, \end{aligned} \tag{2.36}$$

which reverses sign upon reflection from the end mirrors.

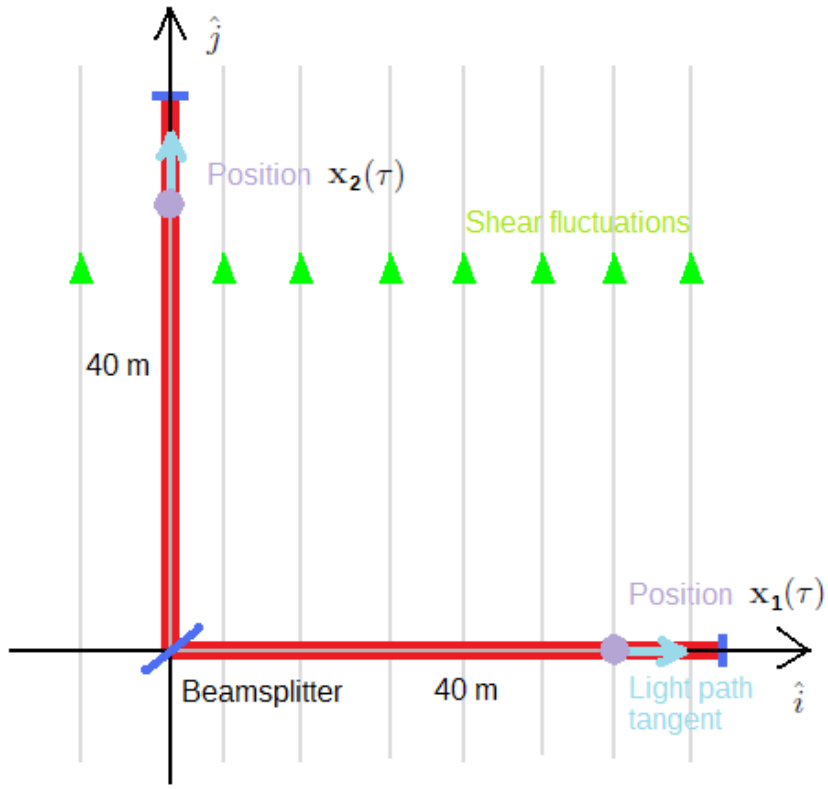

Figure 2.4: Geometrical schematic of the optical paths through each Holometer interferometer. Light incident on the beamsplitter is split into two straight, orthogonally-oriented paths 40 m long, as indicated by the thick red lines. In the presence of exotic spatial shear, represented by the set of vertical grey field lines, one path accumulates purely longitudinal fluctuations while the other accumulates purely transverse fluctuations. Only the longitudinal deviations affect the optical phase difference measured at the AS port.

2.4.2 Coupling to Shear Effect

Adopting the stochastic shear field of eq. 2.27, substituting eq. 2.36 into eq. 2.26 yields a geometrical coupling of

$$\begin{aligned}\Theta(\tau) &= \begin{cases}\hat{\mathbf{j}}\cdot\left(\hat{\mathbf{j}}-\hat{\mathbf{i}}\right)\ , & \tau<T\\ \hat{\mathbf{j}}\cdot\left(-\hat{\mathbf{j}}+\hat{\mathbf{i}}\right)\ , & \tau\geq T\end{cases}\\ &= \begin{cases}1\ , & \tau<T\\ -1\ , & \tau\geq T\end{cases}\end{aligned} \tag{2.37}$$

to exotic shear noise.

2.4.3 Coupling to Rotational Effect

For completeness, this section will further demonstrate that the Phase I Holometer configuration has sensitivity to *purely* the shear effect. Adopting the stochastic rotational field of eq. 2.28, substituting eq. 2.36 into eq. 2.26 analogously yields a geometrical coupling of

$$\begin{aligned}\Theta(\tau) &= \begin{cases}\hat{\boldsymbol{\theta}}\cdot\left(\hat{\mathbf{r}}-\hat{\mathbf{r}}\right)\ , & \tau<T\\ \hat{\boldsymbol{\theta}}\cdot\left(-\hat{\mathbf{r}}+\hat{\mathbf{r}}\right)\ , & \tau\geq T\end{cases}\\ &= 0\end{aligned} \tag{2.38}$$

to exotic rotational noise. The rotational coupling can thus be seen to vanish under purely radial propagation of the light. Appendix A provides an example of an optical geometry which does have non-vanishing sensitivity to the rotational effect.

2.4.4 Auto-Response Functions

The following sections will now calculate the statistical response of each interferometer individually to exotic spatial shear noise.

Autocovariance

Substituting the coupling function of eq. 2.37 into eq. 2.31 yields the autocovariance of each Holometer interferometer,

$$C_{SS}(\tau \,|\, l_P) \;=\; l_P \times \begin{cases} 2L - 3c|\tau| \,, & |\tau| < T \\ c|\tau| - 2L \,, & T \leq |\tau| \leq 2T \\ 0 \,, & |\tau| > 2T \end{cases} \,. \tag{2.39}$$

This function is shown in the top panel of Fig. 2.5 for $l_P = l_p$. At zero delay, the autocovariance reduces to the variance of OPD measurements. The scale of optical path displacement is thus $\sqrt{C_{SS}(0 \,|\, l_p)} \approx 36$ am, a scale far smaller than the 1064 nm wavelength of light in the Holometer. This places the optical response of the interferometers firmly in the linear-response regime, as was discussed in §2.3.1.

Power Spectral Density

Substituting eq. 2.39 into eq. 2.32 then yields the PSD of each Holometer interferometer,

$$\widetilde{C_{SS}}(f \,|\, l_P) \;=\; l_P \left[\frac{8L^2}{c} \operatorname{sinc}^2\left(\frac{\pi L}{c} f\right) - \frac{2(2L)^2}{c} \operatorname{sinc}^2\left(\frac{\pi (2L)}{c} f\right) \right] \,. \tag{2.40}$$

This function, shown in the bottom panel of Fig. 2.5 for $l_P = l_p$, represents an OPD noise of exotic origin for which the Holometer is designed to search.

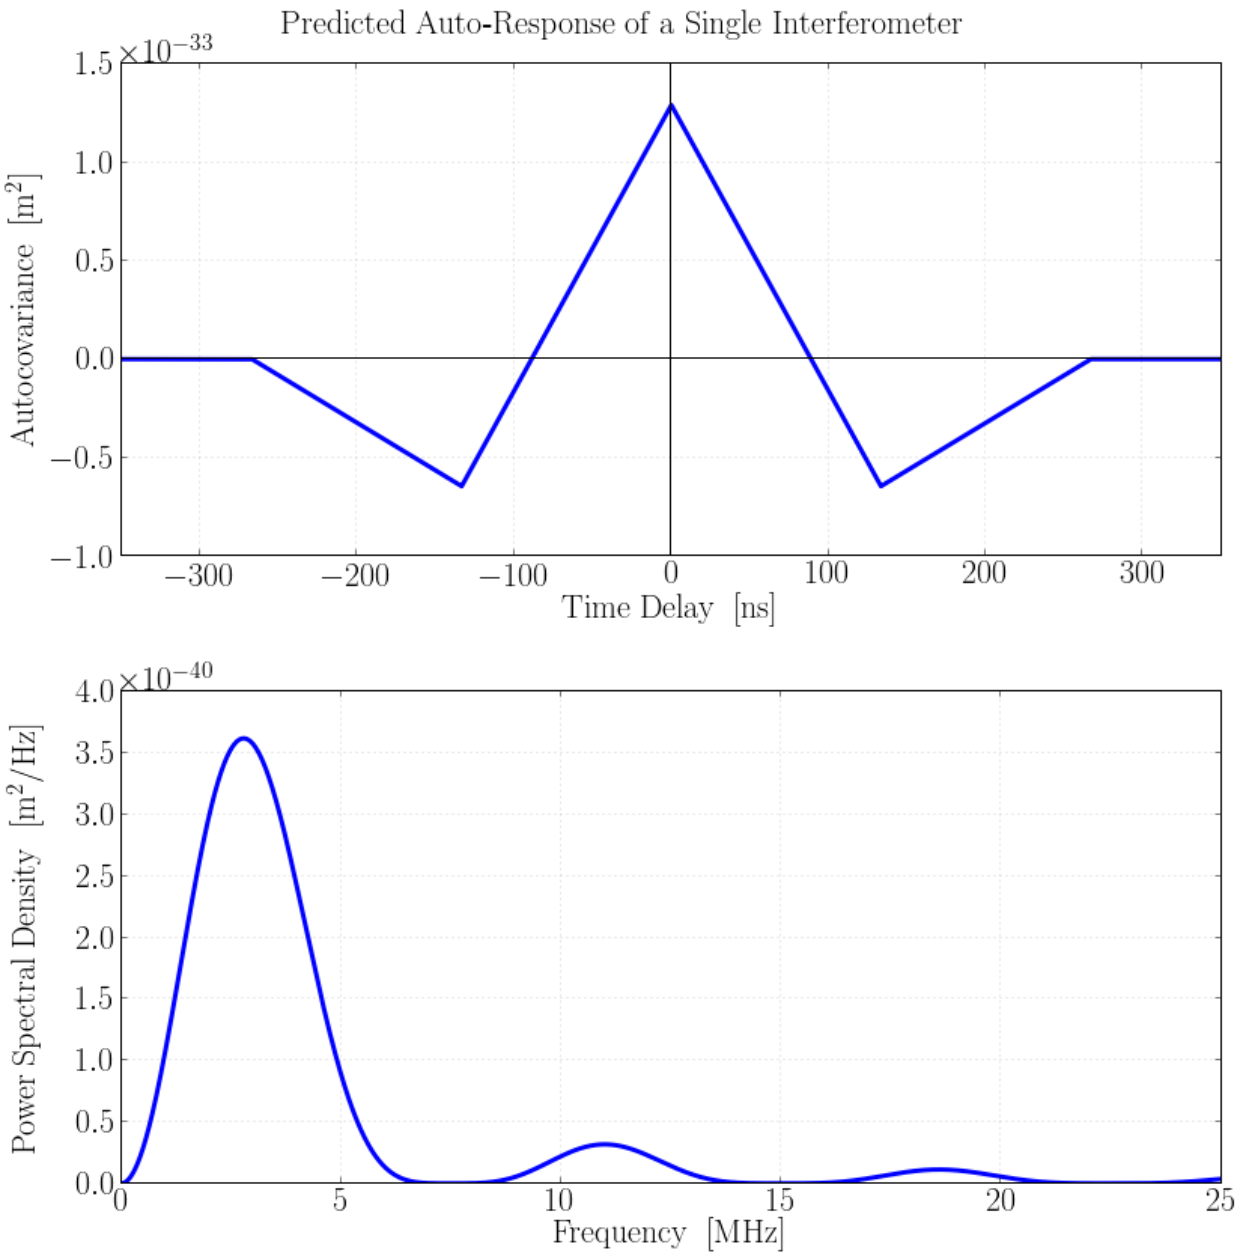

Figure 2.5: Time- and frequency-domain response functions of each Holometer interferometer to exotic spatial shear noise. The autocovariance (top panel) is given by eq. 2.39 and the PSD (bottom panel) is given by eq. 2.40.

2.4.5 Cross-Response Functions

As will be shown in §3.2, the exotic spectral noise power given by eq. 2.40 is a factor of 10^3 weaker than standard optical shot noise, the dominant interferometer noise component at radio frequencies. This precludes the possibility of directly detecting this noise excess in a single instrument, as it requires control of all systematic uncertainties to an untenable 0.1%. The Holometer overcomes this limitation with a dual-interferometer detection strategy which isolates the spectral noise component common to co-located but independent instruments. At radio frequencies, no conventional source of optical noise is expected to correlate between independent systems. However, as was demonstrated in §2.3.3, the exotic spatial noise measured by co-located interferometers must exhibit a high degree of coherence. The following sections will now calculate the statistical response of the interferometers to *common* exotic spatial shear noise.

Cross-Covariance

Substituting the coupling function of eq. 2.37 into eq. 2.33 yields the cross-covariance of each Holometer interferometer,

$$C_{SS'}(\tau \,|\, l_P) \;=\; l_P \times \begin{cases} 2L - 3c|\tau| - 3|\mathbf{x}'| \,, & |\tau| < T - \dfrac{|\mathbf{x}'|}{c} \\ c|\tau| - 2L + |\mathbf{x}'| \,, & T - \dfrac{|\mathbf{x}'|}{c} \leq |\tau| \leq 2T - \dfrac{|\mathbf{x}'|}{c} \\ 0 \,, & |\tau| > 2T - \dfrac{|\mathbf{x}'|}{c} \end{cases} \,. \tag{2.41}$$

This function is shown in the top panel of Fig. 2.6 for a variety of spatial separations, with $l_P = l_p$.

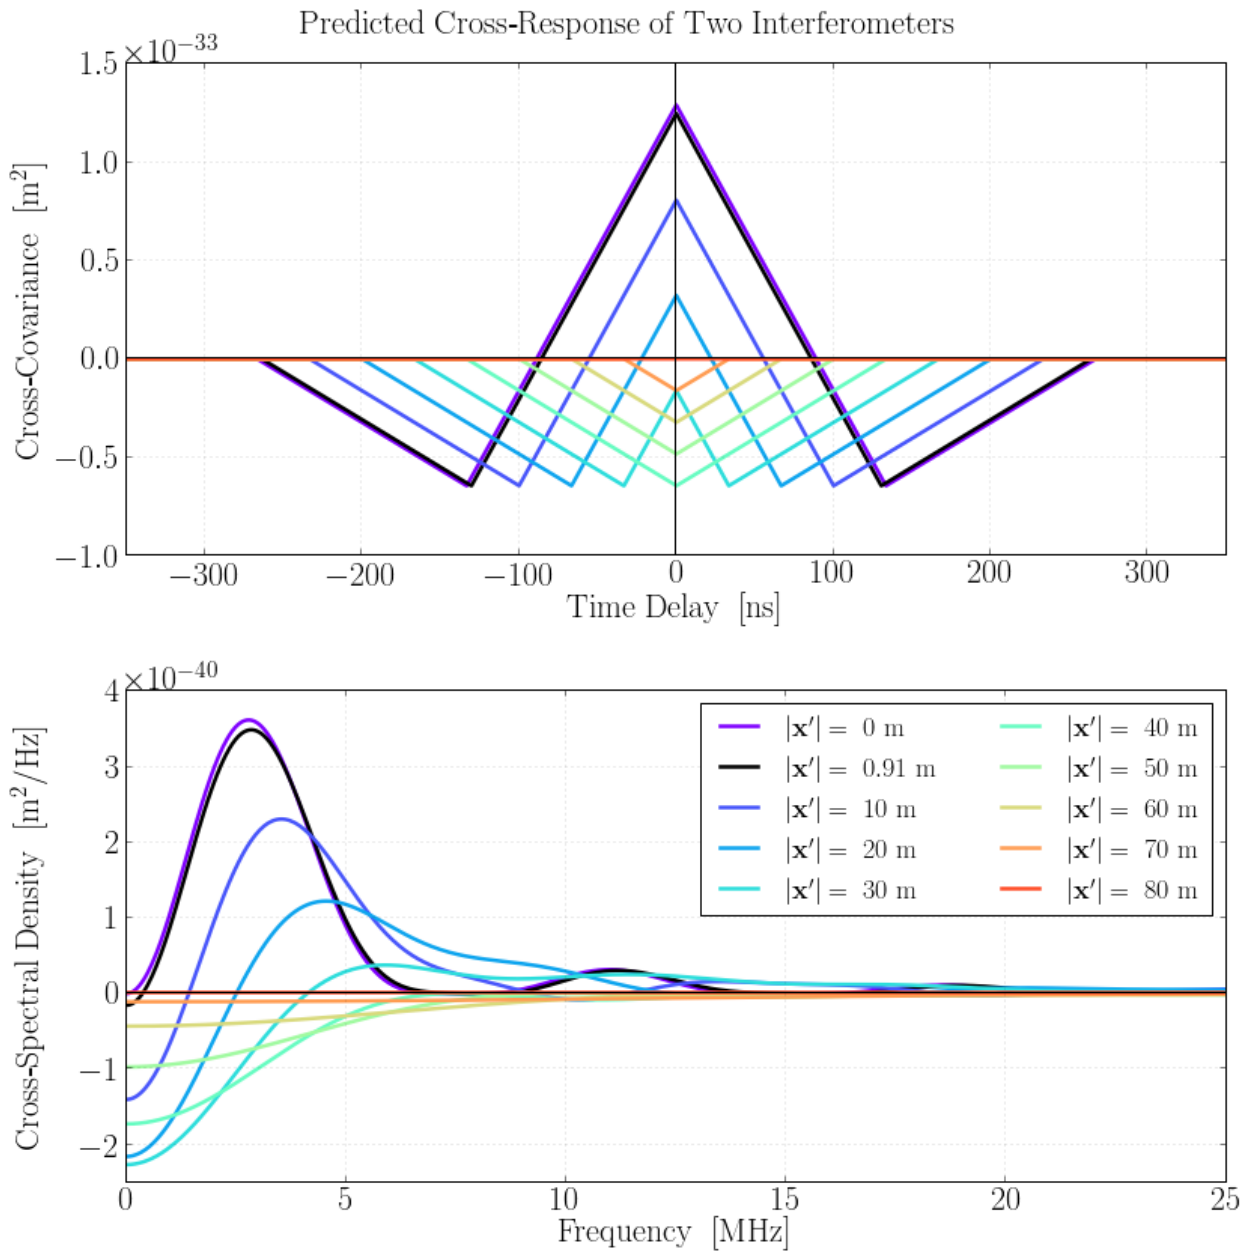

Figure 2.6: Correlated time- and frequency-domain response functions of the Holometer interferometers to exotic spatial shear noise, shown for a variety of spatial separations in a two-dimensional plane. The black curve corresponds to the actual separation of the two interferometer origins, 0.91 m. The cross-covariance (top panel) is given by eq. 2.41 and the CSD (bottom panel) is given by eq. 2.42.

Cross-Spectral Density

Substituting eq. 2.41 into eq. 2.34 then yields the CSD of the Holometer interferometers,

$$\begin{aligned}\widetilde{C_{SS'}}(f\,|\,l_P) &= l_P \times \begin{cases} \dfrac{8\left(L-|\mathbf{x}'|\right)^2}{c}\,\mathrm{sinc}^2\left(\dfrac{\pi\left(L-|\mathbf{x}'|\right)}{c}f\right), & |\mathbf{x}'|<L \\ 0\,, & |\mathbf{x}'|\geq L \end{cases} \\ &\quad - l_P \times \begin{cases} \dfrac{2\left(2L-|\mathbf{x}'|\right)^2}{c}\,\mathrm{sinc}^2\left(\dfrac{\pi\left(2L-|\mathbf{x}'|\right)}{c}f\right), & |\mathbf{x}'|<2L \\ 0\,, & |\mathbf{x}'|\geq 2L \end{cases}\end{aligned} \tag{2.42}$$

which vanishes for spatial light-time separations larger than a measurement duration. This function is shown in the bottom panel of Fig. 2.6 for a variety of spatial separations, with $l_P = l_p$. The black curve, particularly, corresponds to the actual separation of the Holometer interferometers, $|\mathbf{x}'| = 0.91$ m. As can be seen, at this small separation scale, dual-interferometer detection reduces the measurable exotic noise power by only 3% relative to the single-interferometer PSD.

2.5 Previous Experimental Limits

There are several long-baseline gravitational wave detectors with sensitivity to exotic transverse spatial fluctuations. Of these, the GEO600 experiment has the highest sensitivity to this effect[3]. The GEO600 detector is a Michelson-class interferometer comprised of two straight 600-m arms oriented at 90°, with each arm folded once to double the distance traversed by the light. Like the Holometer, the straight-arm design of GEO600 makes it sensitive to purely the shear-field effect due to the near purely radial propagation of its light.

3. Although the LIGO interferometers have a longer baseline, > 99% of the injected optical power promptly reflects back to the beamsplitter from the Fabry-Perot arm cavities, making the effective distance over which spatial fluctuations accumulate far shorter.

No gravitational wave detector is sensitive to the rotational-field effect, as straight radial arms are common to the design of all of these interferometers.

Each pass through the GEO600 interferometer makes an identical measurement as the Holometer, but over a longer baseline $L = 600$ m. The geometrical coupling of the double-pass measurement is simply a repetition of the single-pass coupling (eq. 2.37) for a second measurement duration,

$$\Theta(\tau) \equiv \begin{cases} 1\,, & \tau < T \\ -1\,, & T \leq \tau < 2T \\ 1\,, & 2T \leq \tau < 3T \\ -1\,, & \tau \geq 3T \end{cases}, \tag{2.43}$$

with the propagation time domain extended to $\tau \in [0,\, 4T]$. The PSD of the effect is obtained by numerically evaluating eq. 2.32 for a shear field under the geometrical coupling of eq. 2.43. As a function of normalization, the predicted band-integrated shear noise power is

$$\begin{aligned} P_{\text{model}}(l_P) &= \int_{f_1}^{f_2} \widetilde{C_{SS}}(f \,|\, l_P)\, df \\ &= \frac{l_P}{l_p} \int_{f_1}^{f_2} \widetilde{C_{SS}}\left(f \,|\, l_p\right) df\,, \end{aligned} \tag{2.44}$$

where the limits of integration will be chosen to match the GEO600 detection band.

The GEO600 Collaboration has reported a detection of a broadband strain noise of unknown origin, in excess of conventional noise model estimates, in the 100 Hz-6 kHz band. Fig. 2.7, reproduced from a 2014 seminar given by Hartmut Grote of the GEO600 Collaboration, shows the reported noise excess. The fit of an additional, flat noise term to the

GEO600 strain sensitivity implies an unexplained strain noise excess of

$$\widehat{h} \;=\; 1.25 \times 10^{-22}\ \mathrm{Hz}^{-1/2}\ , \tag{2.45}$$

as indicated by the green curve. Whether of exotic or conventional origin, it places the strongest existing upper limit on possible normalizations of a stochastic shear effect.

This upper limit is obtained by requiring that the in-band noise power of exotic shear fluctuations not exceed the unexplained component of the GEO600 noise budget. The noise power of this component, integrated over the 100 Hz-6 kHz band, is

$$\begin{aligned} \widehat{P}_{\mathrm{excess}} &= \int_{f_1}^{f_2} \left(4L\,\widehat{h}\right)^2 df \\ &= 16L^2\,(f_2 - f_1)\,\widehat{h}^{\,2}\ . \end{aligned} \tag{2.46}$$

In the above integrand, one factor of two accounts for the doubling of strain sensitivity to gravitational waves achieved through the arm folding and a second factor of two converts from one-way to round-trip optical phase convention. The condition

$$P_{\mathrm{model}}\,(l_P) \;\leq\; \widehat{P}_{\mathrm{excess}} \tag{2.47}$$

then implies an upper limit of

$$\begin{aligned} l_P &\leq \frac{16\,l_p\,L^2\,\widehat{h}^{\,2}}{f_2 - f_1}\left(\int_{f_1}^{f_2} \widetilde{C_{SS}}\,(f\,|\,l_p)\,df\right)^{-1} \\ &\leq 240\ l_p \end{aligned} \tag{2.48}$$

on the shear effect normalization. With its sensitivity to *spatially-correlated* position displacements, the Holometer measurement will conclusively determine whether the unexplained noise in GEO600 is attributable to exotic transverse position uncertainty.

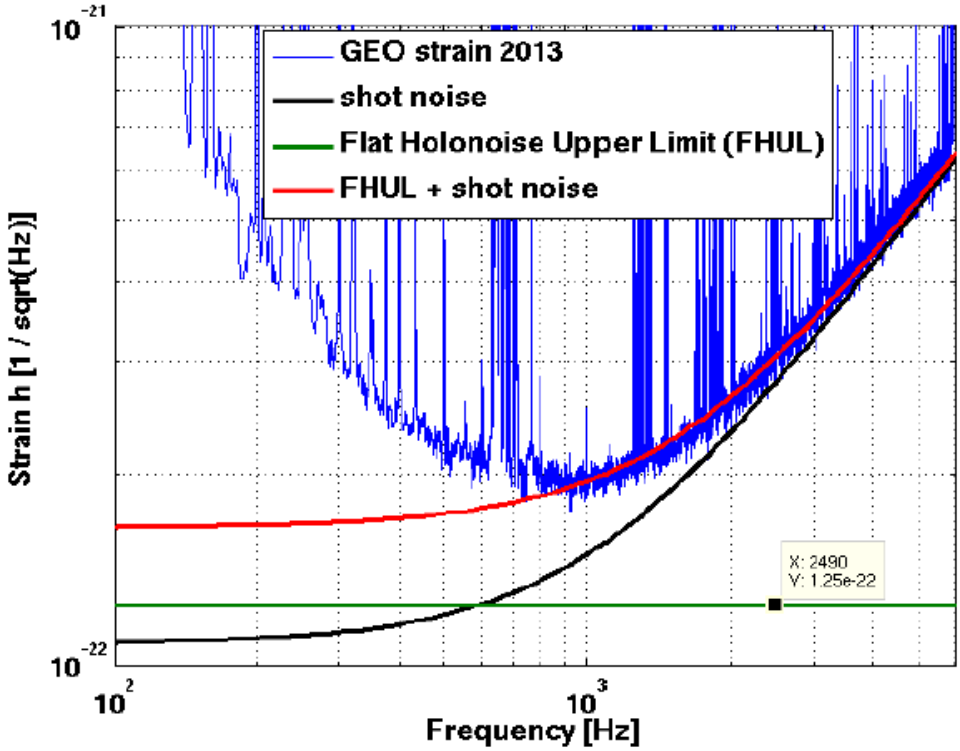

Figure 2.7: GEO600 upper limit on flat unexplained strain noise. Reproduced from a seminar given by H. Grote of the GEO600 collaboration at Fermi National Accelerator Laboratory (April 7, 2014).

CHAPTER 3

THE HOLOMETER DETECTOR

The Holometer detector consists of two identical, independent 40-m interferometers whose outputs are read out synchronously. For historical reasons, the two interferometers will be referred to as the "L" and "T." Although the interferometers are based largely on technology developed for gravitational wave detectors, their design differs in several important respects: a smaller apparatus, higher sampling frequency and bandwidth, and simpler mechanical systems. In other respects, such as optical cavity design and control systems, their design is very similar. The following sections describe the design and implementation of each component Holometer subsystem.

3.1 Physical Environment

The Holometer is situated in an old meson tunnel of the Tevatron at Fermi National Accelerator Laboratory (FNAL). These tunnels are constructed of pre-fabricated reinforced concrete slabs set on a gravel bed at ground level. The floor consists of 3.7-m by 6-m concrete sections 150 mm thick placed end-to-end along the length of the tunnel. The walls are U-shaped concrete sections 2.4 m wide placed on top of the floor sections. The entire structure is buried under a mound of gravel and dirt of minimum thickness 2.5 m, which acted as a radiation shield. The tunnel structure was put in place more than 30 years ago and is now a very stable platform. Fig. 3.1 shows the Holometer site and its location on the FNAL property.

One arm of each interferometer is installed inside the tunnel, bolted directly to the floor, while the other arm extends perpendicular out of the tunnel above ground level. The end stations of the perpendicular arms are bolted onto a newly-constructed 1.9-m by 1.9-m reinforced concrete slab 450 mm thick. The slab rests on three reinforced concrete pillars

Figure 3.1: The Holometer site at Fermi National Accelerator Laboratory (satellite images via Google Maps). The interferometers are situated in an old meson tunnel of the Tevatron.

0.6 m in diameter which extend 1.8 m into the ground. A climate-controlled enclosure constructed over the slab, pictured in Fig. 3.2, protects the electronics and end station vessels.

3.2 Optical System

The experiment is optimized for sensitivity to high-frequency position fluctuations in two co-located, but independent, Michelson interferometers. As demonstrated in §2.3.1, each interferometer converts position fluctuations into a changing rate of photons counted by a detector at the AS port. An exotic effect would appear as a correlated deviation from the mean incidence rate over a time scale $2L/c$. The measurement is limited by Poisson statistics, or "shot noise," in estimating the mean incidence rate. For a mean photon count rate of $\dot{N}$, the optical phase error due to shot noise is $1/\sqrt{\dot{N}}$. The shot noise-limited sensitivity to

Figure 3.2: Climate-controlled enclosure housing the end stations of the above-ground interferometer arms.

differential position displacement is then

$$\begin{aligned} PSD_{\text{shot}} &= \left(\frac{\lambda}{2\pi}\right)^2 \frac{1}{\dot{N}} \\ &= \frac{\lambda c h}{4\pi^2} \frac{1}{P_{\text{BS}}} , \end{aligned} \tag{3.1}$$

where P_{BS} is the power incident on the beamsplitter. Increasing the interferometer power can thus significantly improve the sensitivity of the measurement, which the Holometer achieves through the addition of a power recycling mirror (PRM) at each interferometer input. The PRM forms an effective Fabry-Perot cavity with the end mirrors, resulting in a resonant enhancement of the optical beam power inside the instrument. With an operating power of $P_{\text{BS}} > 1$ kW, each interferometer achieves a shot limited-sensitivity of approximately 5×10^{-36} m^2/Hz.

3.2.1 Power-Recycling Cavity

The power-recycled interferometers are designed to be nearly confocal resonators, folded by the 90° beamsplitter so that each arm forms a flat-curved cavity. The interferometers have an arm length of $L = 39.2$ m, and the end mirror radius of curvature, $R = 75 \pm 1$ m, is chosen in order to place a waist with radius

$$w_0 = \sqrt{\frac{\lambda}{2\pi}\sqrt{2L(2R-2L)}} \approx 3.57 \text{ mm} \tag{3.2}$$

at the position of the PRM. The end mirrors are each located one Rayleigh range away, where the beam half-width has grown to $w_1 \approx 5$ mm. The deviation from a pure confocal configuration satisfies the resonator stability criterion, $R < 2L$.

To prevent inadvertent coherent co-excitations, the cavity is also designed so as to avoid co-resonances of higher-order cavity modes (HOM) with the fundamental mode. In terms of the Rayleigh range,

$$Z_R \equiv \frac{\pi w_0^2}{\lambda} = 37.6 \text{ m} , \tag{3.3}$$

the Gouy phase shift is

$$\phi_{\text{Gouy}} \equiv \arctan\left(\frac{L}{Z_R}\right) = 46.2^\circ \tag{3.4}$$

for one-way propagation down the interferometer arm. Any Hermite-Gauss mode, H_{mn}, attains a round-trip phase excess of

$$\phi_{\text{mn}} = 2(1+m+n)\,\phi_{\text{Gouy}} \tag{3.5}$$

compared to the propagation of unfocused spherical wavefronts. If $2(m+n)\phi_{\text{Gouy}}$ is an integer multiple of 360°, then the H_{mn} mode will be co-resonant with the fundamental H_{00} mode. With 92° of phase separation between each mode order, the lower-order HOM resonances are well-separated from the fundamental. The fourth-order mode wraps back

near the fundamental, but at a sufficient phase separation of 10°.

As the injected laser beam is not perfectly matched to the pure Gaussian H_{00} mode of the cavity, the continuum of noise sidebands present in the laser populate all of the HOM resonances, as well as the the entire comb of H_{00} resonances at integer multiples of the cavity free spectral range. From spectral measurements of the cavity transmission, the estimated interferometer arm length is confirmed by the frequency of the free spectral range, and the predicted Gouy phase is confirmed by the frequencies of the HOM noise-leakage peaks. Under this cavity design, the Holometer interferometers have achieved resonant enhancement factors as high as 3000 (3 kW of storage power).

3.2.2 Antisymmetric-Port Response

Previously, §2.3.1 derived the response of a single-pass Michelson interferometer to time-varying path-length displacements,

$$P_{\mathrm{AS}}(t) \;=\; P_0 \sin^2\left(\frac{\pi}{\lambda}\left(S_0 + S(t)\right)\right) \;, \tag{3.6}$$

where P_0 is the optical power incident on the beamsplitter. Under power recycling, the optical power itself depends on the time-varying length displacements of the cavity, as represented by the promotion $P_0 \to P_{\mathrm{cav}}\left(S(t)\right)$. The dynamics of this dependence are complicated by the storage time of the power-recycling cavity, which sets the time scale on which the optical power can adjust to new path-length displacements.

The storage time imposes a bandwidth limit of approximately 350 Hz on the cavity response. At frequencies $\lesssim$ 350 Hz, the AS-port response to path-length displacements reflects both a change in the Michelson fringe interference offset and in the cavity storage power. At frequencies >> 350 Hz, however, path-length displacements occur on a shorter time scale than the cavity can respond. In this limit, the power-recycled interferometer

responds equivalently to a single-pass Michelson interferometer with optical power $P_0 = \langle P_{\text{cav}}(S(t))\rangle$.

Fig. 3.3 shows the numerically-calculated transfer function of the power-recycled interferometers at several operating offsets, S_0. This transfer function represents the conversion of path-length displacement into optical power at the AS port by the dynamics of the instrument. The Holometer operates at an offset of approximately $S_0 = 1$ nm, and all science and calibration signals are measured at frequencies ≥ 1 kHz. As shown in the figure, at this offset, the deviation from a single-pass Michelson response is $\leq 1\%$ above 1 kHz. For calibration purposes (see §4.2), the instrument can thus be modeled as an equivalent high-power, single-pass interferometer, neglecting this small cavity correction.

3.2.3 Optical Elements

The following sections detail the optics chosen to achieve large resonant enhancement, as well as the infrastructure implemented to prepare, monitor, and control the optical state.

Injection Optics

Prior to injection, the laser beam of each interferometer is conditioned to match the mode of the cavity at the power-recycling mirror. Each interferometer has a table of optics dedicated to beam preparation and to locking the laser frequency to the cavity (see §3.4.4). Fig. 3.4 shows the principal optical components of the laser launch design. Immediately before injection, each table diverts approximately 1% of its prepared beam power onto low-power detectors continuously monitoring the phase and amplitude noise of the laser.

From the table, each injection beam is relayed to the input window of its central vacuum vessel by a set of three flat mirrors mounted to the wall, through 75-mm-diameter aluminum pipes in air. All of the mirrors are manually adjustable, and the last of the three is remotely controlled by dual-axis pico-motor actuation. Two digitally-controlled steering mirrors on

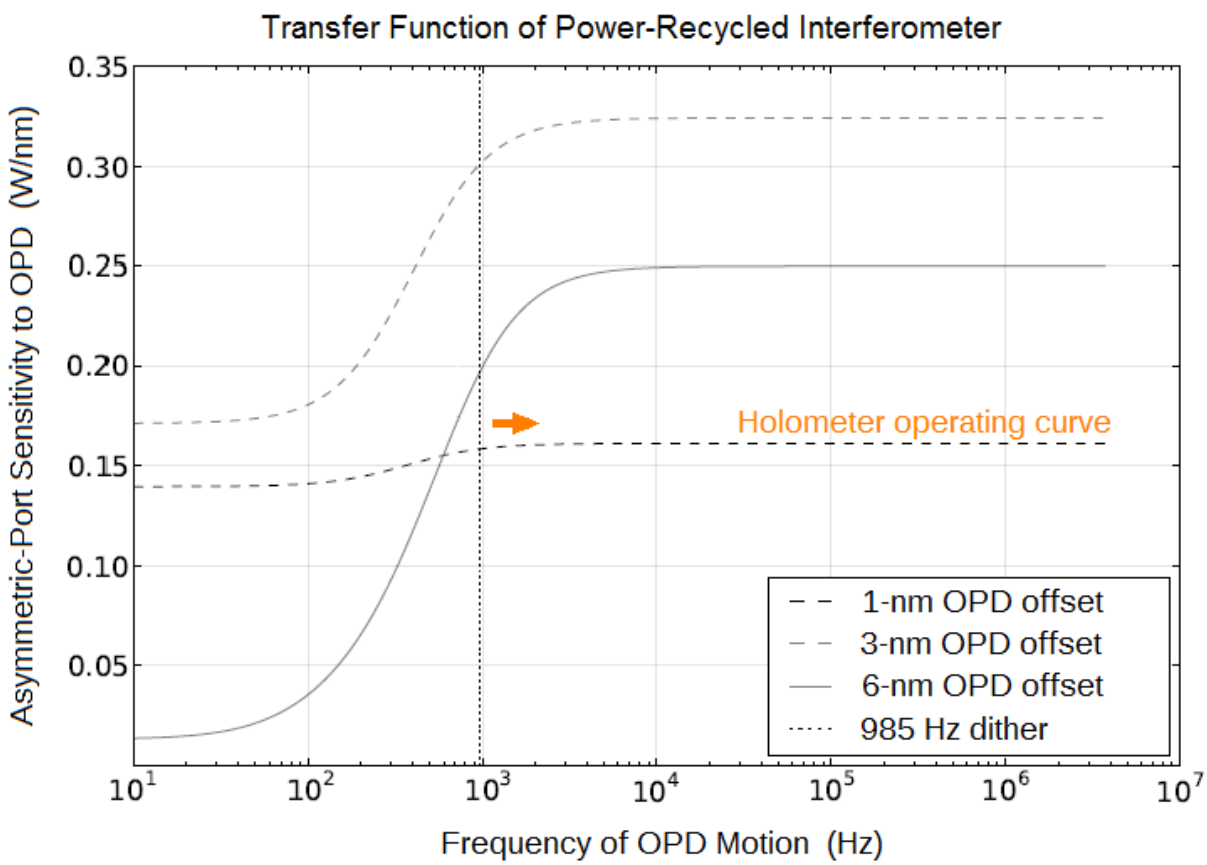

Figure 3.3: Transfer function of the power-recycled interferometers, shown at several OPD operating offsets. A transition between the AC and DC asymptotic limits occurs at the cavity bandwidth, 350 Hz. At frequencies below 350 Hz, the response reflects changes in both the interference offset and the cavity power. At frequencies above the cavity bandwidth, the power-recycled interferometer behaves as a single-pass interferometer of equivalent power. The Holometer operates firmly in the AC asymptotic regime of the 1-nm offset curve.

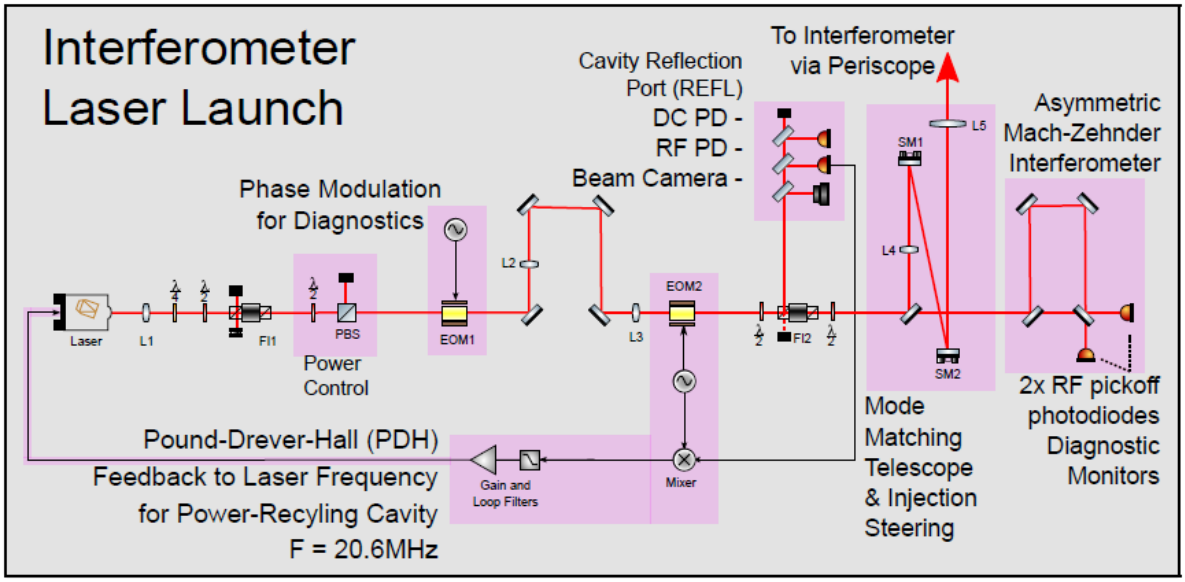

Figure 3.4: Laser launch table for preparation of the injection beam. These optics perform three key functions: mode matching to the cavity, locking the laser frequency to the cavity (see §3.4.4), and continuously monitoring the phase and amplitude noise of the laser.

each table provide fine control for aligning the position and angle of the injection beam to the optical axis defined by the power-recycling mirror.

Power Recycling Mirror

A partially transmissive 2-inch-diameter by 1/2-inch-thick power recycling mirror (PRM) is placed at the symmetric port of each interferometer to intercept the constructively-interfering beam and reflect it back into the instrument. This mirror functions as the input coupling mirror of an effective Fabry-Perot cavity formed with the end mirrors. The PRMs are fabricated from Corning 7980 0A low-inclusion fused-silica substrates and polished to sub-nm flatness by Coastline Optics.

The PRMs are coated by Advanced Thin Films to 80% of their diameter by ion beam sputtering. Spectrophotometric measurements at 1064 nm found the transmission on the reflective surface to be 985.8 ppm and the reflection on the anti-reflective surface to be 13 ppm. Because the transmission through the PRM is the dominant loss in the power-

recycled interferometer, the instrument is operated as a slightly over-coupled optical cavity.

Each PRM is mounted on a pico-motor-actuated stage, which provides digitally-controlled dual-axis steering. These actuators make slow adjustments used for aligning each cavity into a resonant configuration. In each interferometer, the PRM is mounted, along with the beamsplitter, on a common critically-damped two-stage vibration isolation platform with a resonant frequency of 10 Hz.

Beamsplitter

The beamsplitters are 3-inch-diameter by 1/2-inch-thick optics, polished to sub-nm flatness by Coastline Optics, with 5 arc-minutes of wedge. At 45° incidence, the beamsplitters intercept the Gaussian beam with an acceptance of more than five times the waist size. They are fabricated from high-purity, low-OH^--content Heraeus Suprasil 3001 substrate chosen for its low power absorption of 0.3 ± 0.2 ppm/cm, which reduces its susceptibility to thermal lensing effects. With 2.5 kW of storage power, the contrast defect due to thermal lensing in this substrate is expected to be only 10 ± 4 mW.

The beamsplitters are coated by Advanced Thin Films to 80% of their diameter by ion beam sputtering. The reflective surface has been measured by the coating vendor to have 49.991% transmission. The back surface is anti-reflection coated to 60 ppm. The absorption due to impurities in the coatings has been measured via photo-thermal common-path interferometry by Stanford Photo-Thermal Solutions. The absorption on the reflective coating was found to be 0.95 ppm uniformly across the surface of a witness sample, and the absorption on the anti-reflective coating was found to be 1.3 ppm.

Like the PRM, each beamsplitter is mounted on a pico-motor-actuated stage, providing digitally-controlled dual-axis steering for slow alignment adjustments. The beamsplitter is mounted together with the PRM on a common critically-damped two-stage vibration isolation platform with a resonant frequency of 10 Hz.

End Mirrors

The end mirrors are 2-inch-diameter by 1/2-inch-thick optics fabricated from Corning 7980 0A low-inclusion fused-silica substrates by Gooch & Housego. They are coated to 80% of their diameter by ion beam sputtering and have an effective radius of four times the beam size. This large mirror size is chosen to prevent losses in the tails of the Gaussian beams. Using an optical chopper and a lock-in amplifier, the reflective front coating has been measured to have sub-ppm transmission. The back surface is anti-reflection coated in order to allow the instantaneous cavity power to be monitored via the light leakage through the mirror.

In order to select optimally matched pairs of end mirrors, the mirror surface structures were analyzed using both Zernike and Laguerre-Gauss decomposition of metrology data from Zygo interferometry. The actual mirror maps were also used in FINESSE Monte Carlo simulations to determine the fraction of light scattered into higher order modes by mirror surface imperfections, which contributes to the contrast defect power. Pairs of mirrors were selected on the basis of matching in radius of curvature to < 10 cm and with all other surface deviations predicted to produce a contrast defect of < 20 ppm.

Each end mirror is clamped to an aluminum support ring that is driven by three Noliac SCMAP05-12mm piezo-electric transducers (PZTs). The PZTs provide three degrees of actuation over a total range of approximately 10 μm. These actuators are managed by the interferometer control system, and their performance is characterized is §3.4.1. The mirror is held by three clamps spaced 120° apart. To minimize mirror distortion from the clamping, a pad of indium, 5 mm in diameter and 20 μm thick, cushions each point of contact on both sides of the mirror.

Each end mirror-PZT assembly is connected to a 2-kg stainless steel reaction mass, which is itself connected to a supporting mount. The mount has two stages of vibration isolation. The first stage, installed only in the two end stations in the tunnel, where floor vibrations up to 100 Hz are present, isolates the entire mirror-mount structure from the ground with a

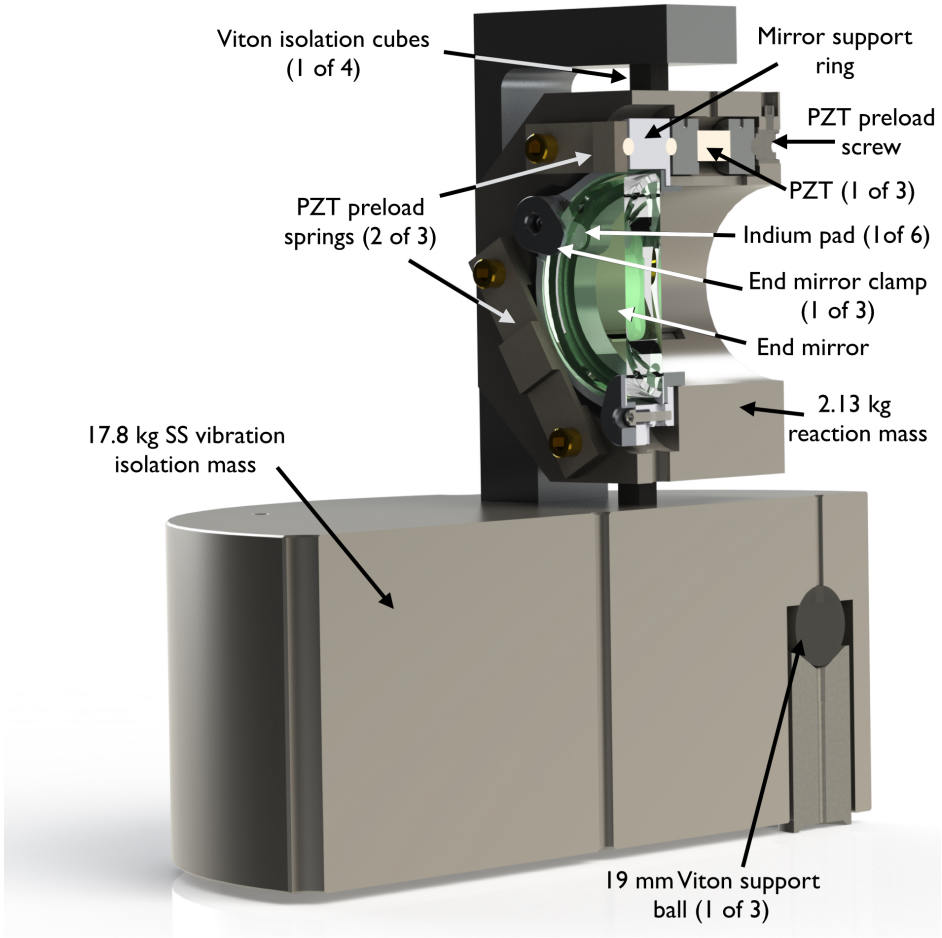

Figure 3.5: Sectioned view of the end mirror vibration isolation mount and PZT actuator. The lower stage passively isolates the entire structure from the ground, while the upper stage passively isolates the end mirror and reaction mass from the support mount.

damped resonance at 11 Hz and a quality factor near one at 37 C, which is maintained by a linear proportional thermostat system. It consists of an 18-kg mass mounted on three Viton balls, whose plane contains the center of mass of the isolated system to decouple rotational and translational motion. The second stage isolates the mirror and the reaction mass from the mount, so that PZT drive motion pushes purely on the reaction mass. This is done to avoid exciting mechanical resonances in the rest of the mount. Fig. 3.5 shows this design.

End Mirror Transmission Optics

The light transmitted through each end mirror is split by a beamsplitter and directed onto a photodiode and a video camera. Each optical set-up is mounted on a 1-foot by 1-foot breadboard bolted directly to the back of its end station cube. The photodiode signal provides a control-signal proxy for the cavity storage power, and the video camera is used to view the position and mode quality of the beam inside the cavity. Manual irises placed between the end mirrors and transmission cameras are used to center the beams on the end mirrors during alignment procedures.

Output Optics

The 3.6-mm interference beam exiting each interferometer at its AS port is telescopically focused down to 1-mm size and directed onto a set of high-frequency signal detectors, low-frequency control system detectors, and a video feed. Each optical set-up is mounted on a 1-foot by 2-foot breadboard supported by 8-inch-diameter tubes bolted directly to the floor. Of the approximately 200 mW of power in the interfering beam, 1% is picked off for control and monitoring purposes. The remaining 99% is directed onto two New Focus 1811 detectors which have been modified to absorb high power (see §3.5.1). These detectors are protected by a digitally-controlled shutter that is triggered to open when the exiting power is at a safe level for the detectors.

The low-power control pickoff is split by a beamsplitter and directed onto a New Focus 2903 quadrant photodetector (QPD), which provides the AS-port control signals, and a video camera, used for monitoring the mode quality and angular alignment of the interfering beam. As the QPD has an adjustable gain, it is used for both high-power interferometer operations (power-recycling mode) and low-power operations (single-pass Michelson mode). Its 2-mm collecting area is divided into four equal-area sections. The sum of the entire receiving area is used to for the DC-readout control signal, while the differences between the quadrants are

used as error signals for the angular alignment control loop (see §3.4.4).

3.3 Vacuum System

The PRM, beamsplitter, and end mirrors of each interferometer are housed under vacuum. This is done to reduce phase noise caused by fluctuations in the index of refraction of air, which would otherwise dominate the noise budget, as well as to avoid burning hydrocarbons onto the optical surfaces exposed to high power. A separate, identical vacuum system is implemented for each interferometer. As every hydrocarbon deposit of 0.1-monolayer thickness on an optical surface contributes 2 ppm of absorption losses, achieving an acceptably slow rate of hydrocarbon deposition requires that the optics be exposed to partial pressures $< 10^{-12}$ Torr. The dual vacuum systems of the Holometer are designed, and have been demonstrated, to exceed this benchmark. A sketch of the vacuum layout is shown in Fig. 3.6.

Each 40-m interferometer arm is constructed of 12 sections of 10-foot-long stainless steel tube with 6-inch ConFlat flanges at the ends. The stainless steel tubes are supported on rollers every 10 feet. A 100-l/s ion pump is installed at the end of each arm via a reducing tee with a 2 3/4-inch ConFlat flange. Another reducing tee with a 4-inch ConFlat flange and gate valve is installed at the center of each arm, where a 4-inch turbo pump can be attached during baking to extract hydrocarbons vaporized from the tube walls. The arm tubing is covered with two layers of fiberglass pipe insulation, separated by a layer of aluminized mylar, for a total thickness of 5 inches. It is pictured in Fig. 3.7. The insulation allows the arm tube to heat to a temperature of 150 C during baking. After assembly, each arm was baked for a minimum of three days by applying 100 A of current through the tube with a turbo pump operating at the center.

In order to prevent multi-path interference from light reflected and scattered by the inner walls of the vacuum tubes, a set of conical baffles is installed in each of the arms, spaced such

Figure 3.6: Illustration of the dual vacuum systems housing the high-power interferometer optics (not to scale). Each interferometer vacuum system has 40-m arms, with 0.91 m of separation between the two beamsplitters.

that neither beamsplitter nor end mirror has a direct line of sight to the tube wall. Each baffle is a truncated cone (half-angle 35°) with an outside diameter of 150 mm, matching the inner diameter of the vacuum tube, and an inside diameter of 110 mm. The baffles are constructed from sheets of thin stainless steel spot-welded into a cone that fits snugly into the vacuum tubes. Prior to assembling the 40-m arms, the baffles were pushed into the individual 10-foot tube segments using a long pole. Each arm contains a total of 31 baffles spaced in a geometrical series derived from the design in the LIGO interferometers[6]. Starting from each end of the arm and working inwards, the first 30 baffles are positioned a distance $l_n = \alpha^n l_0,\ n = 0, 1, \ldots, 14$, from each end, where $\alpha = 1.28$ and $l_0 = 0.5$ m. The final baffle is positioned at the center of the arm.

The central vacuum vessels, each housing the beamsplitter and power recycling mirror of one interferometer, are custom-built stainless steel cylinders 600 mm in diameter and 330 mm in height, with a bottom plate 25 mm thick. Each vessel is closed off with a rubber

Figure 3.7: Installation of fiberglass pipe insulation on the interferometer arms. The reducing tee attachment for a turbo pump is visible at the center of the arm.

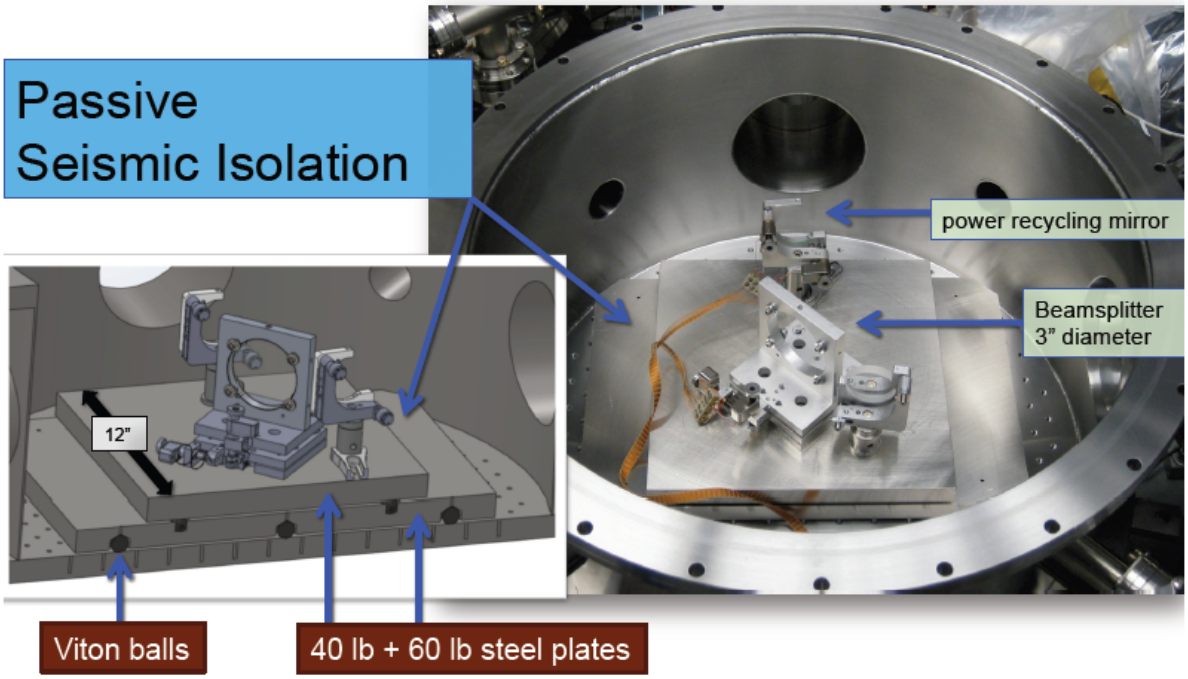

Figure 3.8: Central vacuum vessel housing the beamsplitter and PRM. The inset shows the common two-stage vibration isolation platform.

o-ring-sealed top plate, also 25 mm thick. The positions of the central vacuum vessels are not adjustable. Each vessel stands on three aluminum tubes 200 mm in diameter and 13 mm thick. Under each tube is a 150-mm by 150-mm by 6-mm aluminum support plate, which was positioned on the somewhat uneven floor with a thick layer of epoxy underneath. The tubes and vacuum vessel were then placed on top and the epoxy allowed to harden. Three dog clamps drilled into the concrete floor push down on the bottom plate of the vacuum vessel near the support tube underneath. The lowest horizontal resonance of the mount is calculated to be > 200 Hz, based on the assumed stiffness of the concrete. One of the central vessels is shown in Fig. 3.8.

The end station vacuum vessels, each housing an end mirror, are 10-inch 6-way stainless steel ConFlat cubes purchased from Lesker. A custom-built bottom plate mounted to the bottom ConFlat flange of the cube extends outside the cube itself and serves as its support

base. The cubes stand on three aluminum tubes similar to those of the central vessel, but each of these additionally has a top plate with a 2-mm boss in the center. When the cube is placed on top of the cylinders, these bosses act as compliant elements. As with the central vessel, the the cubes are secured to the floor with dog clamps drilled into the concrete. In this case, however, the compliance of the supporting tubes permits small mechanical adjustments of the pitch and roll orientations of the end station cube

The end station cubes are mechanically isolated from the arm tubes by stainless steel bellows, whose length can change by up to 90 mm to accommodate thermal expansion. The cubes are supported against air pressure on the evacuated bellows by two stainless steel stands bolted to the floor, one on each side of the tube. Each support stand is outfitted with a 5/8-inch 80-threads/inch adjustment screw which contacts the arm tube flange at the center height of the tube. These screws provide for mechanical adjustment of the end station position and orientation. Adjusting both screws commonly lengthens/shortens the bellows, and thus the arm length, and adjusting them differentially rotates the yaw-axis of the end mirror. A 150-mm diameter gate valve, installed between the bellows and cube, permits venting the end station vessel for access to the optics without venting the tube or releasing the tension on the bellows. Fig. 3.9 shows the end station design.

3.4 Interferometer Control System

A sophisticated control system is required to maintain stable linear operation of the power-recycled interferometers and to monitor their performance. This system enforces an interdependent set of optical resonance conditions which allow the instrument to both store high power and achieve interference. The following sections describe the respective hardware and software implementations of the system. The final section then discusses each optical resonance condition and the feedback control loop which enforces it.

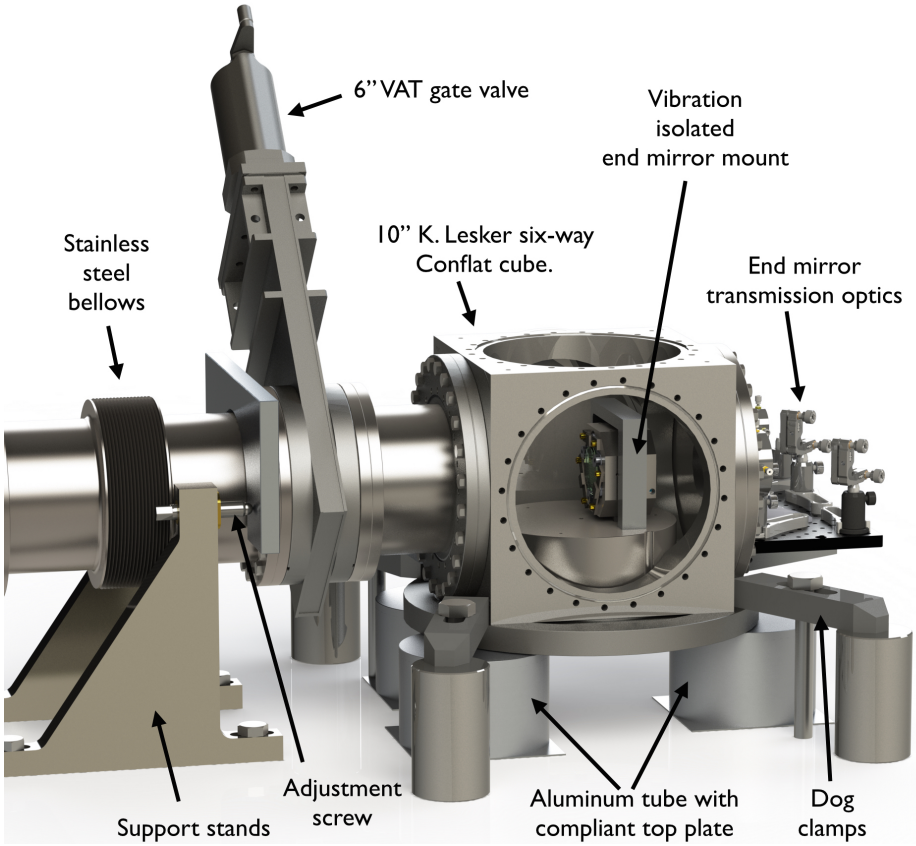

Figure 3.9: End station vacuum assembly with top and side flanges removed. Mechanical adjustment of the end station orientation is provided by the two adjustment screws and the three dog clamps.

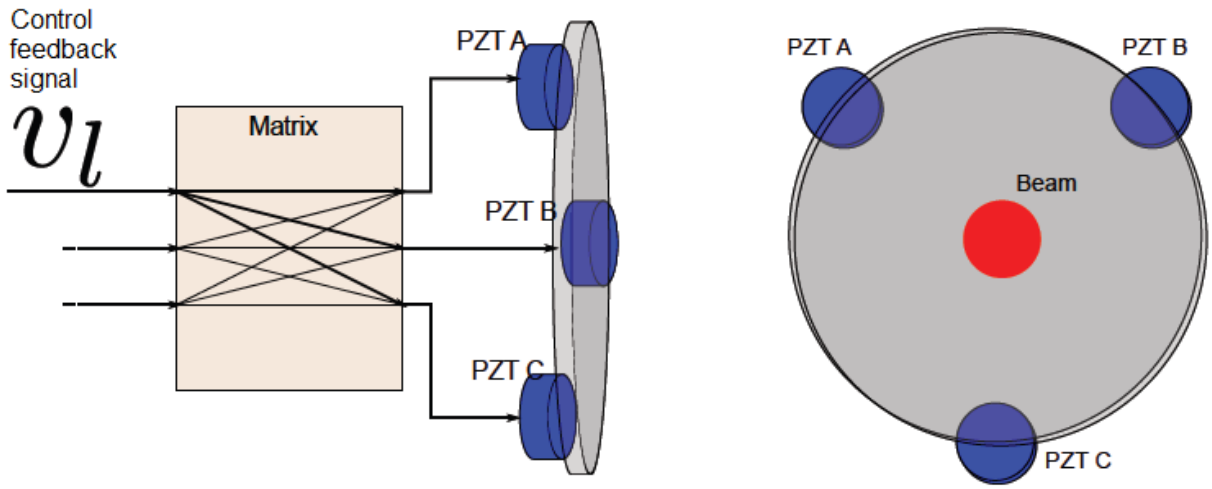

Figure 3.10: PZT-actuated end mirror drive. Each mirror is connected to three PZT actuators capable of moving 10 μm in the direction perpendicular to the reflective surface. The left panel shows a side view and the right panel shows a face-on view of the end mirror.

3.4.1 Piezoelectric Drive System

Each end mirror is actuated with three piezoelectric drivers capable of moving 10 μm in the direction perpendicular to the reflective surface, as was discussed in §3.2.3. The control system of each interferometer engages these actuators to maintain both the longitudinal operating offset and the angular alignment of the instrument. A sketch of the end mirror actuation is shown in Fig. 3.10. The PZTs can drive the end mirrors at frequencies of up to 1.5 kHz, at which point the system becomes limited by mechanical resonances. These resonances must be regularly re-measured, as they are observed to drift up to 10% in frequency over many weeks. The origin of this drift is believed to be thermal.

The mechanical resonances are measured by injecting white noise into the PZT drive of one end mirror while the differential-mode loop (see §3.4.4) is engaged. The transfer function of drive signal into physical motion of the mirror is inferred from the OPD response measured at the AS port. Digital notch filters are implemented in the differential-mode servo loop to eliminate the gain spikes at these resonant frequencies, allowing the total gain of the electronics to be increased. Fig. 3.11 shows the transfer function of the PZT drive of one

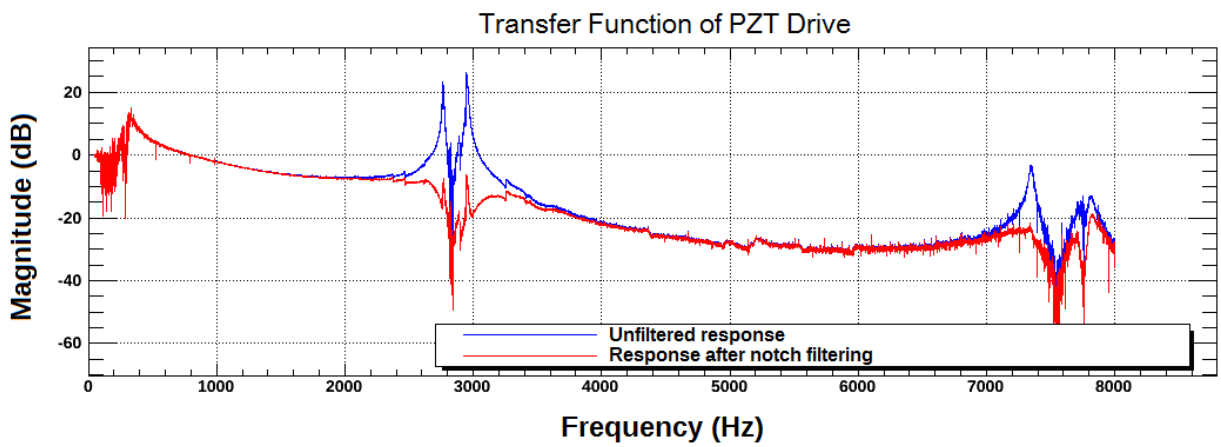

Figure 3.11: Transfer function of the PZT drive of one end mirror. Mechanical resonances are visible as spikes at 3 kHz and 7.5 kHz in the unfiltered response, represented by the blue curve. The red curve shows the response after digital notch filters centered on the resonant frequencies are applied.

end mirror, both with and without notch filtering.

3.4.2 Low-Speed Digitization Electronics

The signals of each interferometer are routed through two National Instruments (NI) PXI-7852R units. Control signals for the feedback loops (see §3.4.4) are routed exclusively through one unit, while auxiliary signals are routed through the other. Each unit consists of an FPGA controller interfaced to an 8-channel analog-to-digital converter (ADC) and an 8-channel digital-to-analog converter (DAC), both of which can sample up to 750 kHz at ±10 V over 16 bits. All of the I/O channels are conditioned in a custom filter buffer board which provides anti-alias filtering of the inputs and DAC-transient-noise filtering of the outputs. In both cases, this conditioning is implemented through analog 40-kHz, 2-pole Butterworth filters.

The input voltage range of each channel is individually adjustable through an AD8253 programmable-gain instrumentation amplifier in the filter buffer board. These analog gains are set using the digital outputs of the PXI-7852R, which powers the filter buffer board via

a DC-DC converter running from a 1-A 5-V line. The input signals can be amplified to ranges of ± 0.005 V through ± 5 V and the output signals can be attenuated to ranges of ± 0.01 V through ± 10 V. These adjustments allow convenient tradeoffs between dynamic range and bit-noise. Each amplifier has an input voltage noise of 10 nV/$\sqrt{\text{Hz}}$ near 1 Hz, with the exception of the gain 1 setting, where the input-referenced noise is 45 nV/$\sqrt{\text{Hz}}$. The analog inputs are differential with common-mode rejection of over 60 dB below 40 kHz, except when used at the gain 2 setting.

3.4.3 Controls Software

The FPGA controllers running the control software function as a micro-controller pipeline feeding the eight input signals through stages of matrix elements and filter banks. The main pipeline runs 1400 digital bi-quadratic filters (BQF) at 32 kHz. Before entering the pipeline, raw input signals to the 16-bit ADC are sampled at 128 kHz and immediately fed into a single BQF for anti-aliasing, then downsampled to 32 kHz. In the pipeline, the 32-kHz signals are routed through fully-configurable matrix elements, mixers, and dividers to synthesize more linear discriminants, demodulate dithers, and transform bases to/from detector, controller, and actuator degrees of freedom. BQF filter banks are applied to whiten signals, impose control policies, and invert actuator responses. Final feedback control signals are then upsampled back to 128 kHz and fed through a smoothing BQF filter, then output to a 16-bit DAC.

A python management script sets up the topology and interfacing between the interferometer state variables and the settings of the FPGA micro-control code. The interferometer state variables are centrally managed by Experimental Physics and Industrial Control System (EPICS) software. This state is accessed via custom Motif Editor and Display Manager (MEDM) software, which provides a graphical interface for users to monitor interferometer operation and adjust control parameters in real time. The MEDM software implements auto-

mated lock-acquisition scripts for high-level interferometer state management. These scripts interface through EPICS to repeat action sequences for the low-level control parameters.

3.4.4 Feedback Control Loops

Operating the power-recycled interferometers requires simultaneously satisfying three optical resonance conditions. Each of these conditions is enforced by a feedback control loop whose implementation is described in the following sections.

Common-Mode Loop

To achieve resonant enhancement of the optical power stored inside the interferometer, the first condition requires that the laser frequency coincide with a multiple of the free-spectral range (FSR) of the Fabry-Perot cavity formed by the power-recycling mirror and the end mirrors. This degree of freedom will be notated as CARM, for common-arm length, as the FSR of the cavity depends on the average of the two interferometer arm lengths. As such, CARM can be actuated by either the laser frequency or the common position of the end mirrors. The periodic FSR boundary sets 3.8 MHz of state-space, while the cavity bandwidth of 350 Hz sets the scale within this space to which the residual CARM noise must be suppressed.

CARM is controlled via a feedback loop which holds the laser frequency on resonance with the cavity. It will be referred to as the common-mode loop. The common-mode loop uses the Pound-Drever-Hall (PDH) method[2] to determine the mismatch between the cavity resonant frequency and the laser frequency. This technique applies phase sidebands to the injected beam by sinusoidally driving an electro-optic modulator. These sidebands are offset from the carrier and thus reflect from the cavity with a different transfer function. The mismatch of reflection coefficient phases produces an amplitude modulation in the reflected beam power, measured by a New Focus 1811 photodetector (unmodified), which is mixed

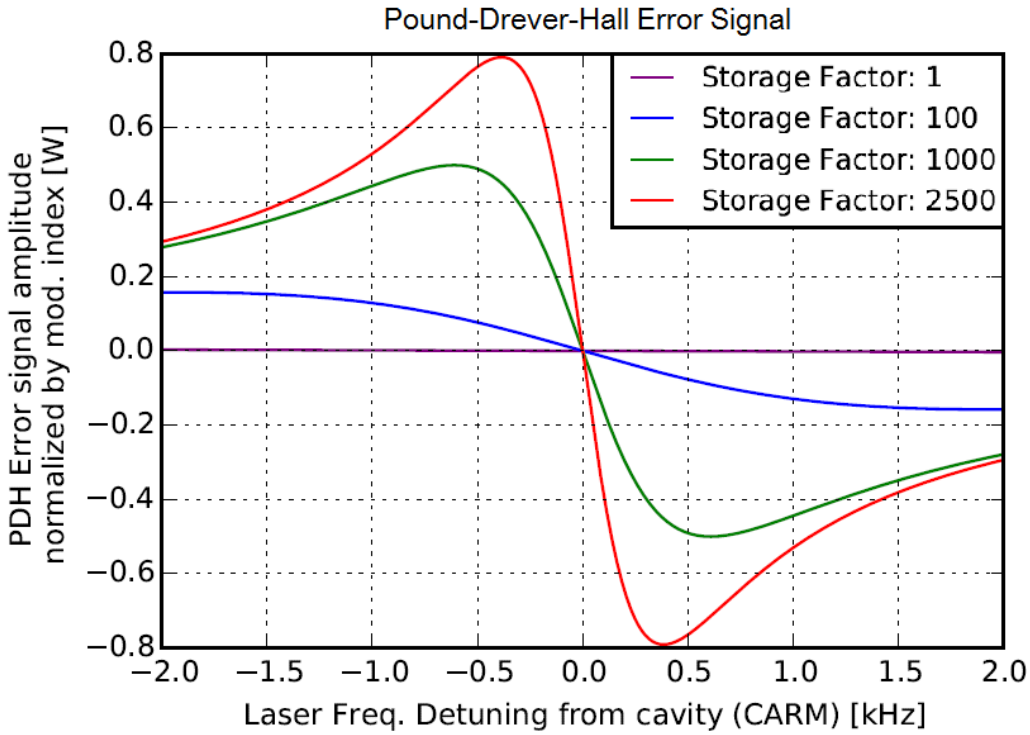

Figure 3.12: Common-mode loop error signals as a function of detuning. The width of the linear-response region around zero detuning is set by the cavity bandwidth (350 Hz), beyond which the response rolls off as a single pole.

in-phase with the drive signal to yield an error signal.

The characteristic shape of the PDH error signal is shown in Fig. 3.12. This signal has a zero where the laser frequency is perfectly matched to the cavity. At low detuning frequencies, the derivative, or gain, is proportional to the cavity storage factor, which itself depends on the OPD offset. The width of the linear-response region is the cavity bandwidth, beyond which the response rolls off as a single pole. The conditioned error signal is fed back to the laser to correct its frequency, which is controlled by an internal PZT-actuated mirror via a servo loop with a unity gain frequency (UGF) of 30 kHz.

Differential-Mode Loop

To maintain stable linear operation of the interferometer at the designed DC output power, the second condition requires that the OPD remain at a constant fringe offset, S_0 (see §2.3.1). This degree of freedom will be notated as DARM, for differential-arm length, as the fringe offset depends on the difference of the two interferometer arm lengths. DARM is actuated by the differential position of the end mirrors. The periodic fringe response sets 1064 nm of state-space. Within this space, the scale to which the residual DARM noise must be suppressed depends on the storage power. To maintain steady-state operation, the output power at the AS port cannot exceed the injected power. For a power storage of N times injection, this requirement imposes the constraint $\sin^2(2\pi S_0/\lambda) < 1/N$. This implies a suppression scale of $S_0 < 5$ nm for a nominal storage factor of $N = 1000$.

DARM is controlled via a feedback loop which holds the OPD at a constant fringe offset. It will be referred to as the differential-mode loop. The non-linearity of the interferometer signals requires that different sources of error signal be used during different stages of the lock acquisition. All of the interferometer signals depend on the total power stored, which is itself dependent on the cavity power loss through the AS port. This causes the output power signal at the AS port to have a characteristic, highly non-linear "volcano" shape arising from the product of the sinusoidal AS-port coupling and the Lorenzian power storage. Thus, instead of using the AS-port power directly, two composite error signals are constructed.

During initial lock acquisition, the error signal is taken as the ratio of input power (calibrated) to storage power (measured), and its upper range is clipped for storage powers < 10 W. The initial lock point is set for a storage power of 100 W, corresponding to a DARM offset of approximately 10 nm, and a 1-Hz single-pole filter is engaged. The digital gain is then increased over a 0.5-s interval to set a UGF of $f_{\mathrm{UGF}} = 500$ Hz at this initial lock point. The loop responds to the increasing gain by actuating DARM by more than one wavelength, ensuring that an interference fringe is crossed. To maintain linearity, the rate of gain increase

is chosen such that the ±10-nm linear operating region of this bounded, non-linear signal is crossed in a longer time than $1/f_{UGF}$. As the DARM offset sweeps near its operating point, the common-mode loop (see §3.4.4) acquires lock first, allowing the power storage signal used by the differential-mode loop to become well defined.

Once the initial lock point is acquired, the operating point is transitioned to 1 kW of stored power through incremental reductions of the DARM offset. During this transition, loop filters are engaged to achieve optimal suppression of the background seismic noise, shown in Fig. 3.13, which is dominated by strong mechanical noises below 100 Hz. This optimization includes an integrator removing all DC offset, strongly shaped filters at 1-100 Hz to suppress coherent mechanical noises, and notch filters to compensate for actuator resonances and prevent instabilities. The loop shaping is engaged in stages, as filter activations inject impulse transients. With filtering optimizations, the differential-mode loop has been demonstrated to suppress ground motion to an RMS scale of < 1 angstrom.

As the power gain of the cavity is increased, the cavity bandwidth decreases, slowing the response of the error signal. At storage powers > 500 W, a whitening filter must be engaged to compensate. The 1-kW lock point corresponds to critical coupling of the recycling cavity. Once it is reached, the control loop is transitioned to instead use the ratio of AS-port output power to storage power as the error signal. This signal is not bandwidth-limited, thus requiring no whitening, and it is insensitive to the mode-matching efficiency at injection. The only degree of freedom affecting this signal is the contrast defect due to mode mismatch and angular misalignment of the returning beams.

Angular Alignment Loop

Finally, to achieve maximal interference at the beamsplitter, a third condition requires the profiles of the two returning beams to tightly overlap. This degree of freedom will be notated as DANGLE, for differential angle, as the beam-profile overlap depends on the difference in

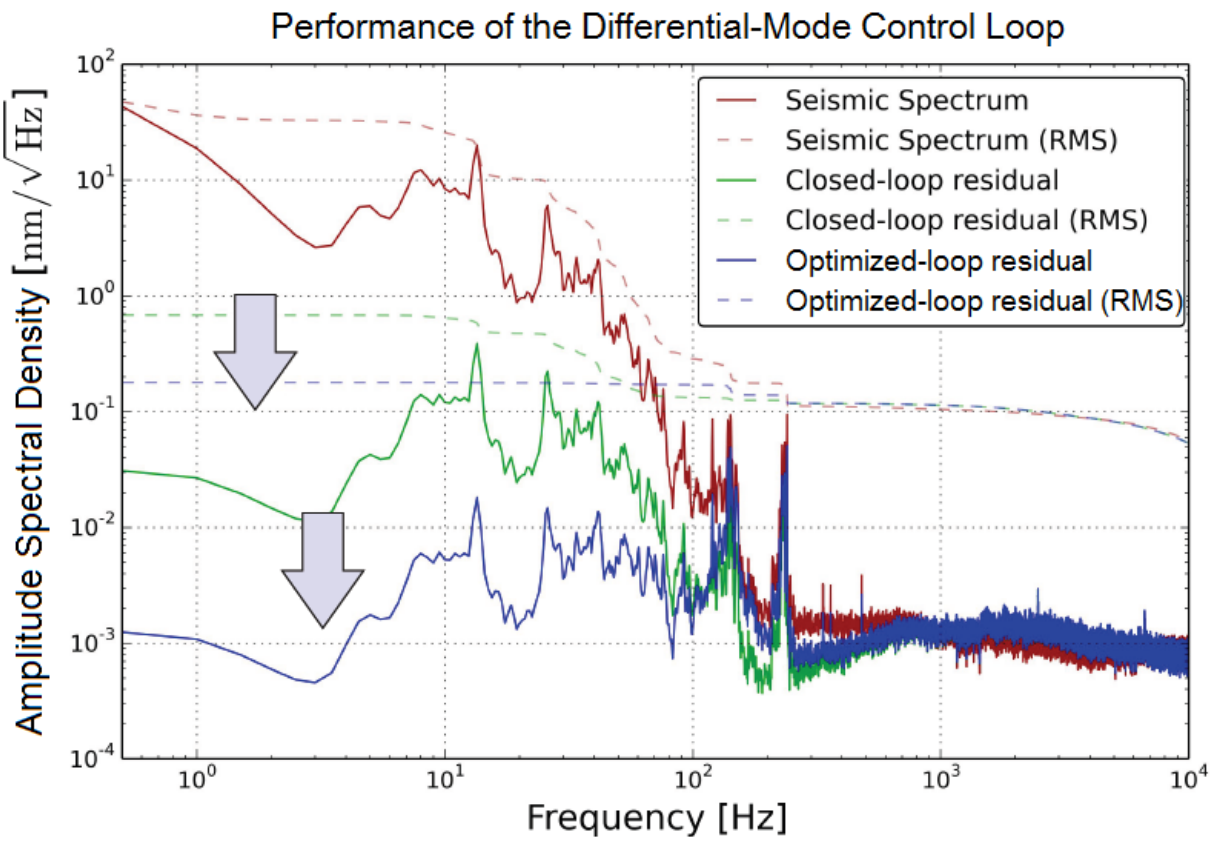

Figure 3.13: Suppression of the seismic noise background by the interferometer control system. The unsuppressed seismic noise spectrum is shown by the red curve. The fully optimized loop, represented by the blue curve, suppresses the ground motion to an RMS scale of < 1 angstrom.

angle of the two end mirrors. DANGLE is actuated by the differential angle of the end mirrors. DANGLE is controlled via a feedback loop which maintains the angular alignment state, and it will be referred to as the angular alignment loop. For small alignment defect power, < 10 ppm, the residual DANGLE noise of the returning beams must be suppressed to < 2 μrad.

The error signal is provided by a New Focus 2903 QPD at the AS port viewing the interference beam. Because of the small DARM operating offset, the QPD is sensitive to the beating of the Hermite-Gauss H_{01} mode of the misalignment with the H_{00} mode of the DC carrier. This loop is operated with a 100-Hz UGF. On one of the interferometers, the mode of the contrast defect light confused this beat signal such that X- and Y-angle misalignments both caused a degenerate QPD response. For this signal, the dithered H_{00} mode at a calibration-line frequency (see §4.3.1) is demodulated with the H_{01} mode of the misalignment to generate a decoupled error signal, but with a much-reduced bandwidth of 20 Hz and a lowered SNR compared to the direct QPD signal.

3.5 Data Acquisition System

Optical path-length measurements separated in time by less than a measurement duration (the round-trip light-crossing time) are predicted to exhibit exotic correlation, as demonstrated in §2.3.4. This imposes an upper limit of $2L/c = 267$ ns on the sampling time of the measurement, corresponding to a minimum sampling frequency of 3.75 MHz. Accordingly, the Holometer samples at 50 MHz to fully resolve the correlation time scale of the effect. This sampling rate is 100 times faster than that of gravitational wave detectors. It carries the considerable advantage of sampling a quantum-limited band, where the contributions of seismic, thermal, and mechanical noise are negligibly small.

However, it also incurs a raw data-stream rate of 5.8 TB/hr, making direct storage to disk untenable. The following sections describe the design and implementation of the Holometer

data acquisition system, a novel high-speed readout and data-processing pipeline. Through its real-time spectral analysis of the 50-MHz-sampled time signals, this system will be shown to reduce the Holometer data storage requirements to a manageable level (< 50 GB/hr) with no loss of information.

3.5.1 High-Speed, High-Power Photodetectors

At the nominal 2-kW operating power of each interferometer, the 50-ppm contrast defect produces an output power of 100 mW. In order to output approximately equal parts defect and signal-carrying power, the OPD offset, S_0 (see §2.3.1), is chosen to produce an output fringe power of 100 mW. The total output power of approximately 200 mW is then split by a beamsplitter and directed onto two New Focus 1811 photodetectors, each of which is modified to absorb 100 mW of DC power. The photodetectors provide both a low-frequency amplification circuit for interferometer control signals and a low-noise, high-frequency amplification circuit for radio-frequency science signals.

Each modified high-power detector contains a 2-mm InGaAs photodiode mounted in a TO-5 package. The photodiode is reverse-biased at 7 V with a LM340 voltage regulator. Low-value resistors are used along the biasing chain, causing a drop in the bias voltage to 6 V at the nominal 100-mA operating photocurrent. These values are chosen because space-charge-related inefficiencies become significant at bias voltages < 5 V and breakdown of the photodiode occurs at a bias voltage of 12 V. To dissipate the 0.6 W of Joule heating, the TO-5 package is wrapped with layer of silicone-based thermal putty and thermally contacted with an aluminum heat sink.

The low-frequency channel uses an INA128 instrumentation amplifier to measure the voltage drop across a 5-Ω resistor placed in series along the photodiode biasing circuit. It provides a gain of 10 V/A and a bandwidth of approximately 150 kHz, limited by the filtering circuitry along the photodiode biasing line. The noise level of the instrumentation amplifier

is 8 nV/$\sqrt{\text{Hz}}$ near 1 kHz, which is subdominant to the digitization noise of the control system.

The high-frequency channel uses a 7.2-kΩ Philips NE5210D transimpedance preamplifier AC-coupled to the photocurrent with a 500-pF capacitor. The 100-mA DC photocurrent is shunted to ground with a 5-μH inductance constructed from two high-current 10-μH inductors wired in parallel with opposite coil orientation to mitigate radio-frequency pickup. Combined with the 60-Ω input impedance of the preamplifier, the inductor and capacitor form a crossover filter centered near 1 MHz. This crossover frequency is chosen to suppress both large-amplitude, seismically-induced optical signals at several 100 kHz, which would otherwise saturate the preamplifier, and unstable amplifier oscillations at several 100 MHz. The input pin of the amplifier is also shorted to ground with 100 pF. This capacitance, combined with the 150-pF capacitance of the reverse-biased photodiode, forms an additional low-pass filter with the 60-Ω input impedance, with a pole frequency near 10 MHz. The high-frequency channel is shot-noise limited at a photocurrent of approximately 5 mA, with a fairly flat amplifier noise level of 0.6 pV/$\sqrt{\text{Hz}}$ from 1-10 MHz.

Transfer function measurements of both photodetector channels are presented in §4.4, where they are used for measurement calibration. A full schematic of the modified photodetector circuit is shown in Fig. 3.14.

3.5.2 High-Speed Digitization Electronics

The high-frequency photodetector readout and the environmental-monitoring signals of each interferometer are routed through two NI PXIe-5122 high-speed ADC units. Each dual-channel ADC can sample up to 100 MHz at ±10 V over 14 bits. One unit exclusively receives the high-frequency readout of the two photodetectors while the other receives the environmental-monitoring signals. Each input channel is anti-aliased via an analog 20 MHz 2-pole Bessel filter. The input voltage range of the two channels on each unit is adjustable

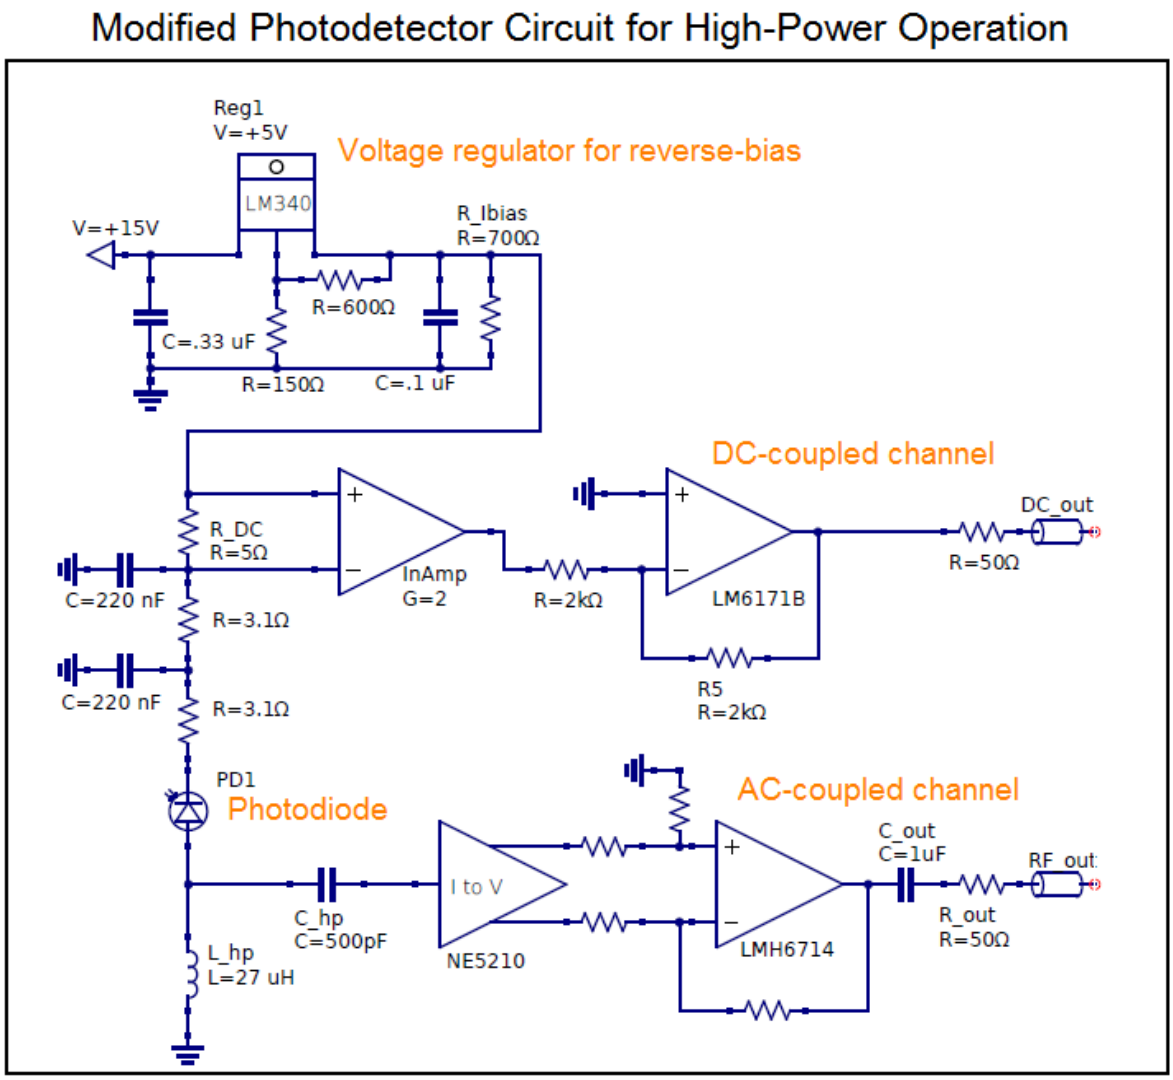

Figure 3.14: Schematic of the modified New Focus 1811 photodetector circuit for high-power (100 mW) operation. It contains both a low-frequency amplification circuit for interferometer control signals and a low-noise, high-frequency amplification circuit for radio-frequency science signals.

through a programmable-gain instrumentation amplifier. The input signals can be amplified to ranges of ± 0.1 V through ± 5 V, with a minimum input voltage noise of 2 nV/$\sqrt{\text{Hz}}$ at the highest gain setting.

Electrical isolation of the ADC units is a critical aspect of the *independent* dual-system design of the Holometer. Each ADC is housed in a separate NI PXIe-1082 chassis connected via a 40-m fiber optic cable to the PCIe backplane of a dedicated workstation, located in a separate building. The central computer receiving the digitized optical signals is the only link between the four chassis. Because direct clock-sharing violates the stringent electrical isolation required of the two detection systems, each ADC unit is synchronized to the other three by phase-locking to a common GPS-provided clocking signal. This signal is generated by an NI PXI-6683H timing card installed in each chassis, with each timing card connected to its own GPS receiver. A schematic is shown in Fig. 3.15. The interchannel decorrelation due to clocking phase noise, assessed in §4.5.3, is shown to be a small effect ($< 5\%$) and is taken into account in the calibration analysis.

3.5.3 Data Pipeline Software

The 32-core central computer runs a custom C++ script which functions as a high-throughput pipeline capable of processing data streams of up to 1.6 GB/s, or 5.8 TB/hr, in real time. The software interfaces with the low-level NI device drivers to provide high-level control of the digitization and synchronization hardware. Before entering the pipeline, the eight input signals to the 14-bit ADCs are anti-aliased via an analog filter and sampled at 100 MHz, then downsampled to 50 MHz. In the pipeline, the 50-MHz signals are divided into accumulation time intervals of a configurable duration, typically 1-2 s, and processed in parallel. The signals in each accumulation time interval are Fourier-transformed in batches of a configurable DFT length, N_{DFT}, via the FFTW algorithm. The frequency-space correlation matrix of each DFT batch is then computed. The signal data in each DFT are weighted by

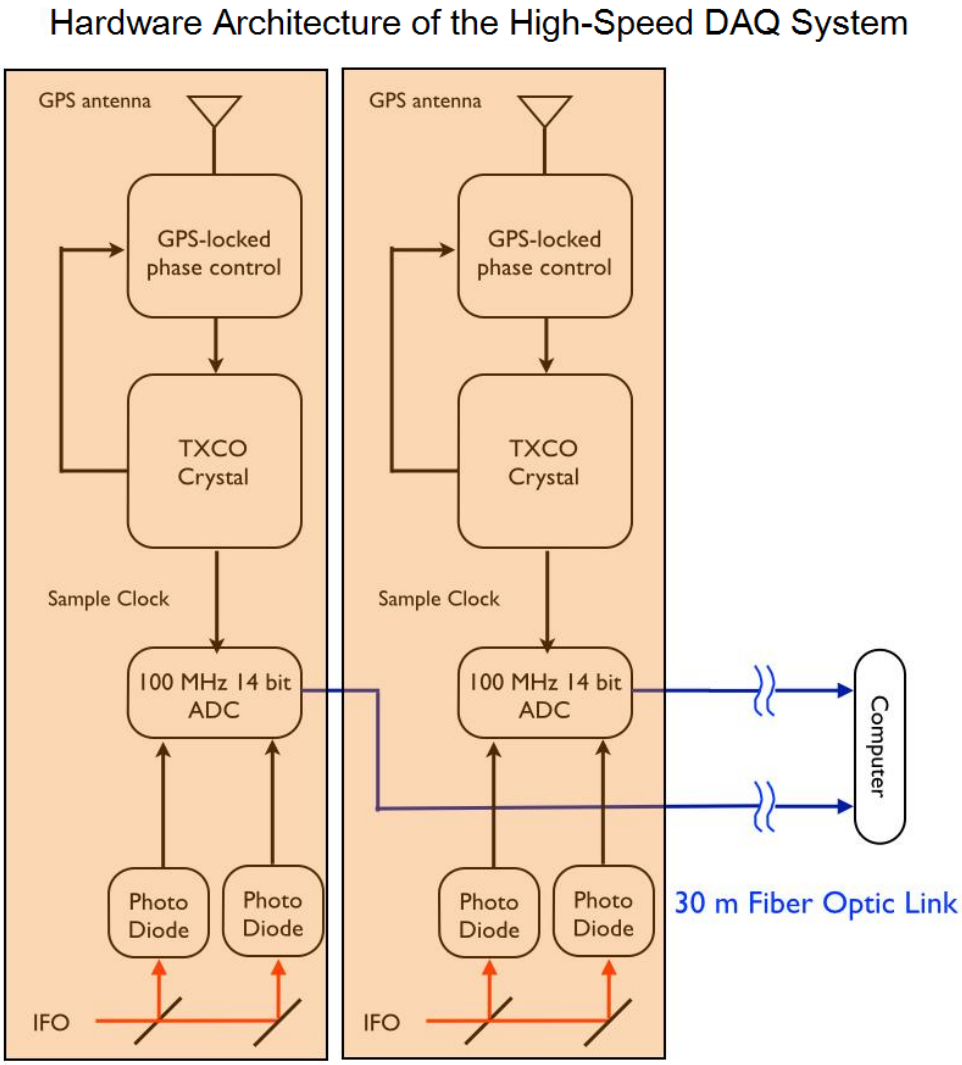

Figure 3.15: Schematic of the high-speed data acquisition system. Each electrically-isolated ADC unit is synchronized to the others by phase-locking its sample clock to a common GPS-provided clocking signal. For simplicity only two of the four PXI chassis are pictured.

a Hann window to minimize aliasing, and successive DFT batches are overlapped by 50% to recover the information lost by the windowing. Nominally, $N_{\mathrm{DFT}} = 2^{17}$ is chosen to achieve a spectral resolution of 570 Hz.

As the data are processed, the correlation matrix of each DFT batch is continually accumulated into a running average over the full accumulation time. The final time-averaged correlation matrix is then written to a 10-TB drive. As the averaged correlation matrix is only written once per accumulation time, this technique achieves a significant data compression factor of > 100 compared to writing the raw time samples directly to disk. Moreover, because the Holometer is searching for a stationary noise background, as opposed to a time-resolved signal, there is no loss of information traded off for this compression.

Achieving this high level of throughput required optimizations of the code at the level of the hardware architecture. During development, profiling revealed that performance was limited not by floating-point operations, but by memory bandwidth between the four non-uniform memory access (NUMA) nodes. The NUMA architecture groups processing units into nodes, each of which has a local memory cache shared by all of the units in the node. Any processing unit can access the memory of any node, but access times for non-local memory are significantly longer, as the memory must be transferred across one or more nodes via lower-bandwidth channels.

Standard memory allocation methods specify the size of the memory block but allow the operating system to assign its location. In this case, the worker threads performing the parallelized spectral analysis were being assigned to effectively random nodes relative to the location of their allocated memory. Optimizations were implemented via the hwloc C++ package which provides programmatic, location-specifiable memory allocation and thread creation. Using this package, each worker thread is bound to a specific core with its memory allocated in the local node cache. Thread assignments are made holistically on the basis of minimizing the total number of internode memory accesses. This programmatic tailoring to

the specific architecture of the system has achieved more than a factor of two increase in throughput.

A remote python management script interfaces with the low-level C++ process via a TCP connection. This custom script provides a graphical interface for users to monitor and control the data acquisition. It renders configurable real-time plots of the accumulating correlation matrix, and it also serves as the access point for retrieving previously-recorded data from the server archive. Any number of remote client applications can simultaneously connect and monitor the system.

CHAPTER 4

MEASUREMENT CALIBRATION

The high-speed data processing pipeline outputs a continuous time stream of 1-s averaged correlation matrices between the AS-port photodetector readouts (see §3.5.3). In order to make physical inferences, these raw voltage data must first be calibrated to physical units of differential length. The following sections now describe this calibration.

4.1 Overview

The time series of 1-s correlation matrices is computed from the digitized high-frequency (HF) readout of each photodetector as

$$\textsc{Csd}\left[\widetilde{V}_i^{\mathrm{HF}}(f,\,t)\,,\,\widetilde{V}_j^{\mathrm{HF}}(f,\,t)\right] \;\equiv\; \widetilde{V}_i^{\mathrm{HF}}(f,\,t)\cdot\left(\widetilde{V}_j^{\mathrm{HF}}(f,\,t)\right)^{*}\,, \tag{4.1}$$

where $\widetilde{V}_i^{\mathrm{HF}}(f,\,t)$ is the Fourier-transformed voltage signal of channel i at time t. The signals are time-dependent because the uncontrolled alignment degrees of freedom (e.g., injection alignment) slowly drift on time scales > 30 s, leading to changes as large as 30% in the interferometer sensitivity over long operation times.

In frequency space, each voltage signal can be expressed as a stationary source of differential path-length noise, $\widetilde{S}_i(f)$, viewed through the time-varying transfer function of the instrument[1],

$$\widetilde{V}_i^{\mathrm{HF}}(f,\,t) \;=\; H_i^{\mathrm{HF}}(f,\,t)\cdot\widetilde{S}_i(f)\,. \tag{4.2}$$

The transfer function, $H_i^{\mathrm{HF}}(f,\,t)$ (units V/m), maps differential length changes in the interferometer into the measured voltage response. It will be referred to as the calibration of

1. Both the response of the interferometers around their operating point and the response of the photodetectors to incident powers < 150 mW are linear.

channel i.

In order to convert the correlation matrix of measured voltage responses to physical units of length displacement, the calibration of each channel must be known. In principle, the calibrations could be directly measured by driving the PZT-actuated end mirrors with a known injection signal, $\widetilde{S}_{\rm inj}(f)$, and measuring each source-normalized response,

$$\widehat{H}_i^{\rm HF}(f,\,t) \;=\; \frac{\widehat{\widetilde{V}}_i^{\rm HF}(f,\,t)}{\widetilde{S}_{\rm inj}(f)}\;. \tag{4.3}$$

While standard practice in gravitational wave detection, direct calibration techniques do not carry over into the Holometer design, as the end mirrors cannot be mechanically driven above kHz frequencies. The unique, high-frequency requirements of the Holometer have led to the development of a novel indirect, but equivalent, method for calibrating the response of a Michelson interferometer at radio frequencies. This technique circumvents the mechanical-drive limitations encountered by direct methods. The following section describes its implementation.

4.2 Indirect Calibration

The transfer function of any linear system can be decomposed into the product of the transfer functions of its individual components. Accordingly, the calibration decouples as

$$H_i^{\rm HF}(f,\,t) \;=\; H_i^{\rm DAQ,HF}(f)\cdot H_i^{\rm IFO}(f,\,t)\;, \tag{4.4}$$

where $H_i^{\rm IFO}(f,\,t)$ (units W/m) is the alignment-dependent, time-varying transfer function of the interferometer and $H_i^{\rm DAQ,HF}(f)$ (units W/V) is the stationary transfer function of the HF signal-readout chain. Previously, §3.2.2 used a numerical model of the Fabry-Perot cavity to demonstrate that the interferometer transfer function is constant to 1% above

1 kHz. Now invoking this result, eq. 4.4 is equivalently represented as

$$\begin{aligned} H_i^{\mathrm{HF}}(f,\,t) &= H_i^{\mathrm{DAQ,HF}}(f)\cdot H_i^{\mathrm{IFO}}(f_0,\,t)\cdot\frac{H_i^{\mathrm{IFO}}(f,\,t)}{H_i^{\mathrm{IFO}}(f_0,\,t)} \\ &= H_i^{\mathrm{DAQ,HF}}(f)\cdot H_i^{\mathrm{IFO}}(f_0,\,t) \end{aligned} \tag{4.5}$$

for all frequencies f and $f_0 \gtrsim 1$ kHz. Eq. 4.5 expresses the underlying principle of indirect calibration. Suppose the end mirrors are differentially dithered at $f_0 \approx 1$ kHz, just below their maximum frequency of linear operation, by a known amount. If the calibration is directly measured at f_0, via eq. 4.3, then eq. 4.5 implies that *only* knowledge of the detection chain response is required to infer the calibration for all higher frequencies.

In reality, the HF signal-readout chain cannot detect the 1-kHz drive, as the photodetector response is high-passed at 1 MHz to suppress lower-frequency seismic noise (see 3.5.1). However, the low-frequency (LF) readout of each photodetector, simultaneously sampled at 32 kHz by the control systems, is sensitive to the injection. The transfer function of the LF system analogously decouples as

$$H_i^{\mathrm{LF}}(f_0,\,t) = H_i^{\mathrm{DAQ,LF}}(f_0)\cdot H_i^{\mathrm{IFO}}(f_0,\,t)\,, \tag{4.6}$$

where $H_i^{\mathrm{DAQ,LF}}(f_0)$ (units W/V) is the transfer function of the LF photodetector channel and the FPGA. Substituting eq. 4.6 into eq. 4.5 then yields an expression,

$$H_i^{\mathrm{HF}}(f,\,t) = \frac{H_i^{\mathrm{DAQ,HF}}(f)}{H_i^{\mathrm{DAQ,LF}}(f_0)}\cdot H_i^{\mathrm{LF}}(f_0,\,t)\,, \tag{4.7}$$

formulated purely in terms of continuously-monitored control signals and the measurable, stationary transfer functions of the signal-readout chains. The following sections now detail the practical measurement of each term.

4.3 Direct Low-Frequency Calibration

The time-varying dynamics of the interferometers are contained purely in the $H_i^{\mathrm{LF}}(f_0, t)$ terms, which are continuously monitored by the control systems via a direct, low-frequency calibration measurement. Its implementation is described below.

4.3.1 Calibration Line Injection and Monitoring

The interferometer control systems are configured to continuously inject a calibration line near 1 kHz via a differential end mirror dither. One interferometer injects a signal at 983 Hz and the other at 984 Hz. Intermodulation studies have placed strong upper limits on the longitudinal impurity of the dithers, as angular motion could introduce coherent systematic biasing of the calibration line measurement. By driving dithers at two frequencies and measuring their intermodulation power with a QPD, these studies limit systematic biasing of the calibration due to angular drive motion to $< 2\%$.

The control systems continuously log the 32-kHz-sampled time series of the diagnostic and interferometer control signals, including the digitally-generated drive signal and the LF readout of the AS-port photodetectors. Since the calibration lines are necessarily measured while the interferometers are locked, the differential-mode control loops act to suppress the drive injection. A loop-gain correction, measured from the coherent time series of the drive and internal test points within the loop, is first calculated to remove the effect of the control loop from each measurement.

With loop effects removed, the data measure the out-of-loop transfer function of the true drive amplitude to the LF readout of each AS-port photodetector,

$$\widehat{H}_i^{\mathrm{LF}}(f_0, t) \;=\; \frac{\widehat{\widetilde{V}}_i^{\mathrm{LF}}(f_0, t)}{\widetilde{S}_{\mathrm{inj}}(f_0)} \,. \tag{4.8}$$

The transfer functions are calculated using complex positive-frequency narrowband filters

on the data, dividing by the complex drive signal and then downsampling to 16 Hz. The narrowband filters act as Nyquist filters for the downsampling and the division normalizes by the drive amplitude and shifts modulation to DC. The time history of the continuously-monitored low-frequency calibrations are shown in Fig.4.1.

4.3.2 Reference to Optical Wavelength

In order to inject a calibrated displacement, the PZT drive of each interferometer must itself be calibrated via the transfer function of drive voltage, $\widetilde{V}_{\mathrm{inj}}$, to the physical length displacement produced by the actuators, $\widetilde{S}_{\mathrm{inj}}$. A simple, highly-precise measurement of this transfer function is made using the laser wavelength as an absolute reference to length. The measurement is made in the single-pass Michelson configuration, where it is simplified by the constant input power and by locking to mid-fringe, where exactly half of the light exits through the AS port. The half-light condition is established by first sweeping DARM through a full fringe, seeing the minimum and maximum extents of the Michelson sinusoidal response. The Michelson interferometer is then aligned so that the contrast defect is at a quadratic minimum and locked to half-power, where any residual variation in contrast defect cannot affect the lock-point offset.

A differential swept-sine signal is injected via the PZT drive and measured at two test points simultaneously to infer the loop gain and the driven sinusoidal optical response at the AS port. From these, the open-loop optical response with respect to PZT drive voltage, $\widetilde{V}_i^{\mathrm{LF}}/\widetilde{V}_{\mathrm{inj}}$, is inferred. Additionally, the linear AS-port detector is known to measure an optical signal proportional to the Michelson AS-port power, given by eq. 2.19. Evaluating the derivative of eq. 2.19 at mid-fringe then provides the open-loop response with respect to true differential length, $\widetilde{V}_i^{\mathrm{LF}}/\widetilde{S}_{\mathrm{inj}}$, whose absolute scale is provided by the laser wavelength. Taking the ratio of the two Michelson open-loop optical responses yields the desired transfer function of the PZT drive, $\widetilde{S}_{\mathrm{inj}}/\widetilde{V}_{\mathrm{inj}}$. Figs. 4.2 and 4.3 show the PZT drive calibrations

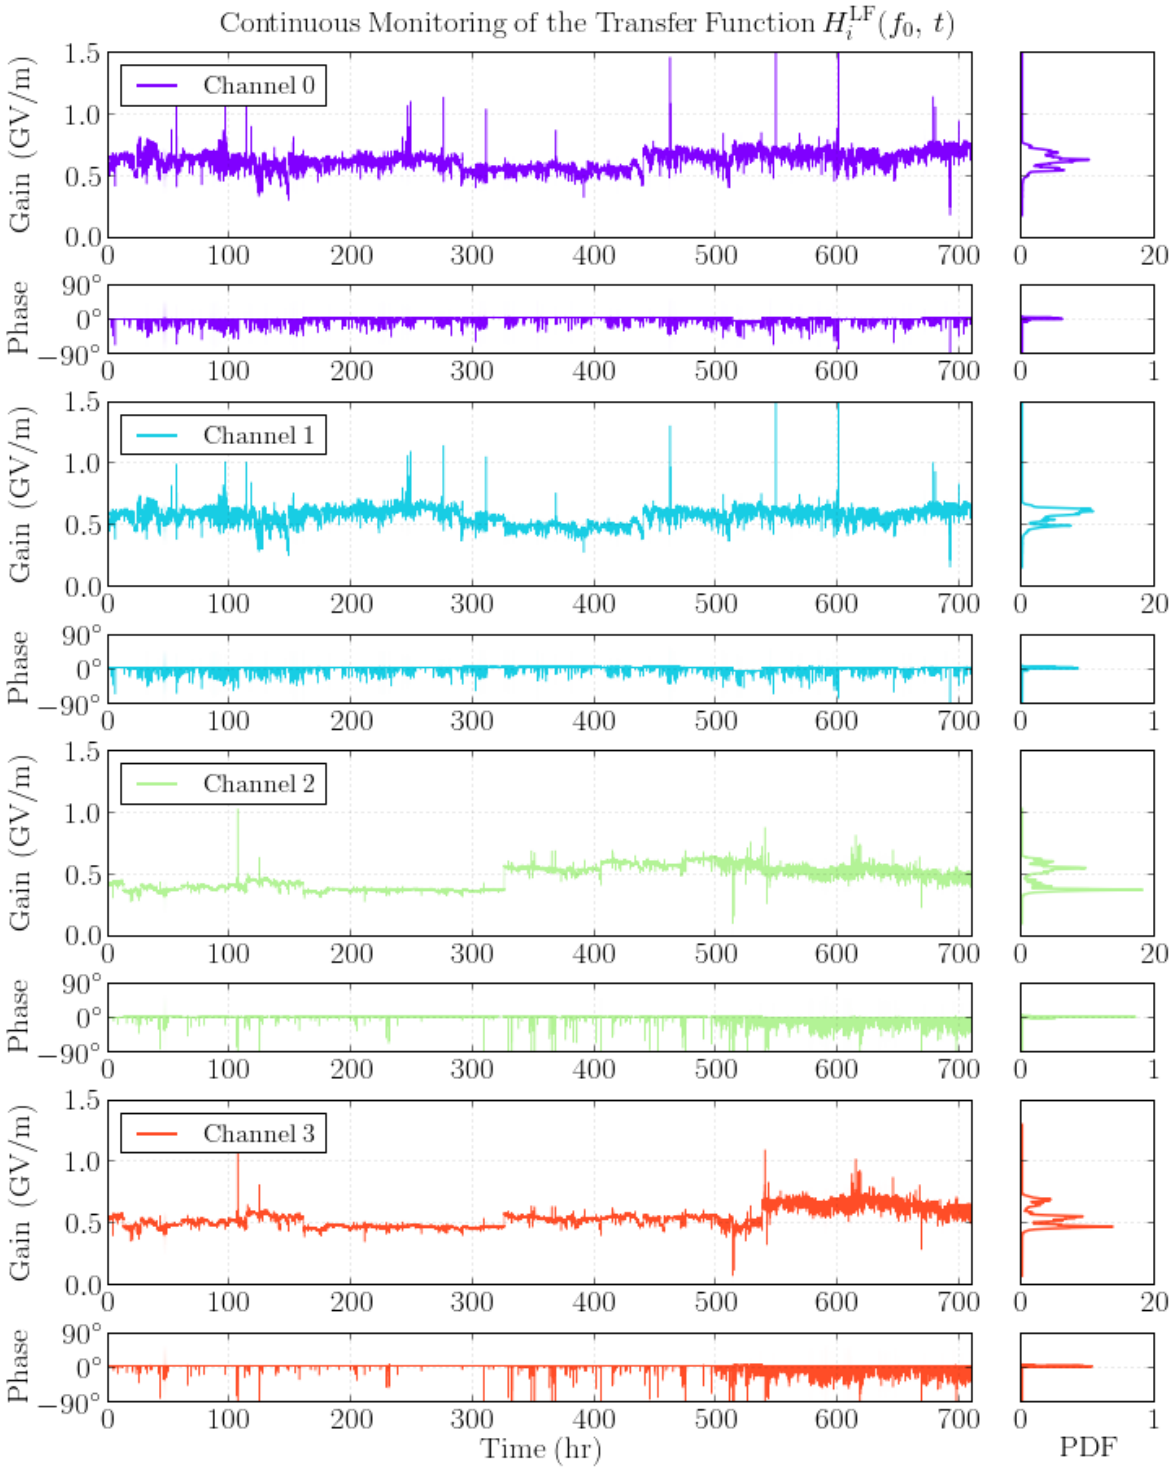

Figure 4.1: Time history of the continuous 1-kHz direct calibrations. The regular, step-like discontinuities in the channel gains reflect periodic realignments of the inteferometers during data collection. The zero phases are indicative of agreement between the single-pass (used for PZT calibration; see §4.3.2) and power-recycled interferometer responses.

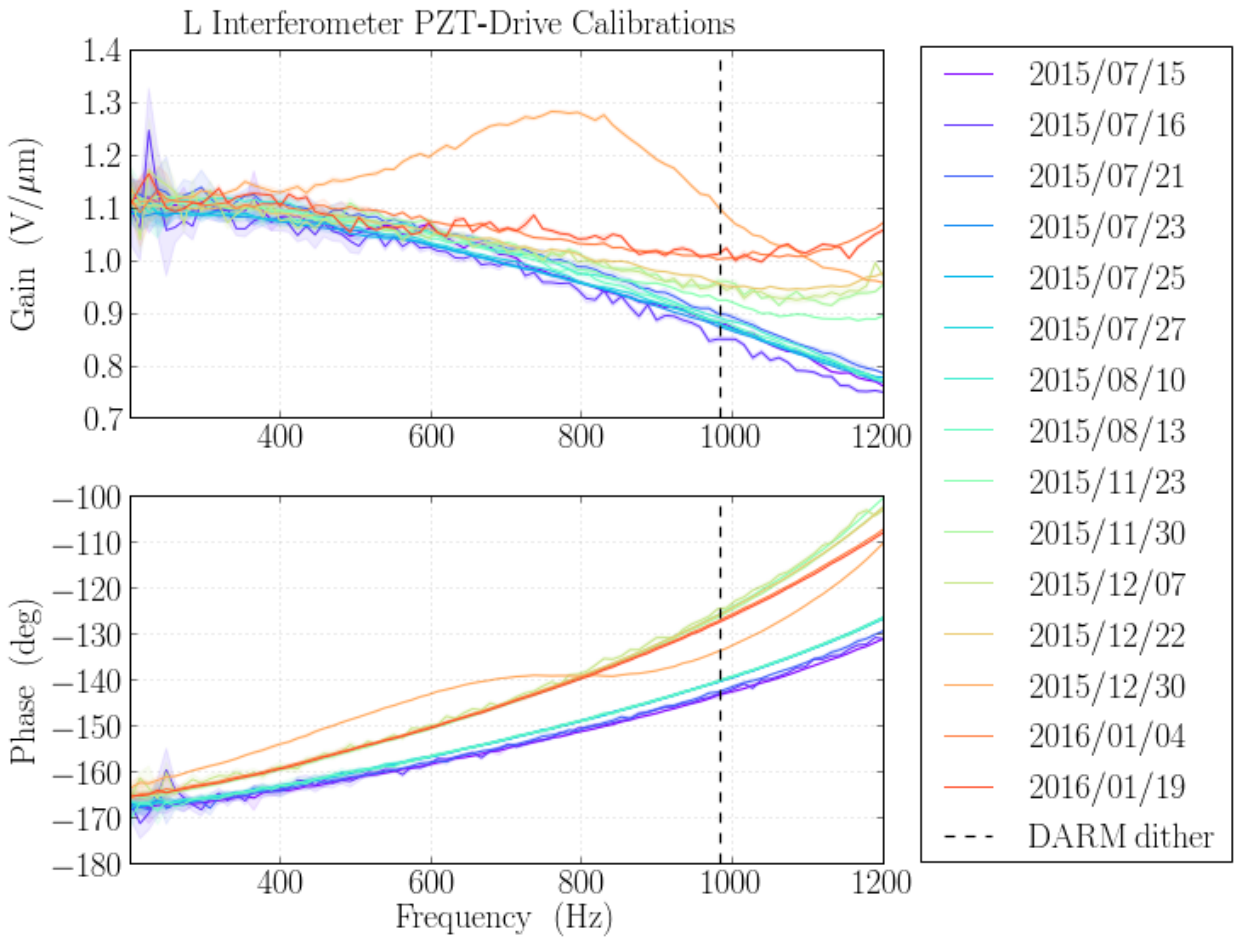

Figure 4.2: PZT-drive calibrations for the L interferometer over the six-month data collection period. The 20% calibration drift between the summer and winter measurements is believed to be caused by the change in ambient temperature. This is discussed further in §4.5.1.

of each interferometer, measured at approximately weekly intervals during data collection. Systematics affecting this calibration will be discussed in §4.5.1.

4.4 Signal-Readout Transfer Functions

The direct, continuously-measured 1-kHz interferometer calibrations are transferred to higher frequencies by the signal-readout transfer functions, $H_i^{\mathrm{DAQ,HF}}(f)/H_i^{\mathrm{DAQ,LF}}(f_0)$. Due to the vast separation in scale between kHz and radio frequencies, a combination of different measurement techniques is required to measure these transfer functions over the entire range.

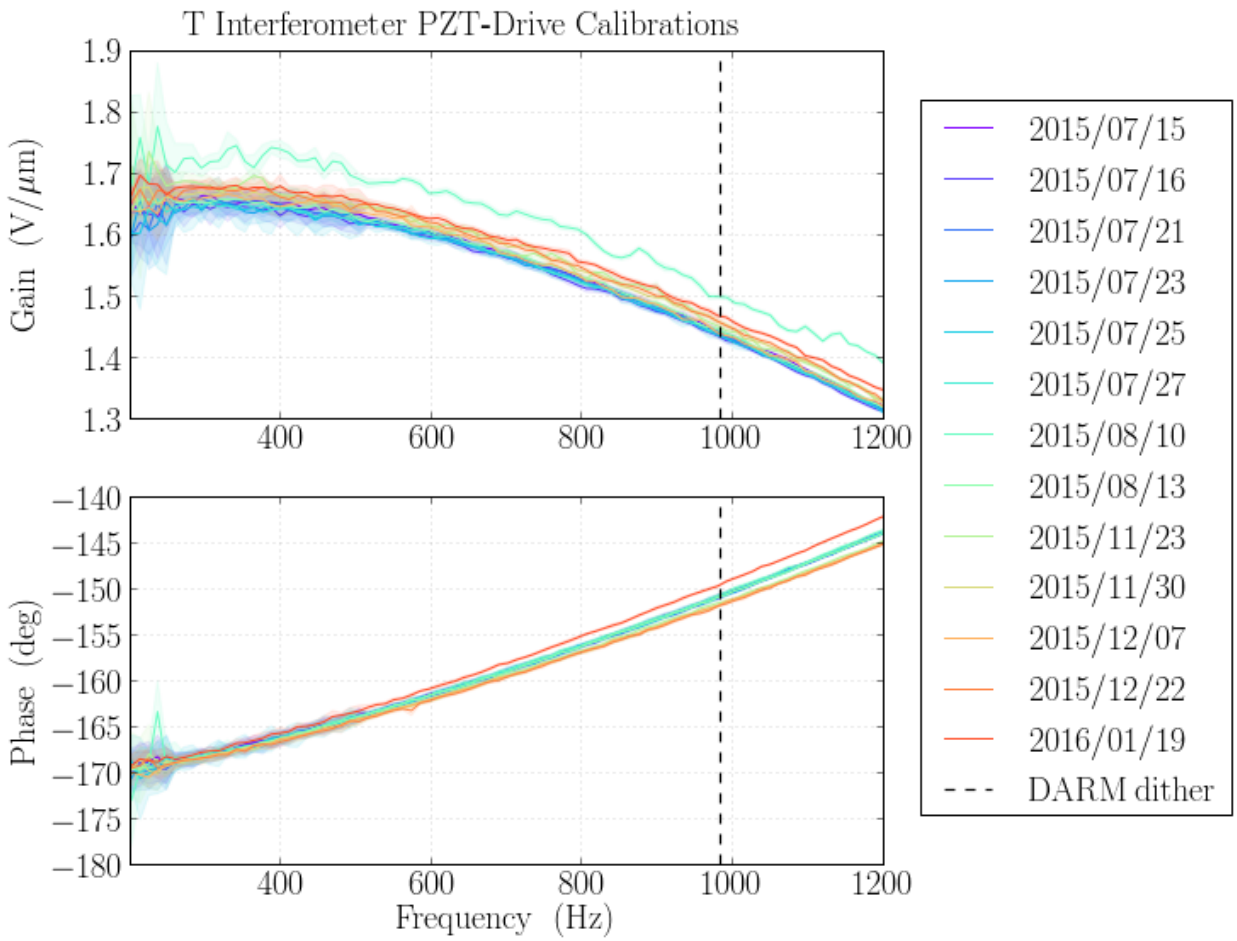

Figure 4.3: PZT-drive calibrations for the T interferometer over the six-month data collection period. Relative drift between the calibration measurements is believed to be caused by changes in ambient temperature. This is discussed further in §4.5.1.

A piecewise division between measurements occurs at 900 kHz, whose location is determined by the practical limitations of the measuring equipment. The following sections sequentially describe each measurement technique, starting at DC and progressing upward in frequency.

4.4.1 Measurement Below 900 kHz

Below 900 kHz, the signal-readout transfer functions are measured via an *ex-situ* lock-in detection technique. Light from a 960-nm LED is focused onto each modified New Focus 1811 photodetector, in turn. The LED power is modulated via a swept-sine signal, $\widetilde{V}_{\text{inj}}(f)$, driven by a function generator. Averaged over many sweeps, the AC- and DC-coupled photodetector channels measure the signals

$$\widetilde{V}_i^{\text{AC}}(f) = H_i^{\text{DAQ,HF}}(f) \cdot G_{\text{align}} \cdot H_{\text{LED}}(f) \cdot \widetilde{V}_{\text{inj}}(f) \tag{4.9}$$

and

$$\widetilde{V}_i^{\text{DC}}(f) = H_i^{\text{DAQ,LF}}(f) \cdot G_{\text{align}} \cdot H_{\text{LED}}(f) \cdot \widetilde{V}_{\text{inj}}(f) \ , \tag{4.10}$$

respectively. In the above equations, G_{align} is an alignment-dependent gain factor and $H_{\text{LED}}(f)$ is the unknown coupling of drive signal to the optical power emitted by the LED. In order to eliminate the dependency on the unknown coupling function, the LED light is split by a beamsplitter and additionally directed onto a Thorlabs PDA 20CS photodetector, whose response, H_{ref}, is known from external calibrations to be flat from DC to 1 MHz. Averaged over many sweeps, it measures the signal

$$\widetilde{V}_{\text{ref}}(f) = H_{\text{ref}} \cdot G'_{\text{align}} \cdot H_{\text{LED}}(f) \cdot \widetilde{V}_{\text{inj}}(f) \ , \tag{4.11}$$

where G'_{align} is the alignment-dependent gain factor for this second detector.

In terms of the measured sensor voltages, each signal-readout transfer function is

$$\frac{H_i^{\mathrm{DAQ,HF}}(f)}{H_i^{\mathrm{DAQ,LF}}(f_0)} = \frac{\widetilde{V}_i^{\mathrm{AC}}(f)}{\widetilde{V}_{\mathrm{ref}}(f)} \cdot \frac{\widetilde{V}_{\mathrm{ref}}(f_0)}{\widetilde{V}_i^{\mathrm{DC}}(f_0)} , \tag{4.12}$$

independent of the photodetector alignment. Moreover, each of the response signals can be locked-in to the drive signal to provide an unbiased measurement. The lock-in is implemented as the measurement

$$\begin{aligned} \frac{H_i^{\mathrm{DAQ,HF}}(f)}{H_i^{\mathrm{DAQ,LF}}(f_0)} &= \frac{\widetilde{V}_i^{\mathrm{AC}}(f)}{\widetilde{V}_{\mathrm{ref}}(f)} \cdot \frac{\widetilde{V}_{\mathrm{inj}}^*(f)}{\widetilde{V}_{\mathrm{inj}}^*(f)} \cdot \frac{\widetilde{V}_{\mathrm{ref}}(f_0)}{\widetilde{V}_i^{\mathrm{DC}}(f_0)} \cdot \frac{\widetilde{V}_{\mathrm{inj}}^*(f)}{\widetilde{V}_{\mathrm{inj}}^*(f)} \\ &= \frac{\mathrm{CSD}\left[\widetilde{V}_i^{\mathrm{AC}}(f)\,,\,\widetilde{V}_{\mathrm{inj}}(f)\right]}{\mathrm{CSD}\left[\widetilde{V}_{\mathrm{ref}}(f)\,,\,\widetilde{V}_{\mathrm{inj}}(f)\right]} \cdot \frac{\mathrm{CSD}\left[\widetilde{V}_{\mathrm{ref}}(f_0)\,,\,\widetilde{V}_{\mathrm{inj}}(f_0)\right]}{\mathrm{CSD}\left[\widetilde{V}_i^{\mathrm{DC}}(f_0)\,,\,\widetilde{V}_{\mathrm{inj}}(f_0)\right]} , \end{aligned} \tag{4.13}$$

with the CSDs measured by the high-speed DAQ system (see §3.5).

To accurately account for post-detection transmission losses, the LF and HF readout signals are both routed through their assigned *en-situ* cabling from the photodetectors to the high-speed ADCs. The input load impedance of the ADCs, 50 Ω for the HF signals and 1 MΩ for the LF signals, is chosen to match the *en-situ* measurement conditions. Although the LF signals are measured by the high-speed ADCs, rather than by the FPGAs, the difference is negligible for frequencies $<<$ 32 kHz, as neither system applies any signal conditioning other than a Nyquist filter.

Sufficiently time-averaged, the CSDs isolate the signal components phase-coherent with the drive, thus removing detector dark-noise bias and any behavior of the LED unrelated to the drive modulation. The resulting signal-readout transfer functions are shown in Fig. 4.4. In addition to the statistical uncertainty, each measurement reflects an additional 4-7% systematic error, whose origin will be discussed in §4.5.2. The flat-response range of the reference photodetector, DC-1 MHz, restricts the use of this technique to frequencies $\leq$ 1 MHz.

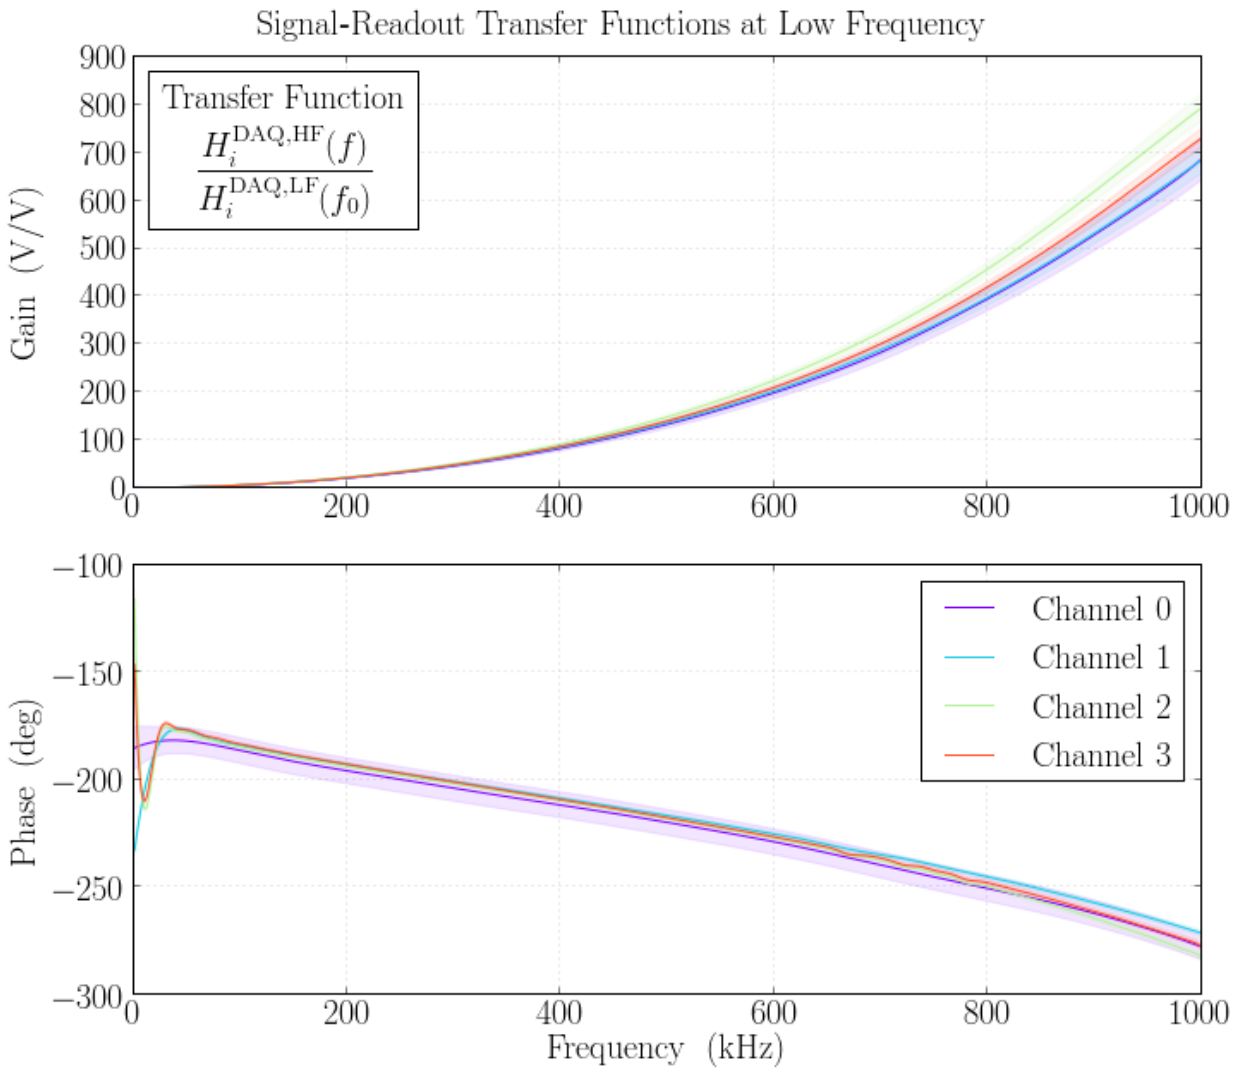

Figure 4.4: Transfer functions of the signal-readout chains measured via lock-in drive detection at low frequency. The gain roll-off at DC is due to the 1-MHz high-pass filter in the AC-coupled photodetector channel, implemented to suppress strong low-frequency optical signals caused by seismic noise (see §3.5.1).

4.4.2 Measurement Above 900 kHz

In principle, with broadband-sensing equipment, the lock-in technique described in §4.4.1 could be used to calibrate the entire measurement band. However, the response of the secondary photodetector rolls off above 1 MHz, necessitating the use of a different gain reference for higher frequencies. A cut-off frequency of $f_1 = 900$ kHz is chosen as the point at which the response of the reference detector has rolled off by 1% relative to its DC value. This transition is implemented as the dual transfer

$$\frac{H_i^{\mathrm{DAQ,HF}}(f)}{H_i^{\mathrm{DAQ,LF}}(f_0)} = \frac{H_i^{\mathrm{DAQ,HF}}(f_1)}{H_i^{\mathrm{DAQ,LF}}(f_0)} \cdot \frac{H_i^{\mathrm{DAQ,HF}}(f)}{H_i^{\mathrm{DAQ,HF}}(f_1)} , \tag{4.14}$$

with the prior lock-in measurement used to bridge the DC-to-AC transfer from f_0 to f_1, followed by a second, AC-to-AC transfer from f_1 to higher frequencies. The gains and phases of the AC-to-AC transfer functions are measured separately, as described below.

Gain Measurement from Optical Shot Noise

The gains of the AC-to-AC transfer functions are measured *ex-situ* directly from optical shot noise. Light from a 1-W incandescent bulb is focused onto each modified New Focus 1811 photodetector, in turn. The AC-coupled photodetector channel measures the signal

$$\widetilde{V}_i^{\mathrm{AC}}(f) = H_i^{\mathrm{DAQ,HF}}(f) \cdot G_{\mathrm{align}} \cdot \widetilde{P}_{\mathrm{shot}} , \tag{4.15}$$

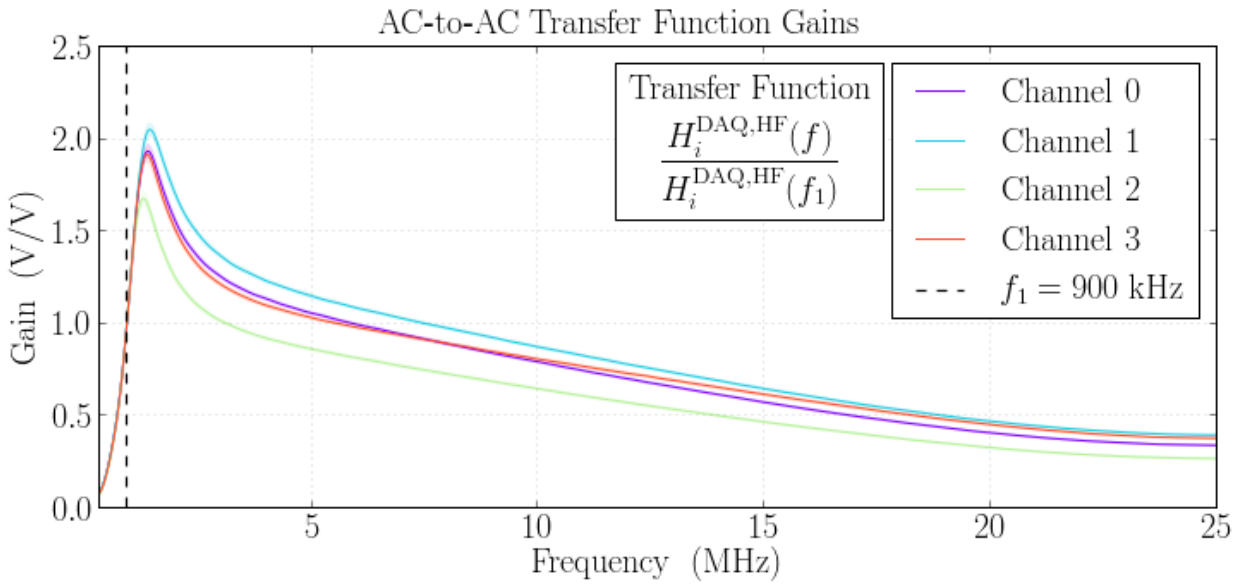

Figure 4.5: Gains of the AC-to-AC transfer functions measured via optical shot noise. Each gain function is referenced to the 900-kHz intermediate transfer frequency.

as the bulb is known to emit intrinsically white shot-noise power, $\widetilde{P}_{\mathrm{shot}}$, above 25 kHz. In terms of the measured sensor voltage, the AC-to-AC gain function is

$$\begin{aligned} \left| \frac{H_i^{\mathrm{DAQ,HF}}(f)}{H_i^{\mathrm{DAQ,HF}}(f_1)} \right| &= \left| \frac{\widetilde{V}_i^{\mathrm{AC}}(f)}{\widetilde{V}_i^{\mathrm{AC}}(f_1)} \right| \\ &= \sqrt{\frac{\textsc{Psd}\left[\widetilde{V}_i^{\mathrm{AC}}(f)\right]}{\textsc{Psd}\left[\widetilde{V}_i^{\mathrm{AC}}(f_1)\right]}}\,, \end{aligned} \tag{4.16}$$

where the PSDs are measured by the high-speed DAQ system via the *en-situ* transmission cabling of each channel.

As the PSD measures the sum of all sources of noise power in the signal, it is susceptible to detector dark-noise bias. Accordingly, each PSD measurement is dark noise-subtracted prior to evaluating eq. 4.16. The resulting AC-to-AC gain functions are shown in Fig. 4.5. In addition to the statistical uncertainty, each measurement reflects an additional 1-3% systematic error, whose origin will be discussed in §4.5.2.

Phase Measurement from Lock-In Drive

Shot noise is an intrinsically white, but incoherent, calibration signal. As such, it provides no phase information. The phases of the AC-to-AC transfer functions are measured via the same lock-in drive technique described in §4.4.1. In this band, however, the reference photodetector has an unknown, frequency-dependent response, $H_{\rm ref} \to H_{\rm ref}(f)$. In terms of the measured sensor voltages, each AC-to-AC transfer function is

$$\frac{H_i^{\rm DAQ,HF}(f)}{H_i^{\rm DAQ,HF}(f_1)} = \frac{\widetilde{V}_i^{\rm AC}(f)}{\widetilde{V}_{\rm ref}(f)} \cdot \frac{\widetilde{V}_{\rm ref}(f_1)}{\widetilde{V}_i^{\rm AC}(f_1)} \cdot \frac{H_{\rm ref}(f)}{H_{\rm ref}(f_1)} \,, \tag{4.17}$$

where the reference detector response now does not cancel. Lock-in detection with the drive signal thus measures

$$\begin{aligned} \frac{H_i^{\rm DAQ,HF}(f)}{H_i^{\rm DAQ,HF}(f_1)} \cdot \frac{H_{\rm ref}(f_1)}{H_{\rm ref}(f)} &= \frac{\widetilde{V}_i^{\rm AC}(f)}{\widetilde{V}_{\rm ref}(f)} \cdot \frac{\widetilde{V}_{\rm inj}(f)}{\widetilde{V}_{\rm inj}(f)} \cdot \frac{\widetilde{V}_{\rm ref}(f_1)}{\widetilde{V}_i^{\rm AC}(f_1)} \cdot \frac{\widetilde{V}_{\rm inj}(f_1)}{\widetilde{V}_{\rm inj}(f_1)} \\ &= \frac{\textsc{Csd}\left[\widetilde{V}_i^{\rm AC}(f)\,,\,\widetilde{V}_{\rm inj}(f)\right]}{\textsc{Csd}\left[\widetilde{V}_{\rm ref}(f)\,,\,\widetilde{V}_{\rm inj}(f)\right]} \cdot \frac{\textsc{Csd}\left[\widetilde{V}_{\rm ref}(f_1)\,,\,\widetilde{V}_{\rm inj}(f_1)\right]}{\textsc{Csd}\left[\widetilde{V}_i^{\rm AC}(f_1)\,,\,\widetilde{V}_{\rm inj}(f_1)\right]} \,, \end{aligned} \tag{4.18}$$

with the CSDs, as before, measured by the high-speed DAQ system via the *en-situ* transmission cabling of each channel.

The phase response measured by eq. 4.18,

$$\mathrm{Arg}\left(\frac{H_i^{\rm DAQ,HF}(f)}{H_i^{\rm DAQ,HF}(f_1)}\right) + \mathrm{Arg}\left(\frac{H_{\rm ref}(f_1)}{H_{\rm ref}(f)}\right) \,, \tag{4.19}$$

is the sum of the desired result and a phase offset due to the reference photodetector. Although the phase offset is not known, it is common to the measurement of every signal-readout channel. Because the CSD is defined as the conjugate-product of two signals, its phase is the *difference* of the individual phases of the two signals. Thus, when the phase measurements of eq. 4.19 are substituted into the full CSD calibration, the common phase

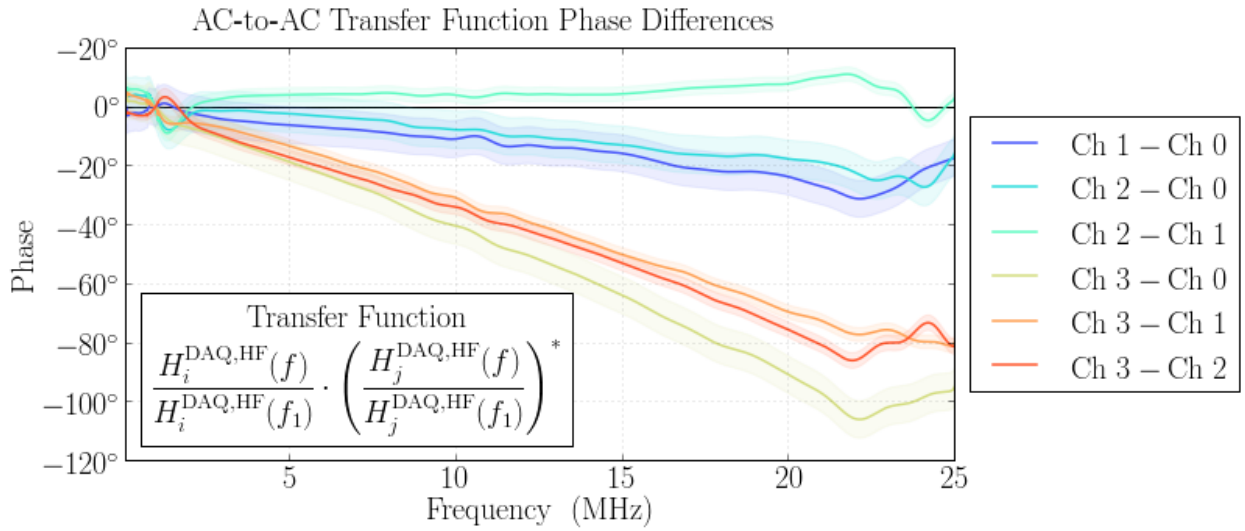

Figure 4.6: Phase differences between the AC-to-AC transfer functions measured via lock-in drive detection. The large phase difference between channel 3 and the other channels is largely due to the response of the 2-pole 20-MHz antialias filter in this ADC channel.

offsets of the two readout channels cancel, leaving only the phase difference of the two AC-to-AC transfer functions. Fig. 4.6 shows the measured phase difference between the AC-to-AC transfer functions of each channel pairing.

4.5 Systematics & Data Quality Control

This section now provides a full accounting of every known systematic effect on the calibration-referred experimental sensitivity. It will be demonstrated that all signal detection and readout systematics are well-controlled over the entirety of the six-month data collection period.

4.5.1 Calibration Drift in the PZT Drives

As discussed in §3.4.1, each PZT-drive system has mechanical resonances above 1.5 kHz that are observed to drift up to 10% in frequency over many weeks. The shifting tails of these resonant peaks are expected to have some impact on the 1-kHz PZT transfer function. Accordingly, the drive calibration of each interferometer were re-measured at approximately

weekly intervals during data collection, as shown in Figs. 4.2 and 4.3. Over periods of several weeks, the calibrations drift by $< 2\%$. However, drifts at the 20% level are seen between the earliest and latest measurements, spaced roughly six months apart. The 12/30/2015 and 01/04/2016 calibrations for the T interferometer were rejected on the basis of producing a 50° phase mismatch, known to be artificial, between the single-pass and power-recycled interferometer responses. The same calibrations for the L interferometer, while drifting in gain, do not exhibit this phase mismatch.

Because the initial and final data, taken in July and January, respectively, also correspond to the annual extremes in ambient temperature, the drift in PZT response is believed (though not proven) to be thermally-related. Although the tunnel and outside enclosure housing the end stations are heated to some degree, annual interior temperature swings of up to 30 C still occur. Regardless of the origin, the large number of regularly-spaced measurements indicate that the time scale of drift is well-resolved by the frequency of re-measurement, implying that any residual biasing of the inferred interferometer sensitivity between PZT recalibrations is $< 2\%$.

4.5.2 Photodetector Dependence on Power

The modified New Focus 1811 photodetectors are observed to exhibit a weak systematic dependence on the incident DC optical power. This dependence affects the responses of both the AC- and DC-coupled readout channels. Although it is not fully understood, it is believed to arise from thermally-induced change in effective resistances inside the circuit at high power. For reasons discussed below, this effect is bounded and then absorbed as an additional, systematic uncertainty of the calibration-referred experimental sensitivity. This treatment is shown to be validated by continuous monitoring of the detector responses during science operation.

Effect on Detector Response

At low frequencies, the left panels of Figs. 4.7 and 4.8 show the systematic variation in the DC- and AC-channel responses of one detector at different levels of DC illumination. This measurement is made using the lock-in drive technique described in §4.4.1, with the addition of a 1-W incandescent bulb with adjustable power supply. The incandescent light is focused onto the modified photodetector, while the reference photodetector is shielded from this light, so that the only condition changing between measurements is the DC power incident on the modified detector. Both detector channels, in general, exhibit some degree of gain compression with increasing DC photocurrent.

This behavior is observed to extend to higher frequencies in the AC-coupled channel. The left panel of Fig. 4.9 shows the AC-channel transimpedance gain of the same photodetector, again measured at different levels of DC illumination. This measurement is made via the optical shot noise technique described in §4.4.2, with the power of the incandescent bulb adjusted to provide varying levels of incident DC power. The transimpedance gain is the square-root of the dark-noise-subtracted PSD, normalized by the photocurrent shot noise, $\sigma_I = \sqrt{2eV/R}$. The DC photovoltage, V, is directly measured via a voltmeter connected to the DC readout channel, whose gain is $R = 10$ V/A.

Bounding Procedure

The above measurements, shown for one detector, are representative of a systematic behavior common to all of the detectors. Although variations as large as 20% occur between the extreme high- and low-power responses, the detectors are operated *en-situ* within a much narrower range of incident powers centered on 100 mW. Over the range of Holometer operating powers, taken as 70-140 mW, the response biasing is much smaller, typically $2-5\%$ per channel. Because the measured detector responses, at a given incident power, have only been found reproducible to the few-percent level, this systematic bias is not modeled, but

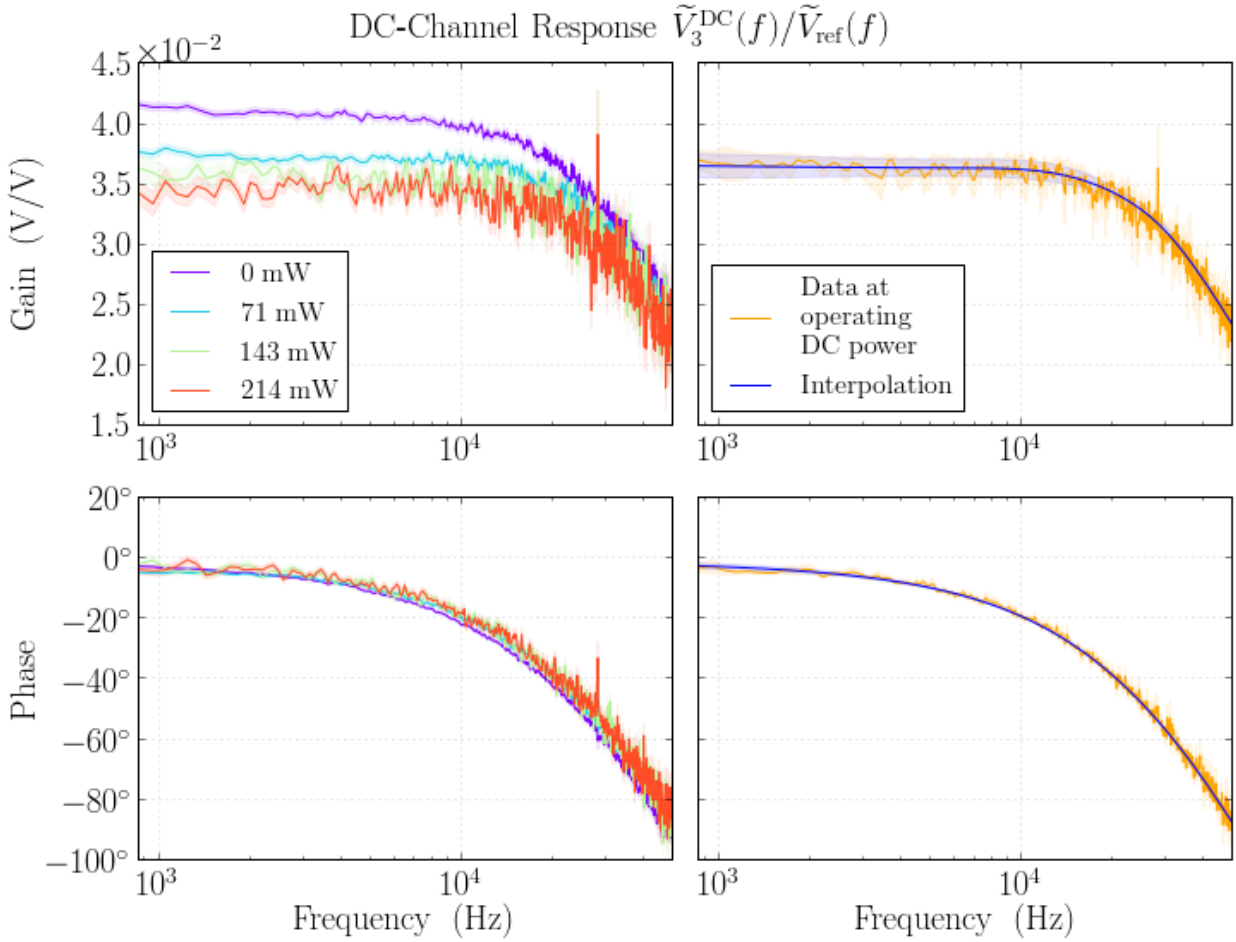

Figure 4.7: Left: DC-channel response of a modified New Focus 1811 photodetector at different levels of DC illumination. Right: Interpolated DC-channel response at operating DC power.

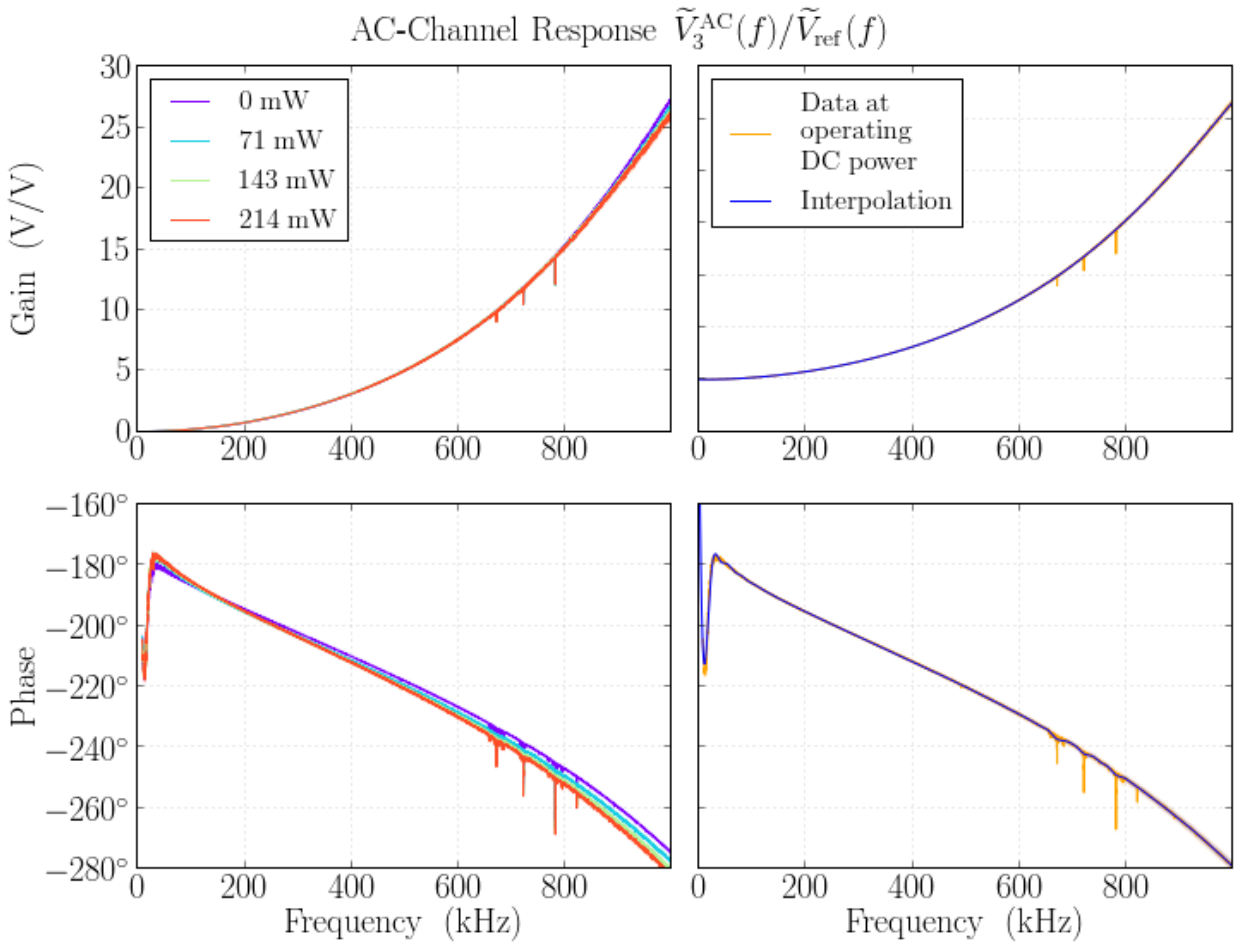

Figure 4.8: Left: AC-channel response of a modified New Focus 1811 photodetector at different levels of DC illumination. Right: Interpolated AC-channel response at operating DC power.

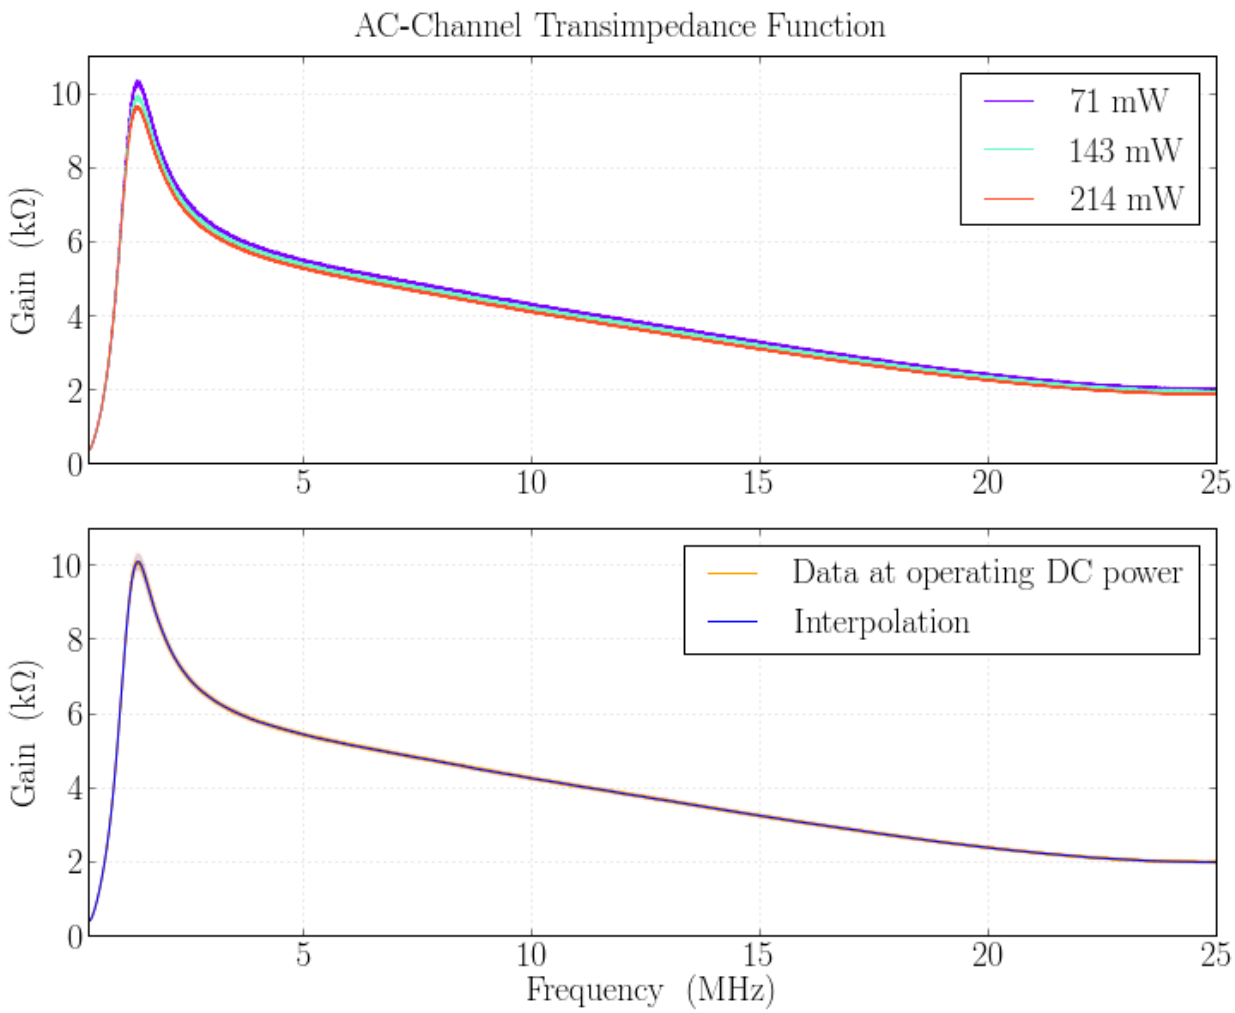

Figure 4.9: Top: Transimpedance gain of the AC-coupled channel of a modified New Focus 1811 photodetector at different levels of DC illumination. Bottom: Interpolated AC-channel transimpedance gain at operating DC power.

is instead bounded within the Holometer operating range and absorbed as an additional, systematic uncertainty of the detector response.

The right panels of Fig. 4.7 show this bounding procedure applied to a DC-readout channel. First, the centroid of the channel responses measured at the two operating power extremes, 70 mW and 140 mW, is calculated. The uncertainty of this centroid is then taken as the statistical measurement error plus a systematic error, summed in quadrature, representing the maximum possible response bias across the operating range. At each frequency, the systematic error is taken as the residual disagreement of the two response measurements, beyond their 1σ statistical error bars. In the figure, the response centroid is denoted by the orange curves, with shaded contours denoting its total uncertainty. The frequency dependence of the centroid is then interpolated via a non-parametric machine-learning technique known as Gaussian process regression[16] to yield the bias-absorbed response, shown by the blue curve.

The right panels of Figs. 4.8 and 4.9 show this bounding procedure analogously applied to the AC-channel measurements. In every case, the detector responses were measured at a minimum of three intermediate powers to confirm that the adopted systematic error does enclose the responses at all in-range incident powers. Applicability of the bias-absorbed responses, by construction, requires that 70-140 mW of exposure be maintained on each photodetector during operation. Enforcing this condition requires continuous monitoring of the power incident on each detector.

Continuous Monitoring of Detector Power

The control systems provide continuous output power monitoring via the DC-channel detector signals. The DC power, P, incident on a detector produces a DC photovoltage of $V = RQP$, where $R = 10$ V/A is the DC-channel transimpedance gain and $Q = 0.71$ A/W is the responsivity of the photodiode. Fig. 4.10 shows the time history of power incident

on each photodetector. Any 1-s CSD matrix whose duration coincides with an out-of-range photodetector power is automatically rejected from analysis. As the operating range allows for up to ±30% drift in the incident power, vetos triggered by this condition are rare (< 1% of the data) and are found to occur almost exclusively as an interferometer is losing lock. As such, any sample bias introduced by this veto condition is negligibly small.

Continuous Monitoring of Detector Performance

The accuracy and stability of the calibrated photodetector responses can be tested *en-situ* using the optical shot noise of the interferometers themselves. At a given shot noise-dominated frequency, chosen as $f_s = 2.3$ MHz, the diagonal elements of the calibrated 1-s CSD matrices are estimators of the instantaneous shot noise-limited sensitivity of each channel,

$$\widehat{N}_i^{\text{shot, AC}}(t) \;=\; \frac{\textsc{Psd}\left[\widetilde{V}_i^{\text{HF}}(f_s,\, t)\right]}{\left|\widehat{H}_i^{\text{HF}}(f_s,\, t)\right|^2} \,. \tag{4.20}$$

If the photodetector is unstable, the true transimpedance gains of its readout channels, $\left|Z_i^{\text{AC, true}}(t)\right|$ and $\left|Z_i^{\text{DC, true}}(t)\right|$, will drift over time relative to the values inferred from the *ex-situ* signal-readout calibrations, $\left|\widehat{Z}_i^{\text{AC}}\right|$ and $\left|\widehat{Z}_i^{\text{DC}}\right|$. Expanded in terms of its dependence on the detector gains, eq. 4.20 is

$$\widehat{N}_i^{\text{shot, AC}}(t) \;=\; \frac{\left|Z_i^{\text{AC, true}}(t)\right|^2}{\left|Z_i^{\text{DC, true}}(t)\right|^2} \cdot \frac{\left|\widehat{Z}_i^{\text{DC}}\right|^2}{\left|\widehat{Z}_i^{\text{AC}}\right|^2} \cdot N_i^{\text{shot, true}}(t) \,, \tag{4.21}$$

where $N_i^{\text{shot, true}}(t)$ is the true shot noise-limited sensitivity. As can be seen, any deviation of the true detector gains from their *ex-situ*-measured values systematically biases the calibration-referred experimental sensitivity.

At the same time, the continuously-monitored DC photovoltages, $V_i^{\text{DC}}(t)$, provide a

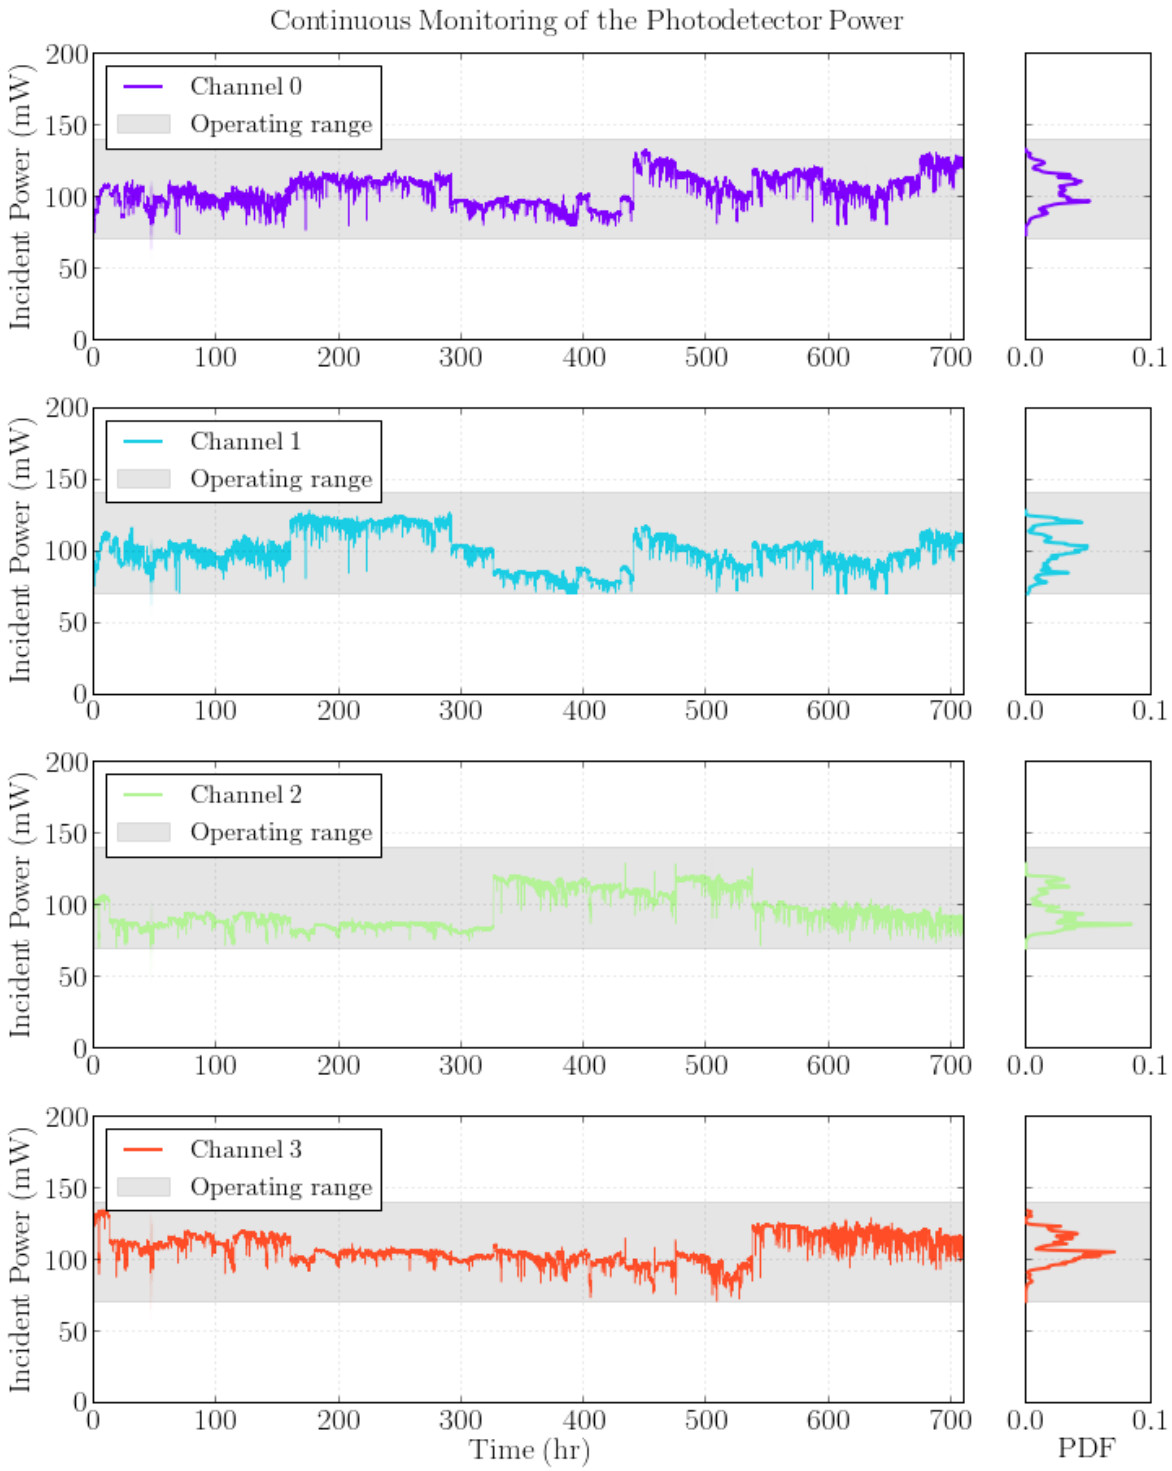

Figure 4.10: Time history of the optical power incident on the modified New Focus 1811 photodetectors. The incident power is directly inferred from the continuously read-out DC voltage of each detector. The regular, step-like discontinuities reflect periodic realignments of the inteferometers during data collection.

second, direct estimator of the shot-noise-limited sensitivity,

$$\widehat{N}_i^{\text{shot, DC}}(t) \;=\; \frac{2e\left|\widehat{Z}_i^{\text{DC}}\right| V_i^{\text{DC}}(t)}{\left|H_i^{\text{LF}}(f_0,\, t)\right|^2}\,, \tag{4.22}$$

which is independent of the AC-channel response. In eq. 4.22, the numerator is the voltage noise power in the DC channel (units V^2/Hz) and the denominator is the directly-measured LF calibration (units V/m). Expanded in terms of its dependence on the DC-channel detector gain, eq. 4.22 is

$$\widehat{N}_i^{\text{shot, DC}}(t) \;=\; \frac{\left|\widehat{Z}_i^{\text{DC}}\right|}{\left|Z_i^{\text{DC, true}}(t)\right|} \cdot N_i^{\text{shot, true}}(t)\,. \tag{4.23}$$

As can be seen, photodetector instability produces a *different* systematic bias in this estimator than in the calibration-referred estimator (eq. 4.21).

The different biasing of the two shot noise-limited sensitivity estimators decouples photodetector performance from the intrinsically time-varying AS-port power. The ratio of the two estimators yields the test statistic

$$\frac{\widehat{N}_i^{\text{shot, AC}}(t)}{\widehat{N}_i^{\text{shot, DC}}(t)} \;=\; \frac{\left|Z_i^{\text{AC, true}}(t)\right|^2 \Big/ \left|Z_i^{\text{DC, true}}(t)\right|}{\left|\widehat{Z}_i^{\text{AC}}\right|^2 \Big/ \left|\widehat{Z}_i^{\text{DC}}\right|}\,, \tag{4.24}$$

representing the instantaneous fractional deviation of the true photodetector response ratio from that of the *ex-situ* calibration measurement. If the photodetector response is stable, this statistic will equal a constant in time, and, if the *ex-situ* calibration is accurate, then the value of this constant will equal unity. Eq. 4.24 is sensitive to gain drift of either the AC- or DC-channel response individually and to common drift of both responses, through the quadratic versus linear dependence on the AC- and DC-channel gain biases, respectively.

Fig. 4.11 shows the stability statistic of each photodetector over the full time history

of data collection. The flatness of the time histories are indicative of stable photodetector operation. The mean of each time history is consistent with unity to $< 2\%$, indicating close agreement of the instantaneous detector responses with the *ex-situ* calibrations. Moreover, the test statistic distributions, shown in the right panels, are all 5-10% in width, consistent with the total uncertainty inferred from the *ex-situ* detector response measurements. This statistic thus provides compelling evidence that the modified high-power photodetectors are both highly stable in their operation and well-characterized by the *ex-situ* calibration measurements.

Any 1-s CSD matrix whose duration coincides with unstable photodetector operation, defined as an excursion larger than $\pm 20\%$ in the stability statistic of any detector, is automatically rejected from analysis. Vetos triggered by this condition are exceedingly rare ($<< 1\%$ of the data) and are found to occur almost exclusively as an interferometer is losing lock, where the assumption of shot-noise-dominated signals breaks down. As such, any sample bias introduced by this veto condition is negligibly small.

4.5.3 ADC Timing Instability

As discussed in §3.5.2, the high-speed ADC units operate under a distributed synchronization model. Each 100-MHz ADC sample clock is phase-locked to its own 1-MHz reference clock generated by a GPS timing card. The sample clocks freely run between every 100 samples, and the reference clocks themselves freely run over the 1-s duration between commonly-received GPS pulses. Relative drift of the sample clock rates between phase realignments causes decoherence in time-averaged measurements, resulting in a systematic degradation in sensitivity to correlated effects.

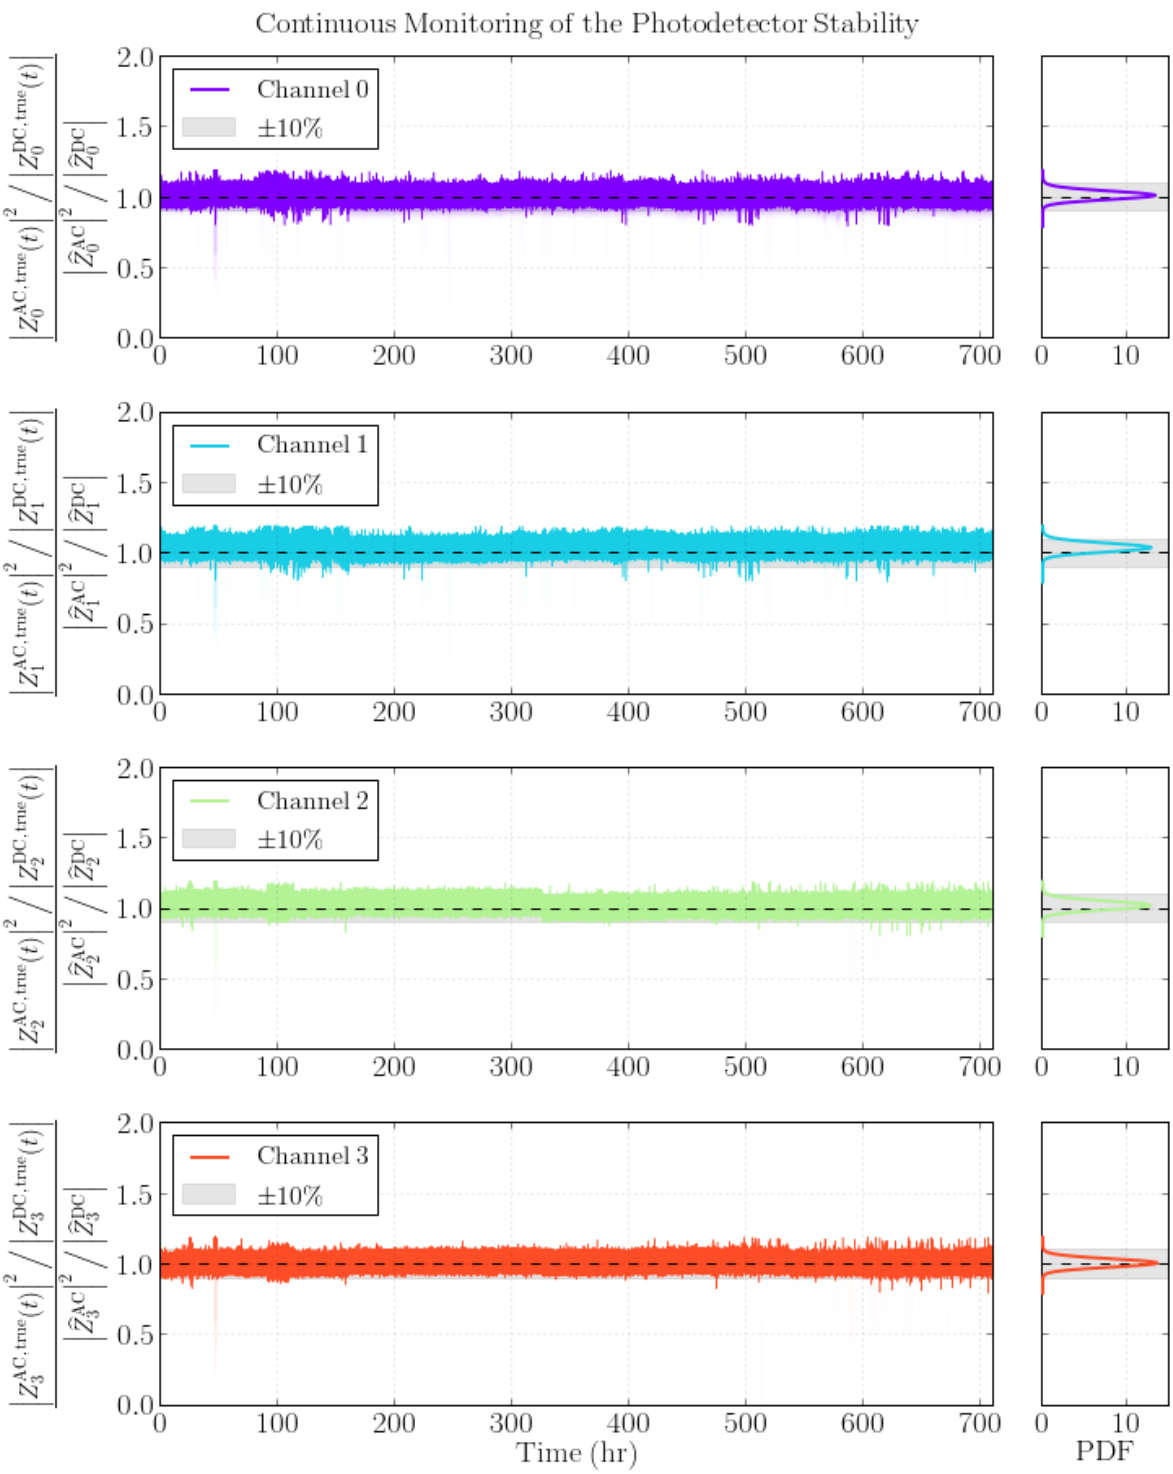

Figure 4.11: Time history of the photodetector transimpedance ratios, a proxy for response stability. The flatness and near-unity values of these time histories are indicative of stable photodetector operation, with the instantaneous detector responses in close agreement with the *ex-situ* calibrations.

Measurement

The degree of clock-rate drift between GPS-disciplined ADC units can be measured by injecting stationary broadband noise from a common source into every channel. Because clock-rate drift occurs on a time scale > 1 s, much longer than the 1-ms DFT duration, a continuous time series of CSD measurements time-resolves the slow drift. To an individual CSD, the drift appears as a constant time offset, τ, in the sampling of the two signals. By the linearity of the Fourier transform, the CSD of a signal, $V(t)$, and a time-shifted version of itself is

$$\begin{aligned} \textsc{Csd}\Big[\mathcal{F}\big[V(t+\tau_i)\big](f),\, \mathcal{F}\big[V(t)\big](f)\Big] &= \textsc{Csd}\left[\mathcal{F}\big[V(t)\big](f)\, e^{i2\pi f\tau}\,,\, \mathcal{F}\big[V(t)\big](f)\right] \\ &= \textsc{Psd}\Big[\mathcal{F}\big[V(t)\big](f)\Big]\, e^{i\phi_\tau(f)}\,, \end{aligned} \tag{4.25}$$

where the CSD phase, $\phi_\tau(f) \equiv 2\pi f\tau$, is determined purely by the time offset.

For each CSD in the time series, a least-squares fit of the parameter τ to the measured phase, $\phi_\tau(f)$, yields the instantaneous time offset between the sample clocks of the two channels. Clock-rate drift manifests as a change in the best-fitting parameter, $\widehat{\tau}$, along the time series. Fig. 4.12 shows this measurement for the two ADC units receiving the interferometer readout signals. The time series of $\widehat{\tau}$ shows that the internal control loop of each ADC allows its sample clock to freely drift until hitting a rail at approximately ± 10 ns, at which point a sudden phase correction is made. Application of the same measurement technique establishes two channels on the same ADC card to be synchronous to < 1 ns.

Clock-Wander Decoherence Model

The sampling time offsets shown in Fig. 4.12 are approximately normally-distributed with a width of $\sigma_\tau \approx 10$ nm. The aggregate statistical effect on a time-averaged CSD measurement can thus be modeled as if each CSD independently draws a random sampling time offset

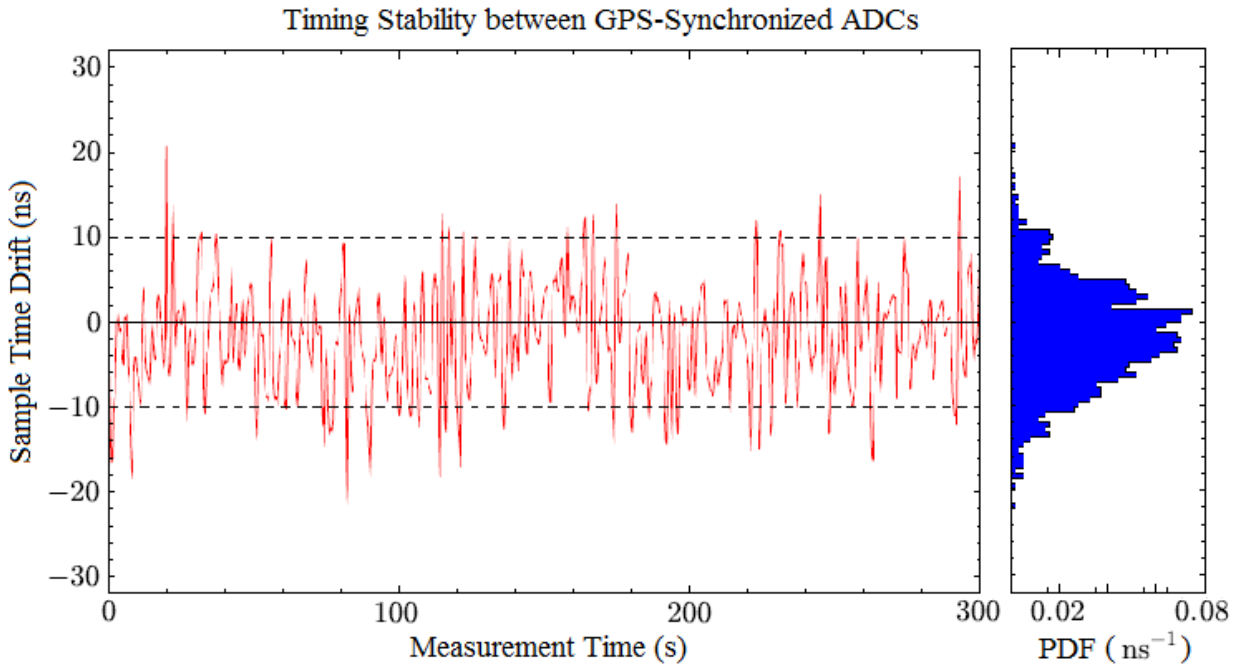

Figure 4.12: Sampling time drift between two GPS-disciplined ADC units measured by coherent broadband noise injection. The time series indicate that the internal control loop of each ADC allows its sample clock to freely drift until hitting a rail at approximately ± 10 ns, at which point a sudden phase correction is made. The distribution of time offsets is asymptotically normal.

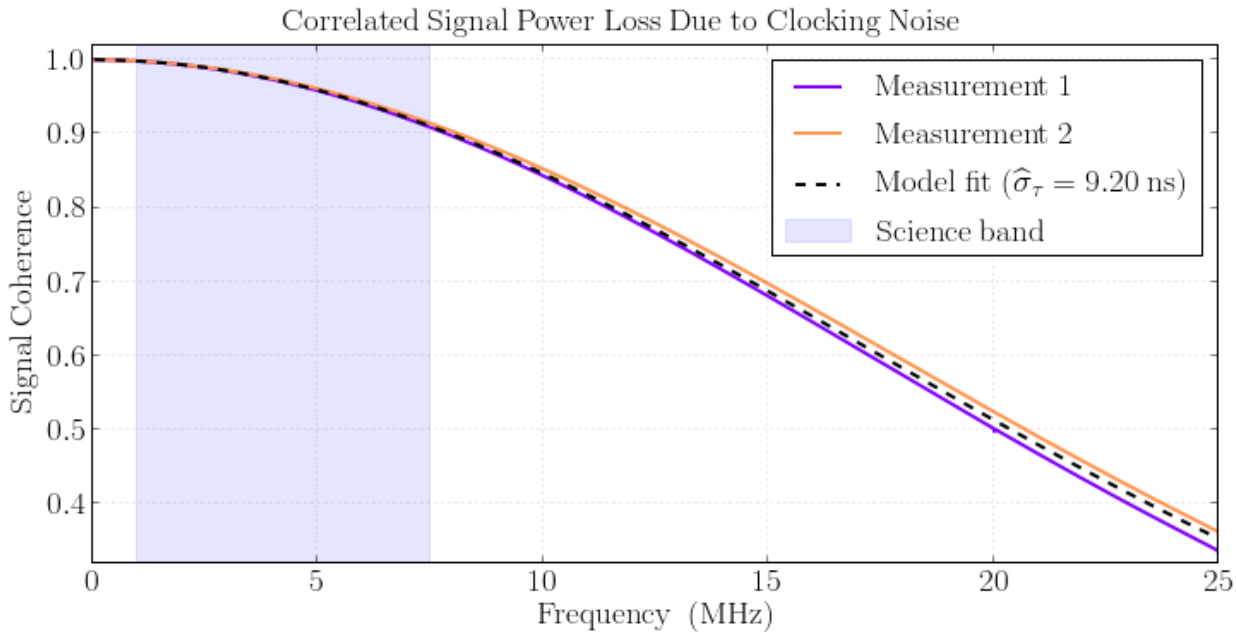

Figure 4.13: Decoherence of time-averaged CSD measurements due to clocking noise between the ADC units. The purple and orange curves denote two three-hour measurements made via coherent broadband noise injection. The dashed black curve represents the fit of the Gaussian decoherence model (eq. 4.26) to the average of the two measurements.

$\tau \sim \mathcal{N}(0,\, \sigma_\tau^2)$. The decoherence of a time-averaged CSD measurement due to clocking noise is then

$$\begin{aligned}
H_{\text{coh}}(f) &= \frac{\left\langle \textsc{Csd}\Big[\mathcal{F}\big[V(t+\tau)\big](f),\, \mathcal{F}\big[V(t)\big](f)\Big]\right\rangle_\tau}{\sqrt{\left\langle \textsc{Psd}\Big[\mathcal{F}\big[V(t+\tau)\big](f)\Big]\right\rangle_\tau \left\langle \textsc{Psd}\Big[\mathcal{F}\big[V(t)\big](f)\Big]\right\rangle_\tau}} \\
&= \left\langle e^{i\phi_\tau(f)} \right\rangle_\tau \\
&= \frac{1}{\sqrt{2\pi\sigma_\tau^2}} \int_{-\infty}^{\infty} e^{i2\pi f\tau}\, e^{-\tau^2/2\sigma_\tau^2}\, d\tau \\
&= \frac{1}{\sqrt{2\pi\sigma_\tau^2}} \int_{-\infty}^{\infty} \cos\left(i2\pi f\tau\right)\, e^{-\tau^2/2\sigma_\tau^2}\, d\tau \\
&= e^{-2\pi^2\sigma_\tau^2 f^2}\,, \qquad (4.26)
\end{aligned}$$

where the parameter σ_τ sets the rate of signal coherence loss with increasing frequency.

Fig. 4.13 shows two measurements of the de-cohering of intrinsically coherent broadband

noise (see §4.5.3), each computed from the time-averaged CSD after three hours of continuous integration. A fit of the Gaussian decoherence model (eq. 4.26) to the average of the two measurements, shown in the figure, yields a best-fit loss parameter of $\widehat{\sigma}_\tau = 9.20$ ns. Because the Holometer signal bandwidth is less than 8 MHz, the inter-channel decorrelation due to clocking phase noise is $< 10\%$ at all science frequencies. To account for the systematic loss of correlated sensitivity, the fitted decoherence function is applied as an overall correction to each of the cross-ADC CSD averages.

Continuous Monitoring of Sampling Synchronicity

As correlated sensitivity must be ensured for the entire 700-hour measurement duration, the instantaneous sampling time offset between the two ADC units receiving the science detector readouts is continuously monitored. This is implemented via direct injection of a coherent narrowband optical signal onto the four high-power photodetectors at the AS port. A 960-nm LED installed in each output optics box is split by a beamsplitter and directed onto the two high-power detectors in each box. The 1-mW power of the two LEDs is coherently modulated via a function generator driving a sine wave at 13 MHZ, above the science band. The instantaneous sampling time offset of each calibrated 1-s CSD matrix is then given by the estimator

$$\widehat{\tau} \;=\; \frac{\phi_\tau(f)}{2\pi f}\,, \tag{4.27}$$

evaluated at $f = 13$ MHz.

Fig. 4.14 shows the history of sampling time drift between all four science readout channels. As with the *ex-situ* measurements, the two ADC units remain tightly synchronized within a 10-ns envelope over the entire 700 hours. The fit of a normal distribution to the sample time offsets yields a best-fit width parameter of $\widehat{\sigma}_\tau = 10.40$ ns, consistent with that

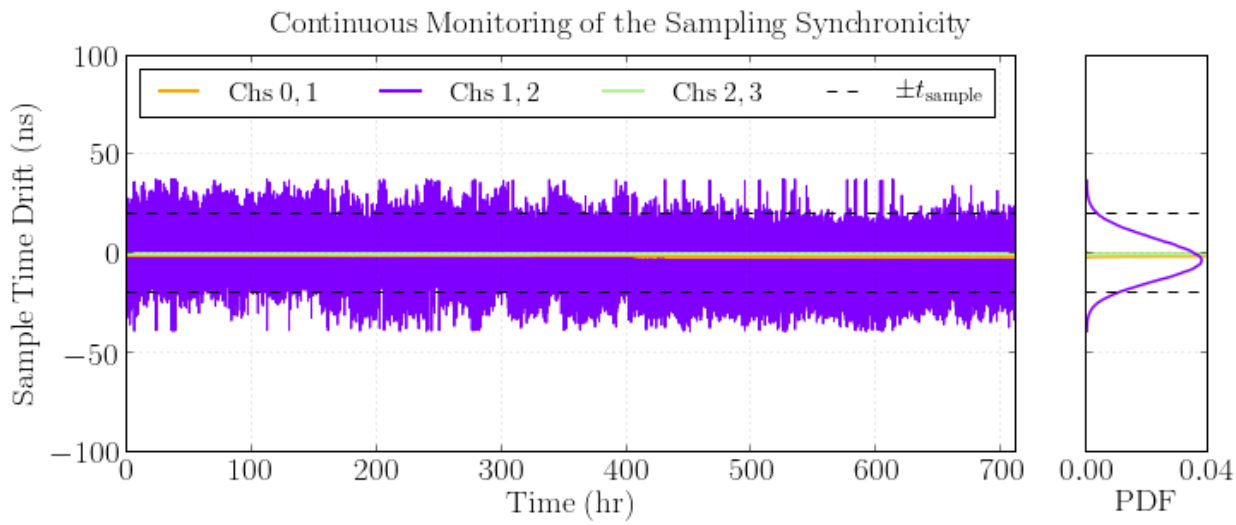

Figure 4.14: Time history of the sampling synchronicity between the interferometer readout channels. This time history demonstrates that the two ADC units remain tightly synchronized within a 10-ns envelope over the entire 700 hours. The fit of a normal distribution to the time offsets yields a best-fit width parameter of $\widehat{\sigma}_\tau = 10.40$ ns and a mean of 2.17 ns, likely reflecting a small phase difference between the two separate sets of cabling connecting the LEDs to the function generator. This small offset corresponds to a phase difference of $< 1^\circ$ at 1 MHz.

of the *ex-situ* decoherence measurements. The mean of the distribution is offset from zero by 2.17 ns, likely reflecting a small phase difference between the two separate sets of cabling connecting the LEDs to the function generator. In any case, this offset produces a phase difference of $< 1^\circ$ at 1 MHz. In principle, the CSD phase information could be used to re-phase each 1-s correlation matrix across the entire 700 hours, thereby recovering the decoherence loss. However, in this analysis, the decoherence is simply absorbed as a (small) loss, as the $< 5\%$ gain in sensitivity was not determined to merit the implementation time.

Any 1-s CSD matrix whose duration coincides with a loss of signal coherence, defined as the cross-ADC LED-drive coherence falling below 50%, is automatically rejected from analysis. Vetos triggered by this condition are rare ($< 1\%$ of the data) and are found to occur almost exclusively as an interferometer is losing lock, when the signal coherence falls due to a strong transient increase in the interferometer noise power. As with the previous vetos, any sample bias introduced by this condition is negligibly small.

CHAPTER 5

ANALYSIS & RESULTS

The time series of 1-s correlation matrices are individually calibrated to differential length via the estimator

$$\widehat{CSD}_{ij}(f,\,t) \;\equiv\; \frac{\textsc{Csd}\left[\widetilde{V}_i^{\mathrm{HF}}(f,\,t)\,,\,\widetilde{V}_j^{\mathrm{HF}}(f,\,t)\right]}{\widehat{H}_i^{\mathrm{HF}}(f,\,t)\cdot\left(\widehat{H}_j^{\mathrm{HF}}(f,\,t)\right)^*}\,, \tag{5.1}$$

where $\widehat{H}_i^{\mathrm{HF}}(f,\,t)$ (units V/m) is the continuously-inferred sensitivity of each readout channel. The estimation of the channel calibrations is discussed at length in §4. Post-processing analysis is then performed on the calibrated data to assess its quality and to place physical constraints on exotic spatial shear. The following sections describe each stage of the post-processing analysis.

5.1 Sensitivity Enhancement by Averaging

Unlike gravitational wave detectors, which must time-resolve transient geometrical disturbances in the spatial background, the Holometer is searching for a universal, stationary quantization noise of the background itself. The stationary statistical nature of the effect makes it possible to achieve large sensitivity enhancements through the averaging of independent measurements, each of which can be made at an arbitrarily different time. Averaging the time series of CSD matrices reduces the measurement variance, extracting the faint correlated noise background common to every measurement. The time series of calibrated 1-s CSD matrices are thus averaged over the full 700-hour exposure time as

$$\widehat{CSD}_{ij}(f) \;\equiv\; \sum_{n=0}^{N-1} w_n\,\widehat{CSD}_{ij}(f,\,t_n)\,, \tag{5.2}$$

where t_n represents each measurement time and N is the total number of 1-s CSD measurements. The time-averaged CSD matrix is notated similarly to the 1-s CSD matrices, but with the time dependence removed.

To make eq. 5.2 a maximum-likelihood estimator, each 1-s CSD measurement is optimally weighted by the inverse of its signal-band-integrated variance as

$$w_n = \frac{\left(\sum_l \mathrm{Var}\left[\widehat{CSD}_{ij}(f_l,\, t_n)\right]\right)^{-1}}{\sum_{m=0}^{N-1}\left(\sum_l \mathrm{Var}\left[\widehat{CSD}_{ij}(f_l,\, t_m)\right]\right)^{-1}}\,, \tag{5.3}$$

where the l summation is taken over all frequencies $f_l \in (1\ \mathrm{MHz},\ 7.5\ \mathrm{MHz})$. A truly optimized estimator would assign an inverse-variance weight to each frequency bin individually, under which $w_n \to w_n(f)$. However, because the dominant source of time variation is the interferometer transfer function, which rescales all frequencies equally (see Fig. 3.3), a band-integrated statistic produces an equivalent weighting in practice, with less computational overhead.

The variance of each Hann-windowed, 50%-overlapped CSD measurement is given by the estimator

$$\mathrm{Var}\left[\widehat{CSD}_{ij}(f,\, t)\right] = \frac{\widehat{CSD}_{ii}(f,\, t)\cdot\widehat{CSD}_{jj}(f,\, t)}{0.947\, N_{\mathrm{meas}}}\,, \tag{5.4}$$

where N_{meas} is the number of individual DFT batches in the 1-s time average. The effective number of *independent* measurements is approximately 95% of this value due to the covariance introduced by the 50% overlap of the DFT batches (see §3.5.3). A derivation of this estimator is given in §IV.C of [7].

The covariance between 1-s CSD measurements due to DFT overlap (significant on 1 ms time separations) is negligibly small. Each time-averaged CSD can be thus regarded as the

weighted sum of N independent random variables, whose reduced variance is estimated by

$$\mathrm{Var}\left[\widehat{CSD}_{ij}(f)\right] = \sum_{n=0}^{N-1} w_n^2 \,\mathrm{Var}\left[\widehat{CSD}_{ij}(f,\, t_n)\right] . \tag{5.5}$$

For concreteness of illustration, consider the idealized case in which all 1-s CSD measurements have equal variance σ^2. The weight of each measurement is then $1/N$, and eq. 5.5 can be seen to imply that the variance of the measurement average,

$$\begin{aligned} \mathrm{Var}\left[\widehat{CSD}_{ij}(f)\right] &= \frac{1}{N^2}\sum_{n=0}^{N-1}\sigma^2 \\ &= \frac{\sigma^2}{N} , \end{aligned} \tag{5.6}$$

is reduced by a factor of N relative to the single-measurement variance.

Fig. 5.1 shows the lower triangle of the time-averaged CSD matrix of the four science readout channels. Because the DFT is performed on purely real input data, the upper triangle matrix elements are Hermitian conjugates of the lower elements and are thus redundant. Channels 0 and 1 correspond to the L interferometer signals at the AS port. Likewise, channels 2 and 3 correspond to the T interferometer signals. The (1, 0) and (3, 2) CSDs represent the cross-correlation of each interferometer beam, split onto two photodetectors, with itself. Because each detector measures independent shot noise, this time-averaged measurement extracts purely the thermal and mechanical component of interferometer noise, which is preserved by the splitting of the output beam. Its significance is discussed further in the following section.

5.2 Analysis of Optics Thermal Noise

This section will demonstrate that the etalon correlations present in each single-interferometer CSD arise from thermally-excited vibrations of the end mirrors and beamsplitters. The ad-

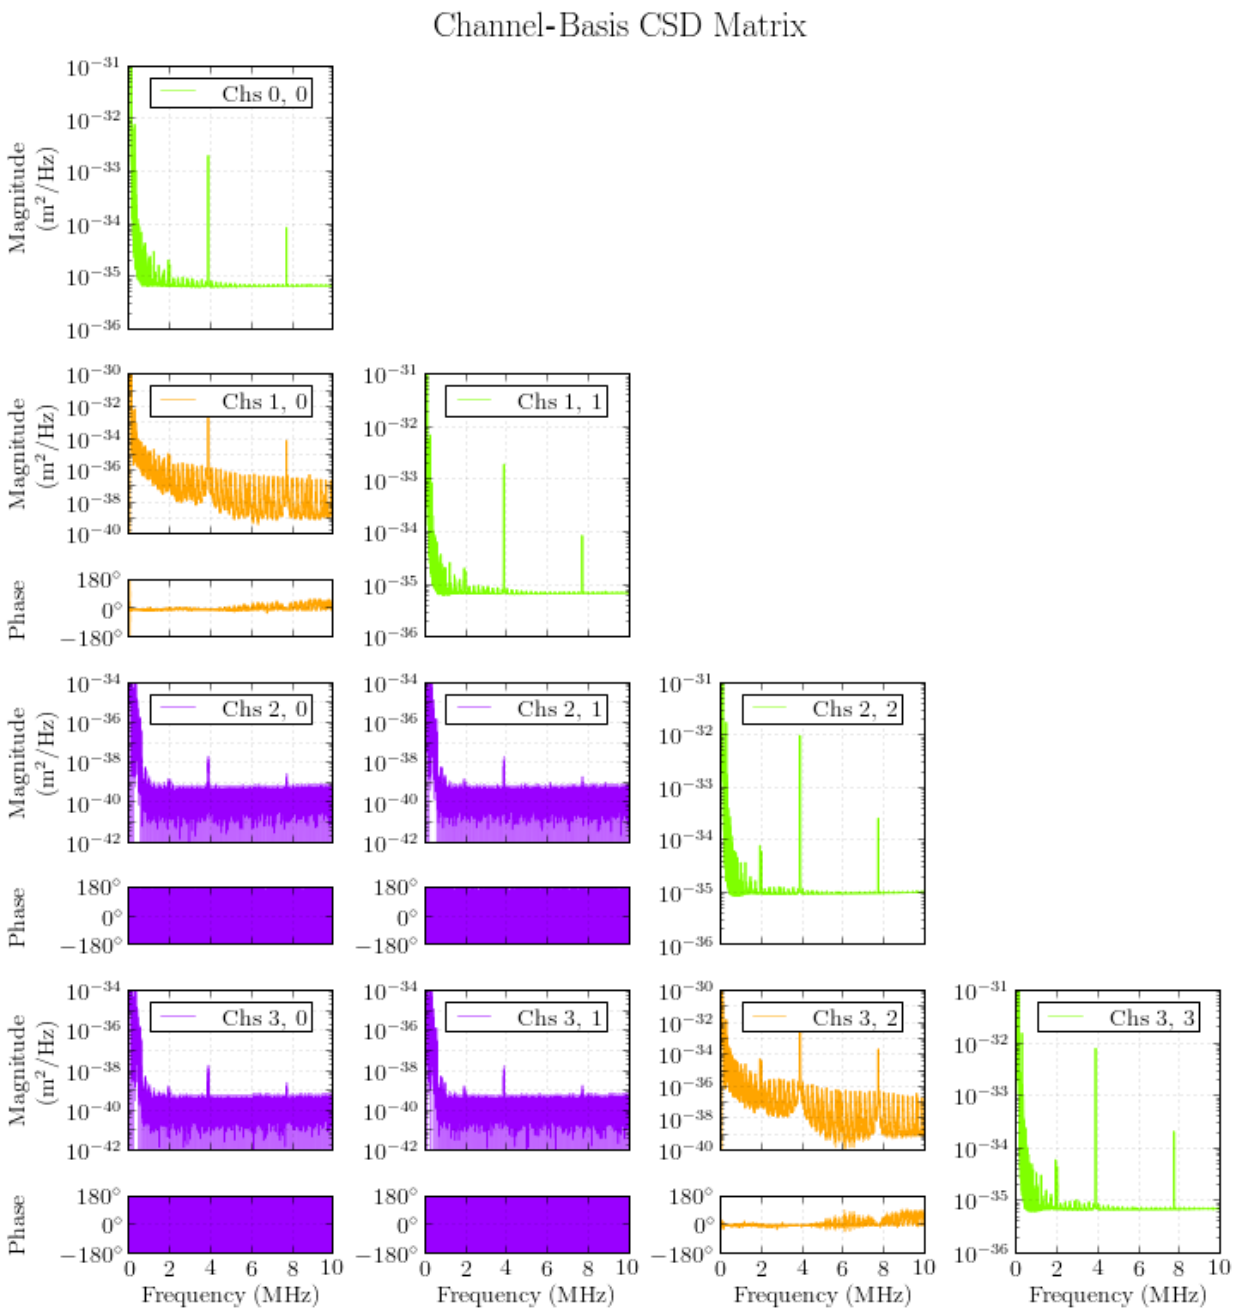

Figure 5.1: CSD matrix measured between the AS-port detector readout channels. Only the lower triangle of spectra are shown, as the upper triangle contains no additional information. Channels 0 and 1 correspond to the L interferometer signals and channels 2 and 3 correspond to the T interferometer signals. The diagonal elements (green curves) represent the PSD of each readout channel. The 2-0, 2-1, 3-0, and 3-1 CSDs (blue curves) represent the cross-interferometer measurements. The 1-0 and 3-2 CSDs (orange curves) represent the cross-correlation of each interferometer beam, split onto two photodetectors, with itself. These terms extract purely the thermal and mechanical noise component in each interferometer.

ditive thermal noise is of small consequence for cross-interferometer CSD measurement, as it does not correlate between the two independent optical systems. However, it is significant in demonstrating the ability of the Holometer to interferometrically detect a true RMS motion, below shot noise, at radio frequencies. The thermal noise is used to validate the calibration of each interferometer to within a factor of two, limited by the unknown mixing of thermally-excited angular and longitudinal noise in the AS-port response.

5.2.1 Spectral Model of Thermal Excitations

Fig. 5.2 provides an enlarged view of the single-interferometer CSDs. As can be seen, each noise spectrum is punctuated with a regular set of spectral lines spaced approximately every 226 kHz. The frequencies of the peak-amplitude lines are consistent with harmonics of the round-trip sound-crossing time in the end mirrors and beamsplitter. Each quartz optic, of thickness $d = 12.7$ mm, has a sound speed of $v_s = 5720$ m/s. As planar vibrational modes in a cylinder have a fundamental frequency of $v_s/2d = 225$ kHz, these lines are identified as thermal excitations of the fundamental planar mode and its higher harmonics. Since the smallest dimension of the optics is their thickness, the other vibrational modes with radial and azimuthal dependence will have different, more closely-spaced excitation frequencies. These modes are interpreted to correspond to the remaining system of peaks seen regularly spaced in Fig. 5.2.

In thermal equilibrium, each vibrational mode of the optical elements has energy $k_bT/2$ at temperature T. Each mode forms a standing wave inside the optics, driving a mean-square surface motion of $\langle x^2 \rangle = k_bT/m\omega^2$, where $m = 30$ g is the effective mass of the mirror and ω is the angular frequency of the normal mode oscillation. The PSD of each resonant response is a Lorentzian curve,

$$L(f) \;=\; h\,\frac{(\Gamma/2)^2}{(f-f_0)^2+(\Gamma/2)^2}\,, \tag{5.7}$$

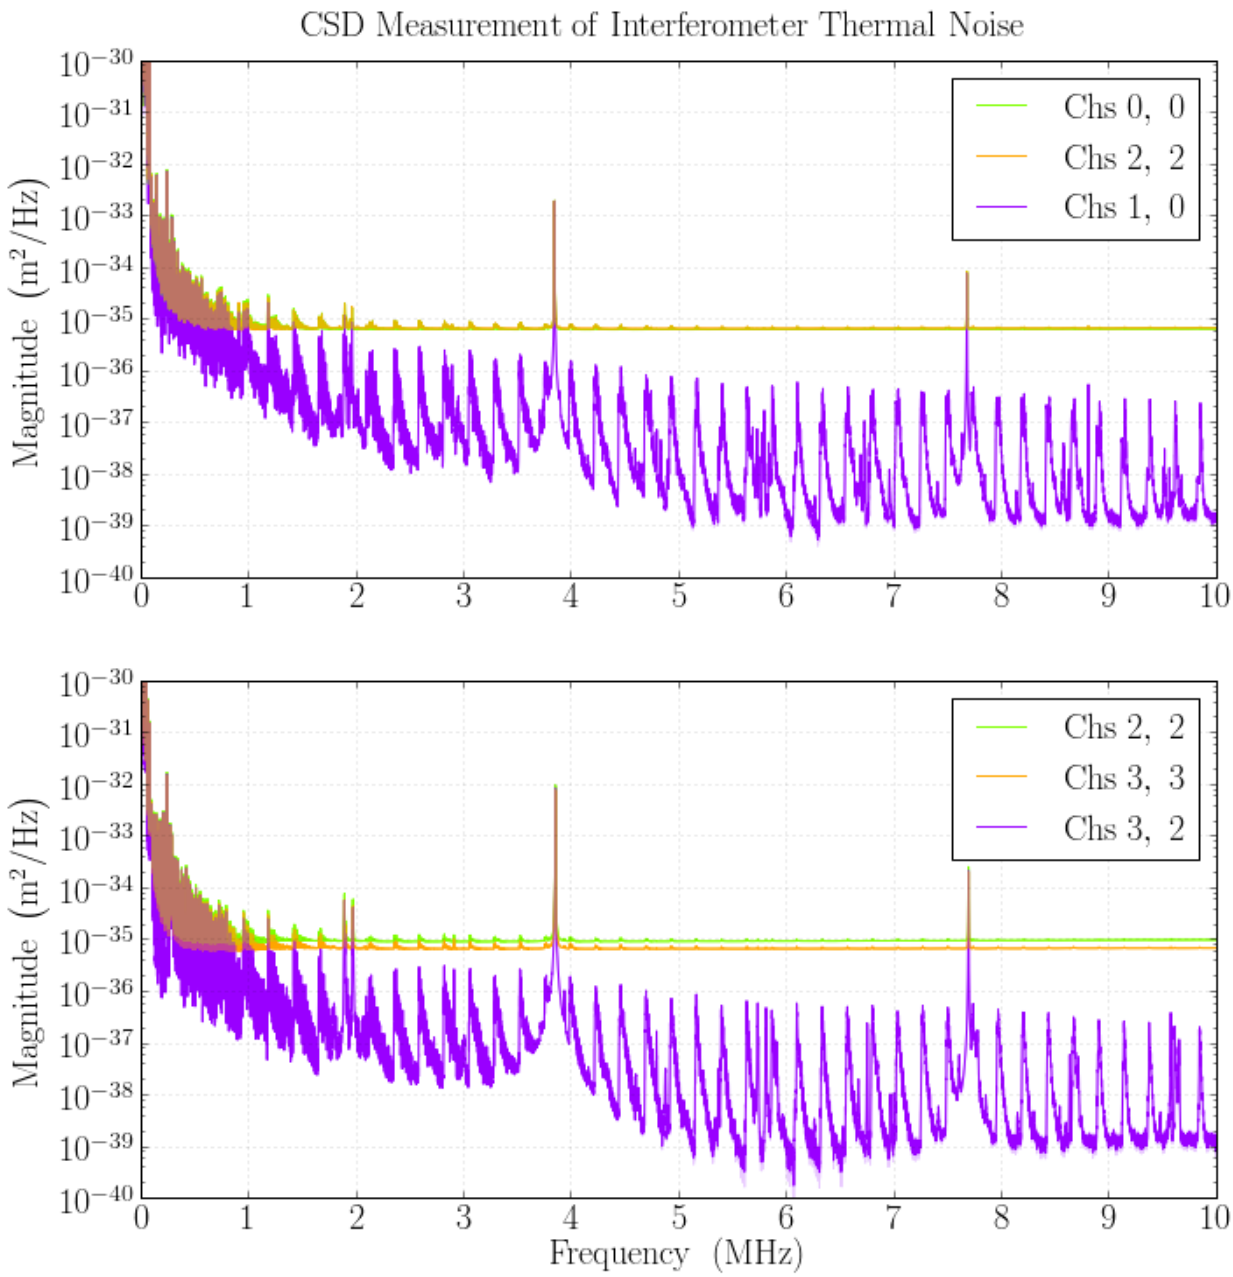

Figure 5.2: Thermal noise of the optics measured via the single-interferometer CSDs. Each noise spectrum is punctuated with a regular set of spectral lines spaced approximately every 226 kHz, consistent with harmonics of the round-trip sound-crossing time in the end mirrors and beamsplitter. The peak-amplitude lines are identified as thermal excitation of the fundamental planar vibrational mode and its higher harmonics. The remaining system of regularly-spaced peaks is identified as other vibrational modes with radial and azimuthal dependence.

where h is the height of the peak, f_0 is the central frequency, and Γ is the width. When the DAQ system is operated at its highest resolution, 23 Hz, each line is individually resolved (the nominal operating resolution is 570 Hz; see §3.5.3). The integral of an individual Lorentzian,

$$\int L(f)\, df \;=\; \frac{h\Gamma}{2\pi}\,, \tag{5.8}$$

is equal to the mean-square longitudinal displacement, $\langle x^2 \rangle$. In principle, these vibrational modes provide a set of direct, independent calibration lines spanning the entire measurement band. However, their application is limited in practice by the unknown mixing of thermally-excited angular and longitudinal noise, as will be discussed below.

5.2.2 Calibration Using the Optics Thermal Noise

Two methods have been used to check the interferometer calibrations against the expected mean-square thermal motion. The first method directly compares the integrated power in a single calibrated line to that which is expected under the thermal model at equilibrium, $\langle x^2 \rangle = k_b T/m\omega^2$. However, because vibrational modes with azimuthal and radial dependence do not align with 100% efficiency to the optical path of the interferometer arm, their plane waves are expected to contribute less than $k_b T/m\omega^2$ to the longitudinal length fluctuation spectrum. In practice, it is difficult to identify purely longitudinal modes due to the complicating effects of the mounting and support structure. While many lines are measured to contribute less than $k_b T/m\omega^2$, no line has ever been found to *exceed* the thermal $k_b T/m\omega^2$ bound, which would be indicative of an overestimated calibration sensitivity. As some lines do nearly saturate the $k_b T/m\omega^2$ bound, the measurements made via this method are concluded to be consistent with the interferometer calibrations.

The second method integrates each single-interferometer CSD over the band of many etalons, subtracting the residual non-thermal background. Unlike the first method, com-

puting the band-integrated noise power does not require the thermal lines to be resolved. The band-integrated noise power serves as a time stability statistic. Although the intrinsic bandlimited thermal noise power is unknown, its inferred value should remain constant in time if the interferometer calibration is stable. However, the etalon correlations in practice measure an unknown linear combination of longitudinal and angular noise power, where the angular noise component is alignment-dependent. As such, time drift in the calibration accuracy cannot be reliably decoupled from angular alignment drift. Over the entire 700-hour span of the data, this statistic varies by a factor of approximately two, which places an upper limit on the possible time drift in calibration accuracy.

5.3 Interferometer Signal Recombination

Because more power exits the interferometers than can be absorbed by a single photodetector (see §3.5.1), each AS-port beam is split by a beamsplitter and directed onto two detectors. The calibrated, time-averaged CSD matrix is naturally calculated in the basis of the four detector readout channels. Its elements,

$$\widehat{CSD}_{ij}(f) \;=\; \widetilde{S}_i(f) \cdot \widetilde{S}_j^*(f) \;, \tag{5.9}$$

are interpreted as conjugate-products of the calibrated readout signals, $\widetilde{S}_i(f)$. In order to extract full statistical power from the cross-interferometer measurement, the CSD matrix must be transformed from this quad-channel representation to one in terms of the two recombined interferometer signals.

The recombined signal of each interferometer is constructed from the weighted average

of its two detector signals,

$$\begin{bmatrix} \widetilde{S}_L(f) \\ \widetilde{S}_L(f) \end{bmatrix} \equiv \begin{bmatrix} w_0\,\widetilde{S}_0(f) + w_1\,\widetilde{S}_1(f) \\ w_2\,\widetilde{S}_2(f) + w_3\,\widetilde{S}_3(f) \end{bmatrix} \tag{5.10}$$

$$= \begin{bmatrix} w_0 & w_1 & 0 & 0 \\ 0 & 0 & w_2 & w_3 \end{bmatrix} \begin{bmatrix} \widetilde{S}_0(f) \\ \widetilde{S}_1(f) \\ \widetilde{S}_2(f) \\ \widetilde{S}_3(f) \end{bmatrix} ,$$

via the linear transformation matrix

$$M \equiv \begin{bmatrix} w_0 & w_1 & 0 & 0 \\ 0 & 0 & w_2 & w_3 \end{bmatrix} . \tag{5.11}$$

The detector signal weights are subject to the normalization constraints $w_0 + w_1 = 1$ and $w_2 + w_3 = 1$.

To make eqs. 5.10 maximum-likelihood estimators, each detector readout signal is optimally weighted by the inverse of its signal-band-integrated variance. The detector signal weights for the L interferometer are thus

$$w_0 = \frac{\sum_l \widehat{CSD}_{11}(f_l)}{\sum_l \widehat{CSD}_{00}(f_l) + \sum_l \widehat{CSD}_{11}(f_l)} \tag{5.12}$$

and

$$w_1 = \frac{\sum_l \widehat{CSD}_{00}(f_l)}{\sum_l \widehat{CSD}_{00}(f_l) + \sum_l \widehat{CSD}_{11}(f_l)} , \tag{5.13}$$

where the l summation is taken over all frequencies $f_l \in (1\text{ MHz},\ 7.5\text{ MHz})$. Similarly, the detector signal weights for T interferometer are

$$w_2 = \frac{\sum_l \widehat{CSD}_{33}(f_l)}{\sum_l \widehat{CSD}_{22}(f_l) + \sum_l \widehat{CSD}_{33}(f_l)} \tag{5.14}$$

and

$$w_3 = \frac{\sum_l \widehat{CSD}_{22}(f_l)}{\sum_l \widehat{CSD}_{22}(f_l) + \sum_l \widehat{CSD}_{33}(f_l)} \ . \tag{5.15}$$

As the dominant source of sensitivity difference between each pair of detectors is simply their difference in incident power, which rescales all frequencies equally, this band-integrated statistic is equivalent in practice to an individualized frequency-bin weighting.

Under the linear transformation M, the 4×4 channel-basis CSD matrix transforms into the 2×2 common-interferometer basis[1] as $\widehat{CSD}(f) \to M\,\widehat{CSD}(f)\,M^{\mathrm{T}}$. The diagonal elements,

$$\widehat{CSD}_{LL}(f) = w_0^2\,\widehat{CSD}_{00}(f) + w_1^2\,\widehat{CSD}_{11}(f) + 2w_0w_1\,\mathrm{Re}\left[\widehat{CSD}_{10}(f)\right] \tag{5.16}$$

and

$$\widehat{CSD}_{TT}(f) = w_2^2\,\widehat{CSD}_{22}(f) + w_3^2\,\widehat{CSD}_{33}(f) + 2w_2w_3\,\mathrm{Re}\left[\widehat{CSD}_{32}(f)\right] \ , \tag{5.17}$$

represent the PSDs of the reconstructed interferometer signals. Similarly, the off-diagonal

1. The other two basis vectors spanning this space are the weighted *differences* between detector signals, which are not of experimental interest.

element,

$$\widehat{CSD}_{TL}(f) = w_2 w_0 \widehat{CSD}_{20}(f) + w_2 w_1 \widehat{CSD}_{21}(f) + w_3 w_0 \widehat{CSD}_{30}(f) + w_3 w_1 \widehat{CSD}_{31}(f) \ , \tag{5.18}$$

represents the CSD of the two interferometer signals.

Because each detector signal is dominated by optical shot noise above 1 MHz, which does not correlate between channels, each CSD matrix element in the interferometer basis can be treated as a weighted sum of independent random variables. The variance estimators of the PSDs are then

$$\mathrm{Var}\left[\widehat{CSD}_{LL}(f)\right] = w_0^4 \,\mathrm{Var}\left[\widehat{CSD}_{00}(f)\right] + w_1^4 \,\mathrm{Var}\left[\widehat{CSD}_{11}(f)\right] + 2 w_0^2 w_1^2 \,\mathrm{Var}\left[\widehat{CSD}_{10}(f)\right] \tag{5.19}$$

and

$$\mathrm{Var}\left[\widehat{CSD}_{TT}(f)\right] = w_2^4 \,\mathrm{Var}\left[\widehat{CSD}_{22}(f)\right] + w_3^4 \,\mathrm{Var}\left[\widehat{CSD}_{33}(f)\right] + 2 w_2^2 w_3^2 \,\mathrm{Var}\left[\widehat{CSD}_{32}(f)\right] \ , \tag{5.20}$$

and the variance estimator of the cross-interferometer CSD is

$$\mathrm{Var}\left[\widehat{CSD}_{TL}(f)\right] = w_2^2 w_0^2 \,\mathrm{Var}\left[\widehat{CSD}_{20}(f)\right] + w_2^2 w_1^2 \,\mathrm{Var}\left[\widehat{CSD}_{21}(f)\right] + w_3^2 w_0^2 \,\mathrm{Var}\left[\widehat{CSD}_{30}(f)\right] + w_3^2 w_1^2 \,\mathrm{Var}\left[\widehat{CSD}_{31}(f)\right] \ . \tag{5.21}$$

Considering the idealized case of equal detector sensitivities, it is clear that eq. 5.21, representing the average of four independent measurements, implies a variance reduction of 4 relative to each channel-basis CSD.

Fig. 5.3 shows the reconstructed interferometer-basis CSD matrix. The cross-interferometer spectrum detects strong correlated noise below 1 MHz, whose power exceeds the predicted exotic noise power by > 3 orders of magnitude. As it exhibits near-perfect coherence with environmental monitors of the laser amplitude and phase noise (see §3.2.3), this low-frequency noise is positively identified as external contamination injected by the lasers. On this basis, the band below 1 MHz is rejected from analysis.

At each frequency, the degree of correlation between the two interferometer signals is quantified by their coherence, defined as the ratio of the CSD magnitude and the geometric mean of the two PSDs. The cross-interferometer coherence is shown in Fig. 5.4. As discussed above, the band below 1 MHz is dominated by environmental noise injected by the lasers and is thus rejected from analysis. Above 1 MHz, however, the AS-port detector signals do not exhibit measurable coherence with monitors of the laser noise and the radio-frequency environment. Indeed, as Fig. 5.4 shows, the noise power of the two interferometer signals in the 1-10 MHz band is found to be uncorrelated to < 1 part in $10,000$.

5.4 Spectral Leakage Estimation

In anticipation of modeling the measured cross-interferometer CSD, the spectral leakage due to the DFT is first calculated. Spectral leakage introduces covariance between neighboring frequency bins. Although the modeling could, in principle, account for these correlations, its implementation is greatly simplified when each frequency bin can be treated as an independent measurement. The high-speed processing pipeline applies a standard Hann window to the DFT data,

$$w_n = \sin^2\left(\frac{\pi(n+1)}{N_{\mathrm{DFT}}+2}\right) , \quad n = 0,\, 1, ...,\, N_{\mathrm{DFT}} - 1 , \tag{5.22}$$

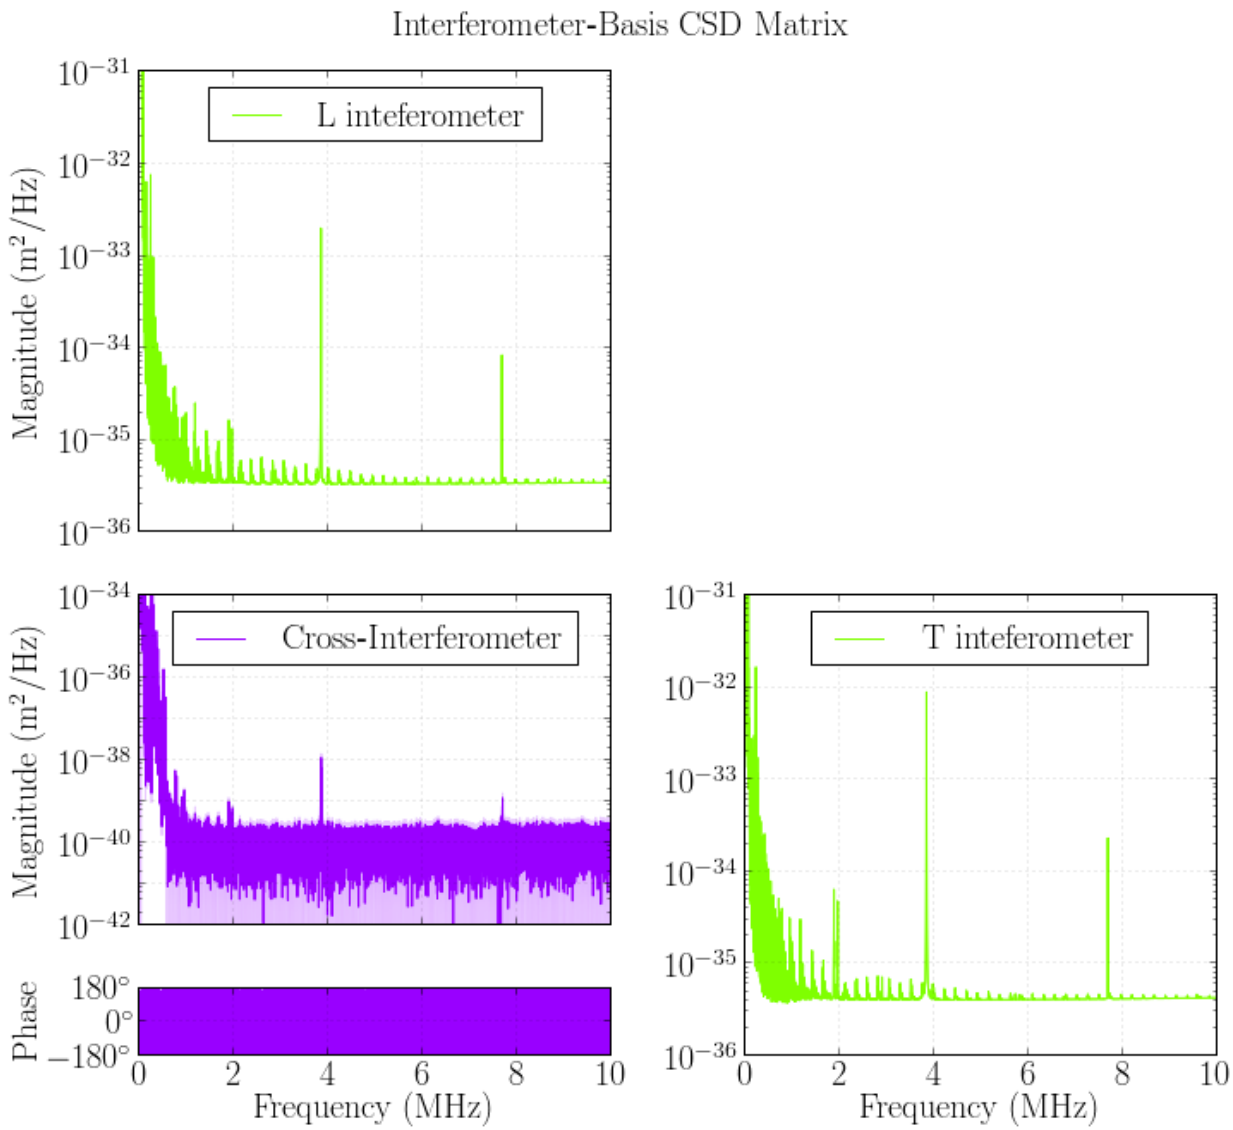

Figure 5.3: CSD matrix measured between the AS-port interferometer signals. The diagonal elements (green curves) represent the PSD of each interferometer signal. The cross-interferometer spectrum (blue curve) detects strong correlated noise below 1 MHz, whose power exceeds the predicted exotic noise power by > 3 orders of magnitude. As it exhibits near-perfect coherence with environmental monitors of the laser amplitude and phase noise (see §3.2.3), this low-frequency noise is positively identified as external contamination injected by the lasers. On this basis, the band below 1 MHz is rejected from analysis.

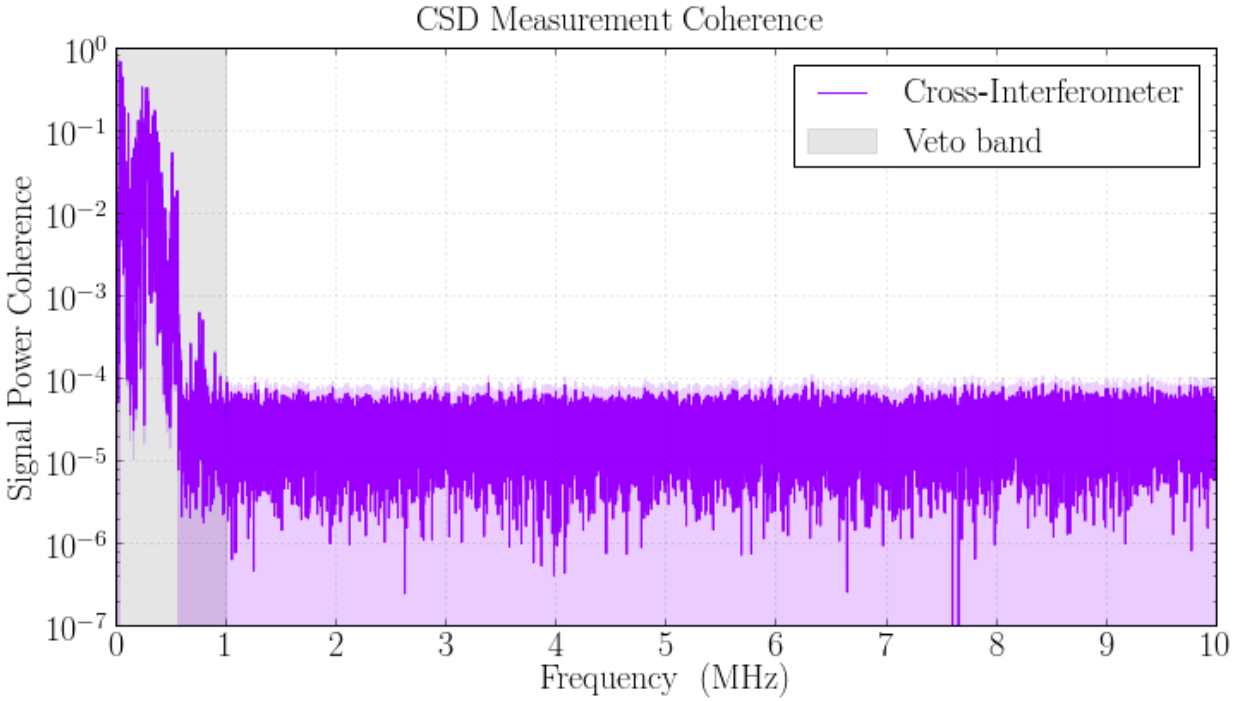

Figure 5.4: Coherence of the cross-interferometer CSD measurement. The band below 1 MHz is rejected from analysis due to strong environmental noise injected by the lasers (see the text). Above 1 MHz, the AS-port detector signals are free from all known forms of radio-frequency environmental contamination. The coherence shows the noise power of the two interferometer signals in the 1-10 MHz band to be uncorrelated to < 1 part in $10,000$.

whose equivalent noise bandwidth of 1.5 bins (570 Hz) sets the frequency-space correlation scale. A set of independent frequency measurements can be generated by averaging every contiguous set of $N >> 1.5$ bins, producing a rebinned CSD of lowered frequency resolution. Because the new bins are much wider than the 570-Hz correlation scale, the residual covariance between their edges is diluted to a negligible level, allowing each rebinned CSD measurement to be regarded as independent.

Accordingly, every $N = 653$ contiguous frequency bins are averaged to produce a rebinned cross-interferometer CSD of 250 kHz resolution,

$$\widehat{CSD}_{TL}\left(\overline{f}\right) = \frac{1}{N} \sum_{m \in 250\text{ kHz}} \widehat{CSD}_{TL}(f_m) \,, \tag{5.23}$$

whose individual frequency measurements are independent. The variance of each rebinned measurement is estimated by

$$\begin{aligned} \mathrm{Var}\left[\widehat{CSD}_{TL}\left(\overline{f}\right)\right] &= \frac{1}{N^2} \sum_{m=0}^{N-1} \sum_{n=0}^{N-1} \mathrm{Cov}\left[\widehat{CSD}_{TL}(f_m),\, \widehat{CSD}_{TL}(f_n)\right] \\ &= \frac{1}{N^2} \sum_{m=0}^{N-1} \sum_{n=0}^{N-1} \sqrt{\mathrm{Var}\left[\widehat{CSD}_{TL}(f_m)\right]} \sqrt{\mathrm{Var}\left[\widehat{CSD}_{TL}(f_n)\right]}\, r_{m-n} \,, \end{aligned} \tag{5.24}$$

where

$$r_{\Delta n} = \frac{\left| \sum_{m=0}^{N-1} w_n^2\, e^{i2\pi m \Delta n / N_{\mathrm{DFT}}} \right|^2}{\left| \sum_{m=0}^{N-1} w_n^2 \right|^2} \tag{5.25}$$

is the frequency-bin correlation function of the window (see [7]). Numerically evaluating eqs. 5.23 and 5.24 for the window function given by eq. 5.22 yields the rebinned CSD mea-

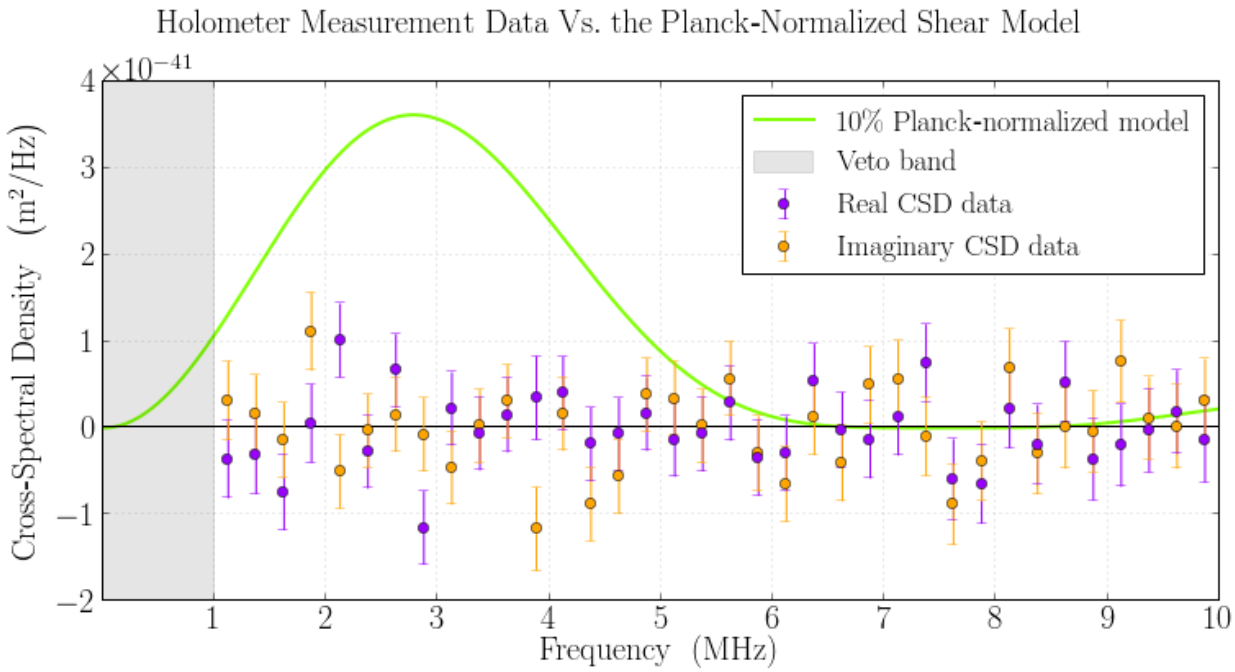

Figure 5.5: Cross-interferometer CSD measurements rebinned to 250 kHz resolution. Each rebinned frequency measurement is statistically independent. The CSD is decomposed into its real (blue data points) and imaginary (orange data points) components. Any exotic spatial correlation would appear purely in the real component, as the predicted spatial fluctuations are detected perfectly in-phase. The green curve shows the exotic shear noise model with a normalization of $l_P = l_p/10$.

surement shown in Fig. 5.5, whose 250-kHz wide bins are statistically independent. The rebinned CSD data are modeled in the following section.

5.5 Constraints on Exotic Spatial Shear

The shear-field model of exotic transverse position uncertainty has one free parameter, the commutator normalization, l_P. This scale is interpreted to set the time between successive stochastic fluctuations, establishing the characteristic "step size" of the random walk. A Bayesian analysis of the cross-interferometer CSD measurement will now place new constraints on the normalization of spatially-coherent stochastic shear. As was shown in §2.4.3, the Phase I measurement has no sensitivity to rotational-field effects.

In Bayesian model comparison, the inferential force of new evidence, in this case, the CSD measurement, is quantified by the Bayes factor. It represents the probability ratio of a model of normalization l_P to a model of null normalization, conditioned on the data. Under the null-normalization model, no deviations from classicality occur. Given a set of N frequency-space measurements, represented as $\widehat{CSD}$, the Bayes factor is

$$B\left(l_P\right) = \frac{P\left(\widehat{CSD} \,\middle|\, l_P\right)}{P\left(\widehat{CSD} \,\middle|\, 0\right)} . \tag{5.26}$$

In general, the above distributions are obtained by marginalizing over the remaining free parameters of each model, which need not be the same. However, since, in this case, there are no additional free parameters, the Bayes factor reduces to the classical likelihood ratio. The Neyman-Pearson lemma demonstrates that the likelihood ratio is the most powerful of all hypothesis test statistics.

For independent measurements, the conditioned probabilities of the measurement set decouple into simple products of the conditioned probabilities of the individual measurements as

$$B\left(l_P\right) = \frac{\prod_{n=1}^{N} P\left(\widehat{CSD}(f_n) \,\middle|\, l_P\right)}{\prod_{n=1}^{N} P\left(\widehat{CSD}(f_n) \,\middle|\, 0\right)} . \tag{5.27}$$

Each measurement probability, given a "true" normalization, is the probability that the deviation of the measured value from the expected value is purely due to the uncertainty of the measurement. For measurements with normally-distributed uncertainty, as is the case for the CSD, each measurement probability is given by a normal distribution centered on

the model prediction (eq. 2.42),

$$
\begin{aligned}
P\left(\widehat{CSD}(f)\,\middle|\,l_P\right) &= \mathcal{N}\left(\widetilde{C_{SS'}}\left(f\,|\,l_P\right),\ \mathrm{var}\left(\widehat{CSD}(f)\right)\right) \\
&= \frac{1}{\sqrt{2\pi\ \mathrm{var}\left(\widehat{CSD}(f)\right)}}\exp\left(-\frac{\left(\widehat{CSD}(f)-\widetilde{C_{SS'}}\left(f\,|\,l_P\right)\right)^2}{2\ \mathrm{var}\left(\widehat{CSD}(f)\right)}\right),
\end{aligned}
\tag{5.28}
$$

with measurement variance.

Substituting eq. 5.28 into eq. 5.27, the Bayes factor is then

$$
\begin{aligned}
B\left(l_P\right) &= \frac{\exp\left(-\frac{1}{2}\sum_{n=1}^{N}\frac{\left(\widehat{CSD}(f_n)-\widetilde{C_{SS'}}\left(f_n\,|\,l_P\right)\right)^2}{\mathrm{var}\left(\widehat{CSD}(f_n)\right)}\right)}{\exp\left(-\frac{1}{2}\sum_{n=1}^{N}\frac{\left(\widehat{CSD}(f_n)-\widetilde{C_{SS'}}\left(f_n\,|\,0\right)\right)^2}{\mathrm{var}\left(\widehat{CSD}(f_n)\right)}\right)} \\
&= \exp\left(-\frac{1}{2}\left(\chi^2(l_P)-\chi^2(0)\right)\right),
\end{aligned}
\tag{5.29}
$$

where

$$
\chi^2(l_P) \equiv \sum_{n=1}^{N}\frac{\left(\widehat{CSD}(f_n)-\widetilde{C_{SS'}}\left(f_n\,|\,l_P\right)\right)^2}{\mathrm{var}\left(\widehat{CSD}(f_n)\right)}
\tag{5.30}
$$

is the χ^2 statistic representing the goodness of fit of the model to the data. Fig. 5.6 shows the Bayes factor, or likelihood ratio, of commutator normalizations inferred from the 250-kHz resolution cross-interferometer CSD measurements (see Fig. 5.5). Because the model prediction, $\widetilde{C_{SS'}}\left(f\,|\,l_P\right)$, is purely real, only the real component of the CSD data is included in the modeling. Discarding the imaginary component, which contains purely incoherent statistical noise, reduces the measurement variance by an additional factor of two.

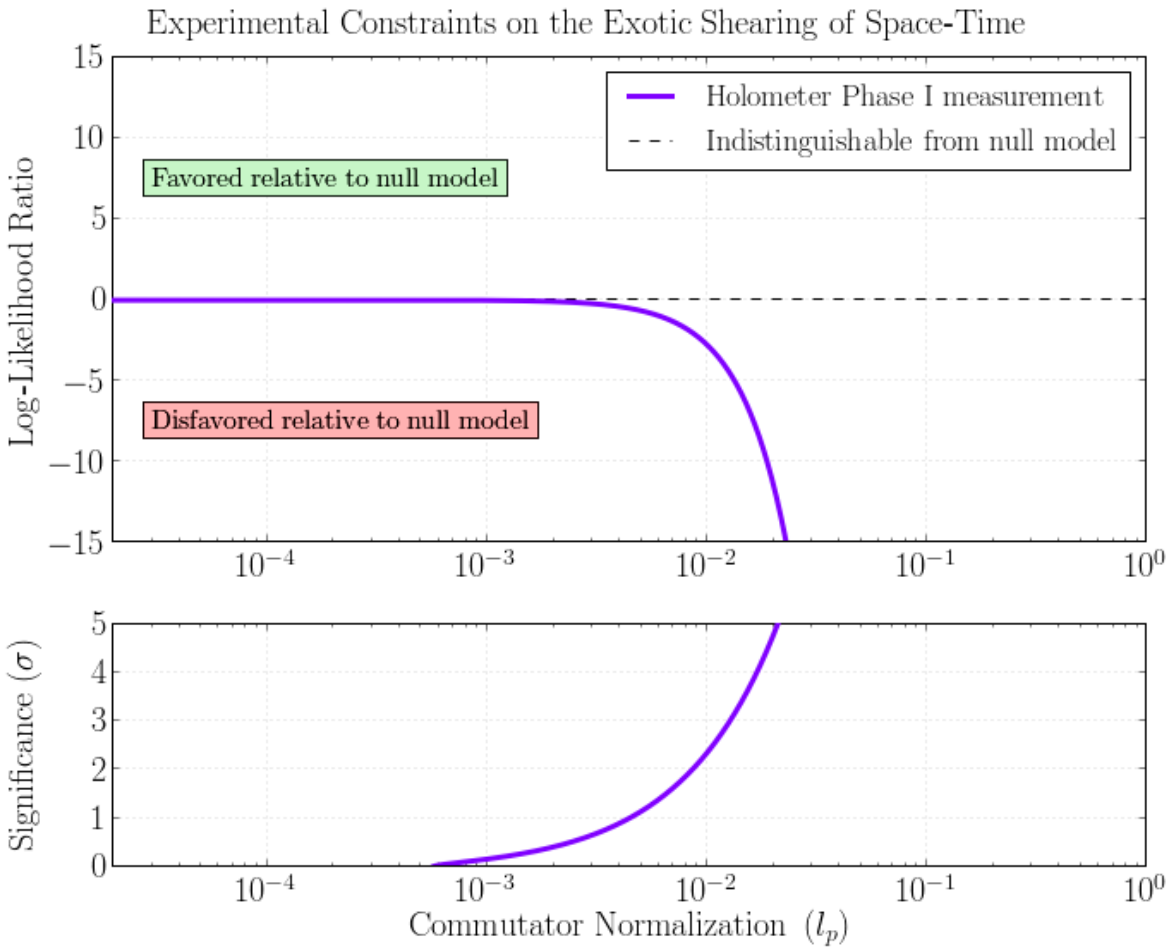

Figure 5.6: Top: Log-likelihood ratio of transverse position commutator normalizations relative to the zero-normalization model. Positive values are indicative of a detection, while negative values are indicative of an exclusion. Values equal to zero imply that the CSD data cannot statistically distinguish the normalization value from a truly zero value (i.e., the models are equally probable under the limited precision of the data). The downward turning of the log-likelihood ratio places an upper limit of $l_P < l_p/100$ on possible shear normalizations. Bottom: Statistical significance of the CSD data as a function of commutator normalization. This shows the Holometer measurement to have $> 2\sigma$ sensitivity to correlated spatial shearing at a scale 1% of a Planck length.

The statistical power reflected by a particular value of the likelihood ratio can be understood in terms of the probability of a type I statistical error, a "false positive," having occurred. The type I error probability, also known as the significance level, is the probability that the null hypothesis is excluded by random chance. In order to calculate this probability, the distribution of the underlying test statistic must be known. Although the distribution of the likelihood ratio (eq. 5.29) is difficult to analytically determine, Wilks' theorem establishes that the test statistic

$$\begin{aligned} D &\equiv -2\ln\left(B\left(l_P\right)\right) \\ &= \chi^2(l_P) - \chi^2(0) \end{aligned} \tag{5.31}$$

is asymptotically χ^2-distributed with one degree of freedom. For a given normalization l_P, the significance level of the measurement is thus

$$\begin{aligned} P(\text{type I}) &= 1 - \text{CDF}\left(D \,\middle|\, \chi_1^2\right) \\ &= 1 - \text{CDF}\left(\chi^2(l_P) - \chi^2(0) \,\middle|\, \chi_1^2\right) \ . \end{aligned} \tag{5.32}$$

Recognizing that D and Z^2 are identically distributed as χ_1^2, where $Z \sim \mathcal{N}(0,\ 1)$ is the z-score of a Gaussian-distributed variable, the above significance can be equivalently expressed in terms of the z-score $Z = \sqrt{D}$, as shown in the bottom panel of Fig. 5.6.

5.6 Physical Implications

The constraints shown in Fig. 5.6 exhaustively exclude spatially-coherent shear fluctuations of space-time. The Holometer measurement disfavors commutator normalizations to a full two orders of magnitude below the Planck length at $\gtrsim 3\sigma$ significance. Otherwise stated, the Holometer has measured an exact symmetry of the emergent Planckian space-time to $\sim 1\%$

precision. Consequently, this measurement excludes exotic spatial shearing as a possible source of the unexplained strain noise reported by GEO600 (see §2.5).

Although necessary, exclusion of a shear-field effect is not sufficient to exclude exotic transverse position uncertainty entirely. As discussed in §2.3.2, the Hogan model does not fully constrain the form which the stochastic vector field can assume. It was demonstrated that both a shear and rotational field satisfy the general requirements of the model. A complete experimental program requires testing all possible forms of the Planckian spatial fluctuations.

Phase II of the Holometer experiment, scheduled to begin construction in summer 2016, is a second interferometric search redesigned for sensitivity to exotic *rotational* fluctuations of space-time. As demonstrated in §2.4.3, the Phase I Holometer optical design has no sensitivity to a rotational effect. Now that a shear-field effect has been experimentally excluded, the second measurement will be known *a priori* to provide a pure measure of rotational spatial noise. Appendix A provides an overview of the Phase II design and a calculation of its predicted spectral response.

APPENDIX A

HOLOMETER PHASE II DESIGN

Phase II of the Holometer experiment is a planned search for a rotational-field effect. The following sections describe the new geometry of the interferometer light paths and calculate the statistical response of the instruments under the Phase II design.

A.1 Optical Geometry

Under the Holometer Phase II reconfiguration, the east arm of each interferometer will be bent to the north at its midpoint, creating an arm with 20 m of eastward propagation followed by 20 m of northward propagation. Fig. A.1 shows a rendering of this modification. Again adopting the beamsplitter-centered coordinate system introduced in §2.3.2, the classical round-trip light path through the east arm becomes

$$\mathbf{x}_1(\tau) = \begin{cases} c\tau\hat{\mathbf{i}}\,, & \tau < l/c \\ l\hat{\mathbf{i}} + (c\tau - l)\hat{\mathbf{j}}\,, & l/c \leq \tau < 2l/c \\ l\hat{\mathbf{i}} + (3l - c\tau)\hat{\mathbf{j}}\,, & 2l/c \leq \tau < 3l/c \\ (4l - c\tau)\hat{\mathbf{i}}\,, & \tau \geq 3l/c \end{cases}, \tag{A.1}$$

where the arm-segment length $l = 20$ m.

The time derivative of eq. A.1 implies the unit tangent vector,

$$\hat{\mathbf{T}}_1(\tau) = \begin{cases} \hat{\mathbf{i}}\,, & \tau < l/c \\ \hat{\mathbf{j}}\,, & l/c \leq \tau < 2l/c \\ -\hat{\mathbf{j}}\,, & 2l/c \leq \tau < 3l/c \\ -\hat{\mathbf{i}}\,, & \tau \geq 3l/c \end{cases}. \tag{A.2}$$

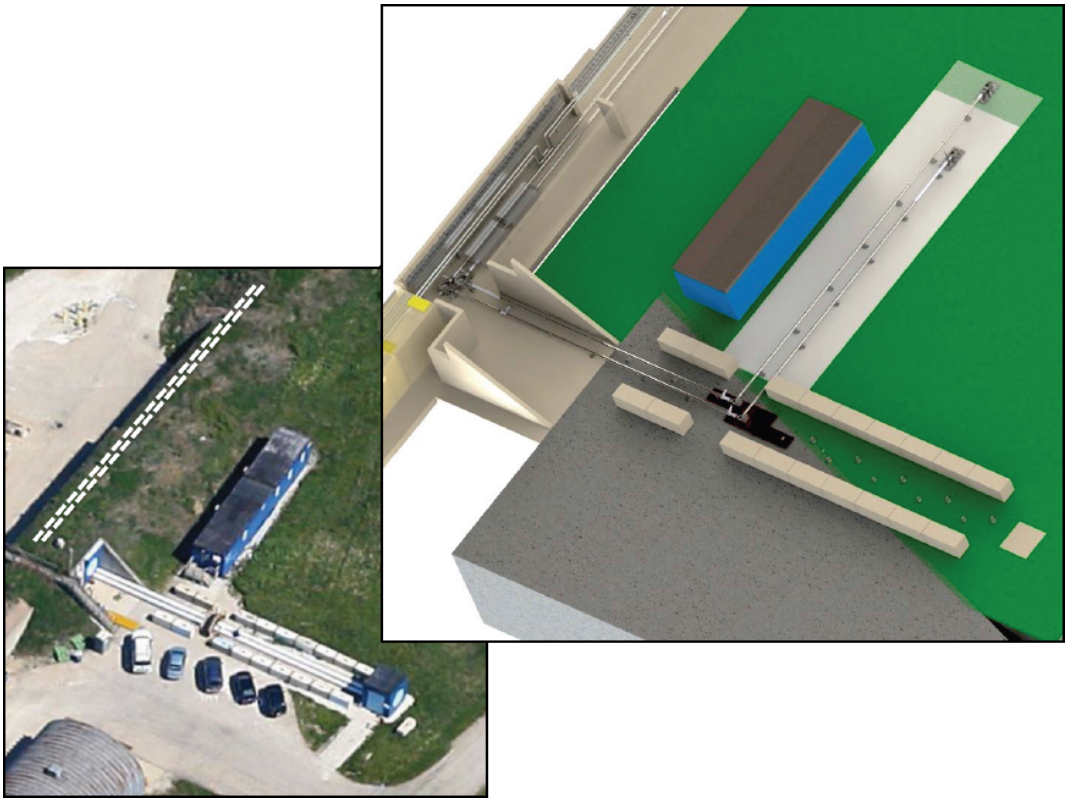

Figure A.1: Rendering of the Phase II reconfiguration of the Holometer interferometers (satellite image via Google Maps). The east arm of each interferometer will be bent to the north at its midpoint, creating an arm with 20 m of eastward propagation followed by 20 m of northward propagation. As the bent arms have some non-zero angular propagation relative to the beamsplitters, this configuration has sensitivity to exotic rotational spatial fluctuations.

Equivalently, in polar coordinates, this vector is

$$\hat{\mathbf{T}}_1(\tau) \quad = \quad \begin{cases} \hat{\mathbf{r}}\ , & \tau < l/c \\ \dfrac{(c\tau - l)\hat{\mathbf{r}} + l\hat{\boldsymbol{\theta}}}{\sqrt{l^2 + (c\tau - l)^2}}\ , & l/c \leq \tau < 2l/c \\ -\dfrac{(3l - c\tau)\hat{\mathbf{r}} + l\hat{\boldsymbol{\theta}}}{\sqrt{l^2 + (3l - c\tau)^2}}\ , & 2l/c \leq \tau < 3l/c \\ -\hat{\mathbf{r}}\ , & \tau \geq 3l/c \end{cases} , \tag{A.3}$$

which now has some non-zero angular component along the "bent" portion of the path. As the north arm is unchanged from Phase I of the experiment, the unit tangent vector of its round-trip light path,

$$\hat{\mathbf{T}}_2(\tau) \quad = \quad \begin{cases} \hat{\mathbf{r}}\ , & \tau < 2l/c \\ -\hat{\mathbf{r}}\ , & \tau \geq 2l/c \end{cases} , \tag{A.4}$$

remains purely in the radial direction.

A.2 Response to Rotational Effect

Adopting the stochastic rotational field of eq. 2.28, substituting eqs. A.3 and A.4 into eq. 2.26 yields a geometrical coupling of

$$\begin{aligned}
\Theta(\tau) &= \hat{\boldsymbol{\theta}} \cdot \begin{cases} \hat{\mathbf{r}} - \hat{\mathbf{r}} \,, & \tau < l/c \\ \hat{\mathbf{r}} - \dfrac{(c\tau - l)\hat{\mathbf{r}} + l\hat{\boldsymbol{\theta}}}{\sqrt{l^2 + (c\tau - l)^2}} \,, & l/c \leq \tau < 2l/c \\ -\hat{\mathbf{r}} + \dfrac{(3l - c\tau)\hat{\mathbf{r}} + l\hat{\boldsymbol{\theta}}}{\sqrt{l^2 + (3l - c\tau)^2}} \,, & 2l/c \leq \tau < 3l/c \\ -\hat{\mathbf{r}} + \hat{\mathbf{r}} \,, & \tau \geq 3l/c \end{cases} \\
&= \begin{cases} 0 \,, & \tau < l/c \\ -\dfrac{l}{\sqrt{l^2 + (c\tau - l)^2}} & l/c \leq \tau < 2l/c \\ \dfrac{l}{\sqrt{l^2 + (3l - c\tau)^2}} & 2l/c \leq \tau < 3l/c \\ 0 \,, & \tau \geq 3l/c \end{cases}
\end{aligned} \tag{A.5}$$

to exotic rotational noise. Substituting eq. A.5 into eq. 2.31 then yields the statistical response of each Holometer interferometer under the Phase II reconfiguration. Although this integral does not have an analytic solution, its numerical solution is shown in the top panel of Fig. A.2 for $l_P = l_p$. Numerically evaluating the Fourier transform of this autocovariance then yields its PSD, as is shown in the bottom panel.

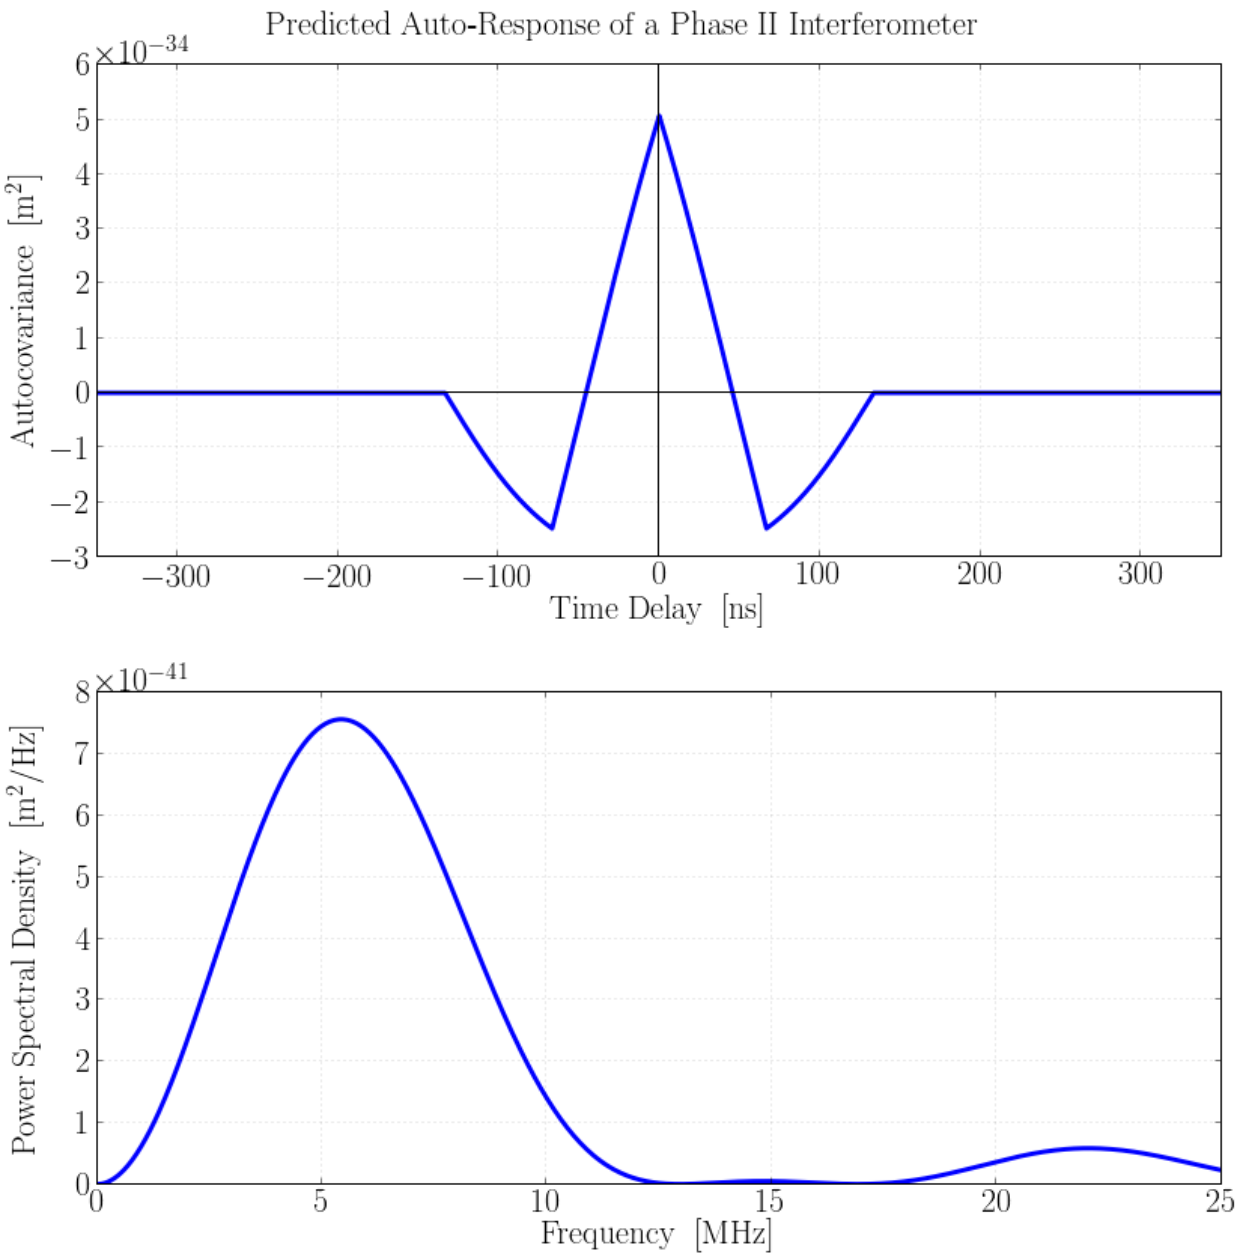

Figure A.2: The time- and frequency-domain response functions of each Holometer interferometer to exotic spatial rotational noise under the Phase II reconfiguration. The integration time to achieve unity SNR is estimated to be only 15 minutes.

References

[1] Abhay Ashtekar. Introduction to Loop Quantum Gravity. *PoS*, QGQGS2011:001, 2011.

[2] E. D. Black. An introduction to pound-drever-hall laser frequency stabilization. *Am. J. Phys.*, 69:79, 2001.

[3] W. A. Christiansen, Y. J. Ng, D. J. E. Floyd, and E. S. Perlman. Limits on spacetime foam. *Phys. Rev. D.*, 83(8):084003, April 2011.

[4] W. A. Christiansen, Y. J. Ng, and H. van Dam. Probing Spacetime Foam with Extragalactic Sources. *Physical Review Letters*, 96(5):051301, February 2006.

[5] M. R. Douglas and N. A. Nekrasov. Noncommutative field theory. *Rev. Mod. Phys.*, 73:977, 2001.

[6] E. E. Flanagan and K. S. Thorne. Noise Due to Backscatter Off Baffles, the Nearby Wall, and Objects at the Far End of the Beam Tube; and Recommended Actions. Technical Report LIGO-T940063-00-R, LIGO, 1994.

[7] F. J. Harris. On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform. *IEEE Proceedings*, 66:51–83, January 1978.

[8] C. J. Hogan. Indeterminacy of holographic quantum geometry. *Phys. Rev.*, D78:087501, 2008.

[9] C. J. Hogan. Measurement of quantum fluctuations in geometry. *Phys. Rev.*, D77:104031, 2008.

[10] C. J. Hogan. Interferometers as probes of planckian quantum geometry. *Phys. Rev.*, D85:064007, 2012.

[11] S. Hossenfelder. Minimal length scale scenarios for quantum gravity. *Living Rev. Rel.*, 16:2, 2013.

[12] Ohkyung Kwon and Craig J. Hogan. Interferometric Tests of Planckian Quantum Geometry Models. *Class. Quant. Grav.*, 33(10):105004, 2016.

[13] R. Lieu and L. W. Hillman. The Phase Coherence of Light from Extragalactic Sources: Direct Evidence against First-Order Planck-Scale Fluctuations in Time and Space. *ApJ*, 585:L77–L80, March 2003.

[14] T. Mohaupt. Introduction to string theory. In D. Giulini, C. Kiefer, and C. Lämmerzahl, editors, *Quantum Gravity: From Theory to Experimental Search*, volume 631 of *Lecture Notes in Physics*, pages 173–251, Berlin; New York, 2003. Springer.

[15] Y. J. Ng, W. A. Christiansen, and H. van Dam. Probing Planck-Scale Physics with Extragalactic Sources? *ApJ*, 591:L87–L89, July 2003.

[16] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

[17] E. S. Perlman, Y. J. Ng, D. J. E. Floyd, and W. A. Christiansen. Using observations of distant quasars to constrain quantum gravity. *A&A*, 535:L9, November 2011.

[18] E. S. Perlman, S. A. Rappaport, W. A. Christiansen, Y. J. Ng, J. DeVore, and D. Pooley. New Constraints on Quantum Gravity from X-Ray and Gamma-Ray Observations. *ApJ*, 805:10, May 2015.

[19] J. Polchinski. *String Theory*. Cambridge University Press, 1998.

[20] R. Ragazzoni, M. Turatto, and W. Gaessler. The Lack of Observational Evidence for the Quantum Structure of Spacetime at Planck Scales. *ApJ*, 587:L1–L4, April 2003.

[21] C. Rovelli. *Quantum Gravity*. Cambridge University Press, 2004.

[22] J. H. Schwarz. Introduction to superstring theory, 2000.

[23] Thomas Thiemann. *Modern Canonical Quantum General Relativity*. Cambridge University Press, 2008.

[24] S. Weinberg. *The Quantum Theory of Fields*. Cambridge University Press, 1996.

[25] F. Wilczek. Quantum field theory. *Rev. Mod. Phys.*, 71:S85, 1999.