Controlling and Calibrating Interferometric Gravitational Wave Detectors

Thesis by
Craig Cahillane

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2021
Defended January 26, 2021

Craig Cahillane
ORCID: 0000-0002-3888-314X

ABSTRACT

In September 2015, the Advanced LIGO detectors made the first direct detection of gravitational waves from a binary black hole merger [1]. Since then, around fifty total gravitational wave detections have been reported by Advanced LIGO and Advanced Virgo over three dedicated gravitational wave observation times, known as observing runs.

Observing run three (O3) ran from April 2019 to March 2020, with higher sensitivity and more stable operation of the Advanced LIGO detectors [2]. In the first half of O3, thirty-nine gravitational wave events were detected [3], as opposed to eleven in all of observing runs one (O1) and two (O2) [4]. The higher rate of detections is due primarily to the increased detector sensitivity to gravitational waves.

Although the Advanced LIGO detectors are more sensitive to gravitational waves than any detector in history, they have not yet achieved design sensitivity. Work continues to push the detectors to their fundamental limit of sensitivity. The work in this thesis partially covers the effort to improve the sensitivity of the LIGO Hanford detector prior to O3.

Calibration of the Advanced LIGO interferometer is the conversion of raw detector data into gravitational wave strain data. This process is crucial to an accurate and precise understanding of astrophysical sources of gravitational waves. The calibration uncertainty pipeline for characterizing the strain uncertainty during O1 and O2 is discussed in detail [5].

This thesis covers topics in long-baseline interferometric gravitational wave detector technology, including an overview of the performance of the detector in O3, commissioning tasks done to increase the sensitivity of the detector for O3, overall calibration uncertainty in the gravitational wave data, and methods for robust estimation of spectral quantities from LIGO data.

ACKNOWLEDGEMENTS

My path through grad school was anything but straightforward. I thank everyone who helped me along the way.

Thank you Alan for giving me the opportunity to work in LIGO from O1 to O3, for your hands-on mentorship in the the early years, and your patience and counsel in the later years.

Thank you Rana for the opportunity to join your group halfway through grad school, for showing us how to control ASC, and for the stories about the old days.

Thank you Max, Tom, and Surabhi for your companionship throughout grad school and for making the third floor of West Bridge fun. Thank you Darkhan and Sudarshan for inviting me into the calibration group at Hanford my first time there. Thank you Evan and Kiwamu for mentoring a young grad student and showing me the path I eventually took for myself.

Thank you Andrew for making the coffee lobster, buying all those boxes, and for teaching me how to take a transfer function. Thank you Johannes, Aaron, and Brittany for bringing light to the basement. Thank you Kevin for holding me to a higher academic standard.

Thank you Daniel for reading my baffling alogs, creating a positive commissioning environment, and steering me on the right course. Thank you Keita for letting me fix frequency noise during run time. Thank you Jeff for talking about calibration with me. Thank you Sheila and Jenne for your incredible mentorship, patience, and friendship. Thank you to the visiting scientists who took interest in my work and invited me into theirs, including Peter, Stefan, Robert, Rich, and Calum.

Thank you Aidan, Anchal, and Jaime for getting coffee in the park. Thank you Gabriele for sharing an office, your data, and your knowledge with me. Thank you Koji for for all the black glass beam dumps, late night OMC measurements, and for teaching me about intensity noise and HOMs.

Thank you Dan for the late night commissioning times and inviting me into the Finesse group. Thank you Anna, Paul, and Andreas for providing valuable feedback on the interferometer simulation inner workings.

Thank you TJ and Mary for letting me sleep in your house for a month. Thank you Gautam for being my go-to friend for all interferometer ideas, co-captaining the Strong Fielders softball team, and for living with me through a pandemic.

Thank you to my mom and dad, Susan and Mark, for the many packages of socks and years of support. Thank you to my brother Dennis who has mentored me and always been my friend. Thank you to Georgia who I met along the way. I love you.

PUBLISHED CONTENT AND CONTRIBUTIONS

[1] **Craig Cahillane**, G. Mansell, et al. “Laser frequency noise requirements for next generation gravitational-wave detectors”. *Optics Express* (in prep). C.C. produced the frequency noise budget and coupling function. Some of this work is reproduced in Chapter III of this thesis.

[2] A. Buikema, **Cahillane, C.**, et al. “Sensitivity and performance of the Advanced LIGO detectors in the third observing run”. *Physical Review D* (2020). doi: 10.1103/PhysRevD.102.062003
C.C. was a member of the four-person paper writing team for this Advanced LIGO colloboration-reviewed paper. C.C. produced several of the results presented, including the intensity, frequency, and H1 correlated noise budgets, as well as the arm power measurement. This work is reproduced in Chapter III of this thesis. ISSN: 24700029.

[3] **Craig Cahillane**. “Median averaging for cross spectral densities”. *(in prep)* (2020).
C.C. derived analytic results for median-averaged CSDs for use with glitchy long-term LIGO data. This work facilitated Chapter V of this thesis, and is reproduced in Chapter VI of this thesis.

[4] **Craig Cahillane**. “beamtrace: python3 Gaussian laser beam ABCD matrix propagator” (2019).
PyPi library: https://pypi.org/project/beamtrace/
C.C. wrote this python library for Gaussian beam propogation for research at the LIGO Hanford detector. This library is used for Appendix B.4.2 and Appendix C.2 of this thesis.

[5] **Craig Cahillane**. “nds2utils: convenient user interface for the python nds2 LIGO data-acquisition client” (2019).
PyPi library: https://pypi.org/project/nds2utils/
C.C. wrote this python library to facilitate data collection for research at the LIGO Hanford detector.

[6] E. D. Hall, **Craig Cahillane**, et al. “Systematic calibration error requirements for gravitational-wave detectors via the Cramér–Rao bound”. *Classical and Quantum Gravity* 36.20 (2019).
doi: 10.1088/1361-6382/ab368c
C.C. verified the Jacobian numerical results of this paper analytically.

[7] B. P. Abbott et. al. (LIGO Scientific Collaboration). “Calibration of the Advanced LIGO detectors for the discovery of the binary black-hole merger GW150914”. *Phys. Rev. D* 95 (6 2017).
doi: 10.1103/PhysRevD.95.062003
C.C. produced the calibration uncertainty budgets for Advanced LIGO’s first detection of gravitational waves, GW150914.

[8] **Cahillane, Craig**, J. Betzwieser, et al. “Calibration uncertainty for Advanced LIGO’s first and second observing runs”. *Phys. Rev. D* 96 (10 2017). doi: 10.1103/PhysRevD.96.102001

C.C. conceived of the calibration uncertainty pipeline, produced calibration uncertainty budgets for Advanced LIGO’s first and second observing runs, and wrote the manuscript. This work is reproduced in Chapter IV of this thesis.

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

Number *Page*

LIST OF TABLES

Chapter 1

INTRODUCTION

In the early 20$^{\text{th}}$ century, Albert Einstein revolutionized physics with the theories of special and general relativity. Einstein cut ideas like absolute universal length, absolute universal time, and simultaneity in favor of a universal speed of light. Relativity unified space and time into a singular "spacetime." Travel through space and time is was no longer independent, and observers traveling through space relative to one another must experience different passage of time.

General relativity redefined gravity as *curvature in spacetime* arising from the presence of mass and energy within that spacetime. The "force" of gravity between two bodies was not a force at all, but two objects following a "straight line", or *geodesic*, through a curved spacetime. General relativity resolved issues with Newtonian gravity, including correctly predicting the precession of Mercury's orbit and gravitational lensing by the sun. The development of general relativity revolutionized physics and astrophysics, provided a new framework for understanding the universe on a large scale, and kicked off the field of cosmology.

One key prediction of general relativity was the existence of waves in spacetime, known as gravitational waves. Gravitational waves were a natural consequence of Einstein's equations describing spacetime curvature, but Einstein predicted these waves were far too weak to ever be detected by humanity.

In 2015, Advanced LIGO made the first-ever detection of gravitational waves from a binary black hole merger [1]. Since then, Advanced LIGO's sensitivity to gravitational waves has increased even further, resulting in 39 detections in the most recent observing run [3].

This thesis will focus on the efforts to characterize and improve the sensitivity of the LIGO Hanford Observatory leading up to its third observing run (O3), with topics in precision detector calibration, noise mitigation, and novel detector measurement techniques.

1.1 What is a gravitational wave?

Gravitational waves are the propagating wave manifestation of a fundamental force of nature. The "electric charge" of gravity is mass, which can only be positive, not both positive and negative. However, gravity is a much weaker force than electromagnetism, and can only emit quadrupole radiation, as opposed to electromagnetism's dipole radiation.

A gravitational wave is often described as a "ripple" in spacetime. As heavy objects move through the universe, they interact with spacetime, creating curvature such that the motions of nearby objects through spacetime appear distorted.

If two extremely heavy objects begin orbiting one another very quickly, spacetime curvature near this orbit becomes extremely strong and changes rapidly. A significant amount of energy in the fluctuating spacetime propagates away to infinity in the form of oscillating spacetime.

In general relativity, Einstein's field equations relate spacetime curvature to the energy and matter residing within that spacetime:

$$G_{\mu\nu} - \Lambda g_{\mu\nu} = \kappa T_{\mu\nu} \tag{1.1}$$

where $G_{\mu\nu}$ is the Einstein tensor describing spacetime curvature, Λ is the cosmological constant, $g_{\mu\nu}$ is the local spacetime metric, $\kappa = 8\pi G/c^4$ is the Einstein gravitational constant governing energy coupling to spacetime curvature, and $T_{\mu\nu}$ is the stress-energy tensor describing the matter and energy within a spacetime.

In the weak-field limit, where there is no matter or energy, the stress-energy tensor $T_{\mu\nu} = 0$. Then the Einstein equations can be reduced to a wave equation and solved for small perturbations in spacetime. These solutions to Einstein's equations are known as gravitational waves.

For a wave traveling transverse to the z direction, the gravitational wave tensor $h_{\mu\nu}$ is

$$h_{\mu\nu} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \tag{1.2}$$

The coefficients h_+ and $h_\times$ correspond to the two polarizations of gravitational waves, and refer to the way they affect spacetime.

The effect of a gravitational wave can be seen in the way it affects the distances of two objects resting in spacetime. In the lab frame, a gravitational wave can be said to create a length change between any two points in space, with the sign of the

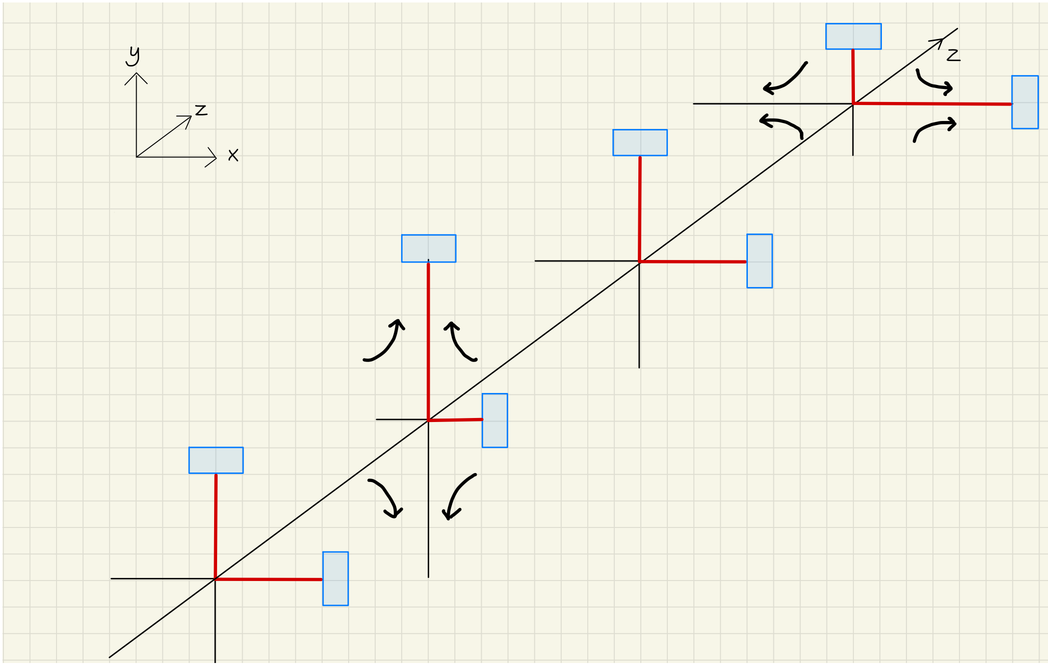

Figure 1.1: Effect of a passing linearly polarized h_+ gravitational wave traveling in the z direction on two test masses on the x- and y-axis. In the lab frame, one axis is "stretched", the other is "squeezed", producing an effective differential length change.

change depending on the polarization and orientation relative to the wave. For two points a distance L apart, the length change ΔL is

$$\Delta L = hL \tag{1.3}$$

where h is the gravitational wave *strain*. Figure 1.1 shows the differential length change effect of a gravitational wave with strain $h \sim 0.5$ on two tests masses. This is the principle upon which gravitational wave detectors are based.

Gravitational waves are produced when any masses accelerate through spacetime, like in the orbit between two objects. In reality, spacetime is a "stiff" medium, or gravity is a "weak" force: only the most massive objects in the universe can make an appreciable dent in spacetime, and only the most massive, most energetic orbits can create significant gravitational waves.

Gravitational waves spread from their source over all space, losing amplitude inversely proportional to their distance from their source. The strongest gravitational waves that reach Earth are all from extremely distant, rare, ultra-powerful astrophysical collisions.

For these reasons, the strongest gravitational waves reaching Earth have a strain on the order of $h \sim 10^{-21}$. For reference, a human cell is 10^{-4} m, the size of an atom 10^{-10} m, the size of the nucleus of that atom is 10^{-15} m. Gravitational waves incident on our $L = 4$ km long detectors will be $hL \sim 10^{-18}$. This is why the sensitivity of the Advanced LIGO detectors to gravitational waves are huge technological feat of engineering, and the gravitational wave data a valuable new font of information on the depths of the universe.

1.2 Sources of gravitational waves

Gravitational waves provide a new unique source of information about the darkest, most massive objects, and most energetic events in the universe. Events normally inaccessible through light, such as binary black hole mergers, supernovae core bounce, or the Big Bang, can be directly observed via GWs.

Mergers of black hole and neutron star binaries are some of the most powerful events in the universe, but are completely invisible to observers on Earth except through the gravitational wave signature they produce [1, 6, 7]. The detections of binary black hole mergers are the first direct observational insights into the physics of massive binary systems. The formation rates of stellar-mass black holes have been more accurately estimated than ever before, as well as the spin parameters of both the inspiraling and final black holes .

Binary neutron stars also offer insight into extreme events of spacetime, including the influence of matter on GW emission [8–10]. Tidal disruption breaks apart the neutron star pair prior to merger, causing irregularly in the inspiral and merger than can provide useful information on the type of matter that makes up a neutron star. For binary neutron stars, multimessenger astronomy has already begun with the detection of prompt electromagnetic followup to a GW merger, which proved GWs travel very near to the speed of light.

Unequal-mass binaries consisting of a neutron star and black hole are also possible, and candidates have already been detected [11]. These are especially interesting because of the orbital precession physics possible, the higher-order multipoles of the GWs detected, and the rates of formation of small black holes and large neutron stars near the so-called “mass-gap”.

Supernovae are the explosive death of stars about $10\times$ more massive than the sun, and the birth of neutron stars or black holes, but the mechanism that powers the explosion is not well-understood. Light from the supernovae comes from the exploding surface of the star, but the gravitational waves from the supernovae would come from the core, and with it valuable information about the formation of the core and the nature of the explosion.

The *stochastic gravitational wave background* is the random noise of the universe. The stochastic background is formed from the sum of all unresolved binary inspirals in the distant universe. Cosmic gravitational wave backgrounds have the potential to provide information directly from the Big Bang, shortly after the era of inflation when the universe expanded exponentially. The current limit to the direct observations from the early universe come from the cosmic microwave background, which occurred during a era called recombination when the universe cooled enough so the first atoms could form, around 370000 years after the Big Bang.

Some of these phenomena cannot be detected with light, or have questions that observation via light cannot answer. Gravitational waves offer a new way of observing the universe, of listening to the universe by measuring spacetime reverberating with the echoes of unimaginably powerful events from billions of years ago.

1.3 Detectors

The future of gravitational wave astronomy and astrophysics relies on the continued improvement of the sensitivity of gravitational wave detectors. Detector scientists, known as commissioners, at the Advanced LIGO detectors are working to achieve the maximum sensitivity possible with the current detectors. Figure 1.2 il-

Table 1.1: Signal-to-noise ratio (SNR) the first gravitational wave detection, GW150914, would have had in current and future GW detectors. Figure 1.2 shows the current and future noise curves beside the characteristic strain of GW150914. Future detector noise curves are reported from gwinc [12].

Detector	SNR	SNR ratio / Observing Run 1
Advanced LIGO Observing Run 1	23	1.0
Advanced LIGO Observing Run 3	34	1.5
Advanced LIGO design	53	2.4
Advanced LIGO A+ upgrade	101	4.5
LIGO Voyager	250	11.1
Cosmic Explorer 1	1232	54.5
Cosmic Explorer 2	2273	100.5

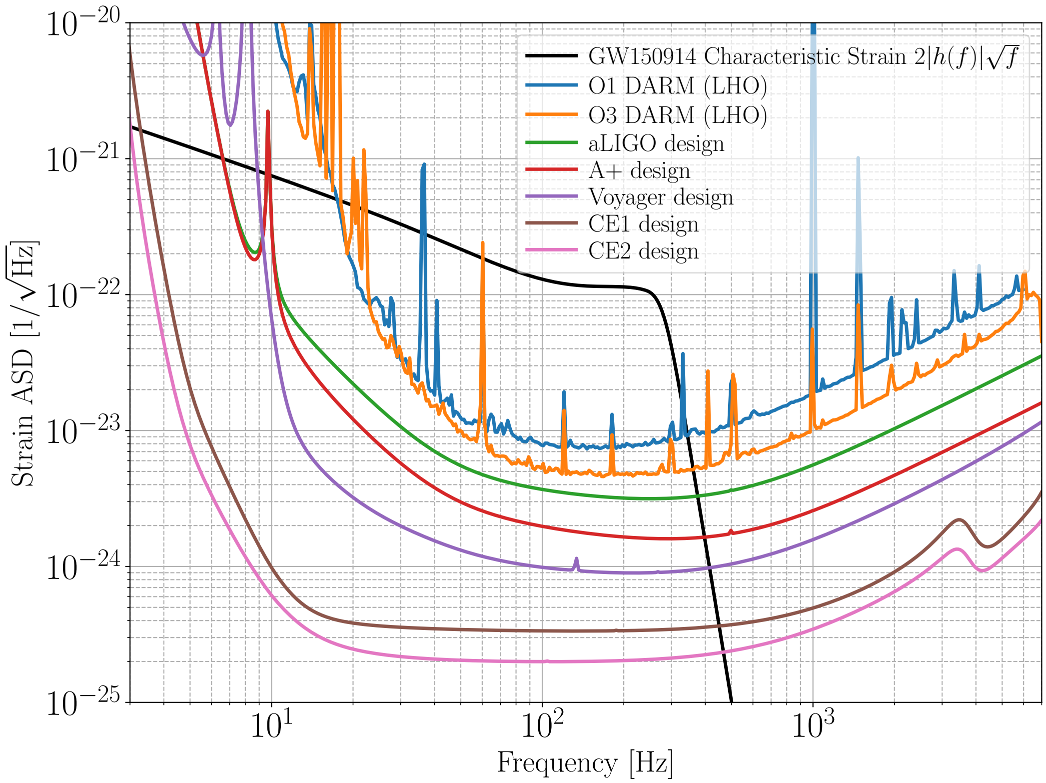

Figure 1.2: Noise curves of current and future gravitational wave detectors, compared with the characteristic strain of the first GW detection, GW150914. The first GW signal measured is shown as the black curve, while each of the colored noise curves represent the measured noise of a current detector or projected noise of a future detector. The lower the noise, the more sensitive the detector is to GWs. By lowering the noise of the detector, a loud signal like GW150914 can be better resolved, and more precise information can be learned from the signal. Table 1.1 calculates the signal-to-noise ratio (SNR) that a signal like GW150914 would have in each detector. Estimated future noise curves from this plot are produced by pygwinc as of January 2021 [12].

lustrates the projected improvement of sensitivity for ground-based long-baseline interferometers in the United States.

As sensitivity of detectors is improved, both the rate of detections will increase, and the signal from current detections will be better resolved. From the increased number of detections, we can learn about black hole and neutron star astronomy, including binary formation rates, and galaxy formation.

From the clearer signals, we can better resolve the physical parameters of the mergers like the masses, spins, distance, orbital plane inclination, and sky location. Also, we can better test general relativity in the most extreme regions of spacetime where two black holes are merging into one, where extremely powerful spacetime curvature itself causes further curvature. Precision measurements of the Hubble constant are also possible.

With a more sensitive detector, weaker sources of gravitational waves could be detected, including gravitational waves from a single, rapidly spinning neutron stars, gravitational waves from exploding supernovae, or the random background of gravitational waves from binary black holes and neutron stars too far away to resolve individually.

This thesis will focus on efforts to achieve the sensitivity acquired by Advanced LIGO for observing run three (O3), with topics in precision calibration of the Advanced LIGO detectors' response to gravitational waves, characterization of noise sources, and new techniques for measurements of fundamental quantities of the interferometer.

Chapter 2

ADVANCED LIGO DETECTOR DESIGN AND O3 UPGRADES

The design of the Advanced LIGO detectors is the culmination of decades of gravitational wave (GW) detector research [13–18]. The Advanced LIGO detector design was finalized in 2010, with several major upgrades compared to the initial LIGO design [19, 20]. The biggest upgrades from initial LIGO to advanced LIGO are the addition of the signal-recycling mirror [21, 22], the introduction of folded recycling cavities for better laser beam geometry stability [23], a novel seismic isolation and quadruple pendulum suspension for the heavier main optics [24–27], lower optical coatings thermal noise [28–32], higher input laser power [33–37], the addition of auxiliary green lasers for achieving detector operation [38–40], and the implementation of DC readout [41] and an output mode cleaner [42].

The success of these upgrades made Advanced LIGO the most sensitive gravitational wave detectors in history [43]. Since then, many upgrades have further increased the astrophysical range of the instrument [2].

In this chapter we will briefly review the basics of Advanced LIGO gravitational wave detection, including the fundamentals of GW detection, the detector topology, optic design, and seismic isolation. We also review some of the hardware upgrades implemented between prior to the start of O3 relevant to this thesis, including increased laser power, squeezed light injection, and replaced core optics.

2.1 Detector topology

The core of the Advanced LIGO detector is dual-recycled Fabry-Perot Michelson interferometer. The arms of the interferometer are formed by 4 km long Fabry-Perot cavities, featuring an input test mass (ITM) and end test mass (ETM). The power recycling cavity (PRC) is formed between the power recycling mirror (PRM) and ITMs. The signal recycling cavity (SRC) is formed between the signal recycling mirror (SRM) and the ITMs. The Michelson is the formed by the beamsplitter and the ITMs. Figure 2.1 shows the simplified layout of the major interferometer components.

A gravitational wave incident on the detector interacts primarily with the high power laser light inside the 4 km long arms. The GW induces a tiny phase shift

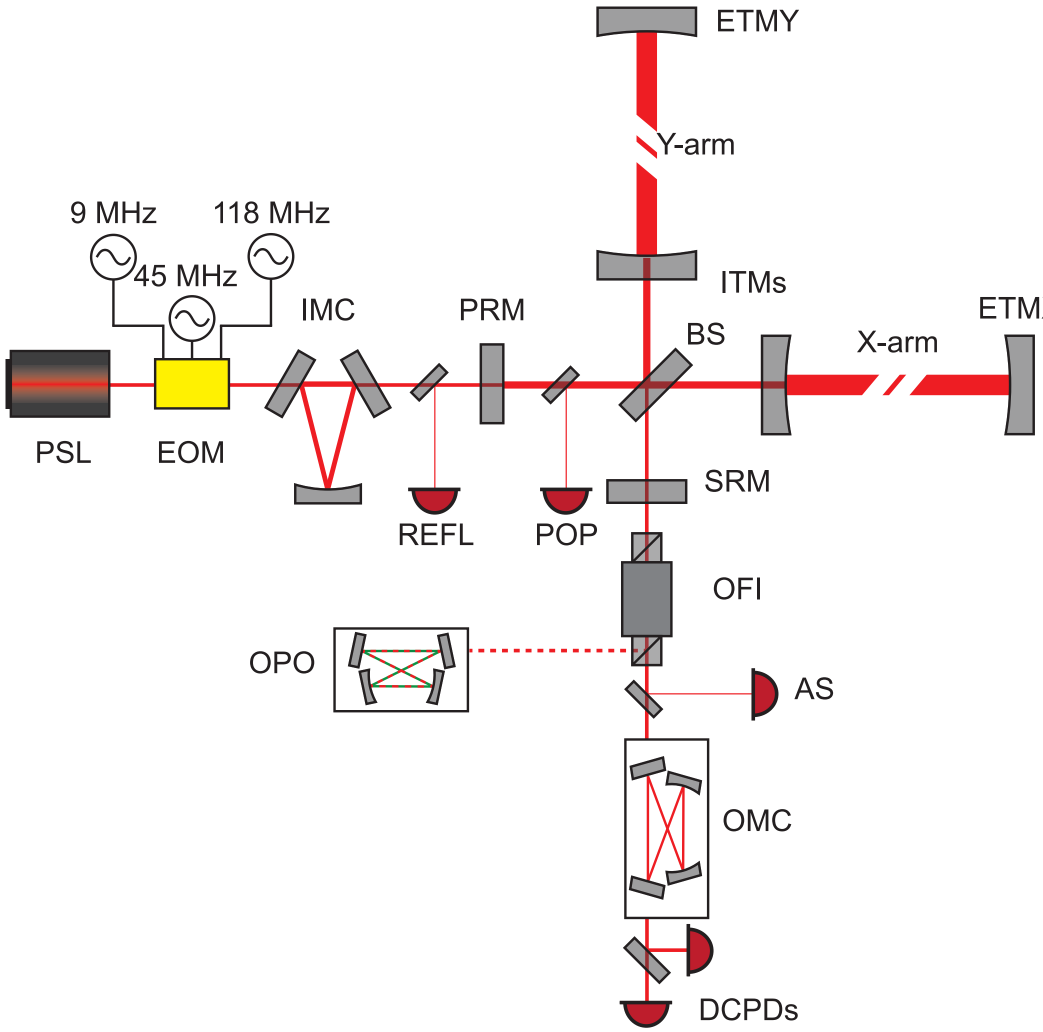

Figure 2.1: Simplified diagram of the optical layout of LIGO Hanford for O3 [2]. The core optics that form the interferometer itself are the end test masses (ETMs), the input test masses (ITMs), the beamsplitter (BS), the power-recycling mirror (PRM), and signal-recycling mirror (SRM). The pre-stabilized laser (PSL) injects laser light with radio-frequency (RF) sidebands at 9 MHz, 45 MHz, and 118 MHz to be used for sensing length changes inside the interferometer. The photodetector sensors REFL, POP, AS, and DCPDs detect power fluctuations due to interferometric length changes. The input mode cleaner (IMC) transmits only carrier laser light and stabilizes the laser frequency to a suspended cavity. The output mode cleaner (OMC) transmits only light carrying the gravitational wave signal, and reflects away all other light. The output Faraday isolator (OFI) keeps light reflected from the OMC from re-entering the interferometer. The optical parametric oscillator (OPO) injects squeezed quantum vacuum into the antisymmetric port of the interferometer.

on the laser light as it propagates down the arms. The tiny phase shift is amplified by three factors:

1. high input laser power,
2. multiple reflections inside the arms,
3. the length of the arm cavities.

The high input laser power increases the light that is available to be scattered into the GW signal. The multiple reflections inside the arm cavities increases the interaction time of the laser with the GW. The length of the arm increases the effective displacement of the arm length due to GWs (Eq. 1.3).

The X-arm and Y-arm are separate cavities, but the Michelson interferometer specializes in discriminating between common and differential arm motion (see Section B.3). Because the common arm motion (CARM) light is reflected back to the symmetric port where the laser entered the beamsplitter, and the differential arm motion (DARM) light is transmitted to the antisymmetric port through the beamsplitter, CARM and DARM are most natural length degrees of freedom eigenmodes.

The power recycling mirror increases the effective input power of the laser by constructively interfering the laser coming back from the beamsplitter with the new input light. The PRC forms a *coupled-cavity* with the arms. A coupled-cavity is two Fabry-Perot cavities in series, where the end mirror of one cavity is the input mirror for the second. The coupled-cavity formed by the PRC and arms is called the *common-arm length*, or CARM cavity. The formation of the CARM coupled-cavity enables the ultra-stable, high resonating laser power inside the Advanced LIGO detectors.

The signal recycling mirror broadens the bandwidth of the detector. The SRC also forms a coupled cavity with the arms, called the differential-arm length, or DARM. Any differential motion in the arms will scatter some light into the DARM coupled-cavity, and therefore the SRC. The SRC is tuned to be *antiresonant* for the main carrier light, so the GW audio sidebands in the SRC are preferentially transmitted through the cavity to the GW signal photodetectors. This setup is known as *resonant sideband extraction*, and is discussed in detail in Section 3.6.

Advanced LIGO is run in DC readout configuration [20, 41]. DC readout employs a small intentional differential offset in the arm lengths to provide DC power on the DCPDs to beat with the gravitational wave signal. Appendix B derives the basics of a simple Michelson operating with DC readout.

The input and output mode cleaners (IMC and OMC) are both in place to "clean" the laser light of *higher-order modes*, which are different spatial modes of the main carrier laser beam. The interferometer geometry, i.e. the radii of curvature of the mirrors and the lengths of the cavities, is designed to contain solely main carrier light. Higher-order modes are created from imperfections in the geometry, carry no signal, but contribute noise. A mode cleaner is a cavity designed to transmit the main carrier laser mode of the interferometer, and reflect all other modes away from the signal photodetectors.

Finally, the optical parameter oscillator (OPO) is a bowtie cavity designed to generated squeezed quantum vacuum for injection into the antisymmetric port of the interferometer [44]. The key component to the OPO is the nonlinear crystal inside the cavity converts the green "pump" photons into two entangled infrared photons. The entangled photons are equivalent to squeezed light. The arrival times of the entangled photons on the photodetector are correlated, lowering the variance of the Poissonian process which describes shot noise. The phase of the squeezed light is controlled such that the quantum shot noise is minimized.

2.2 Gravitational wave signal

The first consideration for detector design is the response to incident gravitational waves. A passing gravitational wave modulates the spacetime metric between any two free masses. Gravitational radiation has quadrupole polarization, and can be broken down into its plus h_+ and cross $h_\times$ orthogonal components.

Figure 1.1 shows the "L-shape" Michelson interferometer response to h_+ GW radiation. A laser is reflected off the test masses to detect their differential motion. Section B.3.2 derives a simple Michelson's response to gravitational waves.

The L-shape interferometer is sensitive only to one polarization of GWs, which we usually define as h_+. A triangle-shaped set of three interferometers is another fundamental design that is sensitive to both polarizations h_+ and $h_\times$. The triangle-shaped detector is employed in future detector designs like Einstein Telescope [45] and space-based interferometers like LISA [46].

The Advanced LIGO detector design is based on the simple Michelson [22, 47]. Both detectors preferentially transmit only differential arm motion, including the

GW signal, out of the antisymmetric port. Most laser power is reflected by the interferometer, keeping the antisymmetric port relatively clean.

The GW signal out of the Advanced LIGO interferometer is amplified due to the high laser power resonating in the Fabry-Perot arms (see Section B.4). The signal bandwidth is broadened by the placement of the signal recycling mirror. Section 2.1 will discuss Advanced LIGO's topology in more detail.

From our derivation of the simple Michelson response in Section B.3.2, Eq. B.45, the response of the interferometer to gravitational waves is proportional to

1. the input power P_{in},
2. the laser frequency ω_0,
3. the arm length L.

Increasing any, naively, will increase the fundamental sensitivity limit of the detector. However, GW signal response is not the only consideration for a sensitive detector design.

2.3 Noise

The second consideration for detector design is the quantification of noises which can mask the gravitational wave signal. *Fundamental noise sources* are those which define the limits of sensitivity to GWs, based on the detector design. *Technical noise sources* are those assumed to be negligible in the detector design, but can be difficult to mitigate in reality.

2.3.1 Quantum noise

Fluctuations of the vacuum electric field at the interferometer readout port impose a fundamental limit to the interferometer sensitivity [22, 48–50]. Quantum noise appears as *shot noise* and *quantum radiation pressure noise*.

2.3.1.1 Shot noise

Shot noise arises from Poisson fluctuations in the arrival time of photons at the interferometer output. The power detected on the photodetector is made up of a finite number of photons which arrive randomly and independently of one another, leading to a detected white noise proportional to the total power P_{dc} on the photodetector:

$$\sqrt{S_{P,shot}(f)} = \sqrt{2\hbar\omega_0 P_{dc}} \qquad \left[\frac{\mathrm{W}}{\sqrt{\mathrm{Hz}}}\right]. \tag{2.1}$$

The quantum nature of the shot noise arises from the fact that the power detected at an audio frequency ω is the result of the main electric field carrier frequency of the laser $\vec{E}_0$ beating with electric field fluctuations at an audio frequency away $\vec{e}(\omega)$. Even with a perfectly stable laser $\vec{E}_0$, the quantum vacuum fluctuations at $\vec{e}(\omega)$ would produce power fluctuations from the $\vec{E}_0^*\vec{e}(\omega)$ terms.

Shot noise shows up as a sensing noise in all photodetectors in Advanced LIGO, and dominates the high-frequency region of the DARM spectrum. As the input power is increased, the DARM signal-to-shot-noise ratio increases $\propto \sqrt{P_{input}}$.

2.3.1.2 Quantum radiation pressure noise

Quantum radiation pressure noise (QRPN) is displacement noise arising from amplitude fluctuations of the electric field in the arms. These amplitude fluctuations produce a fluctuating momentum on the optics via radiation pressure, inducing displacement noise.

The amplitude fluctuations are quantum in nature due to the quantum vacuum at the antisymmetric port of the interferometer beamsplitter [48]. The quantum vacuum conspires to create anti-correlated intensity fluctuations entering each arm by interfering with the main laser power.

Intensity fluctuations in the arms would not affect the phase-quadrature interferometer readout of gravitational waves, except for the coupling of radiation pressure. The intensity fluctuations create a “back-action” force on the mirror, which is displaced according to the compliance of mirror pendulum.

The coupling of quantum amplitude fluctuations to phase fluctuations $\mathcal{K}$ is described

$$\mathcal{K} = \frac{8P_{bs}\omega_0}{mL^2\omega^2(\omega_c^2 + \omega^2)} \tag{2.2}$$

where P_{bs} is power on the beamsplitter, m is the mirror mass, L is the arm length, ω_0 is the laser frequency, ω_c is the arm pole describing the number of reflections inside the Fabry-Perot cavity, and ω is the signal frequency. For a conventional interferometer, the displacement due to QRPN can be described [47]

$$\sqrt{S_{x,\mathrm{QRPN}}} = \sqrt{\frac{L^2 h_{\mathrm{SQL}}^2}{2}\mathcal{K}} \tag{2.3}$$

$$= \frac{1}{mL\omega^2}\sqrt{\frac{32 P_{bs}\hbar\omega_0}{\omega_c^2+\omega^2}} \tag{2.4}$$

where h_{SQL} is the standard quantum limit of conventional interferometer strain sensitivity.

From Eq. 2.3 we see that as input power P_{bs} is increased, QRPN also increases. QRPN is attenuated by the compliance of the mirror pendulums, and so is more important at low frequencies. In Advanced LIGO during O3, QRPN never dominates the gravitational-wave spectrum, as angular and length controls noise is much higher.

2.3.2 Thermal noise

Thermal noise refers to the actual displacement in the mirrors induced by thermal fluctuations in the atoms making up the test mass suspension, substrate, and optical coating cause displacement noise in DARM [28–30, 51]. Generally thermal noise increases with mechanical loss or loss angle, as related by the fluctuation-dissipation theorem [52–54].

The fluctuation-dissipation theorem is a general result showing that thermal fluctuations are equivalent to power dissipated in a mechanical or electrical system. For some observable q of a system with admittance $Y(\omega)$ and temperature T, the power spectral density S_q is

$$S_q(\omega) = \frac{4k_B T}{\omega^2}|\mathrm{Re}[Y(\omega)]| \tag{2.5}$$

where ω is the frequency and k_B is Boltzmann's constant. Eq. 2.5 may be used to characterize a wide variety of systems, from Johnson-Nyquist voltage noise in a resistor [55, 56] to optical coatings Brownian displacement noise on the Advanced LIGO core optics [51].

If a fluctuation $F = F_0\cos(\omega t)$ is applied to a lossy oscillating system, the system coordinate $q = q_0\cos(\omega t + \phi(\omega))$ will respond with a phase lag $\phi(\omega)$. The coordinate will have some velocity $\dot{q} = \dot{q}_0\sin(\omega t + \phi(\omega))$. The admittance transfer function $Y(\omega)$ can be written in terms of the driven force and coordinate velocity:

$$Y(\omega) = \frac{\dot{q}}{F}. \tag{2.6}$$

The real part of the admittance is called the *conductance*, and can be expressed

$$\mathrm{Re}[Y(\omega)] = \frac{\langle F\dot{q}\rangle}{\langle F^2\rangle} = \frac{2W_{\mathrm{diss}}}{F_0^2}. \tag{2.7}$$

where the extra factor of two comes from the time average of $\langle F_0 \cos(\omega t)^2\rangle$. Eq. 2.7 gives the second expression for the fluctuation-dissipation theorem [31, 51]:

$$S_q(\omega) = \frac{8k_B T}{\omega^2}\frac{W_{\mathrm{diss}}}{F_0^2} \tag{2.8}$$

The phase lag $\phi(\omega)$ is related to the power dissipated per driven oscillation W_{diss} [54, 57]:

$$\begin{aligned} W_{\mathrm{diss}} &= \langle F\dot{q}\rangle \qquad (2.9)\\ &= F_0\dot{q}_0\langle\cos(\omega t)\sin(\omega t + \phi(\omega))\rangle\\ &= F_0\dot{q}_0\langle\cos(\omega t)[\sin(\omega t)\cos(\phi(\omega)) + \cos(\omega t)\sin(\phi(\omega))]\rangle\\ &= \frac{1}{2}F_0\dot{q}_0\sin(\phi(\omega))\\ W_{\mathrm{diss}} &\approx \frac{1}{2}F_0\dot{q}_0\phi(\omega) \qquad (2.10)\end{aligned}$$

where in the last line we assume $\phi(\omega) \ll 1$. The *quality factor* $Q = 1/\phi$ is another common measure of the loss in an oscillator.

For LIGO test masses, the fluctuating observable we care about is the optic displacement x. The dominant cause of displacement fluctuations is due to mechanical loss in the optic coatings. For a single coating with thickness d, the dissipated power and coating displacement noise $S_x(\omega)$ due to thermal fluctuations can be calculated [31, 58, 59]:

$$W_{\mathrm{diss}} = \frac{F_0^2(1+\sigma)(1-2\sigma)d}{\pi w^2 E}\phi\omega \tag{2.11}$$

$$S_x(\omega) = \frac{8k_B T(1+\sigma)(1-2\sigma)d}{\pi w^2 E}\frac{\phi}{\omega} \tag{2.12}$$

where σ is the coating Poisson ratio, E is the Young's modulus, w is the beam radius, and ϕ is the mechanical loss angle of the coating.

Because of the direct dependence of the thermal noise on mechanical loss angle ϕ, and the fact that this noise is expected to dominate Advanced LIGO design sensitivity around 50 Hz, much coatings research and development is put into designing and measuring low-loss coatings for future detectors. The actual optical coatings in LIGO are multilayered, switching between two coatings to achieve both high reflectivity and low thermal noise [30].

The test mass quadruple suspension system has been designed to limit thermal noise in the measurement band [27]. The fused silica substrate material is chosen for low mechanical loss and has a small contribution to the thermal noise. A minor contribution to the thermal noise is due to the addition of acoustic mode dampers [60]. The thermal noise contribution from these dampers is estimated to degrade interferometer sensitivity by less than 1%.

Brownian motion of the optic dielectric coatings is the dominant noise in the Advanced LIGO design noise budget from 40 to 100 Hz. Advanced LIGO test masses have titania-doped tantala/silica coatings ($\mathrm{TiO_2}$-doped $\mathrm{Ta_2O_5/SiO_2}$), with 25% titania in the tantala layers and varying layer thicknesses to reduce thermal noise [61, 62]. The coating thermal noise contribution is estimated based on optical measurements of aLIGO end test mass witness samples [32]. The correlated noise measurements in Chapter 5 approach the thermal noise limit as the dominant known noise source around 200 Hz. The coating thermal noise can be reduced with improved low-loss optical coatings or cryogenic optics [63].

Future detectors' design curves rely on improved coatings technology for lower thermal noise. Current research and development is focused on finding and testing better, lower-loss coatings materials. Cryogenically-cooled detectors, such as the Japanese detector KAGRA, employ cooled sapphire optics to reduce thermal noise. Changing the frequency of the laser from 1064 nm to 1550 nm or 2 $\mu\mathrm{m}$ may also provide paths to lower coatings thermal noise.

2.3.3 Seismic noise

Seismic noise is the displacement of the core optics due to the motion of the Earth. The vibrations of the Earth are much larger than LIGO optics can tolerate. Therefore, enormous effort is put into isolating the core optics from the ground vibrations, particularly in the GW sensitive range.

First, the LIGO test masses are suspended from a quadruple stage pendulum chain [27]. The test masses form the bottom stage of the chain. These pendulums are suspended from seismic isolation platforms [64] which themselves are supported by hydraulically actuated pre-isolation structures [65].

This arrangement ensures that the seismic noise contribution at the bottom of the chain sits far below the DARM noise curve. However this seismic noise contribution only accounts for linear coupling to the DARM degree of freedom. Coupling can become nonlinear when low-frequency motion is large, up-converting into the gravitational-wave band. There are circuitous paths by which seismic motion can couple to the interferometer output, such as through angular and length degrees of freedom, or via scattered light. Earthquakes, high microseism, and windy conditions can confuse isolation systems by tilting building floors near wind-driven walls, injecting seismic controls motion that can increase scattered light coupling, cause lock loss, and hinder lock reacquisition.

Improvements to the seismic control scheme granted higher robustness to earthquakes in O3, and helped contribute to the highest overall duty cycle the detectors have had in any observing run [66].

2.3.4 Newtonian noise

Newtonian noise is produced by direct gravitational coupling of test masses to fluctuating mass density fields, such as produced by seismicity and atmospheric pressure fluctuations [67–70]. Newtonian noise, dominated by seismic surface waves called Rayleigh waves, is predicted to limit the design sensitivity of the Advanced LIGO detectors from 10 to 20 Hz [71, 72]. Newtonian noise has not been detected in Advanced LIGO, and is predicted to be below O3 sensitivity levels [73].

2.4 Length sensing and control

In order to achieve sensitivity to gravitational waves, each of the five major length degrees of freedom the interferometer must be constantly sensed and controlled [20]. Pound-Drever-Hall (PDH) locking via radio-frequency (RF) laser sensing provides a robust way resonate laser light inside an optical cavity [74].

PDH locking requires a beatnote between a carrier and sideband which carries the information about that cavity length. PDH locking provides a strong error signal when a cavity is near resonance. A control loop is built around a PDH error signal for each of the degrees of freedom to maintain resonance.

To accomplish this length control of the five degrees of freedom, a double RF modulation scheme is employed. Two RF sidebands at $f_1 = 9\ \mathrm{MHz}$ and $f_2 = 45\ \mathrm{MHz}$ enter the interferometer alongside the carrier. The macroscopic lengths of the cavities are carefully chosen such that the carrier and RF sidebands are resonant where they are required, and antiresonant elsewhere to act as references for the phase changes due to length changes.

In this section we'll overview the detector design considerations for the length sensing and control of Advanced LIGO.

2.4.1 Degrees of freedom

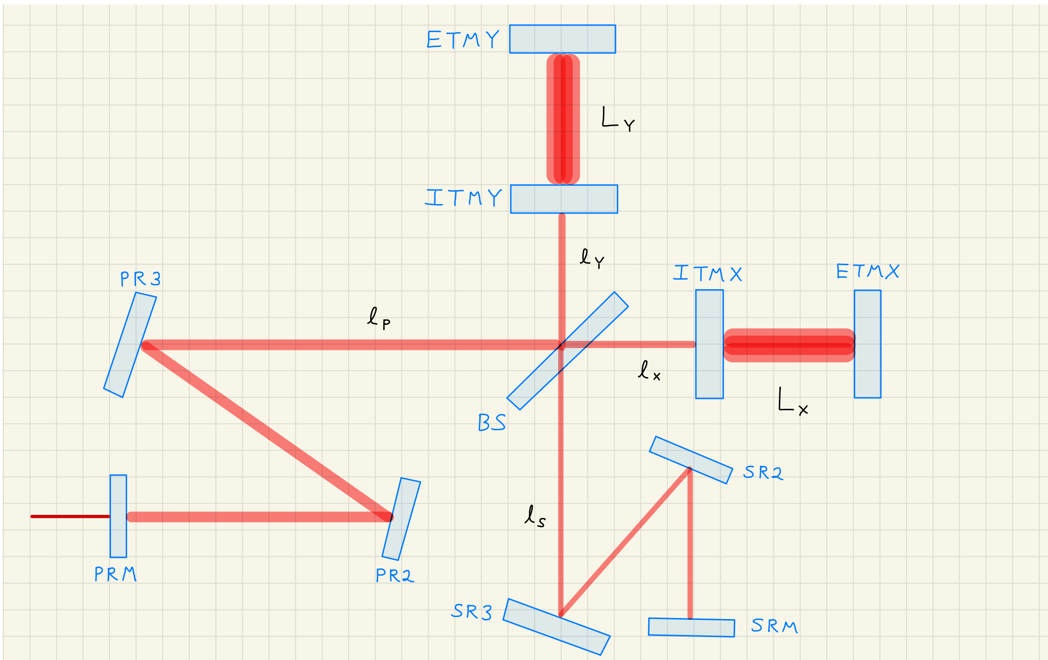

Figure 2.2: Diagram of the Advanced LIGO length degrees of freedom. L_x and L_y are the lengths of the arm cavities, equal to ~ 4 km long. l_x and l_y forms the inner Michelson. l_p includes the entire beam length from PRM to the beamsplitter, and forms part the PRCL length according to Eq. 2.16. l_s includes the entire beam length from SRM to the beamsplitter, and forms part the SRCL length according to Eq. 2.17.

There are five main degrees of freedom that must be controlled to allow LIGO to be sensitive to gravitational waves The five main Advanced LIGO interferometric length degrees of freedom are

$$L_- = L_x - L_y \tag{2.13}$$

$$L_+ = \frac{L_x + L_y}{2} \tag{2.14}$$

$$l_{MICH} = l_x - l_y \tag{2.15}$$

$$l_{PRCL} = l_p + \frac{l_x + l_y}{2} \tag{2.16}$$

$$l_{SRCL} = l_s + \frac{l_x + l_y}{2} \tag{2.17}$$

Figure 2.2 illustrates these degrees of freedom. There are other degrees of freedom

	Arm cavities	Power recycling cavity	Signal recycling cavity
Carrier	resonant	resonant	antiresonant
9 MHz	antiresonant	resonant	antiresonant
45 MHz	antiresonant	resonant	resonant

Table 2.1: Resonance conditions of the carrier and RF sidebands, copied from [76]. No RF sidebands resonate in the arms, since their purpose is to sense length changes in the corner degrees of freedom PRCL, SRCL, and MICH. The carrier laser frequency is antiresonant in the SRC to broaden the detector bandwidth, operating in the resonant sideband extraction scheme (see Section 3.6). 9 MHz resonates in the PRC and is used to detect PRCL. 45 MHz resonates in both the PRC and SRC, and is used to detect SRCL and MICH, SRCL in the I-phase and MICH in the Q-phase (see Table 2.2).

of the interferometer that are unimportant for control, like overall displacement of the entire interferometer.

2.4.2 Macroscopic cavity lengths

The control scheme revolves around setting up robust PDH error signals for each of the five degrees of freedom. The designed cavity lengths enables the control scheme by resonating carrier and sidebands in the correct cavities. A good overview of this process is presented in [20] for the Advanced LIGO length choices, and [75] for the Caltech 40m length choices. Table 2.1 overviews where the carrier and RF sidebands are resonant in Advanced LIGO.

Briefly, the lengths are chosen by first assuming the carrier ω_0 is resonant in the CARM cavity. Then, the PRC length is chosen such that $f_1 = 9.1$ MHz and $f_2 = 45.5$ MHz both resonate alongside the carrier:

$$l_{PRCL} = \left(n + \frac{1}{2}\right)\frac{c}{2f_1} \tag{2.18}$$

where n is an integer. Because $f_2 = 5f_1$ via phase lock, if f_1 resonates in a cavity then so will f_2. The factor of $1/2$ arises because of the π phase flip accrued by the sidebands upon reflection from the arms.

The SRC length is chosen such that f_2 resonants but f_1 does not:

$$l_{SRCL} = m\frac{c}{2f_2}, \qquad l_{SRCL} \neq k\frac{c}{2f_1} \tag{2.19}$$

where m and k are integers. The factor of $1/2$ for the sideband arm cavity reflection still exists, but because LIGO is running in resonant sideband extraction with

Degree of freedom	Sensor
CARM	REFL 9 I
DARM	DCPDs
PRCL	POP 9 I
SRCL	POP 45 I
MICH	POP 45 Q

Table 2.2: Sensors for each length degree of freedom. The DARM degree of freedom uses DC readout, whereas all others use PDH locking RF readout. The RF sensor names come from (1) location, (2) 9 or 45 MHz beatnote, and (3) readout quadrature. Figure 2.1 shows the location of each sensor.

SRC tuning $\phi = \pi/2$, the carrier and sideband phase both accrue an additional $\pi/2$ one-way phase while traveling in the SRC.

Finally, a Schnupp length asymmetry $l_{Schnupp} = l_x - l_y$ in the inner Michelson is selected. The Schnupp asymmetry is chosen such that the 45 MHz sideband f_2 is preferentially transmitted into the SRC. It also enables the PDH sensing of MICH in the Q-phase of 45 MHz sensors.

For Advanced LIGO, the design parameters for the above lengths are [20, 76]

$$L_{arm} = 3994.5 \text{ m} \tag{2.20}$$

$$l_{PRCL} = 57.6557 \text{ m} \tag{2.21}$$

$$l_{SRCL} = 56.0084 \text{ m} \tag{2.22}$$

$$l_{Schnupp} = 0.08 \text{ m} \tag{2.23}$$

2.4.3 Length sensors

The length sensors detect power fluctuations which correspond to length changes in the interferometer. The DARM sensor, the DCPDs, uses DC readout to measure length changes. All other degrees of freedom use Pound-Drever-Hall locking, or RF readout, for length control. Table 2.2 overviews which sensors are used to detect length changes of each degree of freedom for O3.

A PDH signal from the reflection of a cavity is typically sensed in the I-phase, or cosine quadrature. A PDH signal from a length offset, like with MICH and DARM, is typically sensed in the Q-phase, or sine quadrature. Chapter 2 of Martynov [77] and Appendix C of Hall [57] provide good derivations of the heterodyned PDH signal.

DARM is sensed via homodyne of the GW signal with the DARM DC offset light.

The power in the antisymmetric port increases quadratically with the DARM offset:

$$\Delta L_{-} \propto P_{as}^{2}. \tag{2.24}$$

The DARM offset is servoed to maintain $P_{as} \approx 20$ mA on the DCPDs. This linearizes DARM motion δL_{-} such that

$$\frac{P_{as}}{\delta L_{-}}(\omega) \propto P_{as} \tag{2.25}$$

We note here that one cannot improve sensitivity by increasing the DARM offset. The signal from Eq. 2.25 does increase linearly with a DARM offset increase, but the shot noise also increases linearly, as can be seen from Eqs. 2.1 and 2.24.

CARM is sensed from the beatnote between the carrier, which carries the arm length information, and the 9 MHz sideband, which acts as a stable reference. CARM is detected in the REFL port in the I-phase. CARM is technically stabilized to the coupled-cavity consisting of the PRC and arm cavities, and so is sensitive to PRCL as well. The CARM loop servos both to zero indiscriminately. See Section 2.4.4 for a discussion of CARM and PRCL gain hierarchy.

PRCL is sensed from the beatnote between the 9 MHz sideband, which carries the PRCL information, and the carrier, which acts as the stable reference. PRCL is detected in the POP port in the I-phase. As stated before, carrier resonates in the PRC and does carry PRCL information, but this is already sensed and stabilized by CARM.

SRCL is sensed from the beatnote between the 45 MHz sideband, which carries the SRCL information, and the carrier, which acts as the stable reference. SRCL is detected in the POP port in the I-phase. 45 MHz is the only light that resonates in the SRC, and so is most sensitive to SRCL motion. However, 45 MHz is also resonant in the PRC, and carries PRCL information as well. In practice, the PRCL error signal from the 9 MHz is feedforward to SRCL to cancel the PRCL motion in the SRCL loop.

MICH is sensed from the beatnote between the 45 MHz sideband, which carries the MICH information, and the carrier, which acts as the stable reference. MICH is detected in the POP port in the Q-phase. MICH operates on a perfect dark fringe for carrier, so there is no linear response of carrier to MICH motion in the POP port. The Schnupp asymmetry is an effective MICH offset for the 45 MHz sidebands, however, so 45 MHz signal sidebands do appear in the symmetric port.

Degree of freedom	Unity gain frequency
CARM	20 kHz
DARM	60 Hz
PRCL	50 Hz
SRCL	25 Hz
MICH	10 Hz

Table 2.3: Bandwidths of the main LIGO length control loops for O3. The unity gain frequency is the frequency at which the open loop gain magnitude equals one. These values are approximate, as the true values fluctuate as the thermalization of the interferometer changes the optical plant.

2.4.4 Gain hierarchy

The gain hierarchy of the length control loops is also important. For example, the PRCL error signal is made from the 9 MHz beatnote, which senses both CARM motion and PRCL motion. The CARM response to length changes is far more sensitive than PRCL, which would make sensing low frequency PRCL motion difficult in principle. The huge CARM control loop gain relative to the PRCL gain stabilizes carrier to be used as a static reference for the PRCL loop [76, 78].

The control loop gain hierarchy goes CARM, DARM, PRCL, SRCL, and MICH. Other than the CARM loop, the control bandwidths are kept as low as possible to avoid injecting excessive length sensing noise while still controlling displacement noise. Table 2.3 shows the approximate unity gain frequencies of the LIGO control loops in O3.

2.4.5 Feedforward

Controls noise considerations are an important contribution to the final DARM noise. The length sensing noise of SRCL and MICH, and to a lesser extent PRCL, still show up in DARM.

However, these sensing noises are well-monitored via the SRCL, MICH, and PRCL control signals. The control signals are summed into the DARM control signals to preemptively cancel their induced noise. This process is known as *feedforward* [57, 79].

2.5 Lock acquisition

Lock acquisition is the process of bringing the detector into a regime where maximum power buildup is achieved in the arm cavities and all interferometer degrees of freedom are controlled [77, 80]. The locking begins with only 2 W of infrared input power.

First, green lasers at each end station are locked to each arm cavity length. Then, the green transmission beams through each arm are combined with main carrier light to stabilize the PSL to the common arm cavity length and control the differential arm length. The common arm length is moved such that the main infrared laser is *antiresonant* in the arms to avoid arm flashes during corner locking.

Next, PRCL, SRCL, and MICH are locked to the infrared laser via Pound-Drever-Hall error signals in the dual-recycled Michelson (DRMI) configuration. This process involves waiting for "flashes" of resonance in DRMI, then quickly triggering the corner controls. DRMI locking is stochastic, but happens relatively quickly ($<$ 30 seconds) for a well-tuned interferometer settings. The problem is tuning the interferometer settings, including alignment and triggering, to reliable catch the lock.

With DRMI locked, all main degrees of freedom are controlled but there is no infrared light in the arm cavities. To transition to full infrared control, first the PRCL, SRCL and MICH error signals are transitioned from using the first-order radio-frequency sidebands to using the third-order sidebands [81]. This is done because the first-order sideband error signals become zero as the arms are brought from antiresonance to resonance.

Then, the CARM offset reduction stage begins. The green common arm length is brought from infrared antiresonance to the side of an CARM infrared fringe. Here, CARM control is handed off to infrared transmission through the arms.

Next, the infrared light is brought to resonance, where both DARM and CARM are transferred to PDH error signals. The arm angular controls are engaged, and we wait a few minutes here to allow convergence of the slow angular control loops and adjustment of the interferometer to the first stage of thermalization. For the DC readout scheme, a 10 pm DARM offset is applied to allow some carrier light to leave the antisymmetric port and act as a local oscillator for light carrying the gravitational wave signal. The output mode cleaner is locked to this local oscillator beam, further cleaning the mode of the beam to allow only GW signal light onto the output photodetectors.

At this stage the entire interferometer is totally "locked", i.e. it is on resonance and controlled via the main infrared light. Here, the input power is slowly increased from 2 W to 35 W, in steps of 5 W, to avoid rapid changes to the angular optical plants.

Now the interferometer is locked at high power in the “high-noise” state. Here the transition to “low-noise” begins. Low-noise angular and length controls are engaged, feedforward filters are engaged, frequency and intensity noise are suppressed, and squeezed light is injected to achieve maximum sensitivity to gravitational waves. At this point the locking process is complete and the interferometer is ready for observing.

The steps taken to acquire lock are done automatically using a state machine called *Guardian* [82]. Because the locking sequence is not deterministic and can be hindered by poor environmental conditions, there is some variability of the lock acquisition time. The locking sequence takes approximately 30 minutes in good environmental conditions and with good initial alignment. Much of this time is used to allow various slow drift control loops to settle, allow optics to thermalize, and smoothly and reliably move between different control and actuation configurations.

A “lockloss" occurs when the detector falls out of the sensitive linear regime. Locklosses are caused by strong earthquakes, controls and sensor saturations, drifting misalignment, control loop instabilities, and large glitches of known and unknown origin. The cause of lock losses are monitored, and if possible mitigated, to improve detector duty cycle.

2.6 O3 detector upgrades

This section will discuss the instrument upgrades that facilitated the increase in sensitivity and duty cycle for O3, focusing on hardware upgrades to the interferometers. Not all upgrades that were performed will be covered, only those that are immediately relevant to topics in this thesis. See [2] for a full list of upgrades.

2.6.1 Laser power increase

Increasing the laser power reduces instrument noise at high frequency where the sensitivity is shot-noise-limited but comes with operational challenges. Hardware upgrades to the pre-stabilized laser and core optics allowed for an increase in average circulating power in the arm cavities to 201 kW at LHO and 239 kW at LLO for O3 (see Table 3.3). The major technical challenges of operating a high-power interferometer are caused by radiation pressure inducing instabilities in core optic control and absorption of the test masses.

The original aLIGO pre-stabilized laser (PSL) design [36] took the output of a Nd:YAG non-planar ring oscillator (NPRO) operating at 1064 nm and successively amplified the output to above 150 W. The original amplifier chain consisted of a 35 W solid-state amplifier (“front end") followed by a high-power injection-locked ring oscillator. In addition to operational challenges, the high-power oscillator and its high coolant flow produced significant fluctuations of the beam size and pointing angle [83], and thus was never used during an observing run.

For O3 the high-power oscillator was replaced at both observatories with a smaller single-pass solid-state amplifier (neoLASE neoVAN-4S) that requires less coolant flow. The new amplifier produces roughly 70 W of stable output power during the run. After input optics and mode-cleaning cavities, this provides up to 50 W at the power-recycling mirror.

The reduced coolant flow and damping and tuning of problematic optic mounts has reduced the amplitude of angular beam jitter. The higher input power, in addition to the squeezer (Section 2.6.2), lead to improved sensitivity above 100 Hz.

2.6.2 Squeezer

For O3 an in-vacuum squeezer was installed at each site to inject squeezed vacuum into the interferometers and reduce shot noise. A full description of the new squeezer can be found in [44]. In contrast to previous squeezers for gravitational-wave detection [84–86], the squeezed vacuum source (an optical parametric oscillator) is placed inside the vacuum envelope on a separate suspended platform [87]. This reduces squeezing ellipse phase noise and backscattered light noise [88]. The squeezer has been fully integrated into the automated lock acquisition sequence.

Section 2.6.1 discussed increasing the input power to the interferometer, which increases interferometer sensitivity by enhancing the gravitational-wave signal. Injecting squeezed vacuum improves the signal-to-noise ratio by directly decreasing sensing noise. For an entirely shot-noise limited detector, $\sim$3 dB of squeezing is equivalent to doubling the arm cavity power to $\sim$450 kW. With squeezing, the detector sensitivity is brought closer to the Advanced LIGO design sensitivity, which did not include squeezing but specified 750 kW arm cavity power, three times what was achieved in O3. The design vs measured quantum noise in O3 (dashed black line and purple line in Figure 3.1) illustrate the extreme benefit of squeezing.

Above 50 Hz the interferometer sensitivity is increased by 2.0 dB and 2.7 dB at LHO and LLO, respectively. This provides a 12% and 14% increase in binary neutron star inspiral range at each respective site.

Below 50 Hz, injecting frequency-independent squeezed vacuum, as is done during O3, increases the quantum radiation pressure noise. The low-frequency noise at LLO is small enough that this increase in quantum radiation pressure noise is detrimental to sensitivity and binary neutron star inspiral range. The current squeezing level at LLO cannot be further increased without causing a reduction in range [44]. The squeezing angle is set to 7° from the optimal high-frequency configuration. This increases range by reducing low-frequency radiation pressure noise at the expense of a 0.5 dB increase in shot noise at high frequencies. This effect is more fully explored in [89].

Detuning of the signal recycling cavity also produces frequency-dependent squeezing. This effect was used to identify and correct a 2–3 nm detuning in the signal-recycling cavity length locking point at LLO.

2.6.3 Core optic replacement

Several of the core optics were replaced before O3 to improve detector sensitivity, stability, and lock acquisition performance. The motivation and performance benefit of each replacement is presented here.

At both sites the two end test masses were replaced. To improve the lock acquisition sequence via the auxiliary laser system (ALS), the test mass optical coatings reflectivity for green (532 nm) laser light were increased. The green arm cavity finesse increased from 15 to 70 at LHO and to 100 at LLO, providing finer beam quality for locking ALS. This improves the reliability of the early stages of lock acquisition, where control of each arm length is transitioned from green to infrared error signals [80].

The primary reason for replacing the end test masses was to reduce the scatter loss and increase the circulating power. The $\sim$10 ppm reduction in scatter loss has resulted in improved power-recycling gain at both sites. However when increasing the circulating power in the arm cavities, the power-recycling gain has not increased as expected due to nonuniform absorption on the optics increasing scatter losses in the arm cavities. These so-called “point absorbers” stunt the full capabilities of the interferometer to achieve maximum power, instead absorbing high amounts of power and distorting the cavity geometry.

The X-arm input test mass at LHO was replaced before O3 following the identification of a point absorber in the coating. The presence of the point absorber limited high-power operation and coupled jitter noise from the pre-stabilized laser to DARM. The new input test mass shows no significant absorbers. Similar defects have been found on several other test masses currently installed.

The signal-recycling mirror (SRM) at both sites was replaced. The previous SRM was an aluminum and fused-silica composite with a 2" diameter optic that allowed for easy mirror replacement. The composite SRM introduced thermal noise due to internal modes of the composite system with high mechanical loss. The replacement SRM is monolithic fused silica, 150 mm diameter, with no measurable thermal noise contribution to DARM. To maximize the binary neutron star inspiral range, the SRM transmission should be reduced with increasing circulating optical power. For O3, the SRM transmission was reduced from 37% to 32%. The design SRM transmission is 20%.

The reaction masses, which are suspended in a separate pendulum chain behind the end test masses, provide high-frequency actuation via the electrostatic drive [27]. The proximity of the reaction mass to the end test masses can increase the damping noise due to residual gas bouncing between the test mass and reaction mass. This noise is known as squeezed film damping [90]. Before O3 the reaction masses were replaced with annular reaction masses with cored out centers that retained the original electrode pattern. These annular end reaction masses are expected to have reduced the squeezed film damping noise by a factor of 2.5 below 100 Hz [91].

Chapter 3

SENSITIVITY OF THE ADVANCED LIGO INTERFEROMETERS DURING OBSERVING RUN THREE

In this chapter, we will review the sensitivity achieved by the Advanced LIGO Hanford detector during observing run three (O3). First we overview the overall noise performance of the detector. Then we will dive into the technical details of projects the author was focused on, including novel arm power measurement techniques, frequency stabilization improvements and noise budget, intensity stabilization noise budget, auxiliary length control improvements and noise characterization, and feedforward transfer functions from corner degrees of freedom to DARM.

Part of the above is covered in the O3 commissioning paper [2]. Chapter 5 covers related content on the correlated noise measurement.

3.1 O3 overview

3.1.1 Advanced LIGO noise budget

The LIGO detectors are sensitive to gravitational waves via the differential arm (DARM) motion they induce in the interferometer. The sensitivity of a detector is limited by the collection of noises coupled to the gravitational-wave readout. The detector noise is low enough to detect GWs across a broad frequency band from 20 Hz to 5 kHz. To improve the sensitivity of a detector, a source of noise or noise coupling must be identified and mitigated. The *noise budget* is a tool used in this process.

Figure 3.1 shows the LIGO Hanford DARM noise budget for O3. The main measured DARM noise is the blue trace on Figure 3.1. This represents the achieved sensitivity of the detector to gravitational waves. The noise budget collects all known and quantified noise terms onto a single plot. Also included are the DARM noise for O1 and O2, and the Advanced LIGO design sensitivity representing the ultimate sensitivity possible.

There are two overall types of known noises. The first are *fundamental* noise sources. Fundamental noises are the known expected limits of the performance of an interferometer designed like Advanced LIGO. These include noises such as

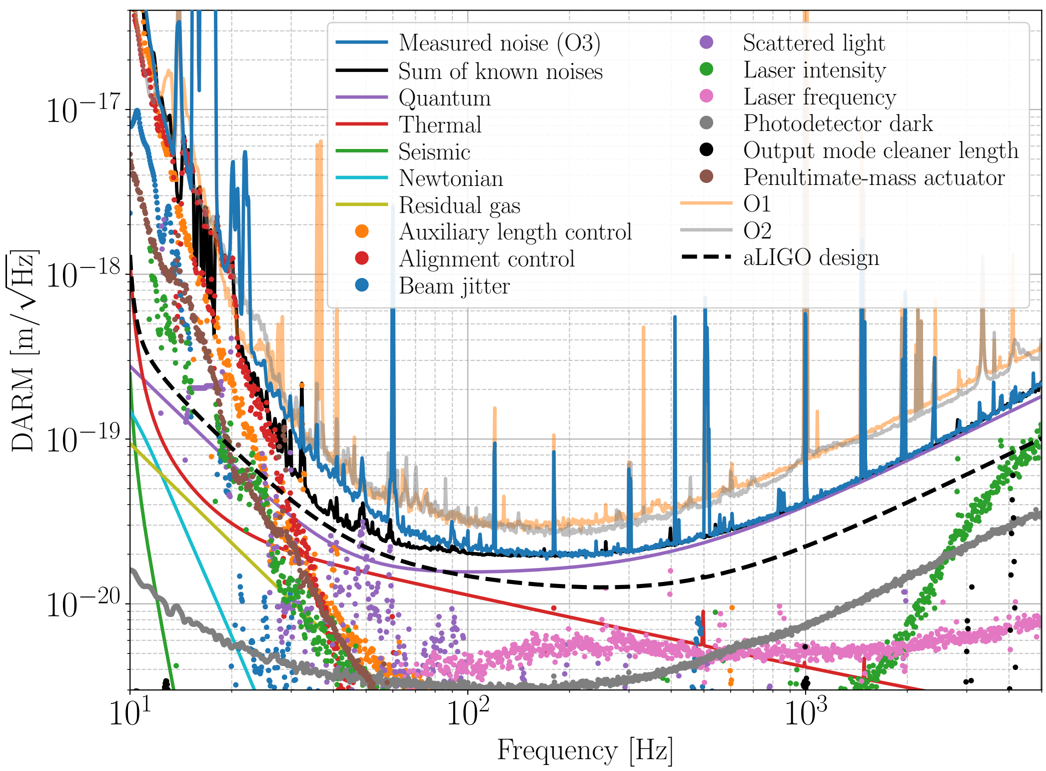

Figure 3.1: Differential arm noise budget for LIGO Hanford in O3 [2]. Also included are the instrument noise floors for previous observing runs, as originally presented in [43] and [92], and the Advanced LIGO design sensitivity [20].

quantum noise, thermal noise, seismic noise, newtonian noise, and residual gas noise. Fundamental noise contribution to DARM are estimated, but typically do not have an independent sensor other than DARM itself. Fundamental noises are plotted as straight lines in Figure 3.1. The sum of fundamental noises is the Advanced LIGO design sensitivity, plotted as the dashed black line in Figure 3.1.

The other type of known noise is *technical* noise. Technical noises are the known, measured limits to DARM from auxiliary aspects of the interferometer These include noises such as length control, angular control, beam jitter, scattered light, laser intensity, laser frequency, photodetector dark noise, and coil driver actuator noise. Technical noise contribution to DARM can be directly estimated by injecting excess noise, measuring the coupling transfer function, and projecting the usual auxiliary noise into DARM. Technical noises should not in principle limit the sensitivity of the detector, but for practical purposes cannot be lowered without additional research. Technical noises are plotted as dots lines in Figure 3.1.

There are some known-unknown noises. These noises often are hard to quantify due to the lack of a independent sensor. Some examples include excess scattered light, higher-order-mode light co-propagating with DARM light, radio-frequency sideband amplitude and phase noise, nonlinear up-conversion from large seismic or angular motion at low frequencies, and charge noise from excess charge on the test masses. Some only periodically appear making projections into DARM even more difficult. Such known-unknown noises are hard to quantify for all times, but are independently checked and mitigated if some excess noise appears in DARM without explanation. Examples of known-unknown noise mitigation are the discharging of the Hanford test masses after a strong nearby earthquake, and the thermal compensation system heating important optics to minimize higher-order-mode content in the interferometer.

Finally, there are some unknown-unknown noises, or "mystery" noise. Mystery noise represents the noise we don't understand that limits DARM. At Hanford around 40 Hz on Figure 3.1 there is a gap between the measured DARM noise in blue and the expected DARM noise in black. This gap represents most of the sensitivity difference between Hanford and Livingston detectors, seen in Figures 3.2 and 3.3. Livingston enjoys lower controls noise at low frequency, higher circulating power, and better observed squeezing, but also observes some mystery noise. Identifying and mitigating mystery noise remains the most important long term task of commissioners.

For Hanford in O3, DARM is largely limited by fundamental quantum shot noise in the 100 Hz to 5 kHz region. Below 30 Hz, the angular controls noise dominates DARM, followed closely by auxiliary length control. The most dramatic improvements made for O3 are due to the injection of squeezed light into the antisymmetric port and the increase of resonating laser power inside the interferometer, both of which lower the quantum shot noise and improve the high-frequency sensitivity to the level seen in Figure 3.1.

Other artifacts include strong lines at various frequencies [93]. The 60 Hz line and harmonics comes from mains power. The 500 Hz line and its harmonics are a combination of all main optic suspension violin modes. The 20 Hz set of lines are alignment dither lines used to hold the beam spots on the optics steady. At 15 Hz, 410 Hz, and 1084 Hz, calibration lines are injected to continuously measure the detector response [94]. Offline subtraction of some of the known lines and noise in DARM help improve offline data analysis [92].

This thesis will focus on the frequency noise (Section 3.4) and intensity noise contribution to DARM (Section 3.5).

3.1.2 Astrophysical range

A useful metric for understanding the sensitivity of a detector is the *binary neutron star inspiral range*, or simply *range*. The range reported is the luminosity distance at which a detector is sensitive to an angle-averaged merger of two $1.4\ M_\odot$ neutron stars for a canonical SNR of 8 [95–97]. The angle average is over the orientation of binary systems and position relative to the detector antenna patterns. The range does not represent a strict maximum distance at which a binary neutron star merger can produce a significant signal. The LIGO Livingston Observatory (LLO) has achieved a binary neutron star range of around 134 Mpc, while the LIGO Hanford Observatory (LHO) has achieved a range of around 111 Mpc. The detector sensitivity to heavier binary black holes extends much further than binary neutron stars.

The range is calculated every minute from the online calibrated strain power spectral density. Figure 3.2 shows the range of each observatory during O3. Figure 3.3 shows two histograms of the binary neutron star range during O3.

3.1.3 Duty cycle

During O3 both detectors were operational a greater percentage of the time compared to the past two observing runs, with LHO and LLO achieving observation duty cycles of 74.6% and 77.0%, respectively, with coincident observation 62.2% of the time. Time not observing is spent either acquiring lock, unlocked and undergoing maintenance, unlocked due to unfavorable environmental conditions (earthquakes, wind, storms), or locked and in a state of commissioning, where improvements are made to the detectors.

"Locking" the detector is the process of achieving laser resonance in every part of the interferometer simultaneously, so the detector is sensitive to gravitational waves. The locking process is automated via Guardian, a state machine whose states are programmed by detector scientists to transfer the interferometer from down to sensitive to GWs [82]. While some states of lock acquisition are faster, the overall the lock acquisition time has not changed significantly from run to run. The rate-limiting steps to lock acquisition are the slow power-up to accommodate thermal changes to the interferometer geometry, and the dual-recycled Michelson lock acquisition time.

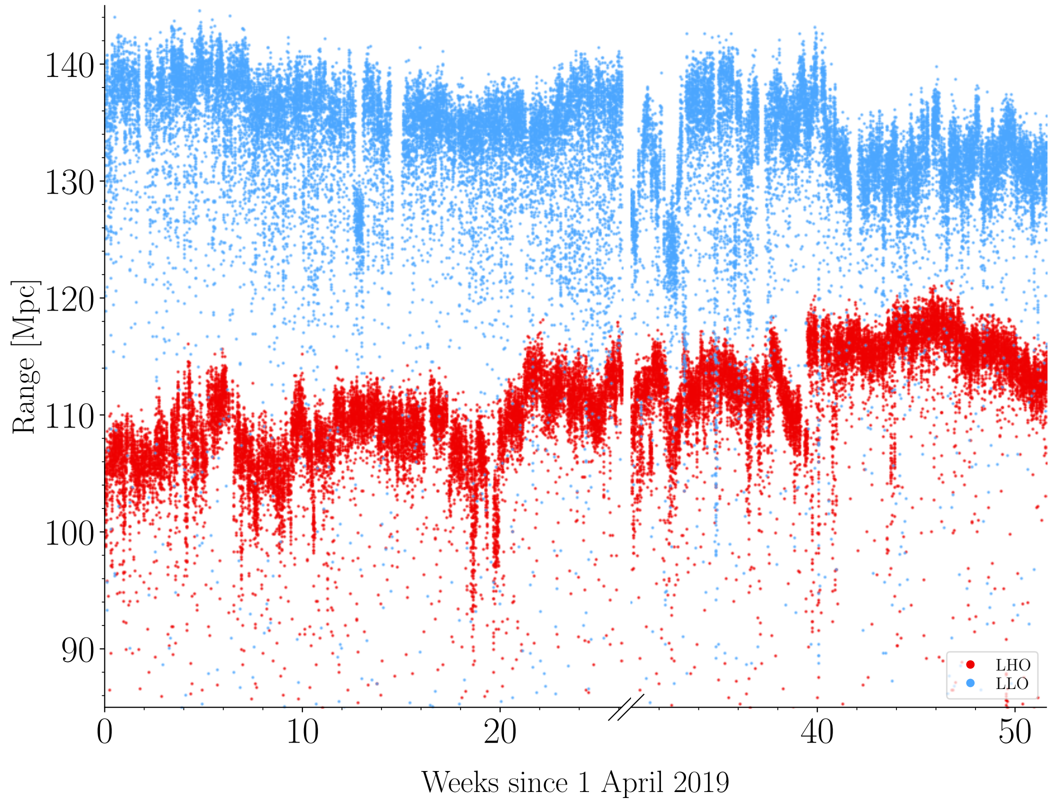

Figure 3.2: Binary neutron star inspiral range of the Hanford and Livingston detectors during O3 [2]. The break in the horizontal axis corresponds to the month-long observing break in October 2019. Brief but significant drops in the range at both sites are caused by instrumental glitches of unknown origin.

Once lock is acquired, the detectors in O3 are more likely to remain locked that in previous observing runs due to improvements in the seismic isolation system, earthquake warnings, and robust detector controls. Table 3.1 quantifies the improvement in average and median lock length and duty cycle over the different observing runs.

Figure 3.4 shows the integrated time-volume sensitivity to binary neutron stars for both sites over the three observing runs. The increase in sensitivity combined with the higher duty factor have resulted in a dramatic increase in the observed time-volume integral, and a roughly proportional increase in gravitational-wave event candidates [98, 99].

3.1.4 Table of O3 Parameters

The remaining sections in this chapter will overview topics in commissioning the Hanford detector undertaken with the goal of understanding and improving the

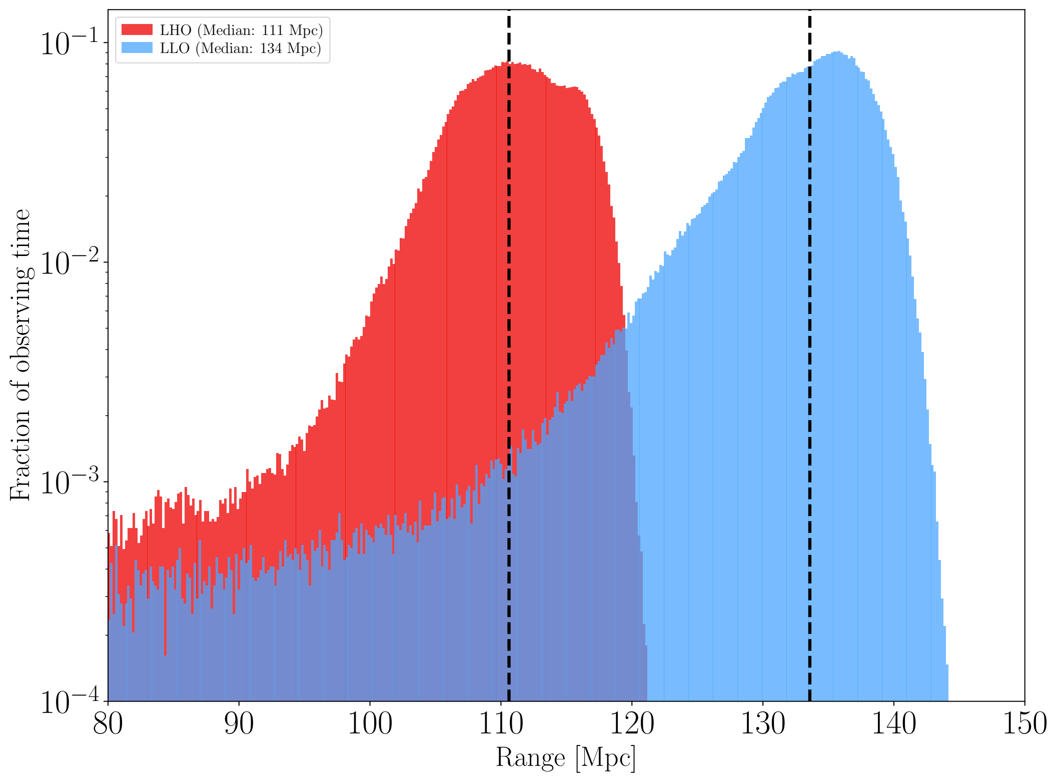

Figure 3.3: Binary neutron star inspiral range histogram of the Hanford and Livingston detectors during O3 [2]. Large brief glitches report a low detector range, skewing the distribution lower. Livingston also experienced daily scattering shelves due to anthropogenic noise, further skewing their range.

detector performance.

Observatory	O1	O2	O3a	O3b
LHO				
Mean (hr)	9.8	9.4	12.4	14.9
Median (hr)	7.2	4.7	8.8	8.9
Duty cycle (%)	62.6	70.6	71.2	78.8
LLO				
Mean (hr)	5.7	5.5	10.2	14.5
Median (hr)	1.9	2.9	6.5	9.3
Duty cycle (%)	55.3	65.8	75.7	78.6

Table 3.1: Mean and median times of low-noise lock segments for each observing run and overall observing run duty factor. Large glitches or unfavorable weather and seismic conditions can knock the interferometers out of lock, reducing the total observing time. In addition to improved sensitivity, both detectors have improved resistance to large disturbances.

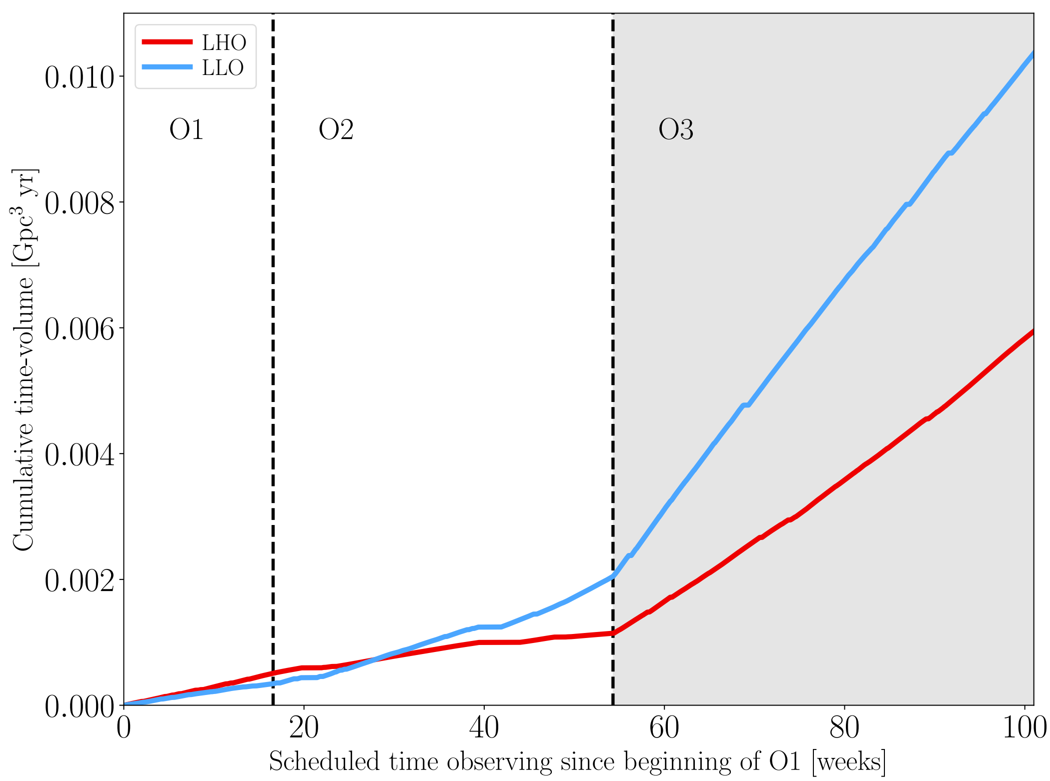

Figure 3.4: Integrated observation time-volume sensitivity over all three observing runs [2]. The observed volume is a sphere with radius equal to the binary neutron star range. The observed time is when the detector was locked and sensitive to gravitational waves. The sharp increase in integrated time-volume is due to the much greater sensitive volume during O3 relative to O1 and O2.

Parameter	Symbol	LHO Value	LLO Value	Units
Squeezing Levels	$\mathrm{dB}_{\mathrm{SQZ}}$	2.0	2.7	dB
1st Modulation Sideband Frequency	f_9	9.100230	9.099055	MHz
2nd Modulation Sideband Frequency	f_{45}	45.501150	45.496925	MHz
3rd Modulation Sideband Frequency	f_{118}	118.302990	118.287715	MHz
1st Modulation Depth	Γ_9	0.135	0.14	rads
2nd Modulation Depth	Γ_{45}	0.177	0.16	rads
3rd Modulation Depth	Γ_{118}	0.012	0.019	rads
ETMX Transmission	T_{ETMX}	3.9	4.0	ppm
ETMY Transmission	T_{ETMY}	3.8	3.9	ppm
ITMX Transmission	T_{ITMX}	1.50	1.48	%
ITMY Transmission	T_{ITMY}	1.42	1.48	%
PRM Transmission	T_{PRM}	3.1	3.1	%
SRM Transmission	T_{SRM}	32.34	32.40	%
Arm Length	L	3994.5	3994.5	m
Power-Recycling Cavity Length	l_{P}	57.7	57.7	m
Signal-Recycling Cavity Length	l_{S}	56.0	56.0	m
Schnupp Asymmetry	$l_{\mathrm{schnupp}} = l_{\mathrm{x}} - l_{\mathrm{y}}$	0.08	0.08	m
Arm Free Spectral Range	f_{FSR}	37.5	37.5	kHz
X Arm Cavity Pole	f_{X}	45.1	44.5	Hz
Y Arm Cavity Pole	f_{Y}	42.7	44.5	Hz
CARM Cavity Pole	f_{CARM}	0.6	0.4	Hz
DARM Cavity Pole	f_{DARM}	411	455	Hz
X Arm Finesse	$\mathcal{F}_{\mathrm{X}}$	415.6	421.3	-
Y Arm Finesse	$\mathcal{F}_{\mathrm{Y}}$	439.2	421.3	-
ETMX Green Transmission	T^g_{ETMX}	7.9	4.8	%
ETMY Green Transmission	T^g_{ETMY}	7.9	5.0	%
ITMX Green Transmission	T^g_{ITMX}	0.96	0.95	%
ITMY Green Transmission	T^g_{ITMY}	1.10	1.11	%
X Arm Green Cavity Pole	f^g_{X}	274.6	175.4	Hz
Y Arm Green Cavity Pole	f^g_{Y}	278.8	186.5	Hz
Input Mode Cleaner Modulation Frequency	f_{24}	24.1	24.1	MHz
Input Mode Cleaner Modulation Depth	Γ_{24}	13	16	mrads
Input Mode Cleaner Round Trip Length	L_{IMC}	32.9434	32.9465	m
Input Mode Cleaner Cavity Pole	f_{IMC}	8625.2	8919.4	Hz
Input Mode Cleaner Finesse	$\mathcal{F}_{\mathrm{IMC}}$	527.5	510.1	-

Table 3.2: Summary optical and physical parameters of the Advanced LIGO interferometers during O3. Measured O3 arm powers and power recycling gain are reported in Table 3.3.

3.2 Arm power measurement

The circulating laser power in the arm cavities governs the optical gain of the interferometer response to gravitational-wave signals. The arm power is difficult to estimate precisely due to large uncertainty in the power on the beamsplitter and optical gain of the arm cavities. Uncertainties are dominated by photodetector calibration and varying interferometer optical losses.

The arm powers in a power-recycled interferometer with a 50:50 beamsplitter should follow

$$P_{\mathrm{arm}} = \frac{1}{2} P_{\mathrm{in}} g_{\mathrm{p}}^2 g_{\mathrm{arm}}^2, \tag{3.1}$$

where P_{arm} is the power in an arm, P_{in} is the input power, g_{p}^2 is the power-recycling gain, and g_{arm}^2 is the arm power gain.

The input power P_{in} is the power incident on the power-recycling mirror, and is estimated from a pick-off just before entering the interferometer. The power on the beamsplitter P_{bs} is estimated directly from a pick-off of the power-recycling cavity. The power-recycling gain is estimated from the ratio of the power incident on the beamsplitter over the input power: $g_{\mathrm{p}}^2 = P_{\mathrm{bs}}/P_{\mathrm{in}}$. Finally, the arm power gain g_{arm}^2 is estimated from the input and end mirror transmissions, as well as the round-trip loss.

Photodetector power uncertainty originates from uncertainty in calibration, losses along beam path combined with beam size mismatch and misalignment. We have assumed a total uncertainty of 5% in power estimated from pick-off photodetectors, P_{in} and P_{bs}. The arm gain g_{arm}^2 at Livingston is assumed to be 265 with uncertainty of 5%. The Hanford X-arm gain is 262, while the Y-arm gain is 276; the 5% gain difference is due to the slightly different transmissions of the input test masses at Hanford. Results are shown in Table 3.3.

A new technique to measure the arm powers using radiation pressure was developed prior to O3 [100]. The length of the signal-recycling cavity (SRCL) is modulated, creating audio sidebands on the carrier laser in the signal-recycling cavity. The audio sidebands enter the arm cavities producing a light power modulation that has opposite sign in each arm cavity, causing a strong signal to appear in DARM via radiation pressure.

The following subsections will overview the fundamental physics of the measurement, the details of the method, and the measurement results.

3.2.1 Fundamentals

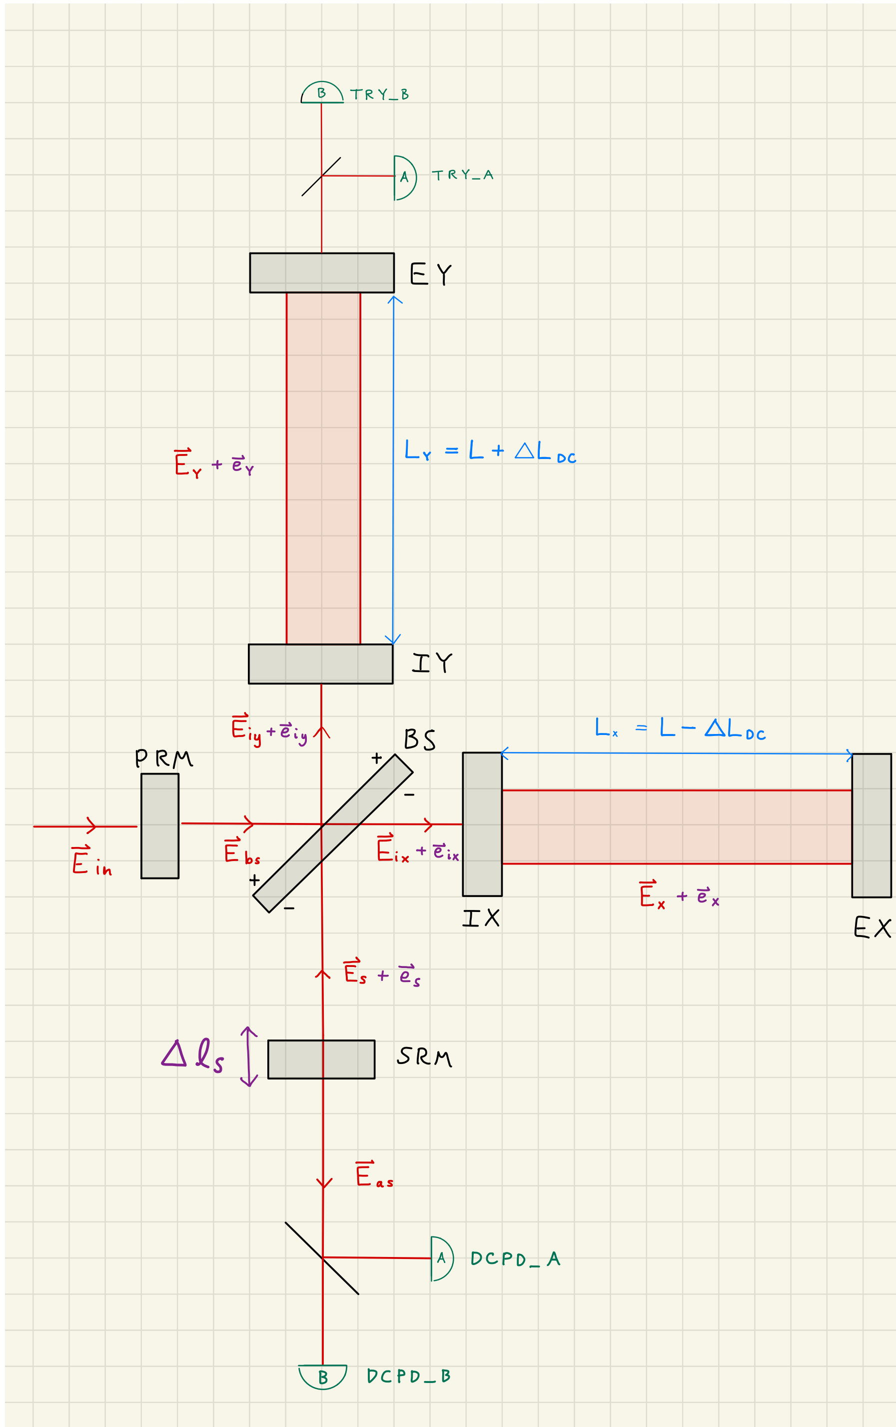

Figure 3.5: Diagram of the signal-recycling cavity length dither to arm power measurement. The SRCL length dither Δl_s modulates the light returning to the arms $\vec{E}_s + \vec{e}_s$, creating a differential arm power modulation due to the phase flip from the beamsplitter reflection.

Here we overview how the SRC length dither causes strong differential power fluctuations in the arms [101]. The measurement technique in subsection 3.2.2 depends on this fundamental physics. Figure 3.5 highlights the important aspects of the measurement.

First, we review the Advanced LIGO interferometer configuration in DC readout. The arms are offset from exact resonance by around $\Delta L_{DC} \approx 10$ pm. This leaks a small amount of carrier light out of the arms, out the antisymmetric port of the beamsplitter into the signal-recycling cavity. This DARM offset light is in the phase-quadrature (see Appendix B).

Some carrier light is reflected off the SRM back toward the beamsplitter. The main carrier light returning to the beamsplitter $\vec{E}_s$ is still in the phase quadrature.

Now, we introduce the SRCL audio sideband dither $\Delta l_s(\omega)$ which phase-modulates the carrier $\vec{E}_s$. This creates the audio sidebands $\vec{e}_s$ in the amplitude quadrature.

The modulated light from the SRM $\vec{E}_s + \vec{e}_s$ returns the beamsplitter, where half is transmitted and half is reflected with phase flip. The phase flip is key to this measurement, as it causes the differential arm power modulation: $\vec{e}_y = -\vec{e}_x$.

The SRC light combines with the input light $\vec{E}_{bs}$ which is also in the amplitude quadrature, and together enter and are enhanced by the arms. The arm power fluctuations are differential, $\delta P_y = -\delta P_x$, which causes a differential length change due to radiation pressure. Figure 3.6 shows the phasors of the light in the SRC and the arms.

The next subsection covers how the arm power can be inferred from this process.

3.2.2 Arm power inference technique

Using radiation pressure coupling to DARM, we can extract the power in the arms by dithering SRCL. This is an overview of the coupling mechanism, reported in [100] and inspired by [101].

3.2.2.1 Definitions

From [101], Equation (14) and (15) report the X-arm power response to a SRCL length dither:

$$\frac{P_x}{l_s}(f) = \frac{8 g_s^2 r_s r_a' \epsilon k}{t_s^2 (1 + s_{rse})} P_x \tag{3.2}$$

$$= \gamma P_x \tag{3.3}$$

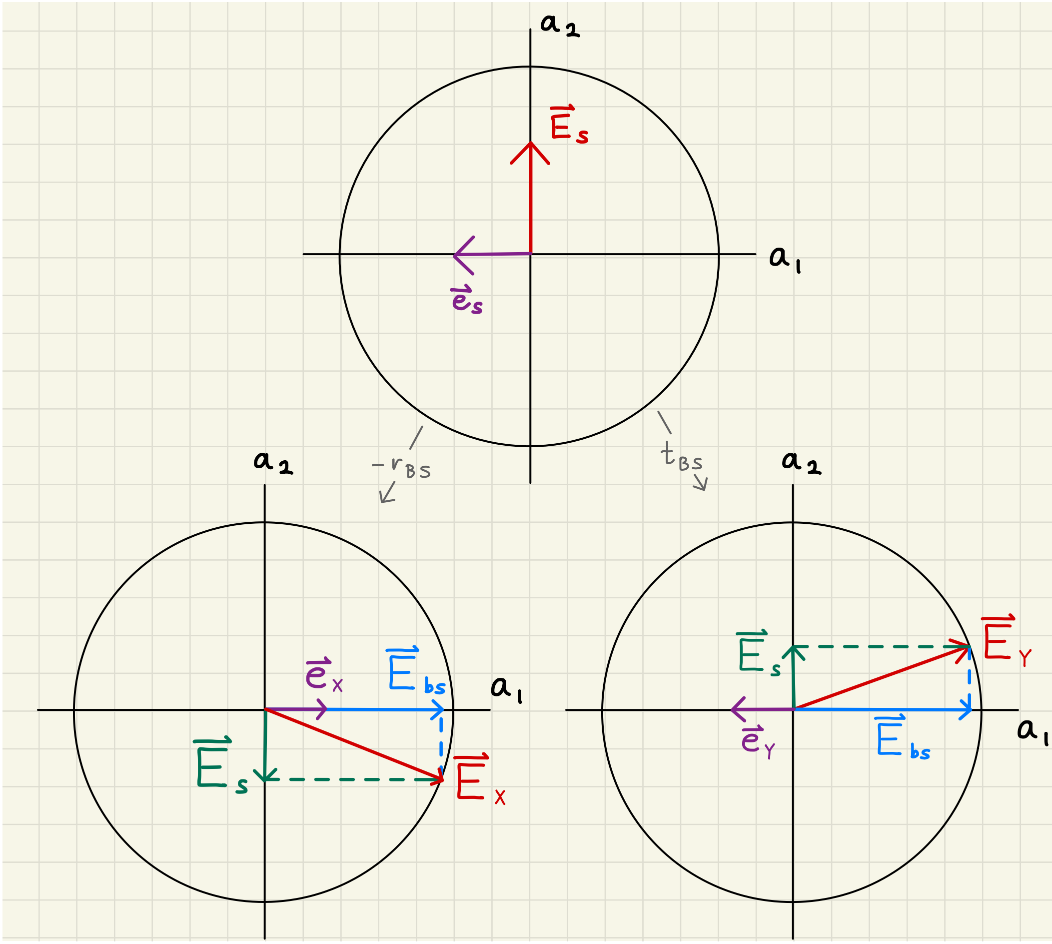

Figure 3.6: Phasors of the signal-recycling cavity length dither to arm power measurement for light in the SRC (top) and the arms (bottom). The top phasor shows the SRCL length dither Δl_s modulates the phase-quadrature carrier light $\vec{E}_s$, creating the audio sidebands $\vec{e}_s$ in the amplitude quadrature. The bottom left phasor is the light in the X-arm, and the bottom right phasor is the light in Y-arm. Due to the phase flip from the reflection off the back of the beamsplitter $-r_{bs}$, the same audio sidebands in each arm $\vec{e}_x$, $\vec{e}_y$ has a different sign. The sidebands modulate the main carrier in the amplitude quadrature $\vec{E}_{bs}$.

where P_x is the power in the X arm, l_s is the signal recycling cavity length, $f = \omega/(2\pi)$ is the audio signal frequency in Hz, g_s is the amplitude signal recycling cavity gain, r_s is the SRM amplitude reflectivity, r'_a is the arm reflectivity derivative with respect to phase, $\epsilon = kL_{\text{offset}}$ is the DARM offset phase in radians, k is the laser wavenumber, t_s is the SRM amplitude transmission, and $s_{rse} = i\omega/\omega_{rse}$ is the resonant signal extraction (rse) DARM cavity pole.

For this measurement, we gather the optical response of the SRCL dither into a factor γ which has units m^{-1}. The power in the Y-arm is the same except for an overall sign flip:

$$\frac{P_y}{l_s}(f) = -\frac{8g_s^2 r_s r_a' \epsilon k}{t_s^2(1+s_{rse})} P_y \tag{3.4}$$

$$= -\gamma P_y \tag{3.5}$$

Compliance of the SRM suspension:

$$\frac{l_s}{F_s}(f) = -\frac{1}{m\omega^2} \tag{3.6}$$

the compliance of the SRM, where F_s is the force applied to the SRM and m is the SRM mass.

SRCL control signal calibration:

$$\frac{F_s}{c_s}(f) = \beta \left[\frac{\mathrm{N}}{\mathrm{cts}}\right] \tag{3.7}$$

where c_s is the SRCL control signal in counts and β is some calibration constant in $[\mathrm{N}/cts]$.

Transmitted power:

$$\frac{P_{txa}}{P_x}(f) = T_{ex}\eta_{xa} \tag{3.8}$$

where P_{txa} is the transmitted power through the X arm falling on the TRX_A photodiode, T_{ex} is the power transmission through ETMX, and η_{xa} is the loss/response/calibration error of the TRX_A photodetector. Each of the four photodetectors (TRX_A, TRX_B, TRY_A, TRY_B) will have slightly different losses (η_{xa}, η_{xb}, η_{ya}, η_{yb}).

Finally, we have relative intensity:

$$RIN(f) = \frac{P(f)}{\overline{P}} \tag{3.9}$$

where $\overline{P}$ is the average power.

3.2.2.2 Constructing the SRCL control to transmitted arm power TFs

We measure the transfer functions from the SRCL control signal $c_s(f)$ to the end station transmitted power P_{txa}, P_{txb}:

$$\frac{P_{tx1}}{c_s}(f) = \frac{F_s}{c_s}(f) \times \frac{l_s}{F_s}(f) \times \frac{P_x}{l_s}(f) \times \frac{P_{tx1}}{P_x}(f) \tag{3.10}$$

$$= \beta \times -\frac{1}{m\omega^2} \times \gamma P_x \times T_{ex}\eta_{x1} \tag{3.11}$$

$$= -\frac{\beta\gamma T_{ex}\eta_{xa}}{m\omega^2} P_x \tag{3.12}$$

We may divide by the average power on the photodetector to get the RIN TF and eliminate dependence on the photodetector calibration, ETM transmission, and arm power:

$$\frac{RIN_{txa}}{c_s}(f) = \frac{\left(\frac{P_{txa}}{c_s}(f)\right)}{\overline{P_{txa}}} = -\frac{\beta\gamma}{m\omega^2} \tag{3.13}$$

This does not help for losses which are different between DC and AC signals e.g. photodetector saturation.

3.2.2.3 Radiation pressure coupling to DARM

From Eqs. 3.2 and 3.4 we see the arm power change from the SRCL dither is differential. The change in arm power will change the radiation pressure force, and arm lengths.

DARM:

$$L_{DARM} = L_x - L_y \tag{3.14}$$

where L_x and L_y are the X and Y arm lengths.

Radiation pressure force:

$$F_i(f) = \frac{2P(f)}{c} \tag{3.15}$$

where $F_i(f)$ is the force on a single optic, and $P(f)$ is the power in the arm.

Quadruple pendulum compliance (force to length):

$$L_i(f) = -\frac{F_i(f)}{M\omega^2} \tag{3.16}$$

where $L_i(f)$ is the displacement of a single optic, and M is the mass of the final stage of the quadruple pendulum.

Combining Eqs. 3.15 and 3.16, and multiplying by two for both the ETM and ITM:

$$L_x(f) = -\frac{4P_x(f)}{Mc\omega^2} \tag{3.17}$$

$$L_y(f) = -\frac{4P_y(f)}{Mc\omega^2} \tag{3.18}$$

Now we find the SRCL to DARM radiation pressure coupling, we combine Eqs. 3.2, 3.4, 3.17, and 3.18:

$$\frac{L_{DARM}}{l_s}(f) = -\frac{8\gamma P_{arm}}{Mc\omega^2} \tag{3.19}$$

where P_{arm} is the static average power in each arm.

Multiplying by the SRM compliance and control signal calibration Eqs. 3.6 and 3.7:

$$\frac{L_{DARM}}{c_s}(f) = -\frac{8\beta\gamma P_{arm}}{Mmc\omega^4} \tag{3.20}$$

3.2.2.4 Inferring the arm power

The arm power is inferred by taking the transfer function from the transmitted arm RIN to DARM while injecting a strong SRCL dither. Combining Eqs. 3.13 and 3.20:

$$\frac{\frac{L_{DARM}}{c_s}(f)}{\frac{RIN_{txa}}{c_s}(f)} = \frac{-\frac{8\beta\gamma P_{arm}}{Mmc\omega^4}}{-\frac{\beta\gamma}{m\omega^2}} \tag{3.21}$$

$$\frac{L_{DARM}}{RIN_{txa}}(f) = \frac{2P_{arm}}{Mc\pi^2 f^2} \tag{3.22}$$

$$\frac{L_{DARM}}{RIN_{txa}}(f) = \frac{\alpha}{f^2}. \tag{3.23}$$

where $\alpha = 2P_{arm}/Mc\pi^2$ is some constant fitting parameter. Solving for the static arm power P_{arm}:

$$P_{arm} = \frac{1}{2}\alpha\pi^2 Mc \tag{3.24}$$

This is the most precise arm power measurement yet devised. Most complexities divide out of this measurement. The uncertainties on input measurement and losses from Eq. 3.1 are avoided.

The uncertainty in the arm power P_{arm} depends entirely on the uncertainty in α, plus the systematics of the measurement itself like biasing from saturated photodetectors, or slow drift of the arm power itself as the interferometer thermalizes during the measurement. The mass of the quadruple pendulum $M = 40$ kg is known to very good precision. This technique also takes advantage of the highly accurate DARM calibration, which is good to $\sim 2\%$.

3.2.3 Measurement Details

Here we list some practical considerations for the arm power measurement.

1. Turn off the SRCL to DARM feedforward before measuring. Feedforward is designed to cancel exactly this kind of radiation pressure signal.
2. Make sure the beams are well-aligned on the transmission quadrant photodetectors. The transmission platforms are known to drift significantly, saturating one of the photodetector quadrants. When a photodetector quadrant is saturated, the AC response is suppressed, causing an overestimate of the arm power. A good indicator of this issue is when the SRCL line shows up strongly in DARM, but you struggle to maintain coherence on the transfer function.
3. Take the measurement a least an hour after the interferometer has locked at full power. The arm power tends to slowly drift after locking from thermalization and spot position changes.
4. Measure in the frequency region where the SRCL to DARM coupling is radiation-pressure dominated, but high enough such that the free mass pendula compliance approximation is valid. Between 10 to 100 Hz is usually sufficient.
5. Avoid measuring on the calibration lines, the alignment-dither lines, and the mains 60 Hz line.
6. When fitting the arm RIN to DARM transfer function, first flatten the TF by multiplying by f^2. This will improve uncertainty and avoid biasing. Typically, measurement systematics will far outweigh fit uncertainty.

3.2.4 Results

Figure 3.7 plots the arm RIN to DARM transfer functions from Eq. 3.21. The transfer function is "flattened" by multiplying by f^2, then fit to the high-coherence points in the radiation pressure regime. Table 3.3 reports the measured arm powers during O3. Measurements derived from signal-recycling cavity length modulation are consistent and more precise compared with measurements inferred from the input power and test mass reflectivity.

The arm powers at Hanford are significantly lopsided: there is a measured 6.7% difference in the arm powers. This was because the two input test masses installed

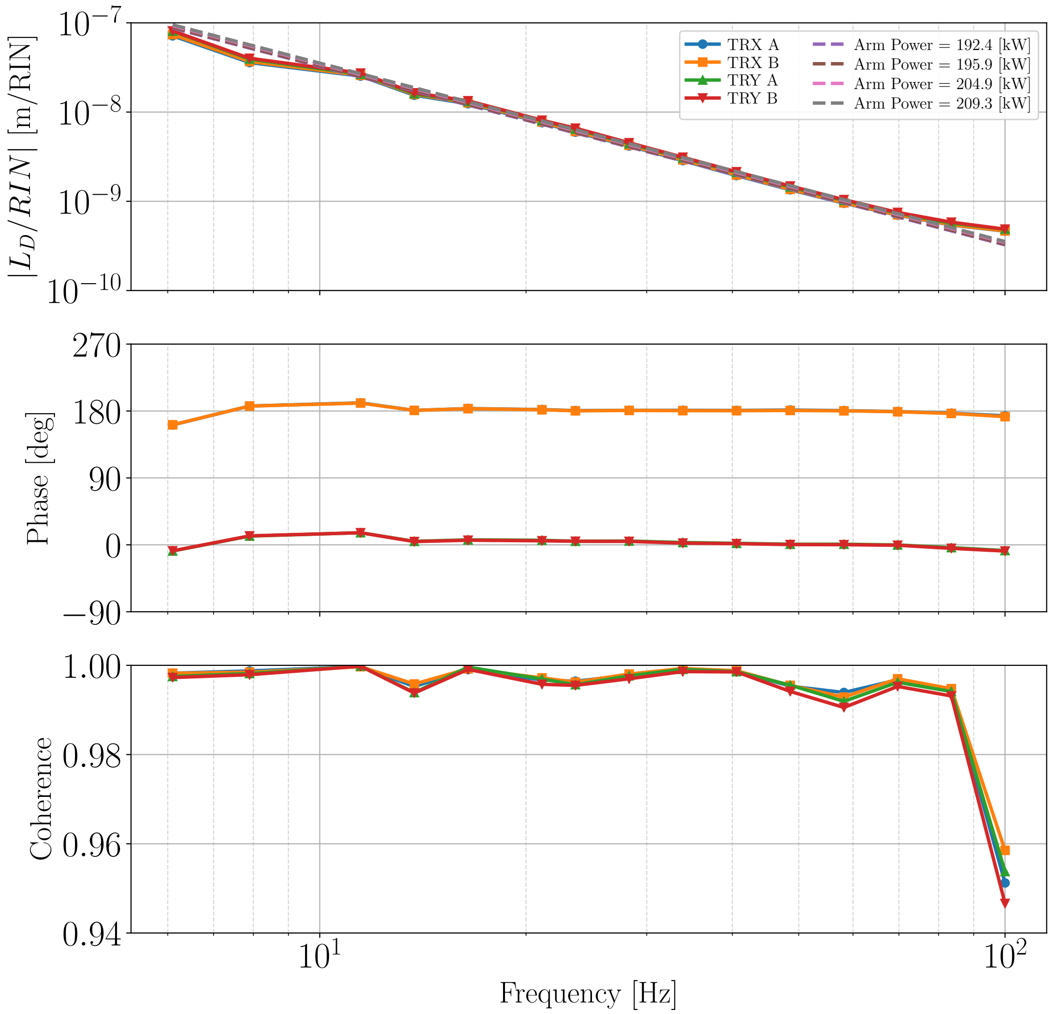

Figure 3.7: Arm power measurements at Hanford via the arm RIN to DARM transfer functions (Eq. 3.21). The magnitude falls like $1/f^2$, as predicted for a radiation pressure regime. Table 3.3 lists the results from this measurement.

at Hanford were not "twins", i.e. they were not made at the same time. ITMX was replaced because of a very large point absorber detected during O2. The new ITMX has a transmission of 1.5%, whereas ITMY has a transmission of 1.42%. This yields an X-arm finesse $\mathcal{F}_x = 416$ and a Y-arm finesse $\mathcal{F}_y = 439$, or an 5.6% difference. This difference in ITM transmission also affected the frequency noise coupling to DARM, see Section 3.4.5.

Power	Symbol	LHO	LLO	Units
Input	P_0	34 ± 2	38 ± 2	W
Power-Recycling Gain	g_p^2	44 ± 3	47 ± 3	W/W
X-arm via Eq. 3.1	P_x	190 ± 14	240 ± 18	kW
X-arm via Eq. 3.24	P_x	194 ± 2	232 ± 15	kW
Y-arm via Eq. 3.1	P_y	200 ± 15	240 ± 18	kW
Y-arm via Eq. 3.24	P_y	207 ± 2	245 ± 5	kW

Table 3.3: Highest measured laser power levels during O3. Input power is estimated via a pick-off from the light incident on the power-recycling mirror. Power-recycling gain is estimated from the pick-off of the power-recycling cavity, using a ratio of power on the beamsplitter and input power. Arm powers are estimated in two ways. The first method is via input power and gain estimates, Eq. 3.1. Arm power uncertainties for Eq. 3.1 are propagated from uncertainty in the input power, power-recycling gain, and loss in the arms. The second method is via radiation-pressure relative intensity noise to DARM transfer function, Eq. 3.24. Arm power uncertainties for Eq. 3.24 are derived from the coherence of the measured transfer function. Typical arm power levels at LLO were about 5% lower over the course of the run.

3.3 Auxiliary length control improvements

The auxiliary length sensing and control (ALS) is integral to the rapid locking of Advanced LIGO [40, 80]. ALS employs two green lasers at each end station, injected onto the back of the ETMs.

Originally, the locking scheme for Advanced LIGO was supposed to hand off laser frequency control from ALS COMM directly to CARM. However, the noise of the as-built ALS system was far larger than requirements largely due to an error in the ETM coatings green transmission. This prevented direct handoff from ALS COMM to CARM, since the ALS COMM noise far exceeded the CARM linewidth of ~ 1 Hz.

The ETM green coatings were fixed on replacement ETMs installed prior to O3. This section overviews the ALS control scheme and the method used to check the performance of the ALS COMM system.

3.3.1 ALS control scheme

The auxiliary length sensing and control (ALS) is integral to the rapid locking of Advanced LIGO [40, 80]. ALS employs two green lasers at each end station, injected onto the back of the ETMs. Each laser is phase-locked to the main pre-stabilized laser (PSL) via two pick-off fibers running from the PSL to each end-

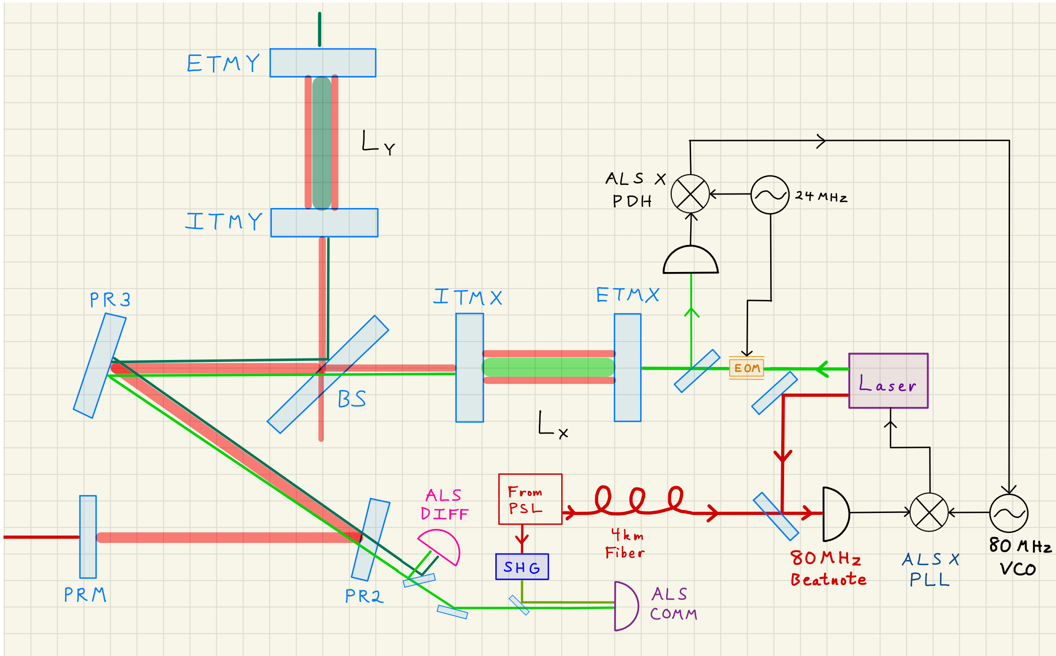

Figure 3.8: Simplified auxiliary length sensing control scheme. The X-arm ALS controls are highlighted here, but the Y-arm employs the same control scheme. The dual-color end-station laser on the right emits both 1064 nm (infrared) and 532 nm (green) light. The green light is locked to the arm cavity, and the transmitted beams are mixed together. The mixed beams beatnotes are sensed and used to control the PSL frequency (ALS COMM) and differential arm length (ALS DIFF). The ALS COMM and DIFF phase-locked loops and VCOs for control are not shown here.

station. Then, each green laser is PDH-locked to its arm cavity.

The green beams transmitted through both arms are routed to an in-air optical table. There, the beams from each arm are combined to form the ALS DIFF signal. The green beam from the X-arm is combined with green PSL light to form the ALS COMM signal.

The ALS COMM signal detects frequency fluctuations between the PSL and the green X-arm. The green X-arm PDH signal follows the arm length, while the PSL is stabilized to the IMC at this point. We lock the PSL to the ALS COMM error signal so the PSL frequency follows the green X-arm length. ALS COMM has a unity gain frequency of around 650 Hz [102].

Then a large ~ 1 kHz offset is placed on the PSL light, so the infrared does *not* yet resonate in the arms. This is done so the dual-recycled Michelson can be locked without "arm flashes", i.e. moments of spurious resonance in the arms which spoil the PRCL, SRCL, and MICH error signals.

The ALS DIFF signal yields information on the differential arm length. This is fed back to one of the ETM suspensions, actuating on the differential arm length itself.

3.3.2 ALS upgrades for O3

Prior O3, the ETMs were replaced to lower scatter losses for infrared. Another improvement of the new ETMs a dramatically increased green reflectivity.

The old ETMs at Hanford had green transmission of 37% for ETMX and 32% for ETMY [103]. The new ETMs both have green transmission of 7.9%. ITMX was also replaced at Hanford, its green transmission went from 1.1% to 0.96%.

The new core optics affected the ALS cavity parameters: the green cavity finesse increased from 15 to 70, green arm poles decreased from $1.3\ \mathrm{kHz}$ to $280\ \mathrm{Hz}$.

The core optics coatings upgrades make the green PDH lock more stable, particularly avoiding mode hopping. They also improve the noise performance of the ALS controls by improving the green PDH noise.

Additionally, major efforts to improve the acoustic noise coupling in the PSL have been undertaken. The benefits of this work can be seen in the lowered peaks apparent in Figure 3.10.

3.3.3 ALS COMM frequency noise measurement

To measure the improvements to the green ALS system, the infrared (IR) light reflected off the X-arm is used as an out-of-loop witness of ALS COMM frequency noise [104–106]. We take advantage of the fact the IR light is stabilized to the X-arm length via the green PDH, the arm cavity is a high-quality cavity, and the RF sidebands and RFPD are already set up for low-noise frequency discrimination.

The procedure for the ALS COMM frequency noise measurement is as follows:

1. Lock the IMC and X-end phase-locked loop.
2. Misalign the Y-arm, beamsplitter, and PRM.
3. Lock the X-arm green PDH signal.
4. Align the cavity for maximum green circulating power.
5. Separately lock the IR light to the X-arm.
6. Align the input optics for maximum IR in the X-arm.

7. Break the IR lock.

8. Lock ALS COMM.

9. Use ALS COMM phase-locked loop offset to find an IR X-arm peak.

10. Sweep the COMM offset over the fringe to calibrate your out-of-loop IR RFPD response.

11. Return to the IR fringe and take a spectrum of the out-of-loop IR RFPD.

The PSL frequency offset from IR resonance in the X-arm is set to be off-resonance by $\sim$ 300 Hz, then moved slowly over the resonance peak. The out-of-loop sensor used in this experiment is REFL A 9I. The calibration procedure relies on our knowledge of PDH error signals. We write the PDH error signal for a simple Fabry-Perot cavity, V_I,

$$V(f) = 2\eta G_{pd}(E_\omega^* E_\Omega - E_0 E_{-\Omega-\omega}^*) \tag{3.25}$$

where η is all the optical losses, and G_{pd} is the gain of the RF photodetector in V/W. We write the reflection of the carrier E_0, E_ω and 9 MHz sidebands $E_\Omega, E_{-\Omega-\omega}$ like

$$E_0 = r_0 E_{\text{in}}, \quad E_\omega = \left(r_0 - \frac{if/f_z}{1+if/f_p}\right) E_{\text{in}}, \quad E_\Omega = ir_\Omega \frac{\Gamma}{2} E_{\text{in}}, \quad E_{-\Omega-\omega} = ir_{-\Omega}\frac{\Gamma}{2} E_{\text{in}} \tag{3.26}$$

where f is the incident carrier laser frequency, f_p is the arm pole, f_z is the arm zero E_{in} is the input field, and $r_0, r_\Omega, r_{-\Omega}$ are the arm reflectivities for the carrier and sidebands. Since carrier is on resonance but the sidebands are off-resonance, we can write these as

$$r_0 = \frac{-r_i + r_e}{1 - r_i r_e}, \qquad r_\Omega = r_{-\Omega} = \frac{-r_i - r_e}{1 + r_i r_e} \tag{3.27}$$

where r_i is the ITMX reflectivity and r_e is the ETMX reflectivity.

The audio reflection zero f_z is merely a scale factor for the zero at DC. It scales the reflected carrier light offset from resonance by a frequency such that the reflected light beyond the cavity pole is near one. The arm pole f_p is derived in Appendix B Section B.4:

$$f_z = \frac{\text{FSR}}{2\pi r_0'} = \frac{c(1-r_i r_e)^2}{4\pi L(1-r_i^2) r_e}, \quad f_p = \frac{\text{FSR}}{2\pi}\log\left(\frac{1}{r_i r_e}\right) \tag{3.28}$$

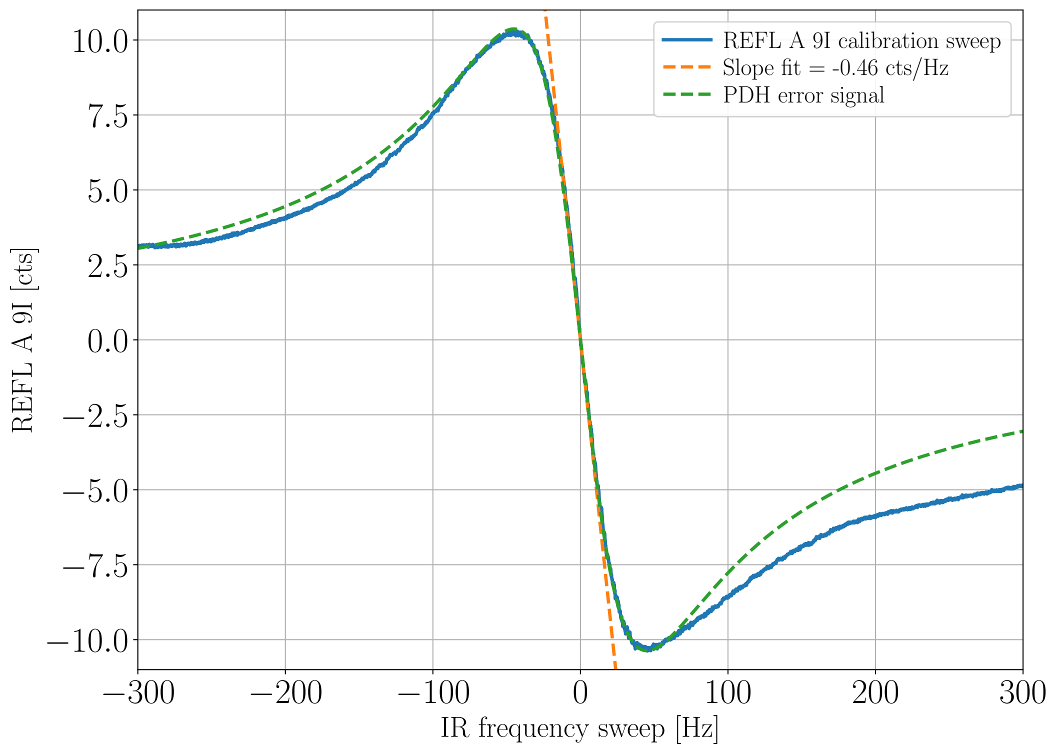

Figure 3.9: Pound-Drever-Hall infrared error signal sweep of the REFL A 9I photodetector while ALS COMM is locked. The slight deviation of the measurement from the model is likely due to the drift of the arm cavity during the sweep time. The model of the PDH is Eq. 3.30, and the slope is from Eq. 3.31.

The signal from Eq. 3.25 is in the I-phase. We substitute Eqs. 3.26 into Eq. 3.25 and take the real part to the RFPD I-phase response:

$$V(f) = -\eta G_{pd} r_\Omega \Gamma P_{\text{in}} \frac{2ir_0 f_p f_z + f(f_p - 2r_0 f_z)}{(f_p - if) f_z} \tag{3.29}$$

$$V_I(f) = \Re(V(f)) = -\eta G_{pd} r_\Omega \Gamma P_{\text{in}} \frac{f}{1 + f^2/f_p^2}. \tag{3.30}$$

For small $f \ll f_p$, i.e. close to resonance, the slope of $V_I(f)$ defines the PDH discriminant in V/Hz:

$$V_I(f) \approx \frac{-\eta G_{pd} r_\Omega \Gamma P_{\text{in}}}{f_z} \tag{3.31}$$

Figure 3.9 shows the measured REFL A 9I response to our frequency sweep, alongside the model PDH signal Eq. 3.30 and PDH discriminant Eq. 3.31.

To calibrate the error signal to the arm pole, the peak-to-peak response to the sweep is measured. Eq. 3.30 gives a peak response at $f = \pm f_p$, with value $V_{I,pp} =$

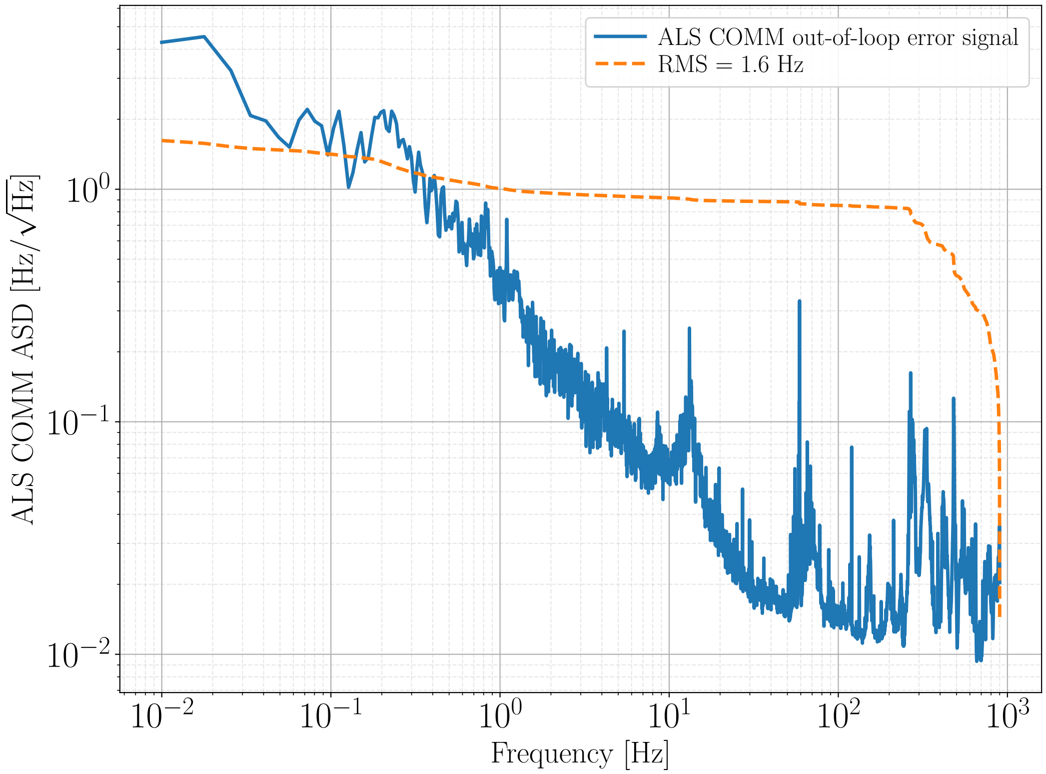

Figure 3.10: ALS COMM out-of-loop error signal, plus cumulative RMS between 0.01 and 1000 Hz. The out-of-loop sensor used is REFL A 9I, with 10 W requested input power P_{in}. The resulting RMS of 1.5 Hz is around an order of magnitude better than in O1 (Figure 4 of [40]).

$\eta G_{pd} r_\Omega \Gamma P_{\text{in}} f_p / f_z$. This gathers all the complexity of the calibration into this peak to peak value. The X-arm IR cavity pole is known to be 45.1 Hz, so the discriminant Eq. 3.31 gives our calibration to the arm pole frequency.

Figure 3.10 shows the results of the ALS COMM frequency noise measurement, with the X-arm IR pole f_p undone and the X-arm calibration applied. ALS COMM is susceptible to many noises, including laser frequency noise, green PDH sensing noise, fiber noise, in-air optics table phase-wrapping noise, acoustic peaks from optics table, and PLL controls noise. Also, the out-of-loop IR sensor has sensing noise as well, which we suppressed by turning up the input power P_in to 10 W to increase signal on REFL A 9I.

The final measured RMS of the ALS COMM frequency noise is 1.5 Hz, about 10 times lower than the measurement performed before O1 (Figure 4 of [40]). This is a result of the enhanced green cavity finesse, and improvements to acoustic coupling on the PSL optics table. With some further improvements, the ALS COMM noise may be reduced to the point where direct CARM handoff is possible, although this is not the rate-limiting step of the locking process.

3.4 Frequency stabilization

Control of the laser frequency is crucial to the optimal performance of the Advanced LIGO interferometer. The laser frequency is just one aspect of the interferometric length sensing and control, and is degenerate with the common-arm (CARM) length via the relation

$$\frac{\Delta L}{L} = -\frac{\delta \nu}{\nu_0} \tag{3.32}$$

where ν_0 is the carrier frequency, $\delta\nu$ is the frequency noise, and $L = (L_x + L_y)/2$ is the common-arm length.

An interferometer enjoys a natural frequency noise mitigation in its *common mode rejection*: motion that is common to both arms tends to be reflected back toward the input laser, whereas motion that is differential preferentially transmitted thought to the antisymmetric port. Changes in the laser frequency must be common to both arms, so to first order most frequency noise is promptly reflected by the interferometer and does not mask gravitational wave signal.

However, to second order imperfections cause *contrast defect* light carrying frequency noise to appear at the antisymmetric port. Worse, frequency noise on contrast defect light appears in the phase quadrature in the antisymmetric port as the gravitational wave signal (see Section B.3.3).

Another natural factor aiding the suppression of frequency noise is the exceptionally low linewidth of the interferometer cavities. The interferometer can be thought of as a highly selective phase-sensitive sieve which only accepts light of the correct frequency. Unacceptable light is promptly reflected and does not make it into the interferometer, making the laser light inside the interferometer incredibly stable: the main laser frequency $\nu_0 \approx 281 \times 10^{12}$ Hz must be stabilized to the interferometer linewidth of around 1 Hz.

The laser frequency stabilization scheme in Advanced LIGO serves two purposes. First, to incrementally stabilize the interferometer length degrees of freedom and achieve resonance in an interferometer with an extremely low linewidth. Second, to suppress frequency noise so that it does not limit sensitivity to gravitational waves. The requirements for stability of the frequency is around 8 orders of magnitude lower than the freerunning NPRO in the detection band (NPRO freerunning noise is $100\ \mathrm{Hz}/\sqrt{\mathrm{Hz}}$ at $100\ \mathrm{Hz}$, requirement is $\sim 10 \times 10^{-6}\ \mathrm{Hz}/\sqrt{\mathrm{Hz}}$).

In this section, we overview the performance of the Advanced LIGO frequency stabilization servo at Hanford for O3, with focus on the frequency noise budget and the frequency noise coupling to DARM.

3.4.1 Control scheme

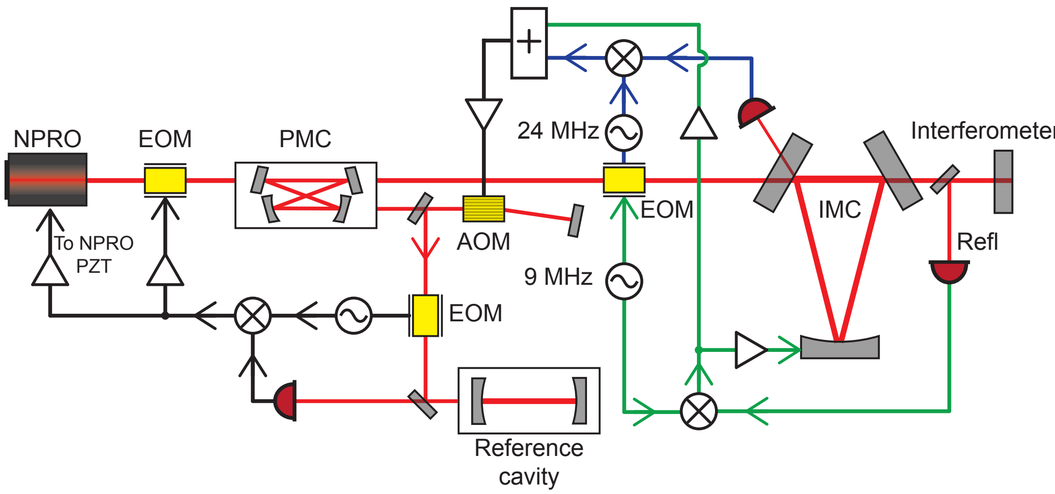

Figure 3.11: Laser frequency stabilization optical diagram. Laser beams are shown in red, electronics are shown in black, blue and green. Reference cavity electronics are shown in black, input mode cleaner electronics are shown in blue, and common mode electronics are shown in green.

There are three hierarchal control loops in the frequency stabilization servo. The first is the reference cavity loop, the second is the input mode cleaner (IMC) loop, and the third to the common-arm (CARM) loop. Because the loop bandwidth required is higher than the data acquisition rate of 16384 Hz, the frequency control loops are mostly analog, with the exception of the slow feedback to the MC2 suspension. Figure 3.11 shows the optical layout of the Advanced LIGO frequency stabilization servo.

3.4.1.1 Pre-stabilized laser and reference cavity

A non-planar ring oscillator (NPRO) solid-state Nd:YAG laser with wavelength $\lambda = 1064\ \mathrm{nm}$ is used to generate the initial seed light due to its exceptional freerunning frequency stability ($(1\ \mathrm{kHz}/f)(10\ \mathrm{Hz}/\sqrt{\mathrm{Hz}})$) with high power output [107]. The NPRO laser is amplified to 80 watts, then locked to a four-mirror pre-mode cleaner (PMC) [36]. In transmission of one of the highly-reflective PMC mirrors, a beam is double passed through an acousto-optic modulator (AOM). The beam is then locked to a fixed-spacer reference cavity with high bandwidth (~ 300 kHz).

Next the beam is phase modulated by an electro-optic modulator (EOM) at 9.1 MHz, 24.1 MHz, 45.5 MHz, and 118.3 MHz [108]. The 9, 45, and 118 MHz sidebands are to be used for main interferometer control, and are the 1st, 5th, and 13th harmonics of each other. The 24 MHz sidebands are used for locking to the input mode cleaner.

3.4.1.2 Input mode cleaner

The beam is then locked to the suspended 16 m input mode cleaner (IMC) with bandwidth ~ 80 kHz [37]. The IMC cleans the main beam of higher-order modes and beam jitter from the main laser, and stabilizes the laser frequency. The free spectral range (FSR) of the IMC is designed to be 9.1 MHz in order to pass the three RF sidebands for main interferometer control. The feedback from the IMC goes to the IMC voltage-controlled oscillator (IMC VCO) which controls the double-passed AOM before the reference cavity.

3.4.1.3 Common-arm cavity

Finally, the beam proceeds to the main interferometer. The power recycling cavity and the arm cavities together form the CARM coupled cavity. The laser is stabilized to CARM with a bandwidth of 20 kHz. The CARM bandwidth cannot be increased past the FSR of the coupled cavity $\mathrm{FSR} = c/(2L) \approx 37$ kHz because of the dynamics of the optical plant: the reflection of the carrier goes through a second resonance and loses 180 degrees of phase, making higher-bandwidth control impossible.

The carrier and 9 MHz sidebands beatnote are used for sensing the CARM degree of freedom [33]. The carrier enters both the power recycling cavity and the arm cavities, returning information on the length of both. The 9 MHz resonates in the power recycling cavity, but not the arm cavities, returning information only on the power recycling cavity length. The beatnote between the carrier and 9 MHz carries the arm length information, which is detected on the REFL photodetectors in reflection of the interferometer.

Figure 3.12 shows the open loop gain measurements of the three hierarchal frequency stabilization loops. The three loops together suppress the NPRO laser frequency noise down to the CARM shot noise level, as characterized in Section 3.4.4.

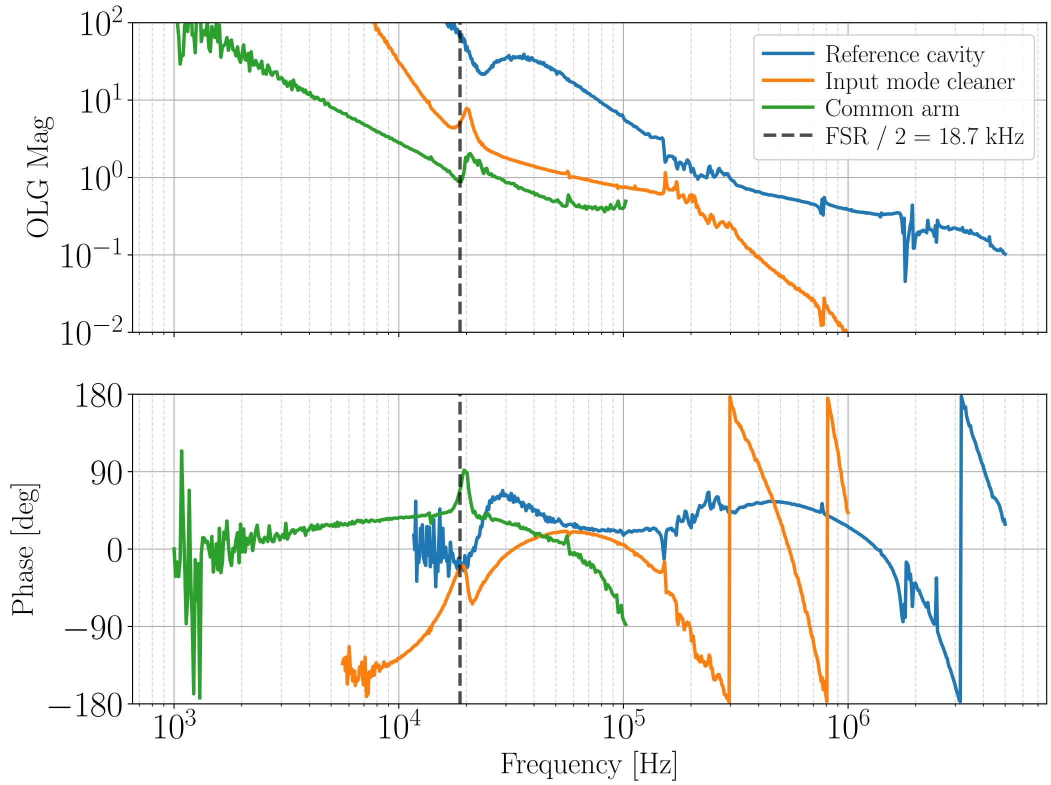

Figure 3.12: Frequency control loop open loop gain measurements. The half free spectral range (FSR) is highlighted because there we see the 9 MHz sidebands plus the audio signal resonate in the arms, affecting the common-arm optical gain around 19 kHz.

The 9 MHz sidebands cause a small dip at around half the free spectral range (FSR). This is because the 9 MHz sidebands are purposefully set to be anti-resonant in the arms, while the carrier frequency is resonant. This places the 9 MHz sidebands nearest resonance in the arm around half of an FSR away from the carrier frequency. If we modulate the main laser frequency at an audio frequency $f \approx \mathrm{FSR}/2$ Hz, then the audio sidebands on the 9 MHz sidebands will resonate in the arms, producing a small effect on the CARM optical gain seen in Figure 3.12.

This can be more easily seen if we think about the Pound-Drever-Hall error signal. The power reflected off the CARM coupled cavity P_{refl}, demodulated at radio-frequency $\Omega/2\pi$ can be written

$$P_{refl,\Omega}(\omega) = 2[-E_0 E_\Omega^*(\omega) + E_\Omega E_0^*(\omega)] \tag{3.33}$$

where E_0 is carrier light reflected off the CARM coupled cavity, E_Ω is reflected RF sideband light, in this case $\Omega/2\pi \approx 9_{\mathrm{MHz}}$, $E_0(\omega)$ is the audio sidebands on the carrier, and $E_\Omega(\omega)$ is the audio sidebands on the RF sidebands.

Normally, when shaking the laser frequency in the audio band $\omega \ll \mathrm{FSR}$, $E_\Omega(\omega) \approx$ const, so the term at $-E_0 E_\Omega^*(\omega)$ is entirely at DC while the error signal at ω is dominated by changes to the carrier $E_\Omega E_0^*(\omega)$. However, if $\omega \approx \mathrm{FSR}/2$, the terms at $\Omega \pm \omega$ from $E_\Omega(\omega)$ start to resonate in the arms, and $-E_0 E_\Omega^*(\omega)$ becomes significant enough to appear in the measured CARM OLG.

3.4.1.4 Mode cleaner length control

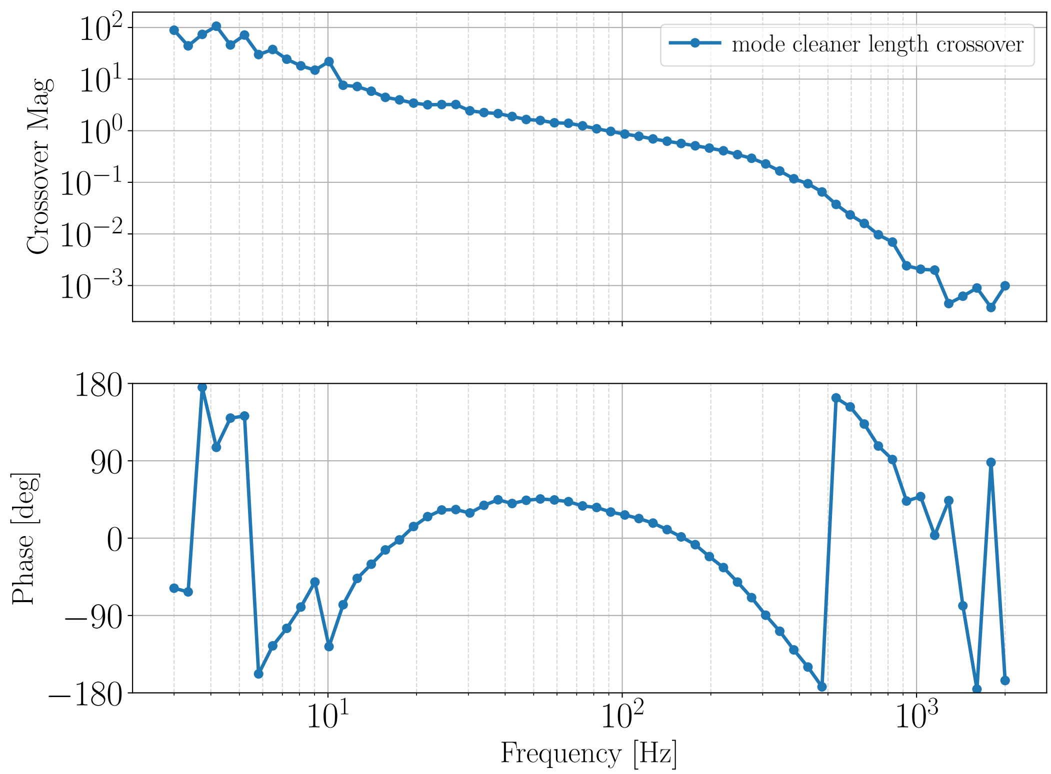

Figure 3.13: Input mode cleaner length to VCO controls crossover measurement. The unity gain frequency at 90 Hz corresponds to the handoff from MCL to VCO controls

The CARM length is the ultimate reference we would like the laser frequency to follow. However, the IMC lies in the path of the laser to clean the beam and stabilize the laser frequency. The beam must be transmitted through the IMC with high efficiency, but to achieve its main tasks the IMC must have a high finesse ($\mathcal{F} \approx$

528). The IMC pole $f_{pole} \approx 8.6$ kHz, which is not wide enough to accommodate all expected CARM length drifts due to e.g. tides.

The solution is to feed the CARM error signal not only to the VCO, but also to the mode cleaner length (MCL). The MCL loop handles the low frequency CARM feedback by adjusting the mode cleaner length to match the CARM length. Because the compliance of the suspension falls off like $1/f^2$, the MCL loop cannot handle high frequency stabilization, so the VCO is relied upon for fast frequency control. The MCL loop is sometimes called the *slow control* and the VCO *fast control.* Figure 3.13 shows the *crossover* measurement of the slow controls over fast controls hierarchy. Here the unity gain frequency of the crossover at about 90 Hz corresponds to the handoff from MCL to VCO controls.

3.4.2 O3 frequency control upgrades

A few upgrades were performed on the frequency control scheme described above for O3. The first was the addition of the REFL B photodetector during the November 2018 vent. The second was the increase of incident power on the IMC REFL photodetector in May 2019 [109].

3.4.2.1 REFL B photodetector

The REFL B photodetector was added in the reflection path from the main interferometer, on the same path as the usual frequency sensor REFL A. REFL B was added because it was suspected that the slewing of the radio-frequency voltage was too fast for the REFL A photodetector electronics, leading to nonlinear response and so-called "fast locklosses". Fast locklosses were when the interferometer would lose lock without an apparent reason in the digital readback signals, prompting suspicion of the analog frequency sensors and controls of malfunction.

To resolve the REFL A fast-slewing problem, the power on REFL A was halved from ~ 10 mW to ~ 5 mW, and the remaining power was directed to a second, identical photodetector REFL B. This halved the slew of the REFL A and B op-amps. Fast locklosses persisted after this change.

But the addition of the REFL B photodetector provides and out-of-loop incident frequency noise sensor. As seen in the frequency noise budget in Figure 3.17, this sensor confirms that frequency noise is CARM shot noise limited in the GW-sensitive band. While locked in low-noise, both REFL A and B are used in-loop to lower the shot noise limit.

3.4.2.2 IMC REFL incident power increase

The IMC REFL photodetector is the sensor used to detect the frequency in reflection of the IMC. Because IMC sensing noise was found to be limiting frequency noise in the ~kHz region, the IMC REFL optical path was reworked to allow more light on the photodiode.

IMC REFL was found to be dark noise-limited, not shot noise-limited, even during high-power operation. "Dark noise" refers to the natural electronics thermal noise, e.g. Johnson noise, which causes voltage noise in otherwise quiescent electronics.

Ideally, the fundamental limit of the sensitivity achievable for an optical cavity is the shot noise limit, where all light on the photodetector contributes to the signal as RF local oscillator for PDH locking, or as the reflected carrier signal from the actual cavity length changes.

In reality, cavity visibility is not perfectly one, and "junk light" carrying no signal, such as higher-order modes, are reflected onto the photodetector. In some cases, the junk light reflected is too high and saturates the photodiode. To solve this, black glass power dumps are placed before the photodiode, and some optical signal is sacrificed for a functional photodiode.

In the case of IMC REFL, only ~ 1.4 mW was reaching the photodetector at full 35 W input power. This level was increased to ~ 9 mW, allowing more signal on IMC REFL and moving from the dark noise-limited regime to the shot noise-limited regime. Coupled with increased CARM loop gain, this reduced IMC sensing-induced frequency noise incident on the interferometer to below CARM sensing noise levels for the entire bandwidth (see Figure 3.17).

3.4.3 CARM calibration

CARM is calibrated to the IMC VCO which controls the double-passed AOM. The VCO control signal is calibrated into units of Hz, which serves as the ultimate reference for the laser frequency [110].

The sensors for CARM are the photodetectors in reflection of the interferometer, REFL A and B. These are RF photodetectors demodulated at 9 MHz to sense the CARM length changes and feed back to the laser frequency.

The usual way we calibrate a control loop is to inject some known quantity, such as a known frequency change in Hz with a VCO, and measure the response in the sensor we care about, such as the volts on the REFL PDs. However, the calibration of CARM is not so straightforward.

This is because of the massive suppression and the hierarchal nature of the full CARM loop. The VCO control signal is made from the sum of the CARM and IMC error signals. The IMC error signal dominates the control signal, because IMC shot noise is much higher than CARM shot noise. The CARM loop gain dominates above 20 kHz, suppressing IMC shot noise down to the level of CARM shot noise level. To do this, the CARM loop inverts the sign of the IMC shot noise it sees and injects that as the VCO control signal.

In short, the VCO control signal is totally dominated by IMC sensing noise, so an injection into the CARM loop would have to be extremely loud to produce an appreciable CARM signal in the VCO control signal. Direct length injections into CARM by, for example, the photon calibrator, make real common length changes in the arms, but do not appear in the REFL error signal because of the huge CARM loop suppression.

I calibrated CARM using the CARM OLG combined with an IMC OLG taken without the CARM feedback. Figure 3.14 shows the full frequency stabilization servo block diagram. We would like to calibrate the CARM plant C in W/Hz.

First, for readability of the following equations I'll make the following consolidations of the block diagram:

$$PCA \to C \tag{3.34}$$

$$IK \to I \tag{3.35}$$

$$HV \to V \tag{3.36}$$

$$SM \to M \tag{3.37}$$

Second, we assume that the reference cavity innermost loop perfectly follows the VCO control signal. In other words, the reference cavity loop bandwidth $\gg 1$ for all frequencies we care about (see Figure 3.12)

The CARM OLG G_{carm} is taken by injecting an excitation x_c at the CARM error signal e_c. Assuming x_c is strong enough to drown out all noise in the loop, we calculate the CARM OLG using Figure 3.14

$$G_{carm} = \frac{CFV + CMIV}{1 - IV}. \tag{3.38}$$

Symbol	Description	Units	ZPK Values
P	IMC transmission pole	Hz/Hz	[], [$8.6e10^3$], 1.0
C	CARM plant	W/Hz	[], [0.6], 3.4×10^{-3}
A	REFL PD sensing chain	V/W	-
	InGaAs photodiode responsivity	A/W	[], [], 0.77 [111]
	Transimpedance	V/A	[], [], 449 [112]
	Demod gain	V/V	[], [], 5.4
F	Fast servo analog board	V/V	-
	Sum node gain	V/V	[], [], $+8$ dB
	REFL IN1 gain	V/V	[], [], $+12$ dB
	REFL Boost 1	V/V	[500], [10], 50
	REFL Compensation	V/V	[4000], [40], 100
	REFL Fast high pass filters	V/V	[$0, 0$], [$5, 5$], $1/25$
	REFL FAST gain	V/V	[], [], $+16$ dB
	IMC IN2 gain	V/V	[], [], -22 dB
I	IMC plant	W/Hz	[], [$8.6e3$], 4.6×10^{-8}
K	IMC REFL PD sensing chain	V/W	-
	InGaAs photodiode responsivity	A/W	[], [], 0.77 [111]
	Transimpedance	V/A	[], [], 378 [112]
	Demod gain	V/V	[], [], 5.4
	IMC IN1 gain	V/V	[], [], $+2$ dB
H	Common servo analog board	V/V	-
	IMC Boost 1	V/V	[$20e3$], [$1e3$], 20
	IMC Boost 2	V/V	[$20e3$], [$1e3$], 20
	IMC Compensation	V/V	[$4e3$], [$4e1$], 100
	IMC FAST gain	V/V	[], [], -18 dB
	IMC Fast daughter board	V/V	[$70e3$], [$140e3, 200e3$], 2.3 [113]
V	Voltage-controlled oscillator	Hz/V	[40], [1.6], $537e3$

Table 3.4: CARM model values for calibrating the CARM path gain via Eq. 3.43. Values are those typical for LIGO Hanford during O3, locked in low noise with input power on the PRM $P_{\text{in}} = 34$ W. The CARM and IMC plants C and I include both the intrinsic optical gain of the cavity and all optical losses, including beam dumps. The responsivity of the InGaAs photodiodes includes the quantum efficiency of ~ 0.9. The measured transimpedance of REFL B was 448 V/A [112]. The calibration of the photodetector signal chains is described in Section C.1. All poles and zeros are listed in Hz. Figures 3.15 and 3.16 shows the CARM path model and IMC path model compared to measurements.

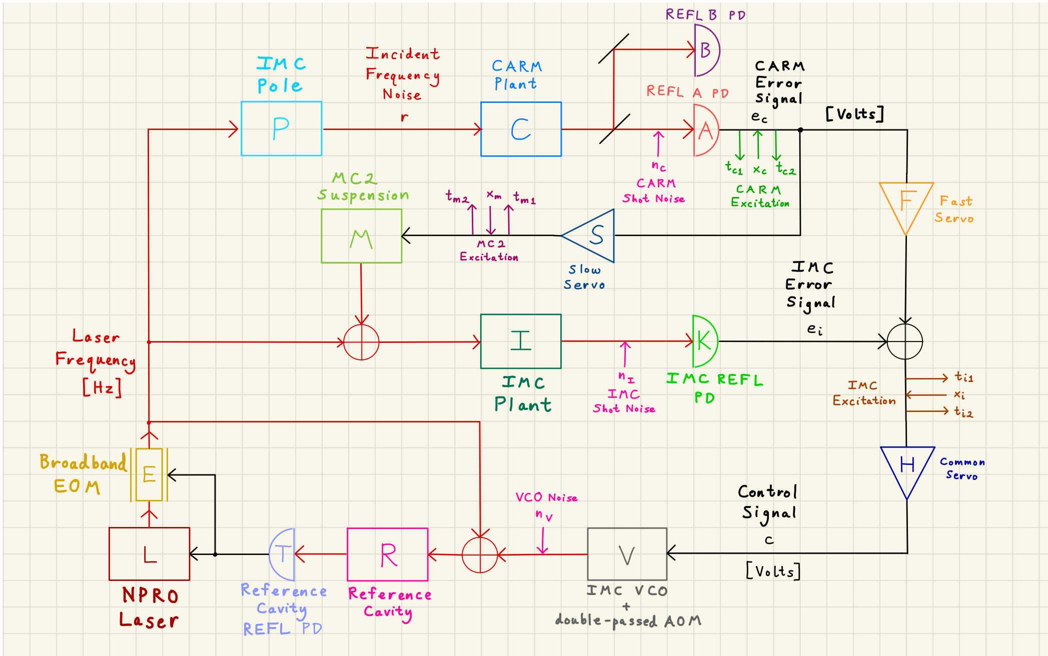

Figure 3.14: Frequency control loop block diagram. The laser L and reference cavity R form the innermost frequency stabilization loop in the bottom left. The input mode cleaner I forms the second loop. The common arm cavity C forms the third loop. The quantity we ultimately care about is the residual frequency noise r incident on the interferometer. The CARM OLG G_{carm} is measured at the CARM excitation point x_c, The IMC OLG G_{imc} is measured at the IMC excitation point x_i, and the MCL crossover G_{mc2} is measured at the MCL excitation x_m.

The denominator represents the IMC loop suppression. The numerator represents the fast and slow paths of the CARM feedback. This quantity is shown in green in Figure 3.12.

Next we measure the IMC loop without CARM feedback:

$$G_{imc} = IV. \tag{3.39}$$

(This measurement is taken in the same configuration as in full lock, e.g. high input power, same analog filters engaged, etc.)

Now we calculate the CARM path gain $\mathcal{G}_{carm}$:

$$\begin{aligned} \mathcal{G}_{carm} &= G_{carm}(1 - G_{imc}) && (3.40) \\ &= CFV + CMIV && (3.41) \end{aligned}$$

The measurements G_{carm} and G_{imc} are done in the frequency range of 1 to 100 kHz. Looking at the MCL crossover in Figure 3.13, in this frequency range the crossover

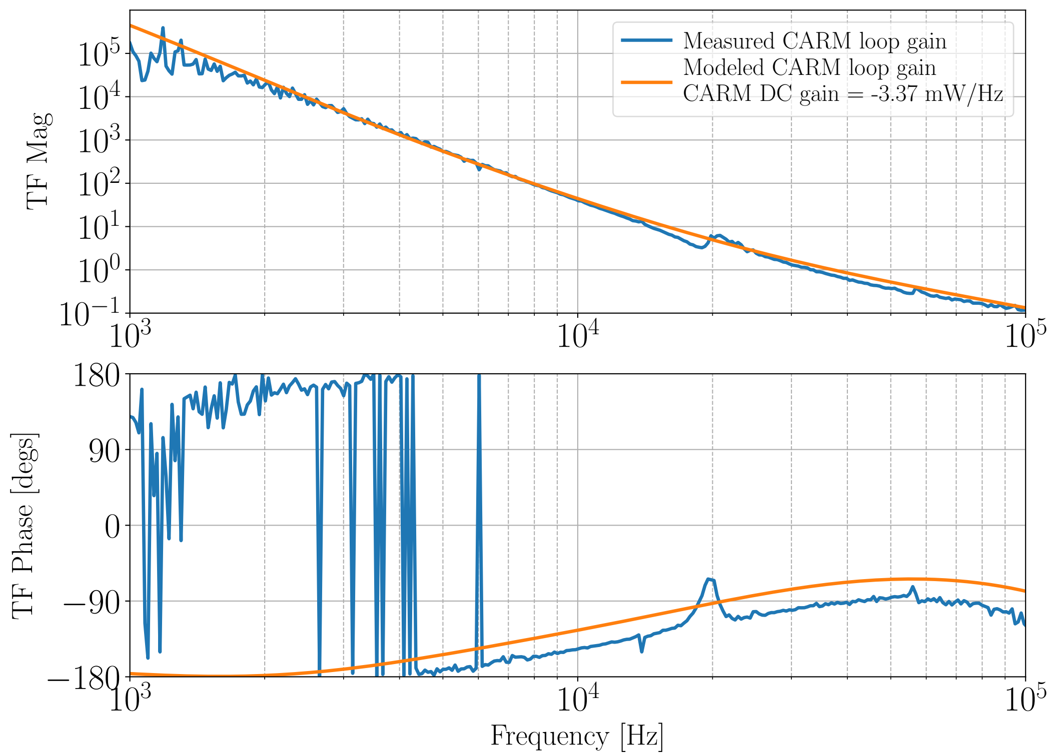

Figure 3.15: CARM path model and measurement comparison. The model of the CARM path gain is computed in Eq. 3.43. The bump at 18 kHz is due to an unmodeled resonance of the 9 MHz sidebands in the arms. This occurs because the 9 MHz carrier is designed to be nearly anti-resonant in the arms, i.e. to resonate in the PRC at a half-FSR frequency. When we shake the laser frequency with an audio frequency near $f = \mathrm{FSR}/2 = 18.7$ kHz, the audio sideband on the 9 MHz will resonate in the arms, increasing the optical gain of the CARM loop near the half-FSR. Table 3.4 shows the values that informed the CARM path gain model.

magnitude is negligible, so we let $M \to 0$ and

$$\mathcal{G}_{carm} = CFV. \tag{3.42}$$

So we've completely eliminated the effect of the IMC on our CARM OLG measurement. Now we must remove the control filters and VCO actuator. This was done with a model of the analog filters.

Expanding our consolidated notation back out:

$$\mathcal{G}_{carm} = PCAFHV \tag{3.43}$$

Table 3.4 shows the values of the CARM model. All are well-known except for the CARM plant C. Therefore we have an overall scale factor in C that we can use to fit our model to the measurement, and measure the CARM plant in this way.

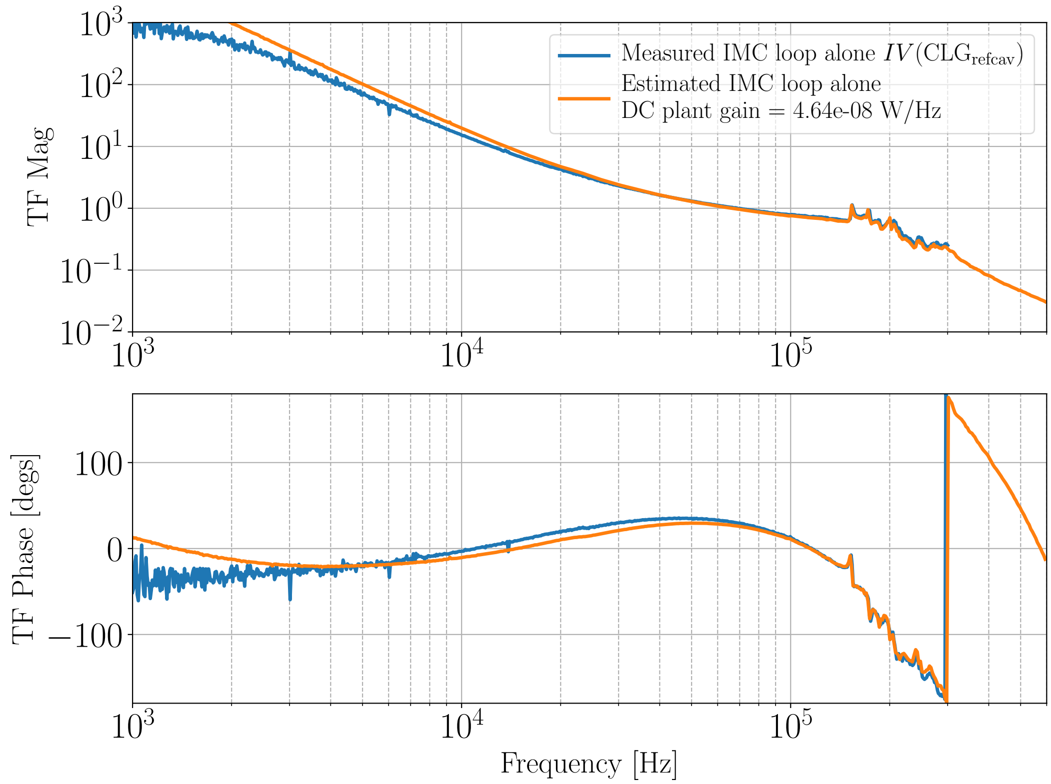

Figure 3.16: IMC path model and measurement comparison. The model of the IMC path gain is computed in Eq. 3.39. Also included here is the measured reference cavity closed loop gain, CLG$_{refcav}$. Table 3.4 shows the values that informed the IMC path gain model.

Using this method, we find the DC CARM plant gain to be around $3.4\ \mathrm{mW/Hz}$. We also model the DC IMC plant gain to be around $4.6 \times 10^{-8}\ \mathrm{W/Hz}$. These plants include the intrinsic optical gain of the cavity plus all optical losses, including excess power beam dumps. Figures 3.15 and 3.16 show the model versus the measurement of these paths.

3.4.4 Frequency noise budget

The frequency noise budget is the characterizes the limit of the CARM stabilization. A couple of main noise sources for frequency noise are the VCO actuator noise, IMC sensing noise, and CARM sensing noise. Figure 3.17 shows the frequency noise budget.

Using the block diagram Figure 3.17, we project noises onto the laser frequency incident on the interferometer r:

$$r = \frac{(AFHVP + ASMIKHVP)n_c + KHVPn_i + Pn_v}{1 - IKHV - CAFHVP - CASMIKHVP} \tag{3.44}$$

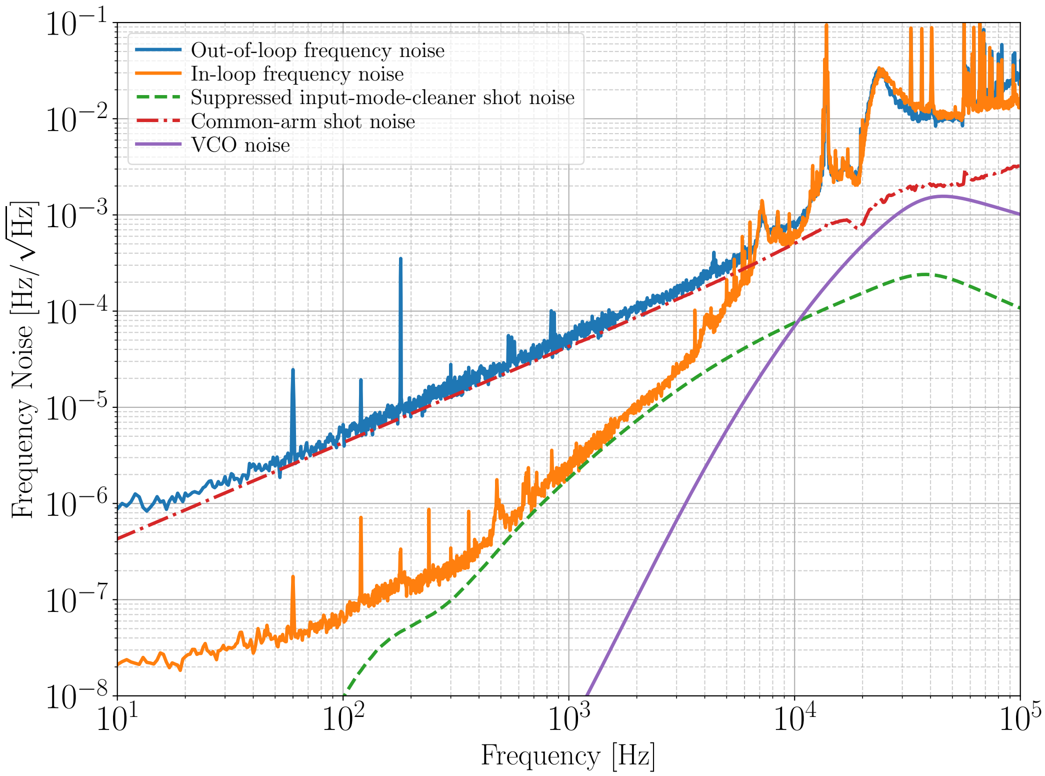

Figure 3.17: Frequency noise budget for LIGO Hanford in O3. In the GW-sensitive frequency band between 10 and 1000 Hz, measured frequency noise is entirely CARM shot noise limited. Improvements to the input mode cleaner sensing noise have rendered its contribution to frequency noise negligible. In low-noise operation, REFL A and B are summed to control CARM, reducing the shot noise limit seen here by a factor of 2.

where n_c is the CARM sensing noise, n_i is the IMC sensing noise, and n_v is the noise of the VCO.

Every individual part of Eq 3.44 is hard to measure and model, so instead we use the components we can easily measure to simplify the algebra. First, the CARM closed loop gain CLG_{carm} can be written

$$\mathrm{CLG}_{carm} = \frac{G_{carm}}{1 - G_{carm}} \tag{3.45}$$

$$= \frac{CAFHVP + CASMIKHVP}{1 - IKHV - CAFHVP - CASMIKHVP} \tag{3.46}$$

Second, the MC2 crossover suppression SUP_{mc2} is written

$$\mathrm{SUP}_{mc2} = \frac{1}{1 - G_{mc2}} \tag{3.47}$$

$$= \frac{1 - IKHV - CAFHVP}{1 - IKHV - CAFHVP - CASMIKHVP} \tag{3.48}$$

We can use the measured CLG_{carm} to significantly simplify the projection of CARM sensing noise n_c:

$$r = \frac{n_c}{C}\mathrm{CLG}_{carm}. \tag{3.49}$$

Since $\mathrm{CLG}_{carm} = 1$ everywhere below 20 kHz, the entire CARM shot noise projection comes to the inverse CARM plant C^{-1}.

We can use the measured SUP_{mc2} to remove the effect of the slow MC2 offloading from the IMC sensing n_i

$$r = \frac{HKPVn_i}{1 - IKHV - CAFHVP}\mathrm{SUP}_{mc2} \tag{3.50}$$

and VCO noise n_v projections

$$r = \frac{Pn_v}{1 - IKHV - CAFHVP}\mathrm{SUP}_{mc2}. \tag{3.51}$$

Eqs. 3.50 and 3.51 leave us with only the suppression from the fast feedback from the IMC and CARM. Eqs. 3.49, 3.50 and 3.51 are all plotted in Figure 3.17.

Other noises not shown in Figure 3.17 include CARM, PRCL, and IMC displacement noise, all of which should be lower than the detected frequency noise. At and below ~ 10 Hz, fringe-wrapping occurs due to relative motion between the interferometer and the REFL photodetectors.

The RMS incident frequency noise is ~ 6 Hz for a bandwidth of 5 to 1×10^5 Hz. The RMS is entirely dominated by noise above 5 kHz, including a 14 kHz peak which may be from the reference cavity first longitudinal mode. The broad hump at ~ 20 kHz is likely due to the odd dynamics of 9 MHz sidebands resonating in the arms leading to a loss of suppression. Unknown are the sources of the forest of peaks above 30 kHz, or the peaks around 7.5 kHz.

3.4.5 Frequency to DARM coupling budget

The final consideration for frequency noise masking GW signals is the measured coupling level from frequency noise to DARM. Frequency noise is common to both arms, and so is largely reflected back toward the input at the beamsplitter. Asymmetries in the interferometer allows frequency noise into the antisymmetric port where DARM is measured. Figure 3.18 shows the frequency noise to DARM transfer function budget.

The asymmetries considered by Izumi [101] and Somiya [114] are straightforward to calculate given the full transfer function from the input to antisymmetric port.

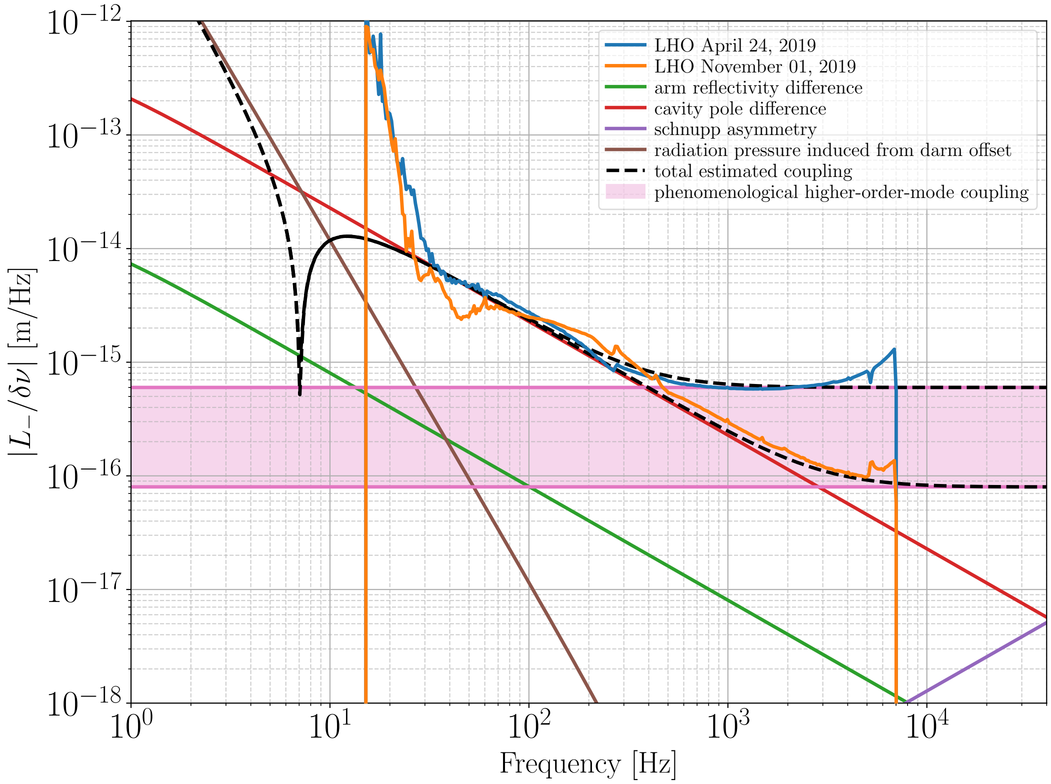

Figure 3.18: Frequency to DARM transfer function budget for LIGO Hanford in O3. From 40 to 400 Hz, the differential arm pole term $\delta\omega_c$ dominates frequency noise coupling to DARM. Below 40 Hz is not well understood, but is likely due to radiation-pressure based coupling through excess frequency noise spoiling the length and angular control loops. Above 400 Hz, there is a variable term is also not well understood, but is commonly attributed to higher-order mode coupling. The content of higher-order modes in the interferometer depends on interferometer geometry, which changes with the thermal state.

Reproducing Eq. 30 of [101] for the frequency noise to DARM coupling:

$$\begin{aligned}\frac{\Delta L_-}{\delta\nu}(f) = &\frac{\pi c}{2\omega_c\omega_0}\frac{1+r_a}{r_a'}\frac{1}{1+s_{cc}} \\ &\times\left[-\delta r_a - \frac{\delta\omega_c}{\omega_c}(1+r_a) + \frac{l_{sch}r_a\omega_c}{c}\left(1-\frac{s_c}{r_a}\right)(1+s_c)\right] \\ &- \frac{16\pi P_a g_s^2 r_a'\omega_0\Delta L_{DC}}{c^2\omega_{rse}s_\mu^2(1+s_{cc})(1+s_{rse})} + k_{HOM}\end{aligned} \tag{3.52}$$

where ω_c is the arm pole frequency, ω_0 is the carrier frequency, r_a is the arm reflectivity, r_a' is the phase derivative of the arm reflectivity, δr_a is the differential arm reflectivity, $\delta\omega_c$ is the differential arm pole, l_{sch} is the Schnupp asymmetry, P_a is the arm power, g_s^2 is the gain of the signal-recycling cavity, ΔL_{DC} is the DARM DC offset, ω_{rse} is the DARM coupled cavity pole, and k_{HOM} is a general higher order mode coupling term. Table 3.5 in the intensity coupling section define and evaluate these expressions. The frequency-dependence of Eq. 3.52 is contained in the s-terms:

$$s_c = i\frac{\omega}{\omega_c}, \qquad s_{cc} = i\frac{\omega}{\omega_{cc}}, \qquad s_{rse} = i\frac{\omega}{\omega_{rse}}, \qquad s_\mu^2 = -\mu\omega^2. \tag{3.53}$$

Summarizing Eq. 3.52, the couplings in the brackets are the three straightforward ways to produce contrast defect in the antisymmetric port. First, the reflectivity difference between the arms δr_a. Second, the arm pole difference between the arms $\delta\omega_c$, sometimes called the arm storage time difference. Third, the Schnupp asymmetry l_{sch} gives a static difference in the inner Michelson length, which produces a very tiny difference in the light travel time. The fourth term is the radiation pressure term due to the frequency noise modulating the phase-quadrature light in the arms due to DARM offset ΔL_{DC}. The fifth term is the least understood and most variable: the coupling due to higher order modes k_{HOM}.

The arm pole term $\delta\omega_c$ dominates at around 100 Hz for Hanford during O3. This was because the two input test masses installed at Hanford were not “twins”, i.e. they were not made at the same time. ITMX was replaced because of a very large point absorber detected during O2. The new ITMX has a transmission of 1.5%, whereas ITMY has a transmission of 1.42%. This yields an X-arm pole $\omega_{cx}/(2\pi) = 45.1$ Hz and a Y-arm pole $\omega_{cy}/(2\pi) = 42.6$ Hz, or differential $\delta\omega_c/(2\pi) \approx 1.2$ Hz. This difference in ITM transmission also affected the arm powers, see Section 3.2.

3.4.5.1 Radiation pressure

The radiation pressure coupling below 30 Hz is not well-understood. The coherence of the measurement dips to around 0.6 to 0.7 as it becomes harder to drive the frequency noise above the DARM noise due to the CARM suppression. Stronger injections were tried, but cause massive upconversion from the many orders of magnitude the injection must cover, which was a problem (see Subsection 3.4.6). Swept sine injections validated the coupling levels measured in Figure 3.18.

The radiation pressure noise due to frequency noise cannot be explained solely by interaction with the quadrature light in the arms. The most likely explanation is the true DARM noise is caused through another path that the excess frequency noise causes, in particular the SRCL path. SRCL is especially capable of creating radiation pressure noise in the arms, a fact we took advantage of in the arm power measurement (Section 3.2). The light in the SRC is all in the quadrature phase, meaning that incident frequency noise will have a strong modulation effect on SRC light.

Working against this is the fact that inside the interferometer the carrier is cleaned of frequency noise by the CARM pole, and the feedforward of the SRCL control signal to DARM compensates for SRCL noise in the DARM spectrum. However, feedforward is tuned to prevent SRCL sensing noise from entering DARM, not displacement noise, and injection of excess frequency noise can spoil the quiescent SRCL state.

This excess frequency to DARM radiation pressure noise may be real at the time of the measurement, but it is possible that the quiescent coupling state is better than what is possible to measure with a strong injection.

3.4.5.2 Higher order modes

Higher order modes (HOMs) of the laser refer to the spatial eigenmodes a laser may have when resonating in a cavity. All LIGO cavities are designed to accept the same TEM_{00} mode, or main carrier light. HOMs are rejected from cavities they are not designed to resonate in. However, HOMs still spawn from the carrier due to mode mismatch, misalignments, and cavity imperfections inside the dual-recycled Michelson. Worse, “point absorbers” on the optics heat the optics irregularly at high power, spoiling the cavity geometry and scattering carrier light into HOMs. We know from the DARM optical spring detuning that there is likely significant HOM content in the SRC, see Subsection 3.6.4. However, we don’t have a sensor or a good model for the level of HOMs in the interferometer.

HOMs can carry frequency noise, and do not experience the usual cleaning effect from the CARM and DARM coupled cavities. Ideally, the output mode cleaner will reflect away most low-order HOMs, but it is possible some high-order HOMs transmit easily through the output mode cleaner, or that there is very large low-order HOM content incident on the output mode cleaner. Another possibility is significant “mode healing” happening in the signal recycling cavity, which is when HOMs are scattered back into the main carrier light by the signal recycling cavity mirrors. This light would not be cleaned by the CARM coupled cavity, but transmit directly through the output mode cleaner.

These different coupling paths will have a different frequency dependence in their coupling to DARM. This is based on the fact that HOMs do not resonate in the arms. Therefore, the frequency and intensity noise the HOMs carry are not suppressed by the interferometer's coupled cavity poles s_{cc} and s_{rse}.

Finesse simulation of simple Michelson suggests that frequency noise on HOMs couple to the antisymmetric port with a flat frequency dependence in W/Hz, same as the carrier frequency noise coupling explored in Subsection B.3.3. If we assume the coupling path is some large HOMs created inside the dual-recycled Michelson, which are then transmitted through the output mode cleaner, this coupling would be flat in W/Hz. Referring watts in the antisymmetric port back to DARM meters, as in Eq. 3.52, would give a frequency dependent coupling $k_{HOM} \propto \kappa_{HOM}(1 + s_{rse})$, with a flat DC coupling and an f-like coupling above the DARM pole.

However, the "mode healing in the SRC" coupling path *would* be cleaned by the DARM pole, because the HOM light returns to carrier light inside the DARM coupled cavity. The antisymmetric port coupling in $\mathrm{W/Hz} \propto \kappa_{HOM}/(1 + s_{rse})$. Referring back to DARM meters would undo the DARM pole, making the coupling flat in $\mathrm{m/Hz} \propto k_{HOM}$.

Finesse simulation of the whole interferometer suggests that the "mode healing" path is most likely. This is the coupling chosen for Eq. 3.52, since it is consistent with measurement and simulation. In the plot, $k_{HOM} \approx (8 - 60) \times 10^{-17}$ m/Hz. However, some measurements have an unexplained uptick at very high frequencies > 5 kHz, which suggests there may be multiple coupling paths of frequency noise through HOMs, or possibly another mechanism entirely.

HOM content in the interferometer can be partially controlled by adjusting the cavity geometry. This is done in Advanced LIGO through the thermal compensation system, which can heat the optics to increase or decrease their radius of curvature. Thermal tuning can be performed on the input and end test masses, and the SR3 mirror.

Through many measurements, we have seen a large variation in the level of frequency noise to DARM coupling at high frequency. Figure 3.18 highlights two extremes of this coupling level, one taken in April 2019 and the other in November 2019. A factor of 10 lowering in the coupling was achieved between these two times through tuning of the thermal state of the interferometer to repair the interferometer geometry, a strong sign that it truly is HOMs causing excess frequency noise coupling at high frequencies. The couplings where monitored via injected frequency noise lines. Other monitor lines were injected for intensity noise, which is also dominated by HOM coupling at high frequency (Section 3.5.3). However, we found that the HOMs coupling frequency noise and intensity noise must be different, i.e. as one coupling increased, the other decreased [115]. A middle-ground thermal setting was found that minimized both couplings as well as possible.

RF sidebands have similar corner-centric coupling mechanisms as HOMs. However, tests where the modulation depth was changed during the injection did not change frequency noise coupling levels [116].

3.4.6 Output mode cleaner dither line

The output mode cleaner (OMC) is a bow-tie cavity on the antisymmetric port of the interferometer. Its main purpose is to reject RF sidebands and HOMs, while passing only carrier light and GW signal.

To accomplish this, the OMC must be locked to carrier. This is done by injecting a dither line at 4.1 kHz on one of the OMC's piezoelectric transducers, then demodulating that line in the OMC DCPD sum to get an error signal for locking the OMC to carrier.

We mention the OMC dither line here because the frequency noise level at the antisymmetric port at 4.1 kHz is now much more important. Excess frequency noise will be downconverted to DC by the OMC dither line if the line is not strong enough.

This was discovered while measuring the CARM loop open loop gain, and was a problem for measuring frequency transfer functions to DARM at low frequency, where upconversion of the frequency noise injection caused downconversion from the OMC dither.

The solution is to ensure the OMC dither line is strong enough for normal low-noise operation such that it is stronger than the frequency noise, and ensure the frequency noise injections at HF are not so strong they upconvert and drown out the OMC dither line. Carefully-chosen swept sine measurements are better for quantifying low-frequency frequency noise coupling.

3.5 Intensity stabilization

The intensity stabilization servo (ISS) is required to limit laser noise masking the gravitational wave signal. Intensity fluctuations from the laser masks the gravitational wave signal via contrast defect and directly through the DARM offset used for DC readout. Section B.3.3, Eq. B.61 shows the intensity coupling for a simple Michelson from first principles.

Intensity fluctuations benefit from the common mode rejection of the interferometer, since any input beam fluctuations will be common to both arms. Intensity noise on the carrier will also be cleaned by the CARM and DARM coupled cavity poles of the interferometer. However, running in DC readout couples intensity noise directly to the antisymmetric port were DARM is sensed.

An analog, two-loop hierarchal intensity stabilization servo is used to suppress the laser intensity incident on the interferometer. The suppressed intensity noise is further cleaned by the coupled cavities of the interferometer such that intensity noise does not mask GW signals. The ISS is DC-coupled in full lock, so the overall laser power entering the interferometer is stabilized. Previous experiments informed the ultra-stable DC-coupled laser intensity control scheme used in Advanced LIGO [34, 35].

This section will overview the Hanford O3 performance of the ISS, discuss the intensity noise budget, and the intensity transfer function to DARM budget.

3.5.1 Intensity control scheme

The freerunning NPRO relative intensity noise (RIN) is measured to be $10^{-5}\ \mathrm{Hz}^{-1/2}$ at $100\ \mathrm{Hz}$, with around $1/f$ dependence [36]. The requirement for RIN incident on the power recycling mirror is $2 \times 10^{-8}\ \mathrm{Hz}^{-1/2}$ at $100\ \mathrm{Hz}$, rising like f [117]. Figure 3.19 shows a diagram of the laser intensity stabilization servo.

The high power laser is first incident on an acousto-optic modulator (AOM), which defracts power based on the input voltage. The beam is then locked to the pre-mode cleaner (PMC), which is on the optical table and passes most intensity noise through the short bow-tie cavity. Transmission through one of the ports of the PMC is then split on two PDs, ISS PDA and ISS PDB. ISS PDB senses intensity noise, which is fed back to the AOM, completing the ISS first loop. ISS PDA serves as an out-of-loop sensor for the ISS first loop.

The main beam propagates forward from the PMC, through the IMC, through the input optics toward the interferometer. A pickoff beam then heads toward an eight-photodetector "ISS array", where $\sim\ 60\ \mathrm{mW}$ is detected. Four of the PDs on the detector form the "ISS inner" signal, which senses intensity noise heading to the interferometer and feeds back to the AOM, completing the ISS second loop.

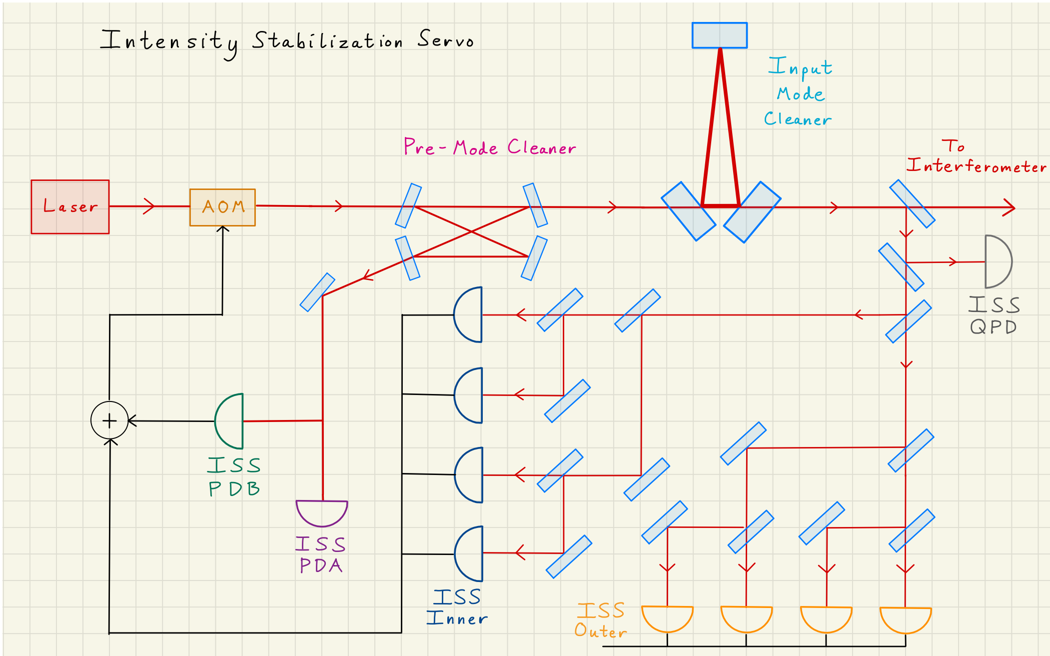

Figure 3.19: Laser intensity stabilization servo diagram. The laser in the top left passes through the AOM, then is locked to the PMC. The ISS first loop is formed from the transmission through one port of the PMC, and intensity noise is sensed by a single photodetector, ISS PDB. The ISS second loop is formed by the beam transmitted through the IMC, picked off on the path to the interferometer, and sent to the eight-photodetector ISS array. Finally, the ISS QPD detects beam jitter heading to the interferometer, which can be misinterpreted as intensity noise.

The other four PDs form the "ISS outer" signal, which serves as the out-of-loop intensity noise sensor. The

The ISS first and second loops are summed together, with the second loop being added into the first loop as an additive offset. Thus the first loop follows the second loop error signal, and both loops suppress the intensity noise sensed by the ISS array. Figure 3.20 shows the open loop gains of the first and second loops of the laser intensity stabilization servo.

The intensity stabilization described here focuses on the audio-frequency intensity noise. However, the ISS is also DC coupled to keep the overall light levels in the interferometer robust to slow drift in the laser output.

One pitfall of the current configuration is the relative polarizations between the transmission ports of the PMC [118]. Figure 3.19 emphasizes that the first and second ISS loops are stabilized to different outputs of the PMC, but these ports have been found to have different polarizations on the PSL optics table. This means

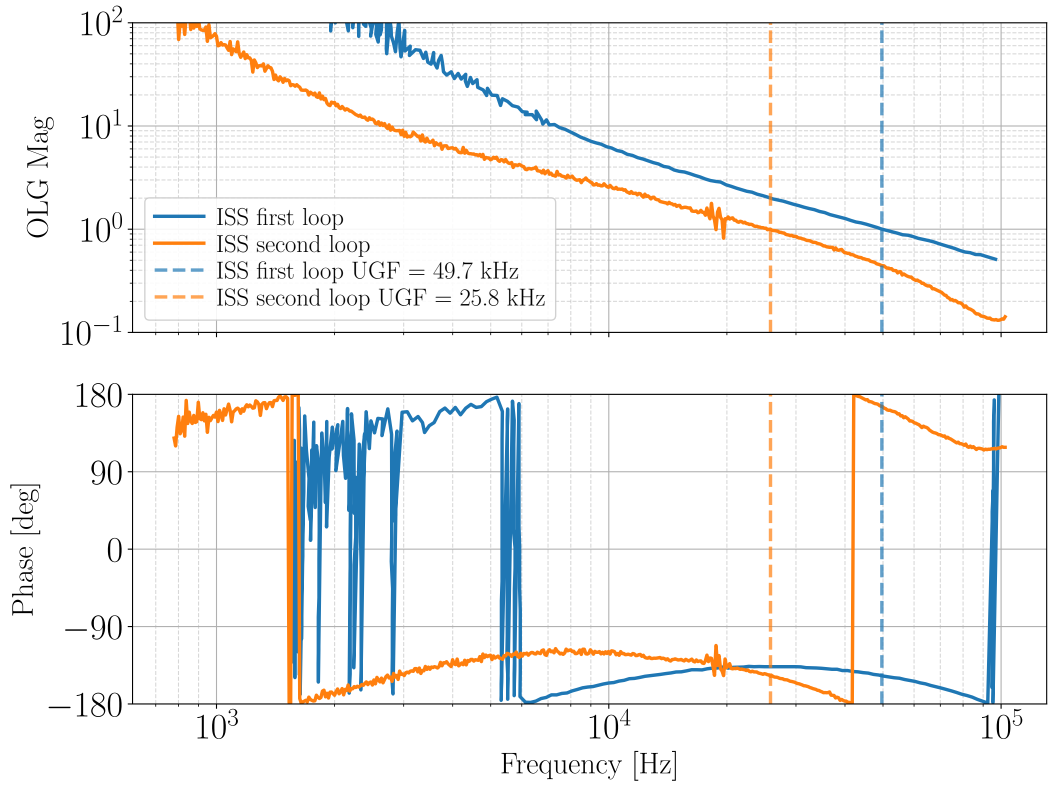

Figure 3.20: Intensity stabilization servo open loop gain measurements. The first and second ISS loops form a hierarchal stabilization scheme, so the second loop may never have a higher unity gain frequency (UGF) than the first loop in the current configuration.

the first and second loops may be sensing and stabilizing slightly different light, which can lead to intensity controls fighting one another.

3.5.2 Intensity noise budget

Figure 3.21 shows the intensity noise budget. The total RIN RMS $\approx 6.5 \times 10^{-6}$. The ISS second loop shot noise limit for intensity noise is achieved for most of the GW detection band. Shot noise is unusually low because of the $\sim$ 30 mW of in-loop light detected. The actual in-loop shot noise will be $\sqrt{2}$ lower than the red trace in Figure 3.21.

The IMC angular controls peaks at 1.1 Hz and 3.4 Hz dominate the intensity noise RMS. Changes in transmitted power due to IMC misalignment are most likely responsible for the registered peaks.

Post-IMC beam jitter measured at the ISS QPD is coherent with the ISS out-of-loop noise in the 10 Hz region. This could be due to jitter in the HAM2 input optics.

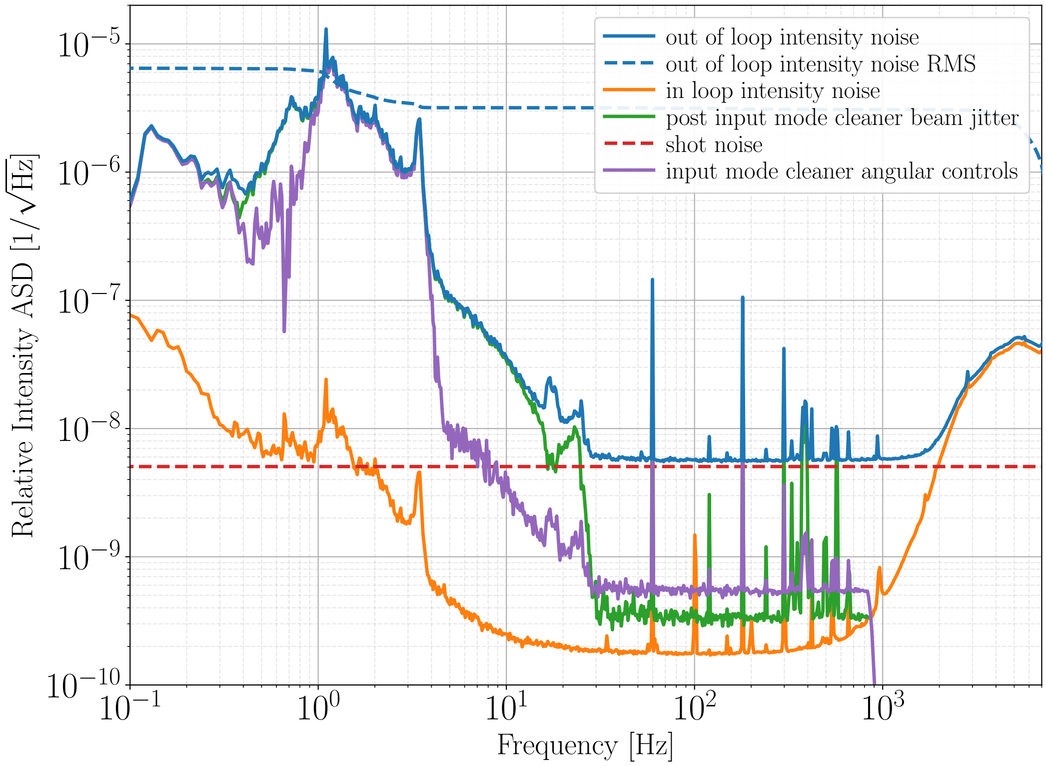

Figure 3.21: Relative intensity noise budget for laser intensity incident on the interferometer for LIGO Hanford in O3. The ISS second loop shot noise limit is achieved between 30 and 1000 Hz. Above 1 kHz the ISS is gain-limited. Between 5 and 30 Hz intensity is coherent with beam jitter detected on the ISS QPD. Angular motion in the IMC dominates the total RMS from peaks around 1.1 Hz and 3.4 Hz.

Another possibility is the ISS QPD and ISS out-of-loop signal are both registering some upconversion from the strong IMC angular peak at 3.4 Hz.

The ISS loop is gain-limited at high frequency. This noise comes close limiting DARM, as seen in the green dots in the DARM noise budget (Figure 3.1). It's not so simple to increase the ISS suppression, as evidenced by the ISS OLGs in Figure 3.20, with unity gain frequencies of $\sim$ 50 kHz and $\sim$ 25 kHz, each with phase margins of around 30°. Higher intensity stabilization unity gain frequencies of ~ 100 kHz where achieved in [34] and [35], but those servos did not have to contend with being DC-coupled to a suspended IMC. The ultimate limit of the ISS stabilization is the AOM actuator. The AOM was measured to have a flat power modulation out to 200 kHz, with a phase lag of 45° at 120 kHz [36].

The gain-limited intensity noise will become more important as squeezing is improved and higher arm powers are achieved, lowering DARM shot noise. More work on the ISS control scheme is required so intensity noise will not limit DARM at frequencies above 1 kHz.

3.5.3 Intensity to DARM coupling budget

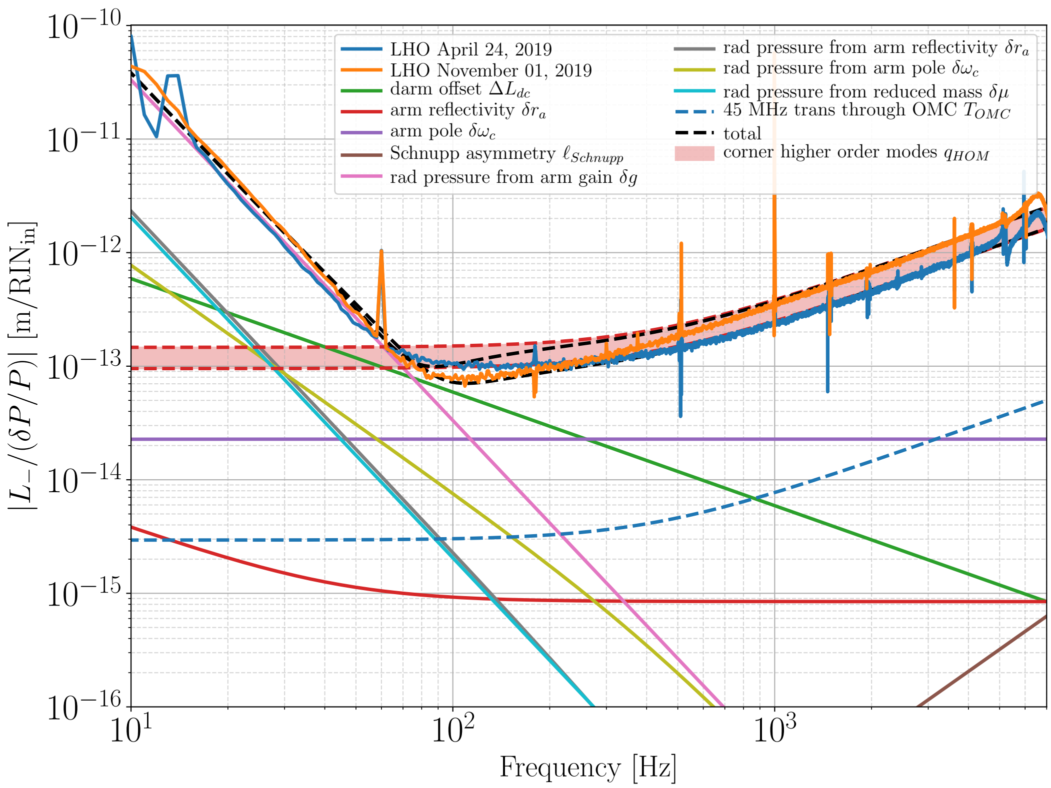

Figure 3.22: Relative intensity to DARM transfer function budget for LIGO Hanford in O3.

The coupling of intensity noise to DARM depends on the level of contrast defect, differential radiation pressure, and RF sidebands and higher order modes (HOMs) through the OMC. Figure 3.22 shows the input relative intensity noise to DARM transfer function budget.

Izumi [101] and Somiya [114] considered the appearance of intensity noise from contrast defect and differential radiation pressure. Reproducing Eq. 26 of [101], but calibrating AS watts into DARM ΔL_- and RAN into RIN gives:

$$\begin{aligned}\frac{\Delta L_-}{\delta P/P}(f) = &\frac{c^2(1+r_a)^2\Delta L_{DC}}{4\omega_0^2}\frac{1}{1+s_{cc}} \\ &+ \frac{\delta r_a}{4r_a'^2\Delta L_{DC}}\frac{1}{1+s_{cc}}\left[\delta r_a(1+s_c) + \frac{\delta\omega_c}{\omega_c}s_c(1+r_a) - \frac{l_{sch}r_a\omega_c}{c}s_c\left(1-\frac{s_c}{r_a}\right)(1+s_c)\right] \\ &+ T_{omc}\frac{\Gamma_{45}^2}{4}\frac{g_{sb}^2 t_{sm}^2}{2g_p^2 g_s^2 r_a'^2\Delta L_{DC}}(1+s_{rse}) \\ &+ \frac{2P_a}{cs_\mu^2}\frac{1}{1+s_{cc}}\left[\frac{\delta P_a}{2P_a} - 2\frac{\delta\mu}{\mu} + \frac{\delta\omega_c}{\omega_c}\frac{g_s s_c(2+s_{rse})}{t_s(1+s_{rse})} - \delta r_a\frac{g_s r_s(2+s_{rse})}{t_s(1+s_{rse})}\right] \\ &+ \frac{q_{HOM}}{2g_p^2 g_s^2 r_a'^2\Delta L_{DC}}(1+s_{rse}) \end{aligned} \tag{3.54}$$

Table 3.5 defines the expressions and values in Eq. 3.54. The frequency-dependence of Eq. 3.52 is contained in the s-terms:

$$s_c = i\frac{\omega}{\omega_c}, \qquad s_{cc} = i\frac{\omega}{\omega_{cc}}, \qquad s_{rse} = i\frac{\omega}{\omega_{rse}}, \qquad s_\mu^2 = -\mu\omega^2. \tag{3.55}$$

To get Eq. 3.54 we had to account for a factor of two from different DARM definitions $\Delta L_-^{\text{Izumi}} \to \Delta L_-/2$ and a factor of two going from RAN to RIN $RAN = RIN/2$, which cancel in the end. We have also written the differential arm power gain $\delta g_{arm}/g$ in Eq. 26 of [101] in terms of differential arm power δP_a,

$$\frac{\delta g_{arm}}{g_{arm}} = \frac{\delta P_a}{2P_a} \tag{3.56}$$

since $P_a \propto g_{arm}^2$.

Summarizing Eq. 3.54, there are nine different expressions for how intensity noise couples to DARM in Eq. 3.54. These are each plotted as a separate line in Figure 3.22. We will go through them line by line.

First, the first line of Eq. 3.54 defines the direct coupling due to the DARM DC offset ΔL_{DC}. Second, the second line defines the coupling due to asymmetries, like arm reflectivity δr_a, arm poles $\delta\omega_c$, and Schnupp asymmetry l_{sch}. Third is the coupling of 45 MHz sidebands directly through the OMC. Fourth are the radiation pressure terms, due to differential arm powers δP_a, optic masses $\delta\mu$, arm reflectivity δr_a, arm poles $\delta\omega_c$. Fifth is the coupling due to higher order modes.

In this case, the differential arm powers that arise from the different finesses in each arm readily explain the intensity noise coupling below 100 Hz. Above 100 Hz, the only plausible explanation is some phenomenological HOM coupling. 45 MHz

Description	Symbol	Value
Carrier frequency	$\omega_0/(2\pi)$	2.82×10^{14} Hz
Arm reflectivity	$r_a = \dfrac{-r_i + r_e}{1 - r_i r_e}$	0.986
Differential arm reflectivity	$\delta r_a = (r_{ax} - r_{ay})/2$	2.5×10^{-3}
Arm reflectivity round-trip phase derivative	$r_a' = \dfrac{t_i^2 r_e}{(1 - r_i r_e)^2}$	268
Arm cavity pole	$\omega_c/(2\pi)$	44.0 Hz
Differential arm cavity pole	$\delta\omega_c = (\omega_{cx} - \omega_{cy})/(4\pi)$	1.5 Hz
CARM coupled cavity pole	$\omega_{cc}/(2\pi)$	0.65 Hz
DARM coupled cavity pole	$\omega_{rse}/(2\pi)$	411 Hz
Schnupp asymmetry	$l_{sch} = l_x - l_y$	0.08 m
Power recycling gain	g_p^2	43 W/W
Signal recycling gain	g_s^2	0.1 W/W
Arm power	P_a	201 kW
Differential arm power	$\delta P_a = (P_{ax} - P_{ay})/2$	-6.5 kW
DARM DC offset	ΔL_{DC}	10 pm
45 MHz modulation depth	Γ_{45}	0.177 rad
45 MHz sideband PRC gain	g_{sb}^2	1.9 W/W
45 MHz sideband SRC transmission	t_{sm}^2	0.21
45 MHz OMC transmission	T_{omc}	100 ppm
Reduced mass of the arm	$\mu = M/2$	20 kg
Higher order mode input RIN to AS watts	q_{HOM}	$(1.3 - 2.0) \times 10^{-5}$ W/RIN
Higher order mode input Hz to DARM meters	k_{HOM}	$(8 - 60) \times 10^{-17}$ m/Hz

Table 3.5: Intensity and frequency noise coupling to DARM parameters from Eq. 3.54 and Eq. 3.52. All differential parameters are divided by two, e.g. $\delta P_a = (P_{ax} - P_{ay})/2$. Round-trip loss in the arms was assumed to be 100 ppm. The same values are used for the frequency noise coupling in Figure 3.18. k_{HOM} is the pink band in that plot.

coupling has the same shape, but even the pessimistic parameters used to estimate this coupling could not come close to the measured coupling ($T_{omc} = 100$ ppm is fairly high).

The coupling mechanisms for HOMs were discussed in Subsection 3.4.5.2. In this case the mechanism is most likely corner HOMs coupling directly through the output mode cleaner, with no DARM pole cleaning. The coupling of intensity noise to the antisymmetric port is flat in W/RIN. Referring back to DARM gives the DARM zero seen in the fifth line of Eq. 3.54.

Similar to frequency coupling, intensity noise coupling was found to vary with the interferometer thermal state. Measurements from April and November of 2019 are plotted in Figure 3.22 showing the $\sim\ 50\%$ difference in coupling above 100 Hz.

We reiterate here that the HOMs carrying intensity noise are not necessarily the same as those carrying frequency noise. This is known because, for some thermal changes in the interferometer, the intensity and frequency couplings are anti-correlated.

3.6 DARM optical plant

The DARM response to gravitational waves sets the ultimate limit for interferometer sensitivity. An enormous response to differential arm motion is necessary to amplify the extraordinarily weak GW signal to a detectable level. The most basic principle of Advanced LIGO interferometer design is maximizing GW response. These are the reasons for the 4 km long arms, the high-finesse arm cavities, the high input laser power, the dual-recycling cavities, and DC readout. All other auxiliary systems are there to enable the interferometer to be maximally sensitive to DARM motion.

The DARM coupled-cavity is formed by the two interferometer arms and the signal-recycling mirror (SRM). The arm cavities resonate the incident carrier light, building up the light in the CARM coupled-cavity to a high level. Then, when DARM motion occurs, light is phase-shifted out of CARM into DARM, showing up at the antisymmetric port of the beamsplitter heading toward the SRM.

The signal-recycling cavity (SRC) is held off-resonance for carrier to shape the overall DARM pole dynamics, so carrier light is preferentially transmitted through the SRC out of the interferometer. This light carries with is the DARM, and GW, signal. The setup where the SRC is held off-resonance for carrier light, i.e. $\phi_{\mathrm{SRC}} = \pi/2$ for carrier, is known as *resonant sideband extraction*, or RSE, and references the broadening of the bandwidth of the detector at the expense of DC gain. Holding the SRC off-resonance, coupled with the fact that the SRC is a low-finesse cavity, makes the SRCL degree of freedom the most difficult to control in Advanced LIGO.

Buonanno and Chen first derived the DARM response and quantum-limited sensitivity of Fabry Perot dual-recycled Michelson interferometers [22]. Later, Ward [119] and Hall [57] modeled and measured the DARM plant dynamics at the Caltech 40 m and LIGO Hanford, respectively.

During O1 and O2, LIGO Hanford operated with significant detuning in the signal-recycling cavity, which produced a DARM optical anti-spring response [5, 120]. The cause of the detuning was not well-understood, nor was is possible to remove completely with SRCL offsets without losing lock. Additionally, the DARM pole value at Hanford was measured to be ~ 10 Hz lower compared to Livingston in O1: Hanford's DARM pole was consistently around ~ 360 Hz, where Livingston's was ~ 370 Hz, much closer to the design DARM pole of ~ 372 Hz [5]. Finally, in O1 the SRC exhibited "mode-hopping" issues, where the SRC would be locked to the correct TEM_{00} carrier mode, then spontaneously switch to a nearby mode, causing locklosses [121]. This was mitigated by aligning the SRC sooner after locking DRMI.

In O3, the SRM transmission was lowered from 37% to 32%, increasing the finesse of the SRC and DARM cavities. During O3, DARM was observed to exhibit both an optical spring and anti-spring. The DARM spring was observed to change sign with the increase of input power, from spring to anti-spring [122], and with the off-on-off test on the SR3 disk heater, from anti-spring to spring to anti-spring again [123]. Again, the DARM pole at Hanford in O3 was low compared to Livingston and design: Hanford DARM pole was 411 Hz [94], Livingston's was 450 Hz, and design was 456 Hz.

This section will report on the latest understanding of the DARM plant at LIGO Hanford in O3. We will also explore the effect of the SR3 disk heater on the DARM plant, and Finesse simulations undertaken to model the effect of higher order modes on DARM detuning.

3.6.1 DARM model

To solve for the fields in a DARM coupled-cavity, one might try initially to write the usual systems of equations for classical plane wave fields. This would get you most of the way to understanding the DARM cavity field dynamics under ideal circumstances. However, for detuned, high-power interferometer the full response becomes more complicated. This is due to the cavity frequency eigenmodes changing from detuning and radiation pressure effects.

The lossless DARM response to GWs is derived from Eqs. (2.20)-(2.24) in Buonanno and Chen [22]. Ward compiled this equation into a convenient expression in Eq. (3.83) of [119]. We will slightly modify the Ward model to give the lossless DARM response in antisymmetric watts per DARM meters:

$$\frac{P_{as}}{L_-}(f) = \frac{\sqrt{2P_{LO}}}{L}\sqrt{\frac{2P_{bs}\omega_0^2}{\omega_c^2+\omega^2}}\frac{t_s e^{i\beta}\left[(1-r_s e^{i2\beta})\cos\phi\cos\zeta - (1+r_s e^{i2\beta})\sin\phi\sin\zeta\right]}{1+r_s^2 e^{i4\beta} - 2r_s e^{i2\beta}\left[\cos(2\phi)+\frac{\kappa}{2}\sin(2\phi)\right]} \tag{3.57}$$

where P_{LO} is the local oscillator due to the DARM offset, L is the arm length, P_{bs} is the power incident on the beamsplitter, ω_0 is the carrier frequency, ω_c is the arm pole, $\omega = 2\pi f$ is the signal frequency, r_s is the SRM amplitude reflectivity, t_s is the SRM amplitude transmission, $\beta = -\arctan(\omega/\omega_c)$ is the phase delay of the arm travel time, ϕ is the SRC detuning angle, ζ is the homodyne angle, and κ is the radiation pressure term for coupling amplitude quadrature fluctuations to phase quadrature inside the arm cavity [47]:

$$\kappa = \frac{8P_{bs}}{ML^2}\frac{\omega_0}{\omega^2(\omega_c^2+\omega^2)} \tag{3.58}$$

Table 3.6 lists the parameters values relevant for Eq. 3.57. We note here that as L changes, P_{as}/L_- remains flat: there is a hidden factor of $1/L$ in the arm pole ω_c.

We have modified Eq. 3.57 vs Ward Eq. (3.83) by adding the prefactor $\sqrt{2P_{LO}}/L$. The factor L comes from converting to DARM meters since $L_- = hL$, P_{LO} is the local oscillator from the DARM offset that beats with the GW signal on the PD, and the factor of $\sqrt{2}$ comes from the quadrature definition, which can be seen in the difference in the prefactor between e.g. Eqs. (2.1) and (2.10) in [22] (see Subsection A.3).

We can see from Eq. 3.57 the importance of the detuning angle ϕ and the homodyne angle ζ. To operate in resonant sideband extraction, $\phi = \pi/2$. This puts the GW signal entirely in the amplitude quadrature upon exiting the interferometer. To detect the GW signal, the homodyne angle $\zeta = \pi/2$. The level of contrast defect in the interferometer sets the homodyne angle, which is hard to estimate in full lock where it matters [124].

In perfect RSE, the κ dependence in the denominator goes to zero, keeping the DARM plant resonance exactly at DC. However, detuning in the SRC will create a DARM optical spring for $\phi < 90°$, or an anti-spring for $\phi > 90°$ [125].

By itself, the DARM spring does not harm the response, and design papers advocate for locking with intentional detuning to increase SNR to binary neutron stars. However, Advanced LIGO is designed to lock with no spring, meaning our SRCL error signal zero-point is detuned from perfect resonance for the carrier for unknown reasons, and must be compensated for with an offset in the SRCL control. The strong, unintentional optical spring, coupled with the low DARM pole, indicates a serious problem with the Hanford SRC, especially when compared to Livingston's consistent DARM plant.

One possibility is excessive losses harm the DARM response. No losses are considered in Eq. 3.57. In the next section we include losses in the arms, SRC, and after exiting the interferometer.

3.6.2 DARM model with losses

Buonanno and Chen also derived the DARM response including loss. Reproducing the DARM response from Eq. (5.6) of [22] here:

$$\begin{pmatrix} b_1^L \\ b_2^L \end{pmatrix} = \frac{1}{M^L}\left[\sqrt{2\kappa}t_s e^{i\beta}\begin{pmatrix} D_1^L \\ D_2^L \end{pmatrix}\frac{h}{h_{SQL}}\right] \tag{3.59}$$

where b_1^L, b_2^L are the amplitude and phase quadratures of the output light from the interferometer, h is the GW signal, h_{SQL} is the standard quantum limit for GW detection:

$$h_{SQL}(\omega) = \sqrt{\frac{8\hbar}{m\omega^2 L^2}} \tag{3.60}$$

Description	Symbol	Value
Arm length	L	4 km
Carrier frequency	$\omega_0/(2\pi)$	2.82×10^{14} Hz
Mass of test masses	M	40 kg
End test mass transmission	T_e	4 ppm
Input test mass transmission	T_i	1.46%
Signal recycling mirror transmission	T_s	32%
Arm cavity pole	$\omega_c/(2\pi)$	44.0 Hz
Input power	P_{in}	34 W
Power recycling gain	g_p^2	44
Power on the beamsplitter	$P_{bs} = P_{\text{in}} g_p^2$	1.4 kW
Local oscillator power	P_{LO}	23.8 mW
SRC detuning	ϕ	90 ± 0.5 degs
Homodyne angle	ζ	89.3 ± 2 degs
Round-trip arm loss	$\mathcal{L}_{rt}$	100 ppm
SRC loss	λ_{sr}	3%
Post-SRM path loss	λ_{pd}	25%

Table 3.6: Typical DARM model parameters for Hanford in O3 from Eq. 3.57. These are known to change with improved interferometer thermal compensation and beam spot positions on the optics. Spot position changes can also affect the angle-to-length coupling in the interferometer [126]. The homodyne angle estimate comes from contrast defect measurements at 2 W input, scaled up to 34 W [124]. The squeezing loss budget for the post-SRM path is documented in [127], which totals 17% loss from the back of the SRM to DCPDs. We need around 25% loss for λ_{pd} for typical DARM measurement fits.

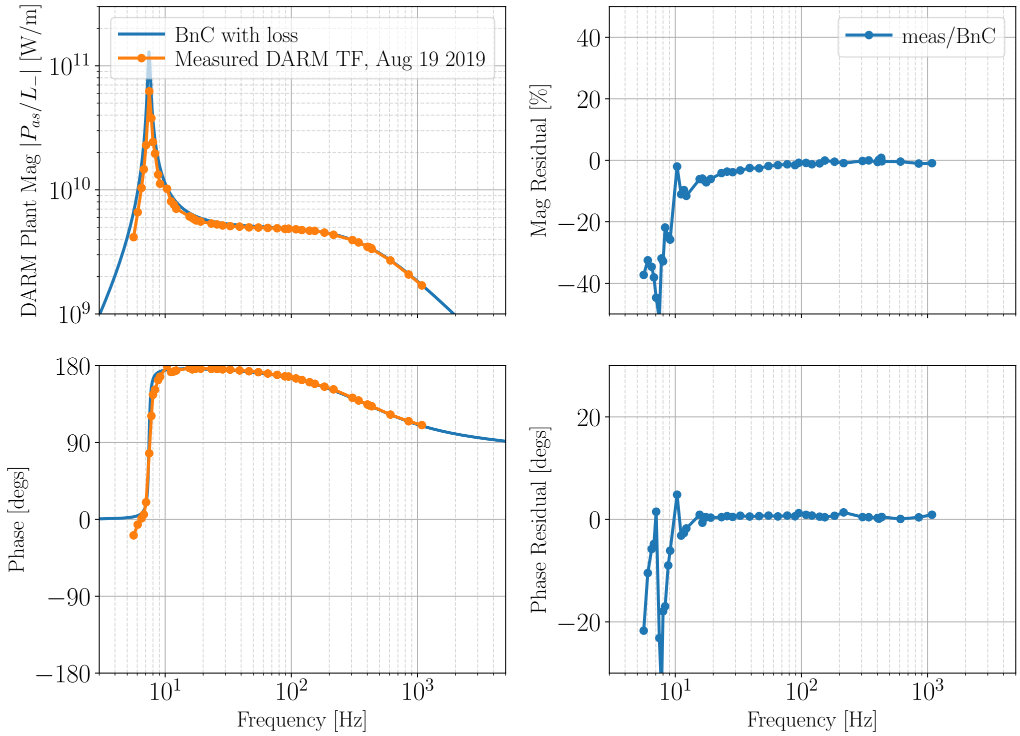

Figure 3.23: DARM model with losses vs LIGO Hanford measurement in August 2019. The DARM measurement represents an extreme DARM plant for illustrative purposes, not the typical response. Here, DARM exhibits a strong, high-frequency DARM spring at 7.5 Hz, and a low DARM pole of 400 Hz. The fit parameters for the model in Eq. 3.66 are $\phi = 89.37°$, $\mathcal{L}_{rt} = 100$ ppm, $\lambda_{sr} = 6.8\%$, and $\lambda_{pd} = 25\%$. All others are the same as Table 3.6.

The expressions for M^L, D_1^L, D_2^L are from Eqs. (5.7) and (5.9) of [22]:

$$\begin{aligned} M^L &= 1 + r_s^2 e^{4i\beta} - 2r_s e^{2i\beta}\left(\cos(2\phi) + \frac{\kappa}{2}\sin 2\phi\right) \\ &\quad + \lambda_{sr} r_s e^{2i\beta}\left(-r_s e^{2i\beta} + \cos 2\phi + \frac{\kappa}{2}\sin 2\phi\right) \\ &\quad + \epsilon r_s e^{2i\beta}\left(2\cos^2\beta(\cos 2\phi - r_s e^{2i\beta}) + \frac{\kappa}{2}(3 + e^{2i\beta})\sin 2\phi\right) \end{aligned} \tag{3.61}$$

$$D_1^L = \sqrt{1-\lambda_{pd}}\Bigg(-(1 + r_s e^{2i\beta})\sin\phi + \frac{\lambda_{sr}}{2} e^{2i\beta} r_s \sin\phi \tag{3.62}$$

$$+ \frac{\epsilon}{4}(3 + r_s + 2r_s e^{4i\beta} + (5r_s + 1)e^{2i\beta})\sin\phi \Bigg) \tag{3.63}$$

$$D_2^L = \sqrt{1-\lambda_{pd}}\Bigg(-(-1 + r_s e^{2i\beta})\cos\phi + \frac{\lambda_{sr}}{2} e^{2i\beta} r_s \cos\phi \tag{3.64}$$

$$+ \frac{\epsilon}{4}(-3 + r_s + 2r_s e^{4i\beta} + (5r_s - 1)e^{2i\beta})\cos\phi \Bigg) \tag{3.65}$$

where λ_{sr} is the signal recycling cavity losses, λ_{pd} is the post-SRM losses on the path to the photodetectors, and $\epsilon = 2\mathcal{L}_{rt}/T_i$ is the fractional round-trip loss in the arms where $\mathcal{L}_{rt}$ is the total round-trip loss in the arms.

Taking the dot product of Eq. 3.59 by the homodyne P_{LO} with angle ζ yields the DARM model with losses

$$\frac{P_{as}}{L_-}(f) = \frac{\sqrt{2P_{LO}}}{L}\sqrt{\frac{2P_{bs}\omega_0^2}{\omega_c^2+\omega^2}}\frac{t_s e^{i\beta}\big(D_1^L \sin\zeta + D_2^L \cos\zeta\big)}{M^L} \tag{3.66}$$

$$\begin{aligned} &= \frac{\sqrt{2P_{LO}}}{L}\sqrt{\frac{2P_{bs}\omega_0^2}{\omega_c^2+\omega^2}}\frac{t_s e^{i\beta}}{M^L} \\ &\times \frac{\sqrt{1-\lambda_{pd}}}{4}\Bigg[-\left(-4+(3+e^{2i\beta})\epsilon\right)\cos\left(\zeta+\phi\right) \\ &+ \Big(-4e^{2i\beta} + \big(5e^{2i\beta}+2e^{4i\beta}+1\big)\epsilon + 2e^{2i\beta}\lambda_{sr}\Big) r_s \cos\left(\zeta-\phi\right)\Bigg] \end{aligned} \tag{3.67}$$

Eq. 3.66 is the DARM model with losses used in 3.23. The losses in the path to the PD λ_{pd} represent an overall scale factor. The losses in the arms ϵ have a very strong effect on the optical gain, and can affect the frequency of the optical spring by reducing the resonant power in the arms. Overall, arm losses act similarly to overall reduction of power on the beamsplitter.

The losses in the SRC λ_{sr} are the most interesting, because they have a frequency-dependent effect on the DARM gain. Increasing the SRC losses lowers the optical spring quality factor, broadening the resonant peak of the optical spring, Increasing the SRC losses also lowers the DARM pole frequency. Because of the frequency dependent nature of the SRC losses, these are the best-known losses. We will explore the SRC losses and DARM pole frequency in the next section.

In the end, the DARM model in Eq. 3.66 is unable to completely model the DARM plant measured at Hanford. The modeled DARM optical spring is always larger than the measured spring. Finesse simulation agrees very well with the DARM model presented, and also overestimates the DARM quality factor. Losses cannot compensate the broad spring and produce a realistic interferometer.

New efforts to model the DARM loop are underway, including exploring how angle-to-length coupling can affect the DARM plant at low frequencies [126].

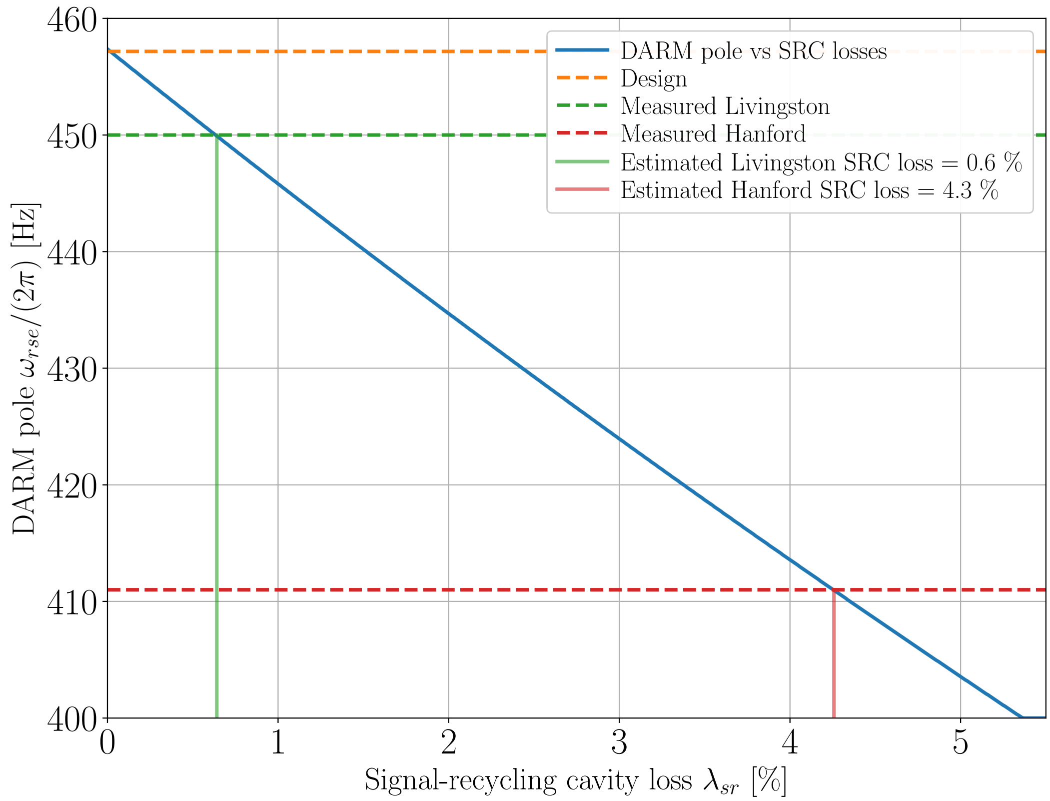

Figure 3.24: DARM pole vs signal recycling loss for LIGO Hanford in O3. Plotted in blue are the roots of setting Eq. 3.61 equal to zero, then solving for the DARM pole frequency. Livingston's DARM pole is consistent with $< 1\%$ losses in the SRC, but Hanford's SRC losses were closer to 3.5% typically.

3.6.3 DARM pole and SRC loss

The DARM pole quantifies the bandwidth of the detector sensitivity to GWs. The denominator M^L of the DARM response in Eq. 3.66 defines the DARM pole. By setting Eq. 3.61 equal to zero, we can calculate the DARM pole in response to changes in losses in the arms and SRC.

SRC losses λ_{sr} have a strong effect on the DARM pole, with excessive losses rapidly lowering the detector bandwidth. Figure 3.24 shows the DARM pole vs SRC losses, including measured DARM poles for both Hanford and Livingston. The Livingston DARM pole was 450 Hz, where the Hanford DARM pole varied between the typical 411 Hz down to 400 Hz at bad times. The corresponding Hanford losses in the SRC were from $4.5 - 5.5\%$.

Thermal compensation changes on the ITMs and SR3 were observed to have a strong effect on the DARM pole [128]. Increases in power also changed the DARM pole, which should not affect the DARM pole strongly. Heating due to power-up from 2 W to 35 W of input power reduced the DARM pole from 420 to 411 Hz over about 40 minutes [129]. This indicates mode mismatch between the arms and SRC is the most likely culprit for initial SRC losses of around 3.5%, which get exacerbated by high power operation spoiling the interferometer geometry.

Arm losses $\mathcal{L}_{rt}$ did not affect the DARM pole much, about 1 Hz of change for 1000 ppm arm losses. The detuning ϕ does affect the DARM pole, but for the levels of detuning measured ($< 1°$ from $90°$) this will not change the DARM pole significantly.

3.6.4 SR3 heater and the DARM optical spring

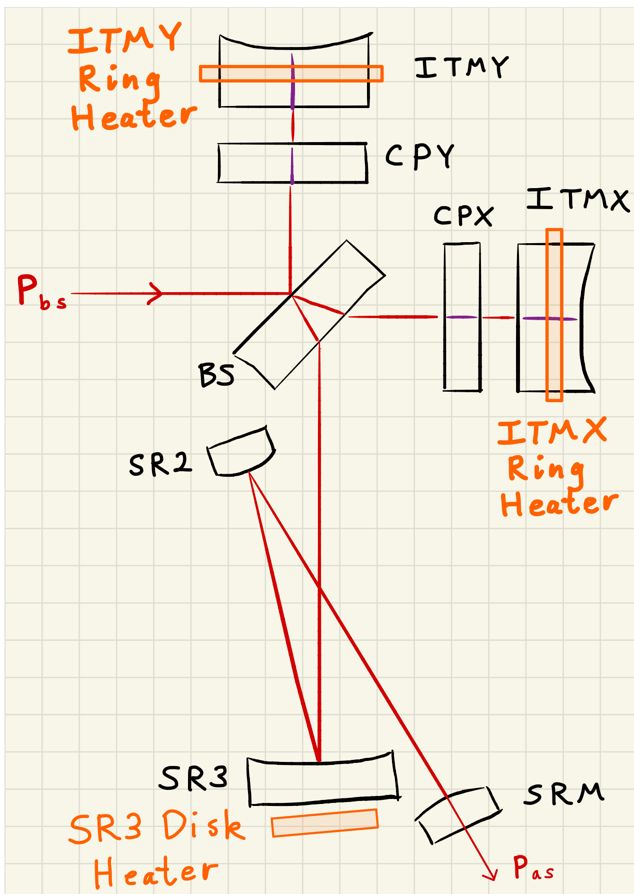

Figure 3.25: Signal recycling cavity optical diagram. The orange boxes are the thermal compensation heating elements. The ITMs have ring heaters surrounding them which can decrease their radius of curvature. SR3 has a disk heater behind it, which also decreases the radius of curvature [130]. CPX and CPY are the transparent compensation plates immediately behind the ITMs, which can lens the beam with the help of the CO2 laser.

Because the SRC is such a problematic cavity, many efforts have been taken to understand its geometry better [131–133]. However, the difficulty is the geometry difference between low and high power. The SRC itself is anti-resonant to carrier by RSE design: therefore it does not have a huge thermal load on any of its optics. The mode matching between the arms changes in lock with different arm powers, producing differential HOMs that combine at the beamsplitter and enter the SRC. If the arm modes become poorly matched to the SRC modes, it can induce excessive scattering from the carrier into HOMs.

The thermal compensation system works to combat the degradation of the optical mode matching inside the interferometer. Figure 3.25 illustrates some of the thermal compensation components. The ITM ring heaters surround the ITMs and adjust the arm mode matching by adjusting the ITM radius of curvature. The SR3 disk heater is capable to adjusting the signal recycling cavity geometry, by directly altering the radius of curvature of SR3 [130]. Not shown are the ETM ring heaters and the CO2 laser heaters incident on the compensation plates CPX and CPY.

3.6.4.1 SR3 heater on-off measurement

Prior to O3, the SR3 heater was engaged to try to improve interferometer sensitivity. The heater was turned on to maximum power of 5 W for three hours, then turned back down to 0 W, while the state of the interferometer was measured.

During this time, we took 18 DARM plant measurements. Figure 3.26 shows three of these measurements illustrating the effect on the DARM plant. Essentially, the DARM optical spring flipped sign from anti-spring, to spring, and back again with the SR3 heater off-on-off test. A small optical gain improvement is also apparent in Figure 3.26.

The optical spring flip indicates the SR3 heater is controlling the detuning ϕ of the SRC. This is interesting because, from Eq. 3.66, the DARM optical spring frequency should *not* depend on SRC losses. A lossy SRC due to e.g. high absorption, will only lower the quality factor of the spring. Therefore, something more complicated is happening due to the HOMs being created in the SRC.

3.6.4.2 Simulation of the SR3 heater

To understand how the SR3 heater changes the detuning of the SRC, we set up a Finesse/pykat simulation of Hanford which changes the SR3 radius of curvature

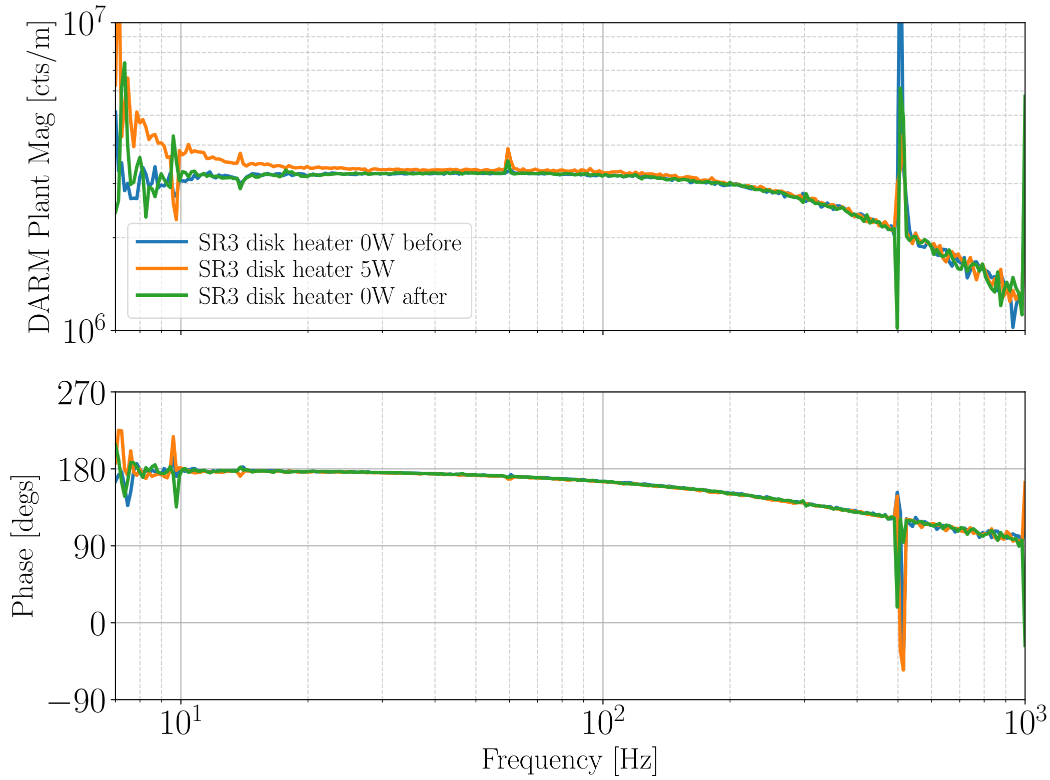

Figure 3.26: DARM optical spring during SR3 heater test in full lock at Hanford. The optical spring begins in the anti-spring state, with weak anti-springs of roughly ~ -2 Hz (blue). Flipping on the SR3 heater generated a spring of around 4 Hz after one hour of thermalization (orange). After turning off the SR3 heater, the anti-spring returned (green).

by a realistic amount in lock [134, 135].

First, we lock our simulated interferometer with similar parameters as Hanford. Next, we change the SR3 radius of curvature within the simulation while monitoring light levels everywhere in the interferometer (Figure 3.27). Then, we freeze the lock at several intervals of the radius of curvature change to sweep the length error signals (Figure 3.28). Finally, we measure the DARM plant transfer function in W/m at the same intervals of the radius of curvature change (Figure 3.29).

Briefly put, the HOMs simulated in this interferometer configuration allow for excessive HOMs to pollute the corner degrees of freedom, particularly MICH and SRCL, such that they lock with significant offsets. All corner degrees of freedom are detected at a pickoff port of the PRC (POP port). HOMs on the carrier actually reaching the POP sensor are relatively small, but HOMs on the ± 45 MHz sidebands, which carry the MICH and SRCL length info, are not small and beat with

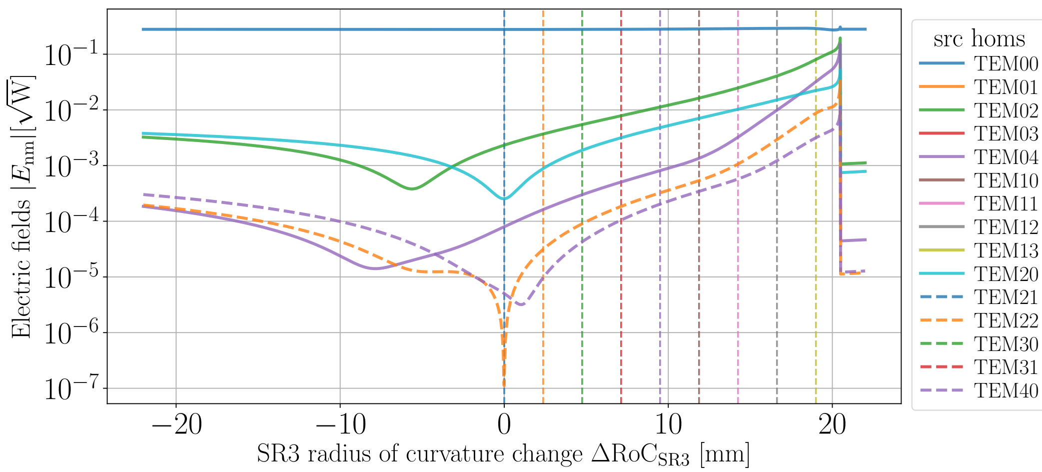

Figure 3.27: Simulation of carrier higher order modes in the signal recycling cavity with an SR3 radius of curvature change. As the SR3 radius of curvature changes, the HOMs in the cavity rise. In the case of extreme curvature, the HOMs approach the levels of the carrier, before reaching a region of lockloss. Plots for the HOMs on the 45 MHz sidebands show a similar effect. The thin dashed vertical lines represent the SR3 curvature change used for the error signal sweeps and DARM plants in Figures 3.28 and 3.29.

the HOMs on the carrier. The HOM effect on the error signal is small relative to the main beatnote signal, but the amount of offset required to induce a strong effect is not large.

When HOMs become significant enough that they affect the SRCL lock offset, the DARM cavity is detuned and an optical spring is exhibited. Figure 3.29 shows how the simulated spring varies with changed SR3 radius of curvature. A measured Hanford DARM plant is plotted for comparison to the simulation results.

This is one likely mechanism for the effect of the SR3 disk heater on the DARM optical spring. The simulation is not perfect: it can be difficult to correctly lock the interferometer and change the parameters "in-lock" while maintaining sensible length signals, we ignore the effect of astigmatism and not all parameters are exactly as measured for LIGO Hanford, for instance the arm powers are slightly low (190 kW), and the radius of curvature change induced on the SR3 in the simulation is the wrong sign compared to what actually occurs in the interferometer. However, it is extremely difficult to properly model the levels of HOMs in all interferometer cavities, and the above results support the idea that the high SRC losses are due to mode mismatch, and can be related to the issues we experience with uncontrolled detuning in the DARM cavity.

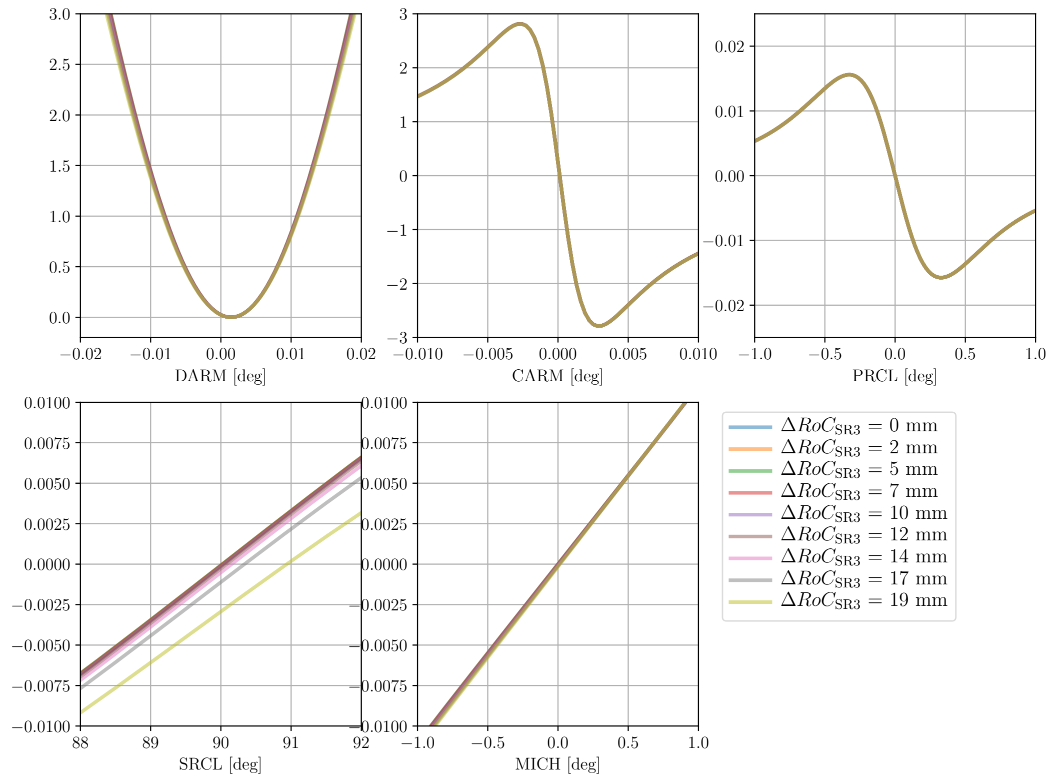

Figure 3.28: Simulation of interferometer error signals with SR3 radius of curvature change. The excessive HOMs in the SRC seen in Figure 3.27 strongly affect the SRCL error signal. The SRCL error signal is made from the 45 MHz sidebands, which carry the SRC length signal, beating with the carrier, which acts as the static reference. The HOMs on the 45 MHz and carrier also reach the sensor, and in high enough levels, spoil the error signal and induce an optical offset.

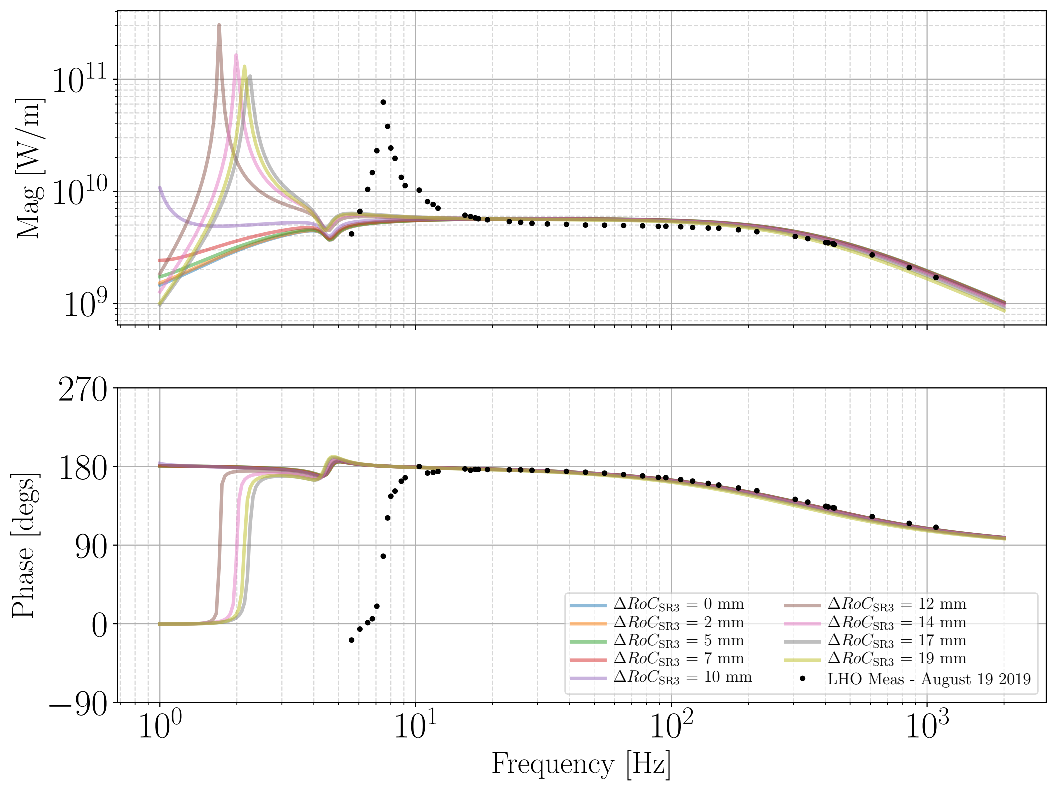

Figure 3.29: Simulation of DARM optical spring with SR3 radius of curvature change. The changing SR3 radius of curvature produces a changing DARM optical spring due to HOMs in the SRCL error signal. Also plotted is a measured Hanford DARM plant for comparison to the simulation.

Chapter 4

CALIBRATION OF THE ADVANCED LIGO DETECTORS

Calibration is the quantification of the Advanced LIGO detectors' response to gravitational waves. Calibration is the final step connecting detector data to true astrophysical strain, and influences all science done with the LIGO detectors.

Gravitational waves incident on the detectors cause phase shifts in the interferometer laser light which are read out as intensity fluctuations at the detector output. Measuring and modeling the detector response to gravitational waves is crucial to producing accurate and precise gravitational wave strain data.

The Advanced LIGO calibration group is responsible for the timely production of accurate strain data for low-latency detection, and quantifying the uncertainty in the calibrated data. The author was responsible for producing calibration uncertainty budgets for O1 and O2, including the first gravitational wave detection, GW150914 [5, 136]. The calibration uncertainty pipeline was improved with the introduction of Markov Chain Monte Carlo (MCMC) methods for fitting calibration models and Gaussian Process Regression (GPR) for quantifying unmodeled deviations. The uncertainty budget method has remained largely unchanged for O3 [94].

In this chapter, we will introduce the calibration process, motivate an accurate and precise calibration, overview the methods currently used to calibrate GW detectors, explore the calibration uncertainty pipeline, and consider future methods for even more precise and accurate calibration.

4.1 Motivation

GW signals are extremely rich sources of information from previously unexplored astrophysical phenomena. Detections from the first three observing runs have vastly altered our understanding of astrophysical binary systems [3, 4].

A miscalibration will produce biased strain data, which biases all downstream products of the data, including astrophysical parameter estimates, tests of general relativity, merger rates and GW backgrounds. Accuracy and precision in GW data is crucial for maximizing the information extracted from GW detections.

4.1.1 Effect of calibration errors on SNR

Calibration errors only effect the SNR of detectors to second order [137, 138]. Thus the likelihood of missed detections due to calibration errors is small. However, calibration error will dramatically affect astrophysics done with the detections.

4.1.2 Optimal calibration

The "optimal" calibration is accurate enough that detected GW data is not biased or limited by calibration uncertainty [138]. Any deeper accuracy of the calibration will be rendered irrelevant by the dominant source of GW data uncertainty, detector noise.

4.1.3 Astrophysical parameter estimation

Calibration model parameters and astrophysical parameters are correlated. The clearest example of this is the positive correlation of the optical gain of the interferometer and the luminosity distance of the source of the gravitational waves: given some GW signal, the larger we believe the optical gain is, the larger the luminosity distance to the GW source. Calibration parameters are considered "nuisance parameters" in the astrophysical parameter estimation process, and are marginalized over, increasing the overall uncertainty of the astrophysical parameters.

For compact binary coalescence GW signals, estimates of the progenitor masses, spins, luminosity distance, orbital plane inclination, final mass, and sky location are derived from the detected waveforms, and each are potentially limited by calibration accuracy [6, 139]. Hall et. al. explored the Cramér-Rao bounds for astrophysical parameter errors due to detector noise, and found requirements for calibration parameters' accuracy such that detector noise dominates astrophysical parameter errors [140]. Vitale et. al. have investigated the potential impact of general calibration errors on parameter estimation pipelines [141, 142]. Vitale et. al. have incorporated physical calibration parameters into the astrophysical parameter estimation pipeline [143].

For Advanced LIGO detections so far with SNRs < 100, detector noise, not calibration uncertainty, limits information from GW detections [144, 145]. As future detectors' noise decreases, for aLIGO design, Einstein Telescope [45], and Cosmic Explorer [146], some very high SNR detections will be made, revolutionizing the science possible [145]. Such high SNR detections may be limited by calibration uncertainty.

4.1.4 New astrophysics

New astrophysics is also being done with O3 detections. A recent neutron-star black-hole candidate merger emitted signature higher multipoles GWs which were detected and fit to high confidence [11]. Higher multipoles are emitted at higher frequencies than the main quadrupole moment. They are also emitted at different angles, offering another way (other than quadrupole polarization) to break the distance-inclination degeneracy problem (see Section 4.2).

The neutron star equation of state describes the properties of the matter in extreme environments like a neutron star. The equation of state affects the high frequency GW signature during merger due to the tidal deformations. Limits on the neutron star equation of state were imposed by the GW170817 BNS detection [9].

Supernovae are theorized to be powered by a “core-bounce” mechanism, which is not well-understood with theory and simulation [147]. GWs offer a way of directly observing the signature of the core-bounce, which could help inform the mechanism by which supernovae explode.

4.1.5 Other astrophysics and cosmology

The rate at which such systems form in the universe can be drawn from detected coalescence events [148–150]. Rate estimates depend on the astrophysical range of the detectors [96]. Rate estimates are particularly vulnerable to calibration errors, since rates are in units of events per time-volume, so any calibration amplitude error gets cubed. As the number of observations increases, rate estimates will become limited by strain amplitude uncertainty.

Testing general relativity has begun with the first detections [7, 151–153]. As the detectors’ sensitivity improves and there are more high signal-to-noise ratio events, calibration uncertainty will limit our test results, and calibration error will bias our test results [154, 155].

Upper limits and observations of sources of continuous gravitational waves, such as rapidly rotating neutron stars, depend on calibration uncertainty [156–158]. Upper limits and observations of the GW stochastic background of unresolvable sources, including the Big Bang, depend on the amplitude calibration uncertainty [159–161].

Using many GW detections to refine estimates of the Hubble constant will be fundamentally limited by calibration uncertainty [162, 163]. There is tension in the

Hubble constant measurement from near- and far-field electromagnetic measurement techniques [164–166]. GWs are useful for measuring the Hubble constant because the luminosity distance can be directly extracted from the detected data. If the GW is a neutron start binary coalescence accompanied by a gamma-ray burst like GW170817, then the host galaxy, and recessional velocity, can be accurately found, making "standard-siren" method of estimating the Hubble constant possible [10]. Methods of estimating the recessional velocity of GW sources using the estimated sky location and galaxy catalogs yield a so-called "dark Hubble" measurement, which requires no electromagnetic follow-up [167]. Spinning neutron start black hole coalescences may also be particularly well-suited for high accuracy distance measurements [168].

4.2 Parameter estimation, self-calibrating signals, and the standard siren

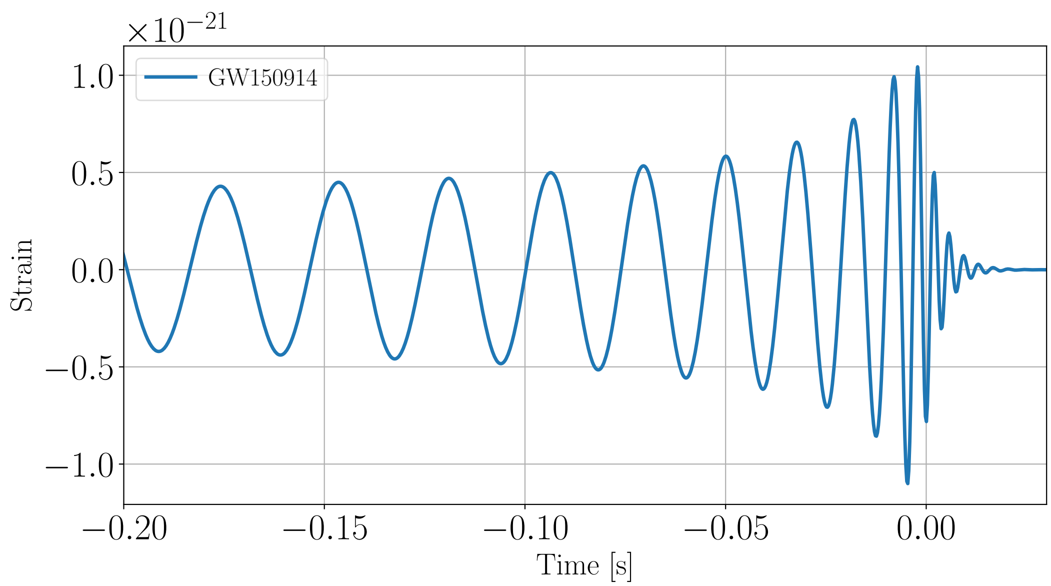

Figure 4.1: Estimated GW150914 strain time series, i.e. waveform, produced using the PhenomD waveforms calculated via PyCBC [169, 170].

A GW signal has a distinctive expected waveform. This can be used to estimate the astrophysical parameters of the binary. Figure 4.1 shows the inspiral and merger waveform detected for GW150914.

A binary merger occurs when two massive objects in orbit inspiral together, reaching relativistic velocities, and violently merge into a single massive object. The binary system emits stronger and stronger GWs as the massive objects move together, orbiting each other more and more rapidly. The increasing frequency and amplitude of the signal produce a distinctive "chirp" characteristic of binary mergers.

The key observation that GW signals are "self-calibrating" was made by Schutz [171]. From the detected GW frequency and amplitude by a network of at least three detectors, Schutz argued that the GW source parameters could be found, and the Hubble constant estimated from them, depending on the accuracy of the detector amplitude and phase measurement.

The "self-calibration" process for a simple, non-spinning GW source goes roughly as follows [172]:

1. The chirp mass $\mathcal{M}$ is found from the inspiral frequency derivative $\dfrac{df_{GW}}{dt}$.
2. The GW amplitude depends on the chirp mass $\mathcal{M}$, sky location (θ, ϕ), orbital plane inclination angle ι, and luminosity distance d_L.
3. With good phase information from the detector network, the sky location θ, ϕ can be determined via triangulation.
4. With polarization information from the detector network, the inclination angle ι may be determined.
5. The last parameter, luminosity distance d_L, is a simple overall scaler to the measured GW amplitude.

From this process, the luminosity distance naturally falls out of the detection parameters without any astrophysical distance calibration. The term *standard siren* refers to binary inspiral's consistent frequency and amplitude dependence on the chirp mass $\mathcal{M}$.

The above process *does* depend on the detector network calibration for accurate GW phase and amplitude information. A phase miscalibration will throw off the sky location triangulation. An amplitude miscalibration will throw off the luminosity distance estimate. This makes an accurate detector calibration imperative for parameter estimation.

In reality, parameter estimation is much more complicated than the picture given above, especially for spinning, precessing binary systems with unequal masses [139, 173]. However, the detector calibration will always represent a fundamental limit to the accuracy of astrophysical parameter estimation.

4.3 Detector calibration fundamentals

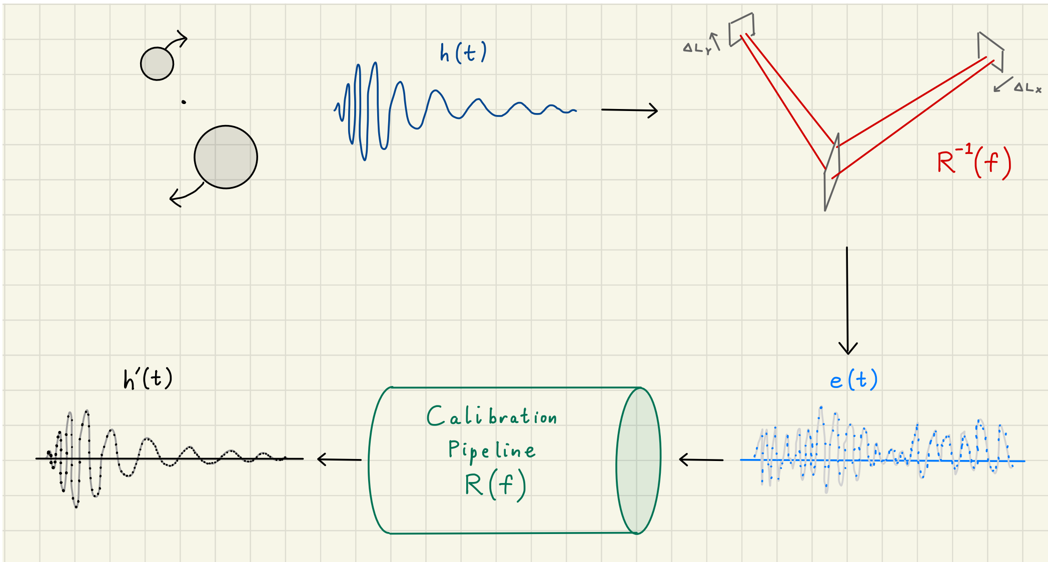

Figure 4.2: Simplified calibration process following the conversion of astrophysical strain $h(t)$ into strain data $h'(t)$.

Figure 4.2 shows the fundamental calibration process. A binary merger produces strong gravitational waves with strain $h(t)$. The detector responds to the incident gravitational wave with response R^{-1}, producing a raw signal $e(t)$. The calibration pipeline takes the raw signal and calibrates it into strain data $h'(t)$, using a measured response function R.

Calibration is the process of measuring the response function R as accurately as possible. Because we cannot generate known terrestrial gravitational waves to calibrate the detector, we use DARM motion according to the relation

$$L_{-} = hL \tag{4.1}$$

where L_{-} is DARM motion, h is incident GW strain, and L is the length of the arms. Subsection B.3.2 derives the GW to DARM transfer function for a simple Michelson.

4.3.1 DARM control loop

DARM is one of the principle degrees of freedom of the interferometer, and must be held on resonance to produce useful data DARM is unstable without a feedback control loop, which suppresses all DARM motion sensed on the DCPDs, include those from GWs. The feedback is routed through the quadruple pendulum position actuators, both the magnetic coil drivers on the upper stages, and the electrostatic drive on the lowest stage. Figure 4.3 shows the interferometer layout and a quadruple pendulum.

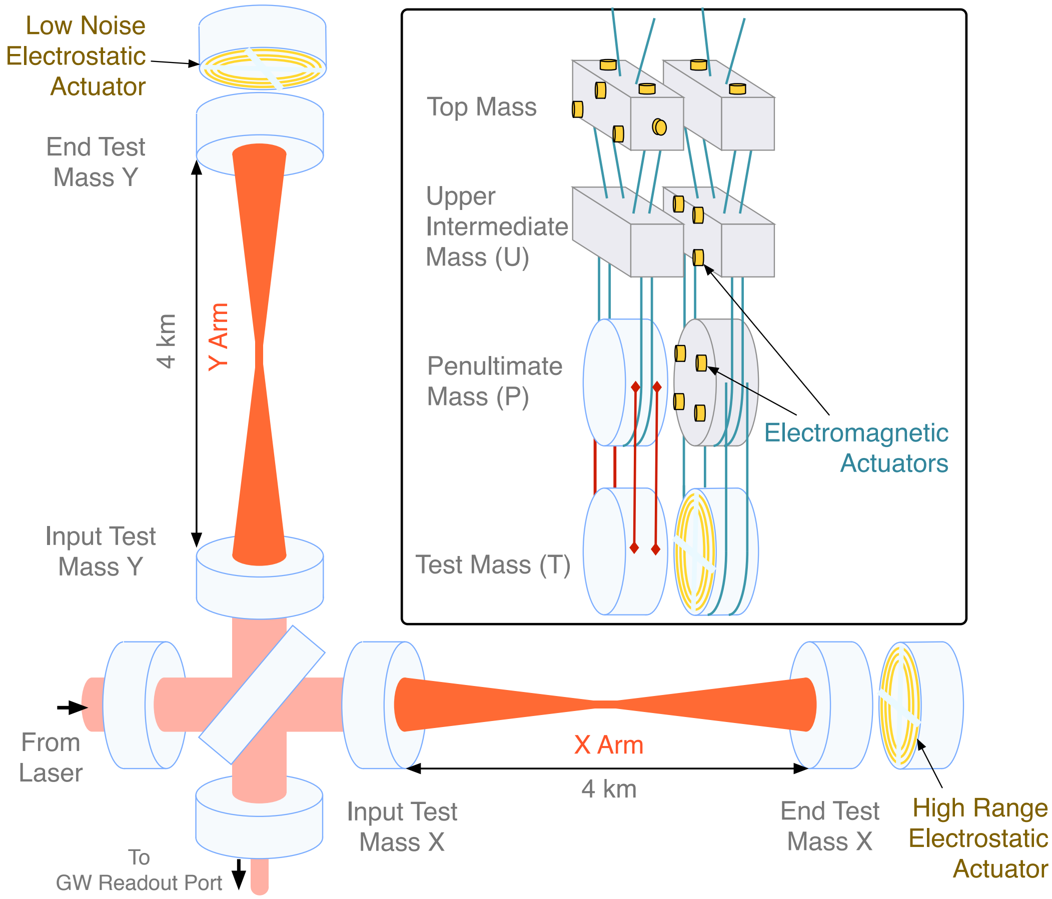

Figure 4.3: Simplified interferometer layout and one of the quadruple pendulum suspension systems for the core optics. For the upper stages of the pendulum, electromagnetic coil drivers are used for length and angular control. An electrostatic drive (ESD) is used to control the test mass position itself [26]. Only one ETM ESD is turned on to control DARM in low-noise lock.

We define three independently quantifiable transfer functions of the DARM control loop, shown schematically in Figure 4.4. The sensing function $C = d_{\mathrm{err}}/\Delta L_{\mathrm{res}}$ defines the measured laser power response to DARM displacement, as well as the data acquisition process, to form the digital error signal d_{err}. Digital filters $D = d_{\mathrm{ctrl}}/d_{\mathrm{err}}$ invert the suspension compliance and shape the loop control signal. The actuation function $A = \Delta L_{\mathrm{ctrl}}/d_{\mathrm{ctrl}}$ moves the optic to cancel any detected DARM displacement within the DARM loop bandwidth.

All transfer functions are complex-valued functions of frequency, with quantifiable magnitude and phase. The digital filters D shape the DARM loop frequency response and are known to negligible uncertainty. The DARM loop transfer functions C and A must be measured and modeled in the frequency domain between 5 and 5000 Hz. Both C and A contribute to the total calibration uncertainty budget.

4.3.2 Calibration pipeline

The error and control signals $d_{\text{err}}, d_{\text{ctrl}}$ are digitally filtered to form a time-series estimate of the GW strain $h(t)$ used for astrophysical searches. The digital filters applied to d_{err} and d_{ctrl} are constructed from models of the sensing function $C^{(\text{model})}$ and actuation function $A^{(\text{model})}$:

$$h = \frac{1}{L}\left[\frac{1}{C^{(\text{model})}} * d_{\text{err}} + A^{(\text{model})} * d_{\text{ctrl}}\right], \tag{4.2}$$

where $*$ indicates convolution in the time domain. The accuracy and precision of the models $C^{(\text{model})}$ and $A^{(\text{model})}$ define the systematic error and statistical uncertainty in the estimated time series $h(t)$.

4.3.3 Response function

We define a transfer function called the response function R,

$$h = R * d_{\text{err}} = \frac{1}{L}\left(\frac{1+G}{C}\right) d_{\text{err}} \tag{4.3}$$

where the DARM open loop gain $G = C * D * A$. Eq. 4.3 illustrates that in the frequency domain, response function error δR is equivalent to the GW strain data error δh and response function uncertainty σ_R is equivalent to the GW strain data uncertainty σ_h. The response error and uncertainty relative to the calibration pipeline model $R^{(\text{model})}$ are quantified as a function of frequency f with time dependence t:

$$\frac{\delta R(f,t)}{R^{(\text{model})}} = \frac{\delta h(f,t)}{h}, \qquad \frac{\sigma_R(f,t)}{R^{(\text{model})}} = \frac{\sigma_h(f,t)}{h}. \tag{4.4}$$

4.3.4 Systematic errors

The values of C and A can drift slowly over time, giving functions of frequency that vary in time $C(f,t)$ and $A(f,t)$. However, our online calibration pipeline digital filters $1/C^{(\text{model})}$ and $A^{(\text{model})}$ are not perfect representations of our understanding of the interferometer. This leads to known systematic errors in our $h(t)$ reconstruction, governed by the sensing and actuation systematic errors $\delta C(f,t)$ and $\delta A(f,t)$. The systematic errors relative to $C^{(\text{model})}$ and $A^{(\text{model})}$ are quantified as

$$\frac{\delta C(f,t)}{C^{(\text{model})}} = \frac{C(f,t)}{C^{(\text{model})}} - 1, \qquad \frac{\delta A(f,t)}{A^{(\text{model})}} = \frac{A(f,t)}{A^{(\text{model})}} - 1, \tag{4.5}$$

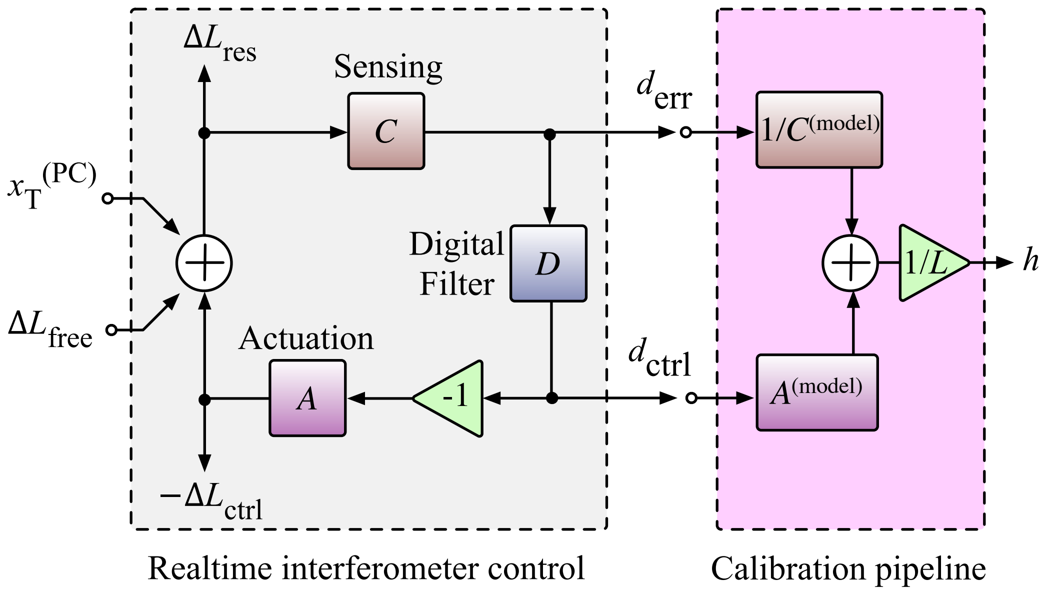

Figure 4.4: The DARM control loop is shown in the grey box on the left. The sensing plant C produces the error signal d_{err} in linear response to residual differential arm motion ΔL_{res}. The digital filters D shape the error signal d_{err} into a control signal d_{ctrl}. Displacement noise from any external source enters the loop as ΔL_{free}. The test mass excitation via the photon calibrator $x_{\mathrm{T}}^{(\mathrm{PC})}$ displaces the test mass above the DARM noise by a precisely known amount. The actuation plant A takes the control signal d_{ctrl} and actuates on the optics by ΔL_{ctrl} to maintain DARM resonance. The pink box on the right shows the calibration pipeline, consisting of an inverse sensing model $1/C^{(\text{model})}$ and actuation model $A^{(\text{model})}$. The output of the calibration pipeline is GW strain data $h(t)$.

where $C(f,t)$ and $A(f,t)$ represent the measured sensing and actuation transfer functions.

Systematic errors δC and δA propagate forward to the relative response function systematic error $\delta R/R^{(\text{model})}$:

$$\begin{aligned}
\frac{\delta R(f,t)}{R^{(\text{model})}} &= \frac{R(f,t)}{R^{(\text{model})}} - 1 = \left(\frac{1+G(f,t)}{C(f,t)}\right) \Bigg/ \left(\frac{1+G^{(\text{model})}}{C^{(\text{model})}}\right) - 1 \\
&= \frac{\left(G^{(\text{model})}\,\dfrac{\delta A(f,t)}{A^{(\text{model})}} - \dfrac{\delta C(f,t)/C^{(\text{model})}}{1+\delta C(f,t)/C^{(\text{model})}}\right)}{1+G^{(\text{model})}}.
\end{aligned} \tag{4.6}$$

4.3.5 Uncertainty

In general, any Gaussian-noise based transfer function follows a joint 2D probability distribution (see Section D.8, Eq. D.32). Eq. D.32 is not necessarily a Gaussian itself, but approaches one in a certain regime of high coherence γ^2 and high number of averaged n.

All calibration transfer functions are swept-sines taken with very high coherence $\gamma^2 > 0.95$, usually $\gamma^2 > 0.99$ over most of the bandwidth, with plenty of averages $n \approx 25$ to ensure the uncertainty of the measurement is squarely Gaussian.

Therefore we can safely use the transfer function uncertainty from Bendat and Piersol Table 9.6 for relative magnitude uncertainty $\sigma^2_{|\hat{H}|/|H|}$ and absolute phase uncertainty $\sigma^2_{\phi_{\hat{H}}}$: [174]

$$\sigma^2_{\frac{|\hat{H}|}{|H|}}(f) = \sigma^2_{\phi_{\hat{H}}} = \frac{1-\gamma^2(f)}{2n\gamma^2(f)} \tag{4.7}$$

4.4 Models

In this section we will overview the models used in the calibration procedure and uncertainty budget creation.

4.4.1 Calibration group DARM model

The calibration group sensing function C finds its basis in the Buonanno and Chen DARM optical plant explored in Section 3.6, Eq. 3.57. However, this function is complex to fully model in real time, and can be simplified to a poles and zeros model without sacrificing much accuracy. The algebra of this simplification calculated in [57] Appendix D, and summarized in [140] Eq. (6), reproduced here:

$$C(f) = \frac{g e^{-2\pi i f L/c}\left(1 + i\dfrac{f}{z}\right)}{1 + i\dfrac{f}{Q_p|p|} - \dfrac{f^2}{|p|^2} - \dfrac{\xi^2}{f^2}} \tag{4.8}$$

where g is the optical gain in $\mathrm{W/m}$, z is the homodyne zero, p is the complex DARM pole with magnitude $|p|$ and quality factor Q_p, and ξ is the spring frequency, related to the phase to amplitude factor κ from Eq. 3.58:

$$z = f_c \frac{\cos(\phi+\zeta) - r_s\cos(\phi-\zeta)}{\cos(\phi+\zeta) + r_s\cos(\phi-\zeta)} \tag{4.9}$$

$$p = f_c \frac{1 - r_s e^{2i\phi}}{1 + r_s e^{2i\phi}} \tag{4.10}$$

$$\xi^2 = f_c^2 \frac{P_{bs} c}{2\pi^3 \lambda_0 f_c^4 m L^2} \frac{2 r_s \sin(2\phi)}{1 - 2r_s\cos(2\phi) + r_s^2} \tag{4.11}$$

where $f_c = \omega_c/(2\pi)$ is the arm pole in Hz.

In the end, we can put Eq. (6) of [140] into a simple poles and zeros form by using the known resonant sideband extraction parameters of the homodyne angle $\zeta =$

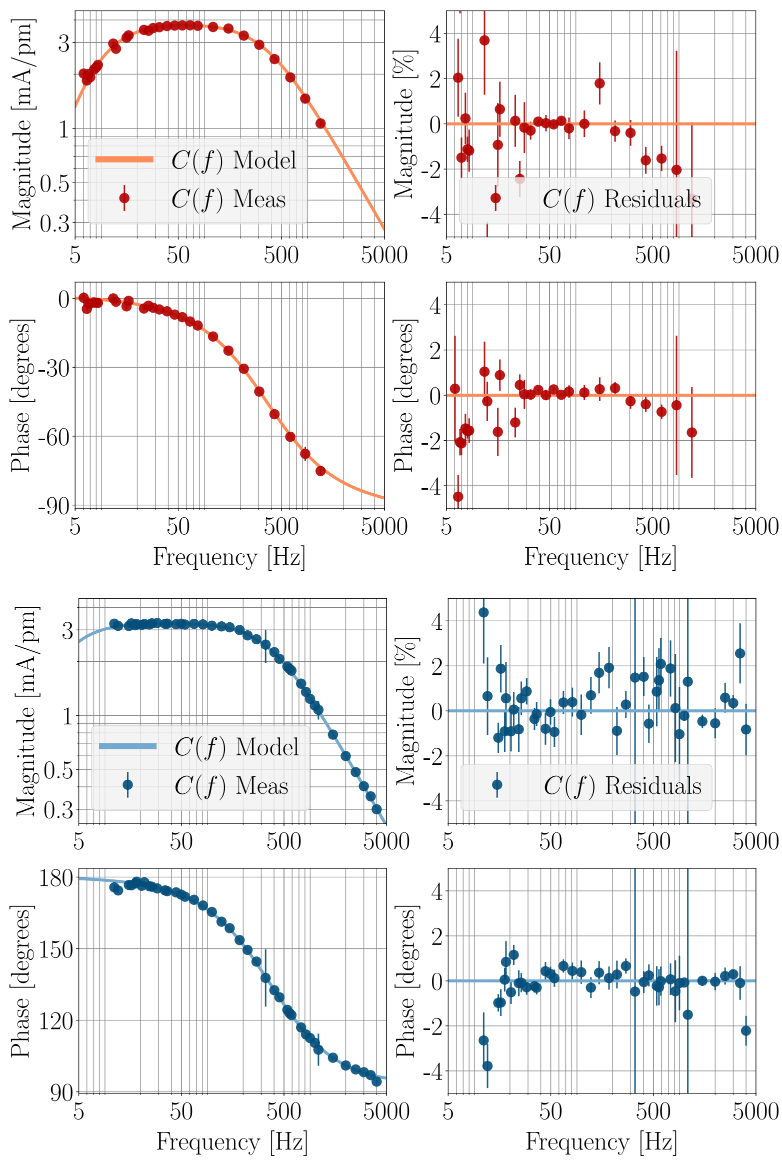

Figure 4.5: Sensing measurements vs sensing model $C(f, t, \vec{\lambda}_C)$ in Eq. 4.13, and their residuals $\delta C(f, t)/C^{(\mathrm{model})}(f, t, \vec{\lambda}_C)$. The Hanford sensing reference measurement from January 4th, 2017 is shown in the four panels in red. The Livingston sensing reference measurement from November 26th, 2016 is in the four panels in blue. The model parameters $\vec{\lambda}_C$ were found via an MCMC. Physically, the magnitude Bode plots represent how many milliamps of current are generated at our transimpedance photodetector per picometer of differential arm motion from 5 to 5000 Hz. The drop in sensitivity at low frequencies shows the effect of anti-spring detuning at both detectors. The 180 degree phase difference between Hanford and Livingston is a sign convention difference between the detectors, most likely from the DARM offset sign.

$\pi/2$ and detuning $\phi \approx \pi/2$, which sets $|p| \approx z$ and $Q_p \approx 1/2$. This cancels the homodyne zero z with one of the DARM poles p, leaving only a factor of $1+if/f_{rse}$ in the denominator.

The above approximation is known as the *single pole approximation*, and refers to the fact that the DARM pole f_{rse} is enough to describe the high-frequency DARM plant dynamics in our current operating scheme. The single pole approximation also ignores the repeating resonances associated with the FSR at extremely high frequencies $f > 10\,kHz$, which is outside the detection bandwidth of the detector.

The spring frequency ξ^2 is strongly dependent on the detuning ϕ. If $\phi = \pi/2$ exactly, then $\xi^2 = 0$. If $\phi > \pi/2$, then $\xi^2 < 0$ and DARM exhibits the optical anti-spring. If $\phi < \pi/2$, then $\xi^2 > 0$ and DARM exhibits the optical spring. We note that this tuning is the opposite sign of the usual two-mirror cavity optical spring tuning [175]: the longer cavity in the SRC $\phi > \pi/2$ produces an anti-spring, not a spring. This is because the SRC is anti-resonant.

In the approximation that $\xi^2 \ll |p|^2$, the terms in the denominator of Eq. 4.8 with $|p|$ become zero. This approximation leaves

$$C \approx \frac{f^2}{f^2 - \xi^2}, \tag{4.12}$$

which can be described as two zeros at $0\ Hz$ and two poles at $|\xi|$. The poles are purely real if $\xi^2 < 0$, and complex if $\xi^2 > 0$.

From these approximations, we simplify Eq. 4.8 further to the calibration group sensing model:

$$C^{(\text{model})}(f, t, \vec{\lambda}_C) = \kappa_C(t) \frac{H_C e^{-2\pi i f \tau_C}}{\left(1 + i\dfrac{f}{f_{rse}}\right)} \frac{f^2/f_s^2}{\left(1 - \dfrac{f^2}{f_s^2} + i\dfrac{f}{f_s Q_s}\right)} \tag{4.13}$$

where H_C is the optical gain in $\mathrm{cts/m}$, τ_C is the delay constant dominated by the light propagation in the arms, f_{rse} is the single DARM pole, f_s is the DARM spring frequency (equal to ξ), Q_s is the DARM spring quality factor, and $\kappa_C(t)$ is the time dependent optical gain factor monitored via calibration lines. f_{rse}, f_s, and Q_s are also monitored. Both f_s^2 and f_s have the same sign, and can be positive or negative according to 4.11.

These parameters are collected into the sensing function parameter vector $\vec{\lambda}_C$:

$$\vec{\lambda}_C = \begin{pmatrix} H_C & f_{rse} & \delta\tau_C & f_s & Q_s^{-1} \end{pmatrix}^T \tag{4.14}$$

Table 4.1: Hanford (left) and Livingston (right) sensing function model parameters $\vec{\lambda}_C$ MCMC fit values and uncertainties for O2. The fits were performed on Hanford's January 4th, 2017 reference measurement and Livingston's November 26th, 2016 reference measurement. The model corresponding to these parameters can be seen in Figure 4.5. The corner plot showing the MCMC results from the Hanford reference measurement is shown in Figure 4.7.

Hanford Parameters	Variable	Value$^{+1\sigma}_{-1\sigma}$	Units
Optical Gain	H_C	$3.834^{+0.003}_{-0.003}$	mA/pm
Coupled Cavity Pole	f_{CC}	360^{+2}_{-2}	Hz
Time Delay	$\delta\tau_C$	$0.6^{+1.3}_{-1.3}$	μs
Optical Spring Frequency	f_S	$6.87^{+0.03}_{-0.03}$	Hz
Optical Spring Inverse Q	Q_S^{-1}	$0.034^{+0.004}_{-0.004}$	-
Livingston Parameters	**Variable**	**Value$^{+1\sigma}_{-1\sigma}$**	**Units**
Optical Gain	H_C	$3.288^{+0.007}_{-0.007}$	mA/pm
Coupled Cavity Pole	f_{CC}	$369.5^{+1.0}_{-0.9}$	Hz
Time Delay	$\delta\tau_C$	$0.84^{+0.13}_{-0.13}$	μs
Optical Spring Frequency	f_S	$2.6^{+0.2}_{-0.2}$	Hz
Optical Spring Inverse Q	Q_S^{-1}	$0.005^{+0.009}_{-0.004}$	-

Figure 4.5 shows the results of the fit of the model in Eq. 4.13 to the reference measurement in O2. Table 4.1 shows the parameters used in Figure 4.5.

4.4.1.1 Sensing systematic errors

Our model of the sensing function $C^{(\text{model})}(f, t, \vec{\lambda}_C)$ is an approximation. The true detector sensing function changes over time and deviates from the sensing model at high frequencies. The sensing model dynamically corrects for $\kappa_C(t)$ with real-time measurement. However, f_{rse}, f_s, and Q_s^{-1} are also changing in time, but were not corrected for in the model in O1 or O2.

The time dependence in f_{rse} was included in the calibration uncertainty budget as a known systematic error, since it was tracked via real-time measurement but could not yet be dynamically corrected for in the model. The time dependence in f_s and Q_s^{-1} results in expanded uncertainty at low frequency. The total systematic error in the sensing function for O1 and O2, $\delta C(f,t)$, was

$$\frac{\delta C(f,t)}{C^{(\mathrm{model})}} = \left(\frac{1 + if/f_{rse}}{1 + if/f_{rse}(t)}\right) \frac{\delta C^{GP}(f)}{C^{(\mathrm{model})}} \, e^{-2\pi i f \delta\tau_C}. \tag{4.15}$$

The first term is the explicit correction for time dependence of the coupled cavity pole, $f_{rse}(t)$. A correction time delay factor $\delta\tau_C$ modifies the original time delay τ_C included in the model.

Further systematic errors originate from the uncorrected time dependence of f_s and Q_s^{-1} or additional unknown systematic errors. Any remaining frequency dependent systematic errors are covered by a Gaussian Process regression $\delta C^{GP}(f)$. Quantifying errors $\delta C^{GP}(f)$ is explained further in Section 4.7.

4.4.2 Long wavelength approximation

Implicit in the DARM modeling everywhere in this thesis is the *long-wavelength approximation*. The long-wavelength approximation assumes that the size of the detector is much less than size of the GW wavelength, $L \ll \lambda_{GW}$. This assumption implies that the GW is "in-phase" across the entire detector, with no complex dynamics due to different phases of GW being incident on different parts of the detector. The long-wavelength approximation is good for low frequency GWs, but at very high frequencies the wavelength approaches the size of the detector and the full treatment is required [176–178].

The total response of the interferometer to GWs is sky location dependent. The systematic errors from the long-wavelength approximation tend to be partially canceled out by the systematic errors from the simple-pole approximation for sensitive sky locations [178]. For a Fabry-Perot Michelson, [178] finds that the approximated response differs from the full response by around 2-3% at 1.2 kHz for reasonable sky locations.

4.4.3 Actuation Model

The Advanced LIGO test masses are suspended via quadruple cascaded pendula [26]. Each suspension stage has independent actuators, as shown in Figure 4.3. The control signal, d_{ctrl}, is digitally distributed as a function of frequency to each stage's actuators via a digital-to-analog converter and signal processing electronics to create the control displacement, ΔL_{ctrl}.

The distribution filters are designed taking into account all actuators' authority to displace the test mass. On the upper intermediate and penultimate stage, the

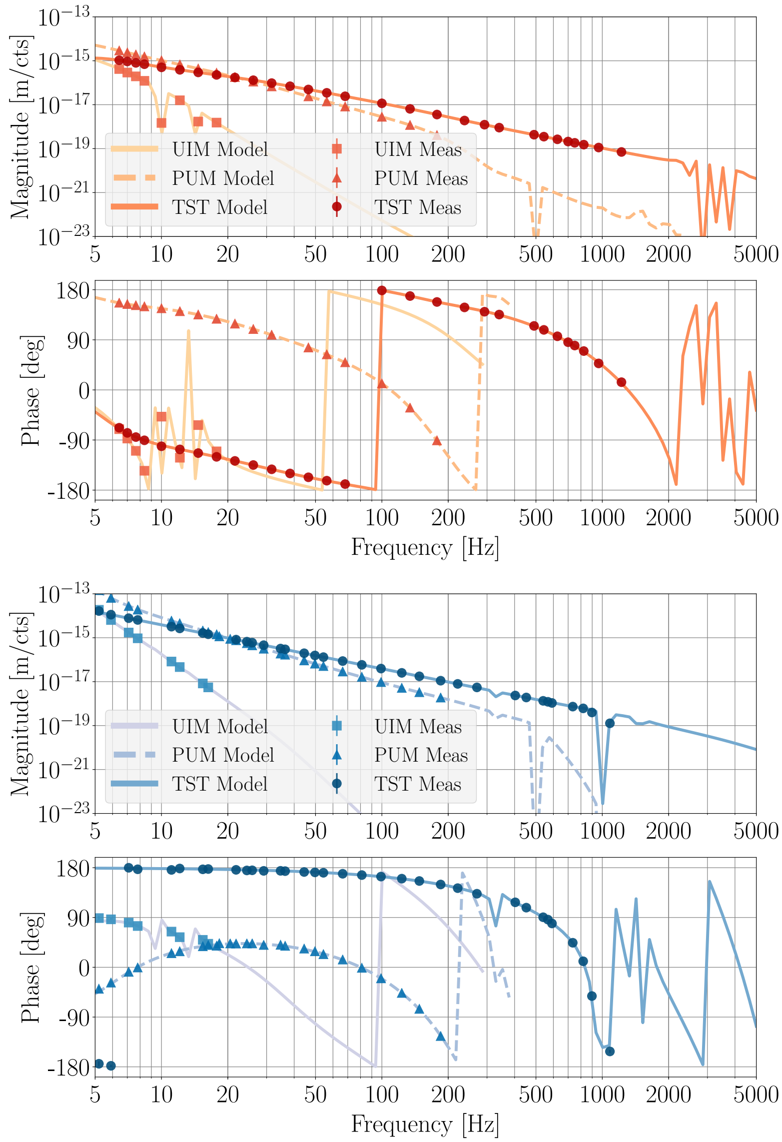

Figure 4.6: Actuation stage measurements and models $[H_i\, A_i(f, \vec{\lambda}_i)]^{(\mathrm{model})}$. Each index i is one of the actuation stages U, P, or T. The Hanford actuation reference measurements from January 4th, 2017 are shown in the two left plots in red. The Livingston actuation reference measurements from November 26th, 2016 are in the two right panels in blue. The model parameters $\vec{\lambda}_A$ for $A_i(f, \vec{\lambda}_i)$ have been found via MCMC. The actuation strength magnitude is in units of meters per d_{ctrl} count. Notches seen in the magnitude plot are purposefully placed to avoid ringing up suspension violin modes at specific frequencies. Each stage's phase is stable for frequencies at which that actuation stage dominates, but then rolls rapidly as it loses authority at high frequencies. For this reason, the UIM and PUM stage phase plots are cut off at 300 Hz and 400 Hz respectively.

digital-to-analog converter drives electromagnets on the reaction stage creating a force on magnets attached to the suspended stage. On the test mass stage, the digital-to-analog converter drives an electrostatic system which creates a force, quadratic in the applied potential, via dipole-dipole interactions between the test mass and a pattern of electrodes on the reaction mass (see Figure 4.3). With a large bias voltage and low control voltage, the requested actuation forces on the electrostatic system are in the linear regime.

The sum of the paths the digital control signal, d_{ctrl}, takes through each stage to displace the test mass, ΔL_{ctrl}, makes up our total actuation model:

$$\begin{aligned} A^{(\text{model})}(f,t,\vec{\lambda}_A) \;=\; & \Big[\kappa_T(t)\, F_T(f)\, H_T\, A_T(f) \\ & +\; \kappa_{PU}(t)\Big(F_P(f)\, H_P\, A_P(f) \\ & +\; F_U(f)\, H_U\, A_U(f)\Big)\Big]\, e^{-2\pi i f \tau_A} \end{aligned} \tag{4.16}$$

where U, P, and T represent the three stages used for control; the upper-intermediate, penultimate, and test mass stages, respectively. Each stage is composed of the normalized electro-mechanical frequency response of the pendulum and its actuators, $A_i(f)$, the digital distribution filter, $F_i(f)$, a scale factor, H_i, and an overall digital delay, τ_A, defined by the common computational delay from each stage. The model time delay τ_A is 45 μs for Livingston and 61 μs for Hanford. $\kappa_{PU}(t)$ is the time dependence of the penultimate and upper intermediate scale factor, and $\kappa_T(t)$ is the time dependence of the test mass scale factor, as calculated in [179].

The penultimate and upper intermediate scale factor $\kappa_{PU}(t)$ is not expected to vary much over time, as it represents the change in the electromagnetic coil actuators' strength. The test mass scale factor $\kappa_T(t)$ does vary significantly over time as the electric charge on the test mass builds up, changing the actuation strength of the electrostatic drive.

The reference scale factor for each stage, H_i, collects scale factors from that of the digital-to-analog converter in V/cts, each stage's drive electronics in A/V or V/V, the actuator itself in N/A or N/V depending on the stage, and the compliance of the suspension in m/N. Time delay correction factors for each stage $\delta\tau_i$ are extracted from measurements as stage-specific corrections to the overall actuation delay τ_A. The electro-mechanical transfer functions, A_i, for each stage are independently measured and included in the model with negligible uncertainty. Remaining scale

Table 4.2: Hanford (left) and Livingston (right) actuation function model parameters $\vec{\lambda}_A$ MCMC fit values and uncertainties for O2. The fits were performed on Hanford's January 4th, 2017 reference measurements and Livingston's November 26th, 2016 reference measurement. The models corresponding to these parameters can be see in Figure 4.6. To get from Newtons/count units in this table to meters/count in Figure 4.6, we multiply by the suspension models which have units of m/N and are known to negligible uncertainty.

Hanford Parameters	Variable	Value$^{+1\sigma}_{-1\sigma}$	Units
Upper Intermediate Gain	H_U	$8.205^{+0.004}_{-0.004} \times 10^{-8}$	N/cts
Upper Intermediate Delay	$\delta\tau_U$	57^{+45}_{-46}	μs
Penultimate Gain	H_P	$6.768^{+0.002}_{-0.002} \times 10^{-10}$	N/cts
Penultimate Delay	$\delta\tau_P$	$0.4^{+0.6}_{-0.6}$	μs
Test Mass Gain	H_T	$4.3573^{+0.0008}_{-0.0008} \times 10^{-12}$	N/cts
Test Mass Delay	$\delta\tau_T$	$2.8^{+0.4}_{-0.4}$	μs
Livingston Parameters	**Variable**	**Value$^{+1\sigma}_{-1\sigma}$**	**Units**
Upper Intermediate Gain	H_U	$7.24^{+0.03}_{-0.03} \times 10^{-8}$	N/cts
Upper Intermediate Delay	$\delta\tau_U$	102^{+56}_{-56}	μs
Penultimate Gain	H_P	$6.41^{+0.02}_{-0.02} \times 10^{-10}$	N/cts
Penultimate Delay	$\delta\tau_P$	$-8.7^{+6.2}_{-6.1}$	μs
Test Mass Gain	H_T	$2.513^{+0.004}_{-0.004} \times 10^{-12}$	N/cts
Test Mass Delay	$\delta\tau_T$	$-4.5^{+1.4}_{-1.4}$	μs

factor and delay parameters dominate the actuation function uncertainty, and are thus collected in the set of actuation parameters:

$$\vec{\lambda}_A = \begin{pmatrix} H_U & \delta\tau_U & H_P & \delta\tau_P & H_T & \delta\tau_T \end{pmatrix}^T. \tag{4.17}$$

The values of these reference parameters λ_A are found in Table 4.2.

Figure 4.6 plots the measured vs modeled actuation functions for every stage of the quad pendulum used for control. Table 4.2 gives the MCMC fit parameters used in Figure 4.6.

4.4.3.1 Actuation systematic errors

The digital filters, F_i, are known a priori, and time-dependent corrections κ_{PU} and κ_T are dynamically corrected for when estimating $h(t)$. The remaining components of the actuation stage model, $[H_i\, A_i]^{(\text{model})}(f, \vec{\lambda}_i)$, may contain systematic errors. We allow for and quantify systematic errors in each actuation stage as

$$\frac{\delta A_i(f)}{A_i^{(\text{model})}} = \frac{\delta A_i^{GP}(f)}{A_i^{(\text{model})}}\, e^{-2\pi i f \delta\tau_i} \tag{4.18}$$

where $\delta\tau_i$ is a time delay phase error on each stage, and $\delta A_i^{GP}(f)$ is the systematic error in scale or frequency dependence from the Gaussian process regression done on each stage's measurement residuals. Systematic error calculations are explained fully in Section 4.7.

4.4.4 Calibration lines and time-dependent factors

The detector is known to vary with time, as losses, alignment, and thermalization of the interferometer affect the response. To capture the time dependence of the calibration during a run, *calibration lines* are applied to the detectors during all observation times. A calibration line is a single-frequency excitation applied to the detector via the photon calibrator and suspension actuators. Using four calibration lines, we are able to capture changes in the detector calibration and partially correct for them in real time.

The calibration lines' response to the applied excitation is recorded in the detector readout d_{err}. These transfer functions are recast into each time dependent parameter, κ_T, κ_{PU}, κ_C, f_{rse}, f_s, and Q_s. The calibration lines are driven with high SNR such that the time-dependent parameter uncertainties are small relative to the parameter values. The calculation of the time-dependent parameters from calibration lines is derived in [179].

4.5 Photon calibrator

The photon calibrator (PCAL) is the ultimate reference for strain data in the interferometer [180, 181]. The laser from two 1047 nm auxiliary lasers are reflected off the end test masses at both sites. The laser intensity is modulated at audio signal frequencies using an AOM, creating a fluctuating radiation pressure force. Radiation pressure pushes on the test masses, creating a true displacement far above DARM sensitivity.

The power incident on the test masses is recorded via two photodiodes calibrated to integrating spheres, one after the AOM before transmission onto the test mass, the other upon reflection off of the test mass. Each photodiode's readout is then digitally recast as a displacement, $x_T^{(PC)}$, which is the amount of PCAL-induced displacement contributing to ΔL_{free}. The full suspension dynamics are incorporated into the transfer function from the PCAL power modulation to the test mass length modulation, giving an accurate frequency response at and below the suspension resonant frequency.

The relative PCAL actuation strength correction factor, $H_{PCAL}(t)$, tracks the actuation strength of the PCAL over time. $H_{PCAL}(t)$ has a value of 1 during times of no clipping, and a value less than 1 during times of clipping. $H_{PCAL}(t)$ has a relative uncertainty of 0.79% over all time. This will affect our total calibration uncertainty budget directly in Section 4.7.

4.5.1 Systematic errors

The effect of the photon calibrator on the test mass is ultimately dead-reckoned from the intensity measurement via integration sphere. The intensity measurements are as accurate and precise as its possible for a laser intensity measurement to be, with precision of around 0.5%. The entire setup also adds uncertainty, yielding a total uncertainty on the PCAL of 0.79% for O1 and O2 [180]. For O3, this number was reduced to 0.41%, largely due to a reduction in the claimed laser power measurement uncertainty from NIST and measured temperature correction factors [181]

The problem with achieving better precision with the photon calibrator is the setup is prone to small systematic fluctuations. Many fluctuations in the photon calibrator are possible. Clipping on the input or output beam is the most serious and common, torquing the test mass instead of longitudinal displacement from PCAL misalignment, temperature changes on the laser change the output laser power, shaking of the PCAL laser or steering mirrors, saturated PCAL intensity servo loops, and drifting alignment onto the test mass are possible. All affect the ultimate calibration of this reference, and some are monitored and corrected for in real time.

Some errors are not possible to know. Moving the integrating spheres between labs is known to change their responsivity ratio [181]. The PCAL laser can introduce elastic deformations on the test mass, which can affect the calibration accuracy above 1 kHz. Dead-reckoning the laser power incident on the test mass is difficult to do with the accuracy and precision required.

A photon calibration team at Hanford monitors the photon calibrator closely to ensure the accuracy of the strain data. Regular maintenance is performed on the photon calibrator, including small alignment adjustments. Any noticeable changes are corrected quickly, times of the change are noted and incorporated into the uncertainty budget.

4.5.2 Other calibration methods

Checks of gross systematic errors in the photon calibrator system have been performed using the free-swinging Michelson and VCO calibration methods in O1 and O2. These found agreement with the photon calibrator to within 10% [136].

A Newtonian calibrator (NCAL) prototype was installed at Hanford prior to O3. The NCAL is a heavy wheel with weights arranged in a quadrupole and hexapole, which spins at 20 Hz near the test mass. This creates a fluctuating gravitational potential with different distance dependence (the quadrupole falls like $1/d^2$, the hexapole like $1/d^3$). This may provide another check of the photon calibrator accuracy.

4.6 Measurements

In this section, we explore how the DARM loop components C and A_i are measured.

The DARM model functions $C(f, t)$ and $A(f, t)$ are measured from swept sine transfer functions of the DARM control loop. A swept sine transfer function is a collection of single frequency excitations applied in successive steps across the relevant frequency band of the detector.

The swept sine transfer functions have the closed loop DARM gain removed to give transfer function measurements of each of the actuation stages and the sensing function. Measurements of the detectors' DARM control loops require the detectors to be running at low-noise observation sensitivity. Once a full suite of reference measurements is taken, the complete response of the detector to GWs can be estimated.

To measure the PCAL to DARM transfer function, a known photon calibrator sine wave excitation $x_T^{(PC)}$ is applied to the detector while the DARM error signal d_{err} is recorded. This excitation is suppressed by the DARM control loop, forming the transfer function

$$\frac{d_{\text{err}}(f)}{x_T^{(PC)}(f)} = \frac{C(f)}{1+G(f)}. \tag{4.19}$$

The measurement suite is a collection of discrete sine waves swept over the frequency range 5 Hz $< f < 1$ kHz. The closed loop gain, $1/[1+G(f)]$, is then measured independently with the standard in-loop suspension actuators at the same frequencies as Equation 4.19. During times of clipping, we underestimate the excitation $x_T^{(PC)}$ by the relative actuation strength $H_{PCAL}(t)$, and must divide $x_T^{(PC)}$ by $H_{PCAL}(t)$ to correct for this. The measured sensing function is then constructed as a function of frequency:

$$C^{(\text{meas})}(f) = H_{PCAL}(t)\,[1+G(f)]\,\frac{d_{\text{err}}(f)}{x_T^{(PC)}(f)}. \tag{4.20}$$

Above 1 kHz, the photon calibrator's signal-to-noise ratio and actuation strength are low. In this region, the open loop gain $G(f)$ is negligible, so

$$\frac{d_{\text{err}}(f)}{x_T^{(PC)}(f)} \approx \frac{C^{(\text{meas})}(f)}{H_{PCAL}(t)}, \qquad f > 1\,\text{kHz} \tag{4.21}$$

We obtain the sensing function at high frequency by performing a long-duration swept sine transfer function measurement. Each single frequency is driven for many hours, and the response is compensated for time dependence using $\kappa_C(t)$.

To measure the three actuation stages, similar swept sine excitations, $x_i(f)$, are applied to each stage at points upstream of the known distribution filters, $F_i(f)$, such that the detector readout measures

$$\frac{d_{\text{err}}(f)}{x_i(f)} = \frac{H_i\,A_i(f)\,C(f)}{1+G(f)} \tag{4.22}$$

where the index i indicates either the upper intermediate U, penultimate P, or test mass T stages. These excitations are then compared to an excitation from the photon calibrator to isolate each actuation plant, as in Eq. 4.19, to form

$$[H_i\,A_i(f)]^{(\text{meas})} = \frac{1}{H_{PCAL}(t)}\,\frac{x_T^{(PC)}(f)}{d_{\text{err}}(f)}\,\frac{d_{\text{err}}(f)}{x_i(f)}. \tag{4.23}$$

The relative magnitude uncertainty and absolute phase uncertainty in a transfer function swept sine measurement point is calculated by Bendat and Piersol Eq. 4.7. The statistical uncertainty in a time-dependent parameter $\sigma_{\kappa_i}(t)$, at any given time, t, is derived from the measured coherence of the calibration lines used to form them (see Equation 4.7, propagated as in [179]). These are used as part of the MCMC and Gaussian process regressions in Section 4.7.

4.7 Calibration error and uncertainty budget

The total calibration uncertainty budget consists of statistical uncertainty and systematic error. Statistical uncertainty is the intrinsic randomness associated with measurements. Systematic error is the bias quantifying the difference between model and measurement.

Our uncertainty budget is numerically evaluated by producing a large number of realizations of the response function. To do this, we first estimate the DARM model parameters using a Markov Chain Monte Carlo (MCMC) method. Next, we stack all measurement residuals and estimate any deviations from the model using a Gaussian process regression (GPR). Then, we sample our MCMC and regression results to form ten thousand resultant response functions. These stacked response functions form the calibration error and uncertainty budget.

4.7.1 DARM model parameter estimation

First, a measurement $\vec{d} = C^{(\text{meas})}(f)$ or $A^{(\text{meas})}(f)$ is obtained as described in Section 4.6. Next, the models $\vec{M} = C^{(\text{model})}(f, t, \vec{\lambda}_C)$ or $A^{(\text{model})}(f, t, \vec{\lambda}_A)$ are fit to the measurement by varying the model parameters $\vec{\lambda} = \vec{\lambda}_C$ or $\vec{\lambda}_A$ via a Markov Chain Monte Carlo (MCMC).

An MCMC algorithm can quickly approximate the posterior probability distributions on the values of the model parameters given a log likelihood function and assumed prior distribution. The log likelihood, $\log \mathcal{L}(\vec{M} \mid \vec{\lambda}, \vec{d})$, is a simple least squares comparison between the model values $\vec{M}(\vec{\lambda})$ given model parameters $\vec{\lambda}$ and measurement data $\vec{d}$ (as described in Section 4.6). All initial parameter estimates in $\vec{\lambda}_C$ and $\vec{\lambda}_A$ were assumed to have flat prior distributions. The maximum a posteriori (MAP) values of the posterior distributions are taken as the best fit values. The ensemble of MCMC distributions are saved to be sampled for the total uncertainty budget in subsection 4.7.3.

The MCMC posteriors are found for both detector's frequency dependent models: $C^{(\text{model})}(f, t, \vec{\lambda}_C)$ and $A_i^{(\text{model})}(f, t, \vec{\lambda}_i)$. The best fit values are reported in Tables 4.1 and 4.2. The plots of the model fits can be seen in Figures 4.5 and 4.6. The one- and two-dimensional posterior distributions for the Hanford sensing model parameters $\vec{\lambda}_C$ are shown in Figure 4.7. The MCMCs were performed using the python `emcee` toolbox [182, 183]. The plot was produced with the `corner` python plotting package [184].

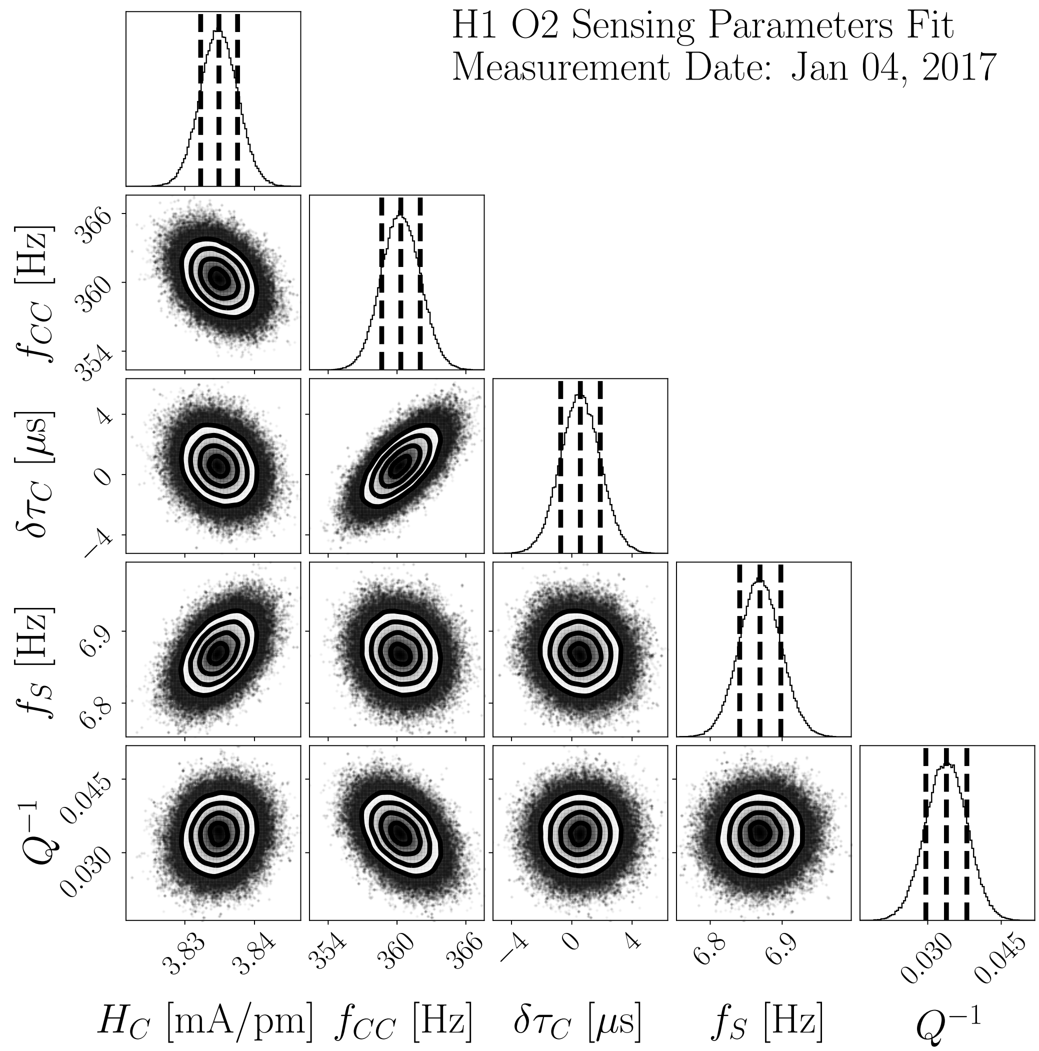

Figure 4.7: Posterior distribution on the Hanford sensing parameters $\vec{\lambda}_C$. Each column represents one of the five sensing parameters: optical gain H_C, coupled cavity pole f_{CC}, time delay correction $\delta\tau_C$, optical spring f_S, and optical spring inverse quality factor Q_S^{-1}. Each point represents a sample in five dimensional parameter space. The diagonal plots represent the variance on each parameter, while the off-diagonal plots show the covariance of each parameter with another. The dashed vertical lines on the diagonal plots represent the median and 1σ values for each parameter.

4.7.2 Quantifying frequency dependent error and uncertainty

Throughout observing runs, collections of detector measurements are taken regularly. Every measurement taken is run through the MCMC method as detailed in subsection 4.7.1. The measurement is then divided by its best fit DARM model to produce a residual, as seen in Equation 4.5.

All of the residuals are gathered together into a collection of all measurements taken over the observing run. These residuals have model-based systematic errors removed, but still contain information about unknown systematic errors. We create a distribution of functions that could describe this residual systematic error, then we incorporate this distribution into the calibration uncertainty budget. To accomplish this, we use a Gaussian process regression [185, 186].

A Gaussian process is a method of producing distributions over random functions. The Gaussian process regression takes in data and a user-defined covariance kernel. The kernel is an estimation of the similarity between any two points in the domain, in our case the log frequency domain $\log(f)$. The regression then trains on the provided data, tunes the covariance kernel hyperparameters to fit the given data, and outputs a Gaussian posterior of potential function fits to the data. This allows an uncertainty budget to be produced for arbitrary frequencies, creating a continuous posterior distribution from discrete data.

From the resulting posterior distribution, we can extract a most probable fit function, known as the mean function. The mean function becomes the systematic error $\delta C^{GP}(f)$ and $\delta A_i^{GP}(f)$ in Equations 4.15 and 4.18. We can also draw frequency dependent uncertainties $\sigma_{\delta C}^{GP}$ and $\sigma_{\delta A_i}^{GP}$ on the systematic error. Posteriors representing $\sigma_{\delta C}^{GP}$ and $\sigma_{\delta A_i}^{GP}$ will be sampled for the total uncertainty budget in subsection 4.7.3.

The main assumptions here are that residual unknown systematic errors in our measurements are Gaussian in nature, and nearby points in the frequency domain have related systematic errors.

The O1 and O2 Gaussian process regression trains on the residual data with the following covariance kernel

$$\begin{aligned} k(\log(f), \log(f')) = \gamma_1^2 &+ \log(f) \cdot \log(f') \\ &+ \left(\gamma_2^2 + \log(f) \cdot \log(f')\right)^2 \\ &+ \gamma_3^2 \exp\left(-\frac{(\log(f) - \log(f'))^2}{2\ell^2}\right) \end{aligned} \tag{4.24}$$

where $\{\gamma_1, \gamma_2, \gamma_3, \ell\}$ are the hyperparameters of the covariance kernel. O3 used a slightly different covariance kernel, but yielded largely similar posteriors [94]. The hyperparameters are tuned by the Gaussian process via gradient descent to best match the training data. This kernel assumes the detector plants' systematic error should be characterized in the log frequency domain, and that the error is relatively smooth and can be captured by a squared exponential and quadratic kernel.

An example collection of measurement residuals for the Livingston detector's sensing function and the resulting Gaussian process regression is shown in Figure 4.8. Here we show the same data from Figure 4.5, but with additional measurements from the entire observation run.

4.7.3 Total calibration uncertainty budget

The total calibration uncertainty budget for any given time is constructed from many sampled response functions $R(f, t)$ from Eq. 4.3. Each sample response function is constructed by sampling from the posteriors of the response function components. The response function components are:

1. The sensing DARM model parameters: $\vec{\lambda}_C = \left\{H_C, f_{CC}, \delta\tau_C, f_S, Q_S^{-1}\right\}$
2. The actuation DARM model parameters: $\vec{\lambda}_A = \{H_U, \delta\tau_U, H_P, \delta\tau_P, H_T, \delta\tau_T\}$
3. The sensing Gaussian process systematic error: $\delta C^{GP}(f)$
4. The actuation Gaussian process systematic errors: $\delta A_U^{GP}(f), \delta A_P^{GP}(f), \delta A_T^{GP}(f)$
5. The time dependent parameters: $\kappa_T(t)$, $\kappa_{PU}(t)$, $\kappa_C(t)$, $f_{CC}(t)$
6. The photon calibrator radiation pressure strength: $H_{PCAL}(t)$

Each of these components to the response have had posterior distributions constructed previously: (1) and (2) from the MCMC ensemble results on the reference

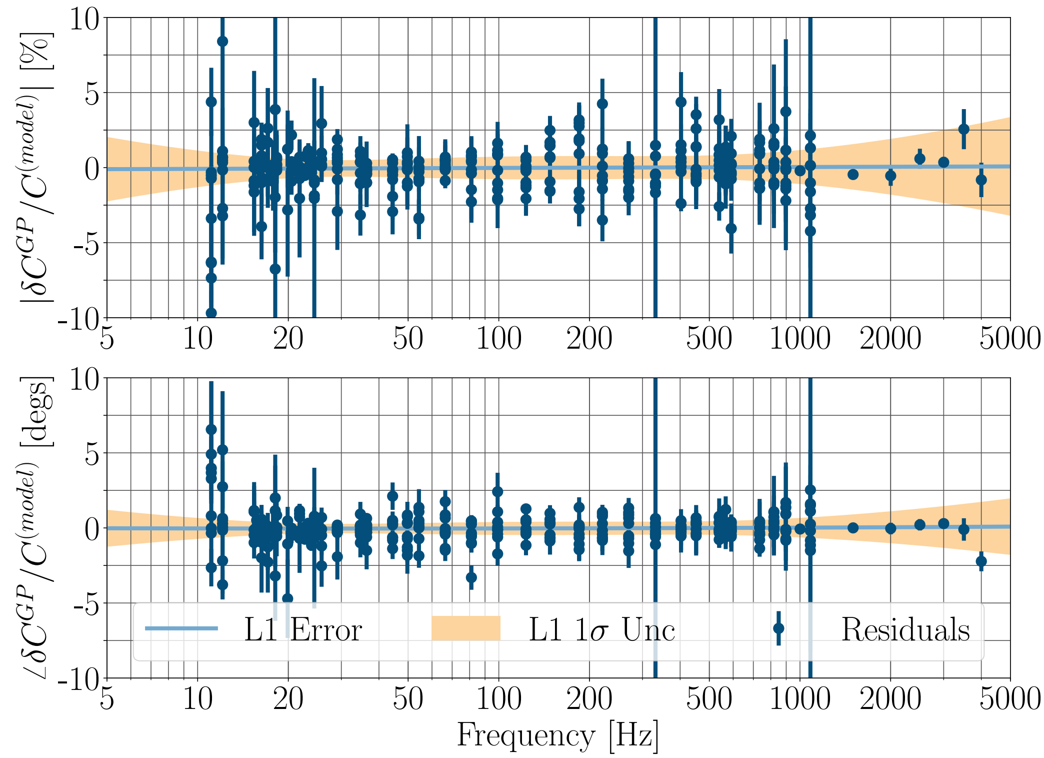

Figure 4.8: Gaussian process regression of Livingston's sensing systematic error $\delta C^{GP}(f)$. The dark blue points are all the sensing measurement residuals, $\delta C/C^{(\text{model})}(f, t, \vec{\lambda}_C)$, taken over the entire observation run. This includes the residuals from the Livingston reference measurement in the far right plots of Figure 4.5. The light blue line is the mean function representing systematic error. The light orange envelope is the 1σ uncertainty on the systematic error.

measurements, (3) and (4) from the Gaussian process regressions on the residuals to incorporate unknown systematic errors, (5) from the calibration line measurements and coherence, and (6) from the 0.79% uncertainty in $H_{PCAL}(t)$ from the photon calibrator paper [180].

Ten thousand samples are drawn from each of these posterior distributions. These samples are used to compute ten thousand sample response functions $R_i(f, t)$ according to Equation 4.3. Each of these response functions is then divided by the nominal response function, $R^{(\text{model})}(f, t)$, which is constructed from the sensing model $C^{(\text{model})}(f, t, \vec{\lambda}_C)$ and actuation model $A^{(\text{model})}(f, t, \vec{\lambda}_A)$. This gives ten thousand relative response functions $R_i(f, t)/R^{(\text{model})}(f, t)$, each of which is plotted in Figure 4.9. The median of this relative response function distribution constitutes the overall systematic error, and the 68th percentile upper and lower contours are the statistical uncertainty, both a function of frequency.

Figure 4.9 shows the calibration uncertainty at the time of the most recent detection, GW170104. Table 4.3 reports the "extreme uncertainty" for calibration between 20-1024 Hz during GW170104. Extreme uncertainty refers to the maximum and minimum of the systematic error $\pm 1\sigma$ uncertainty within a certain frequency band. This quantity is useful for searches requiring single number calibration uncertainty values, and ignore calibration systematic errors or frequency-dependent calibration uncertainty.

4.7.4 Calibration uncertainty for entire observing runs

Calibration error and uncertainty evolves over observing runs, affecting the results of continuous and stochastic gravitational wave searches [156, 157, 159, 160]. To assess the uncertainty of the detectors throughout an observing run, a total calibration uncertainty budget is made for every hour of observing data.

Collapsing the uncertainty budgets along the time axis, the 68th, 95th, and 99th percentile (1σ, 2σ and 3σ) limits are reported. The entire run's calibration error and uncertainty is often reduced to a single statement such as "over the course of an observing run, the 1σ uncertainty is no larger than XX % in magnitude and YY degrees in phase." To do so, the extreme uncertainty is taken in magnitude (XX%) and phase (YY degrees) using the 68th percentile contour over the relevant frequency band.

4.8 Results

The final calibration uncertainty budget for GW170104 is shown in Figure 4.9. The "extreme uncertainties", or the maximum and minimum of error $\pm 1\sigma$ uncertainty, are reported in Table 4.3.

The O1 uncertainty quantification method from [136] reported 10% and 10 degrees uncertainties for GW150914. The O1 calibration uncertainties for all three O1 events are in Table III in [187]. The uncertainty quantification method used for GW170104 was repeated on the O1 events, reported in Appendix A of [5].

Systematic errors are known discrepancies between the detector model and measurement. At low frequency, the systematic error is dominated by the Gaussian process regression on the actuation function residuals. At high frequency, fluctuations in the coupled cavity pole $f_{CC}(t)$, which are not corrected for in the calibration procedure, dominate the error budget.

Uncertainty everywhere is dominated by the Gaussian process regression on both functions. The uncertainty from the MCMC parameter fits on $\vec{\lambda}_C$ and $\vec{\lambda}_A$, and the uncertainty in the time dependent parameters $\kappa_T(t)$, $\kappa_{PU}(t)$, $\kappa_C(t)$, and $f_{CC}(t)$ tend to be about an order of magnitude smaller than the Gaussian process regression results. The 0.79% uncertainty in the photon calibration strength $H_{PCAL}(t)$ contributes only to magnitude uncertainty.

The uncertainty and error for O2 strain data from November 19 through June 19 is shown in Figure 4.10. This percentile plot was created by taking all observing time, producing an uncertainty budget for each hour, then compiling each budget into the percentiles shown. Overall, the detector calibration is stable over time. This consistency is largely due to the correction of the scale factors $\kappa_T(t)$, $\kappa_{PU}(t)$, and $\kappa_C(t)$ in the calibration pipeline models. Uncorrected systematic errors in the cavity pole $f_{CC}(t)$ are particularly visible at Livingston at high frequency.

During some parts of the second observing run, we have found that the reflection photodetector of the PCAL system at the Hanford detector had suffered from clipping. Clipping means that the PCAL laser light incident on the photodetector was slightly off, giving a false low reading of how much power the PCAL was emitting. This means any measurement taken using the reflection photodiode as reference had a systematic error in scale. This includes the scale of any continuously measured time-dependent model parameters which are applied as correction factors for the estimated detector output, $h(t)$. We have quantified this systematic error using the same system's transmission photodiode, and included it as systematic error in the overall response. The systematic error was on the order of a few percent, and can be seen reflected in the upper percentiles of the Hanford uncertainty in Figure 4.10.

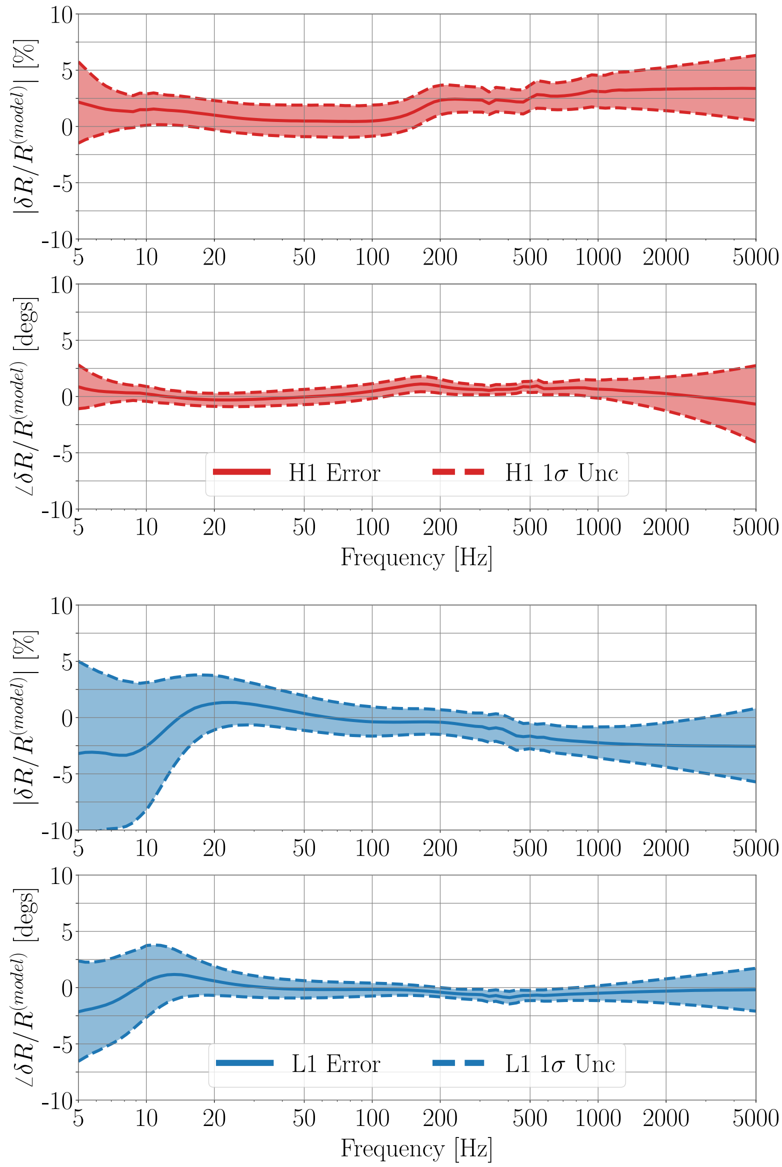

Figure 4.9: Total calibration error and uncertainty budget at the time of GW170104. The uncertainty in the calibrated response function for the Hanford detector is on the top, and for Livingston is on the bottom. The y axis is relative response error $\delta R/R^{(\mathrm{model})}$ and uncertainty $\sigma_R/R^{(\mathrm{model})}$, with magnitude on top and phase on the bottom. The solid line is the median relative response, interpreted as the frequency dependent systematic error on the model response $R^{(\mathrm{model})}$. The dashed lines represent the 1σ uncertainty on this error. Stacking ten thousand drawn response function samples produces the numerical uncertainty budget shown here. The extreme 1σ uncertainties are presented in Table 4.3.

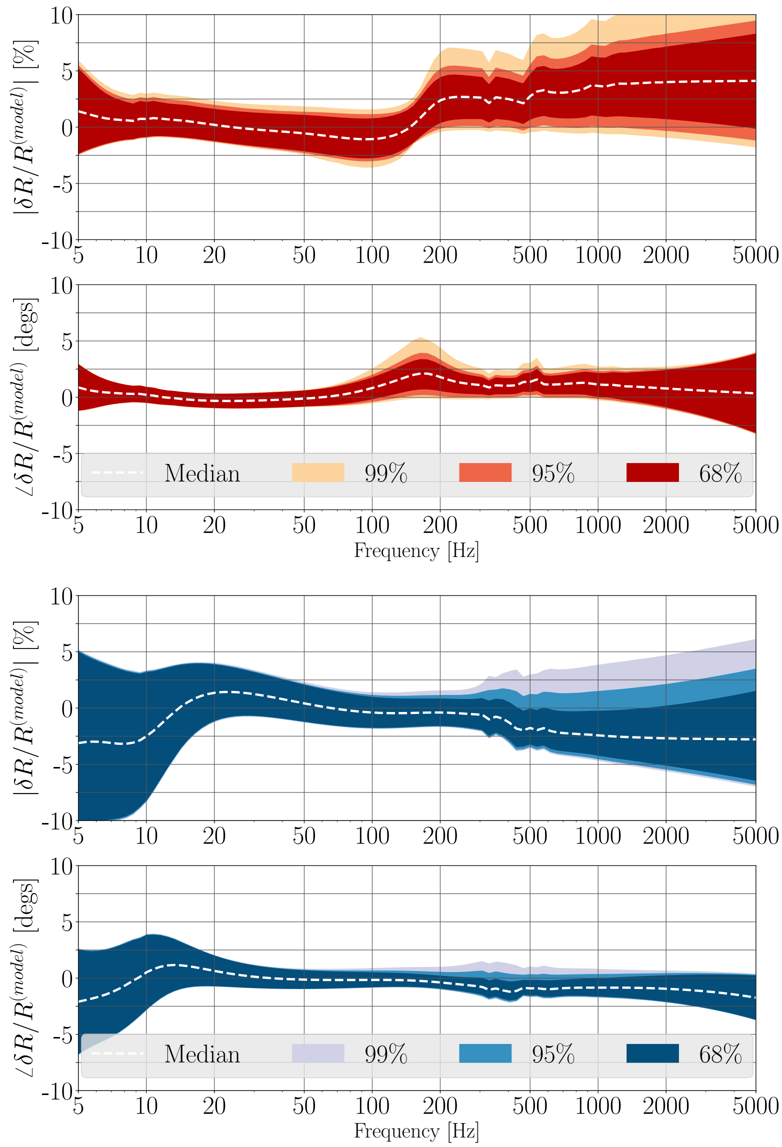

Figure 4.10: Total calibration uncertainty percentiles for observing run two. The percentiles are created for all of O2 data from November 30, 2016 to August 25th, 2017. Hanford's uncertainty plots are the red on the top, and Livingston's are the blue on the bottom. The y axis is relative response $\delta R/R^{(\mathrm{model})}$ magnitude (top) or phase (bottom), stacked for all times in the observing run. The dashed white line is the median relative response, while the colors represent the 1σ calibration uncertainty for 68%, 95%, and 99% of the run's time. The largest changes in the calibration at Hanford were due to clipping of the photon calibrator laser misreporting the strength of our response. The largest calibration changes at Livingston were due to fluctuations in the coupled cavity pole, which changes in time but is not yet corrected for in our calibrated data.

Table 4.3: Below are the extreme calibration uncertainty values for Hanford and Livingston at the time of GW170104 in the 20-1024 Hz frequency range. "Extreme uncertainty" refers to the maximum and minimum of error $\pm 1\sigma$ uncertainty. The plots informing this table can be seen at Figure 4.9

GW170104 Uncertainty	Hanford	Livingston
$+1\sigma$ Magnitude [%]	4.6 %	3.7 %
-1σ Magnitude [%]	-1.0 %	-3.7 %
$+1\sigma$ Phase [degrees]	1.8°	1.9°
-1σ Phase [degrees]	-0.9°	-1.4°

4.9 Fundamental uncertainty limit

In O1 and O2, the relative uncertainty in the photon calibrator actuation strength $H_{PCAL}(t)$ was 0.79% [180]. This has been reduced to 0.41% in O3 [181]. This is the fundamental limit on our uncertainty in the response R and therefore the GW strain data h. The uncertainty in $H_{PCAL}(t)$ is dominated by uncertainty in the laser power and test mass rotation [180]. To push this fundamental limit lower, better measurements of the photon calibrator laser power and test mass rotation must be made, or more precise methods of calibration outside of the photon calibrator may need to be considered.

In the end, the best way to strengthen our confidence in the systematics of the dead-reckoned photon calibrator measurement is to compare other calibration methods to its final results. The Newtonian calibrator is probably the best currently-existing competitor with the photon calibrator [188, 189]. The NCAL has its own share of systematic errors, and is less well-characterized than the photon calibrator. But if these methods agree to some extent, with mostly uncorrelated systematics, we can have far greater confidence that the overall calibration is accurate.

Other methods of calibration were used in initial LIGO prior to the photon calibrator. The free-swinging Michelson is the most familiar, which calibrates DARM to the wavelength of the laser frequency by allowing the end test mass to swing and produce dark flashes in the Michelson.

The VCO calibration calibrates DARM to the green ALS DIFF control signal which is in units of Hz. These Hz are related to meters of DARM motion by $\delta\nu_g/\nu_g = -\delta L_-/L$.

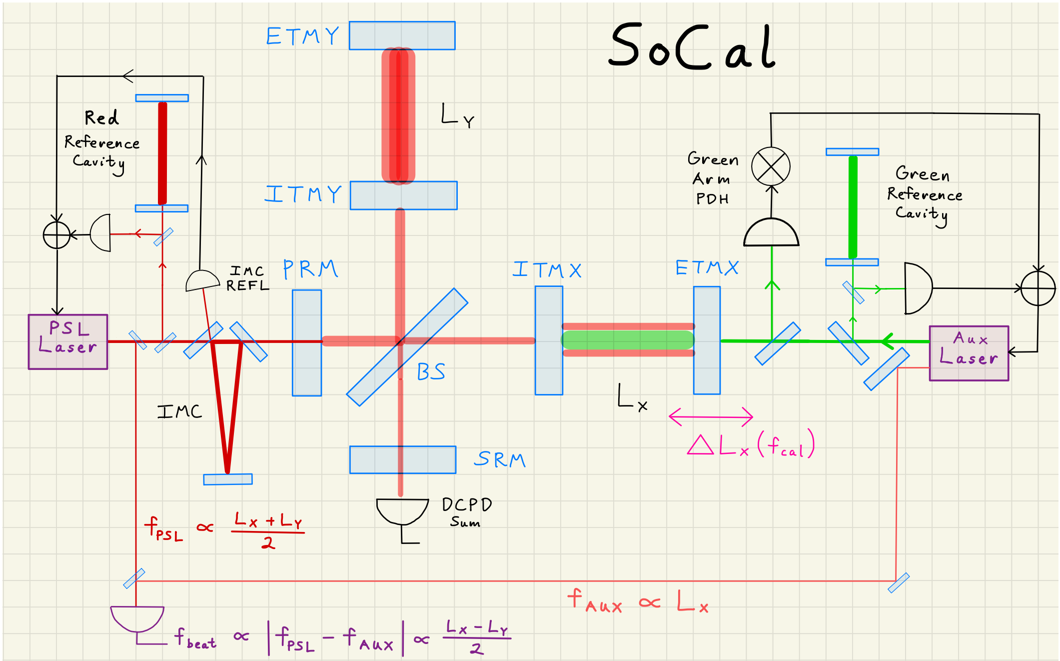

Figure 4.11: Simplistic control scheme for the simultaneous oscillator calibration. A strong calibration signal at f_{cal} is input into DARM, in this case represented as a dither on the X-arm ΔL_x. Both CARM and the green arm PDH feedback the signal to their respective laser frequencies f_{PSL} and f_{AUX}. The calibration signal comes from the frequency beatnote f_{beat} between the PSL f_{PSL} and AUX f_{AUX} lasers.

Both the free-swinging Michelson and VCO calibration suffer from the fact that true DARM displacement cannot be directly injected into a fully locked interferometer. Instead, the quad suspension actuation must be calibrated into meters of motion, which is then used as the calibration injection itself. The suspension can only be calibrated in a high-noise state of the interferometer as well, leading to less precision. This two-step transfer function calibration will incorporate more uncertainty from more measurements than the single-step photon calibration.

4.9.1 SoCal: Simultaneous oscillator calibration

Requirements on the calibration accuracy may become extremely low, near 0.1%, for astrophysical parameter estimation in future detectors such as Einstein Telescope and Cosmic Explorer, where some detection SNRs will be > 100. Another calibration method is under development at the Caltech 40m, colloquially known as SoCal (for "simultaneous oscillator calibration").

Figure 4.11 gives a diagram overview of the system. In short, SoCal uses the infrared laser to sense CARM motion, a green laser to sense a single arm's motion,

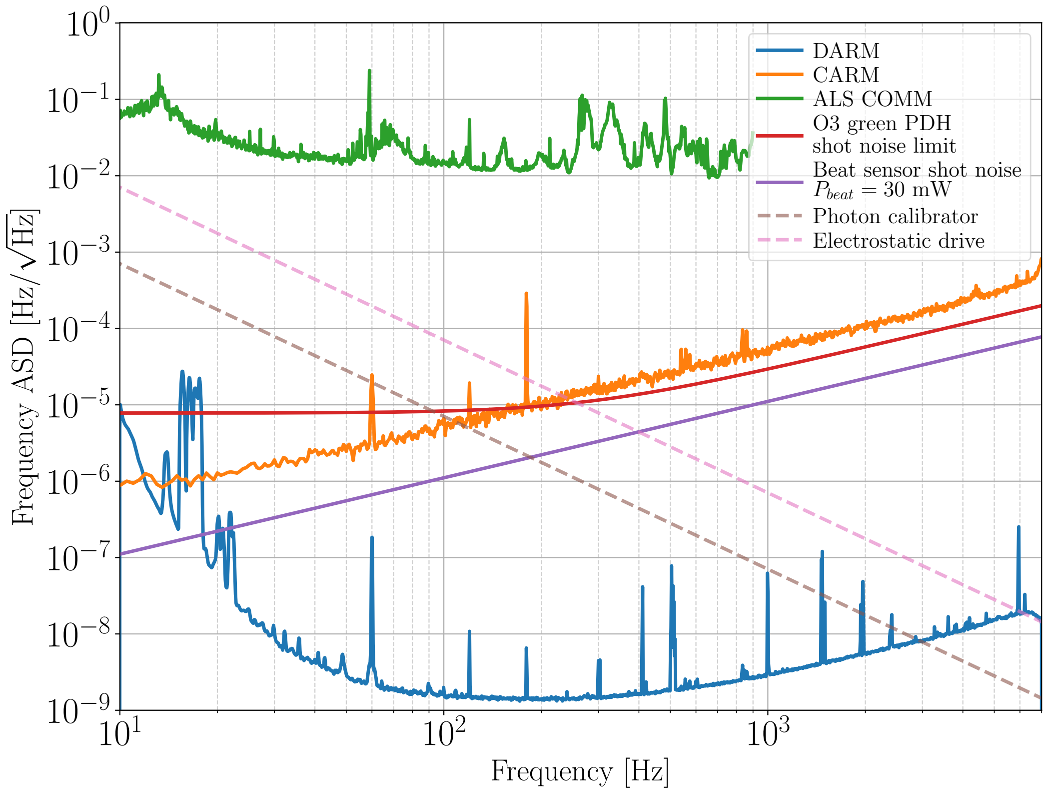

Figure 4.12: DARM, CARM, ALS COMM measured noises, as well as the current ALS X green PDH shot noise limit. The ALS parameters and measurement are taken from Section 3.3 and Figure 3.10, the DARM noise from 3.1, and the CARM noise from 3.17. Currently, the huge dynamic range between ALS performance and DARM renders a scheme like SoCal impossible: DARM is 10^7 times more sensitive that ALS.

and then recovers DARM motion from the beatnote from those signals mixed together. The infrared PSL and auxiliary laser frequencies f_{PSL} and f_{AUX} are locked to the CARM and X-arm interferometer cavity lengths:

$$f_{PSL} \propto \frac{L_x + L_y}{2}, \qquad f_{AUX} \propto L_x. \tag{4.25}$$

The green fringe $2f_{AUX}$ is adjusted such that f_{AUX} is within one FSR of f_{PSL}, then the two beats are mixed on a beatnote sensor and low-passed to produce f_{beat} such that

$$f_{beat} = |f_{PSL} - f_{AUX}| \propto \frac{L_x - L_y}{2} \tag{4.26}$$

which is proportional to the DARM signal in Hz.

If a strong calibration line ΔL_x is injected at a frequency f_{cal}, this will modulate the beatnote f_{beat} with sidebands $f_{beat} \pm f_{cal}$. The usual DARM sensor, the DCPDs, will detect the signal at f_{cal} as well. The DCPDs can then be calibrated into Hz and meters using the output $f_{beat} \pm f_{cal}$.

In reality, we will be detecting the modulated power on the beatnote sensor P_{beat} which comes from the laser fields $\vec{E}_{PSL}$ and $\vec{E}_{AUX}$. Both laser fields will be frequency modulated by the change in arm length ΔL_x at frequency ω_c, but the PSL signal will be modulated by the CARM signal $\Delta L_+ \approx \Delta L_x/2$:

$$\vec{E}_{PSL} = E_p e^{i\omega_p t}\left(1 + ik\Delta L_+ e^{i\omega_c t} + ik\Delta L_+ e^{-i\omega_c t}\right) \tag{4.27}$$

$$\vec{E}_{AUX} = E_a e^{i\omega_a t}\left(1 + ik\Delta L_x e^{i\omega_c t} + ik\Delta L_x e^{-i\omega_c t}\right) \tag{4.28}$$

The modulated power in the beatnote P_{beat} that is proportional to DARM will come from the cross terms between the fields. Ignoring small terms proportional to $\Delta L_x \Delta L_c$ yields

$$P_{beat} = \vec{E}^*_{PSL}\vec{E}_{AUX} + \vec{E}_{PSL}\vec{E}^*_{AUX} \tag{4.29}$$

$$P_{beat} = 2E_p E_a(\cos((\omega_p - \omega_a)t) + 2k(\Delta L_x - \Delta L_+)\cos(\omega_c t)\sin((\omega_p - \omega_a)t)). \tag{4.30}$$

Here we explicitly see the DARM term $\Delta L_- = \Delta L_x - \Delta L_+$ at a frequency ω_c in the sine quadrature of the beat frequency $\omega_{beat} = \omega_p - \omega_a$.

4.9.1.1 Advantages

The advantage of this technique is, to zeroth order, the systematic error is dominated by knowledge the length of the arms and the wavelength of the laser, both of which are known to 1 ppm. The uncertainty can be made negligible compared to the systematic uncertainty by increasing the drive signal or the integration time to achieve a suitable SNR.

Another advantage is the beatnote sensor can be relatively narrow audio-band due to the length of the LIGO arms. Because the $\mathrm{FSR} = 37.5$ kHz, the beatnote frequency f_{beat} can always be made to be within $\mathrm{FSR}/2$ by changing the fringe the green is locked on.

The demodulation of the audio-band calibration line can be done digitally. This removes the noise complexity of an RF detection scheme with crystal oscillator noise.

The SoCal frequency calibration needs only to be performed once, and can be used to more accurately calibrate the more versatile photon calibrator or electrostatic drives. If very large calibration lines are not possible, required SNRs may be achieved by integrating a line at frequency f_i for long times T:

$$\mathrm{SNR}(f_i) = \sqrt{T}\frac{\Delta\nu_{\mathrm{cal}}(f_i)}{\sqrt{S_{\nu,\mathrm{dof}(f_i)}}} \tag{4.31}$$

where $\Delta L_{\mathrm{cal}}(f_i)$ is the modulation line with units of meters, and $\sqrt{S_{L,\mathrm{dof}(f_i)}}$ is the ASD of the noise in the degree of freedom in units of $\mathrm{m}/\sqrt{\mathrm{Hz}}$.

Translating SNRs to uncertainty is done in Appendix E, Eq. E.25. To achieve a relative uncertainty $\sqrt{\mathrm{Var}[\Delta\nu_{\mathrm{cal}}(f_i)]/\Delta\nu_{\mathrm{cal}}(f_i)}$ of 0.1% with an $\mathrm{SNR} = 10$, around $n = 50000$ averages are required. This translates to roughly $T = n/f_b$ seconds, where f_b is the frequency binwidth. If we place our calibration line at frequency $f_{cal} = 50$ Hz, then a frequency binwidth $f_b = 10$ Hz is reasonable, and the calibration can be accomplished in $nT = 5000$ seconds.

However, with a short ASD time of $T = 0.1$ seconds there is a tradeoff with SNR. The signal of the photon calibrator and electrostatic drive for $T = 10\ s$ is plotted in Figure 4.12 [180]. Appendix E explores the line height estimate in depth, including the bias and uncertainty.

4.9.1.2 Challenges

Figure 4.12 shows the noise levels relevant for SoCal. The biggest issues with this scheme are the infrastructure required to make it feasible, and the huge calibration line required to ensure the signal shows up above the CARM noise and green PDH shot noise.

The infrastructure changes would be overall improvements to the ALS subsystem to make it possible to reach the green shot noise limit. Current ALS limits are the VCO which drives the laser frequency control and phase-wrapping from relative motion between the suspended in-vacuum and in-air tables. Putting ALS sensors in-vacuum at the vertex and developing a better green frequency stabilization scheme is required. Also, choosing green reflection parameters with SoCal in mind could push down the green shot noise limit.

Another limit is the beatnote sensor shot noise. The level of light on the beatnote sensor will need to be extremely high to achieve an appreciable SNR, even with an extremely strong calibration signal. Figure 4.12 shows the shot noise for a phase-sensitive PD with 30 mW.

The very large calibration line must make up the factor 10000 difference between DARM and the green shot noise. Also, the line must clear the IMC sensing noise to show up in CARM. Achieving these lines with the current low noise actuators, the photon calibrator and low-noise electrostatic drive, is feasible, but may require switching to the high-noise electrostatic drive setting or coil drivers on the upper stages, which is not the low-noise state we want to calibrate DARM in.

Very large drives cause massive upconversion as sensors and controls saturate. These can not only pollute the spectrum, but cause a nonlinear response in the sensors we are trying to calibrate. Noise must be reduced to the point where the calibration input required is feasible for the interferometer to handle without losing lock. This can be achieved with a lower calibration line amplitude, but a longer integration time.

Systematic errors may arise from auxiliary controls in the interferometer. For instance, angle-to-length coupling will be a problem if the co-alignment of the two beams in the arm is bad. Torque is known to couple to length, and could couple to the two beam differently. For this type of calibration, angle-to-length coupling to DARM, CARM, and ALS will have to be monitored and minimized, so the ALS and CARM signals are following the true lengths of the cavities accurately.

Another source of systematic error is the fact that CARM is stabilized to the sum of common arm length and the PRC. The huge length drive required for this calibration method will show up in the PRCL control. PRCL will have to be notched at the frequency of the drive prior to calibrating there.

Finally, the interferometer is known to change over time, with thermalization and realignment and glitches slightly changing the interferometer response. If long integration times are required, systematics from interferometer fluctuations in optical gain will need to be incorporated into the final uncertainty budget.

4.10 Future Work

As we reduce the calibration uncertainty, properly characterizing systematic errors becomes much more important for precision astrophysics. Any systematic errors left unaccounted for in the calibrated data can result in systematic errors in binary black hole source parameters, compact binary merger rates, or tests of general relativity. Our direct measurements of our detector control loop plants combined with the physics-motivated response function model provide a sanity check that our understanding of the interferometer is close to correct.

4.10.1 Other sources of calibration systematic error

The uncertainty budget does not include error from test mass elastic deformation due to the PCAL laser exciting test mass vibrational modes. Preliminary evidence suggests that above around 3 kHz, elastic deformation has a significant effect on the calibration accuracy. Elastic deformation due to the PCAL must be further understood, monitored, and included in the uncertainty budget directly.

There is a difference between the quadruple pendulum response to an actual gravitational wave versus its response to the photon calibrator. A gravitational wave displaces the entire quad pendulum in the lab frame, whereas the photon calibrator only pushes on the test mass. The effect of this difference on calibrated GW data is on the order of about 1% at 10 Hz, and increases at lower frequencies. This now must be considered quantitatively as uncertainties approach this level.

4.10.2 Conclusions

The uncertainty and systematic error estimates reported in this chapter represent a comprehensive characterization of our Hanford and Livingston detector calibrations for observing run two. In Advanced LIGO's lowest noise region, from about 20 Hz to 1 kHz, the uncertainty in the calibrated data has been reduced from what was previously reported in [136]. The uncertainty estimates for O2 give more refined results, with uncertainty growing at extreme frequency regions below 20 Hz and above 1 kHz, and reduced uncertainty in the low noise frequency region.

Interesting astrophysics exists at high GW frequencies. The equation of state of the neutron star, higher multipoles of GW emission near merger (already detected in [11]), GW "echoes", supernova cores bounces, and unexpected GW detections are all potential exotic phenomena that high frequency GWs carry information about. Calibration accuracy at high frequency will be important for learning about each of these extreme GW regimes.

GW170104's detection and parameter estimation are primarily limited by noise, and not by calibration uncertainty. As Advanced LIGO becomes more and more sensitive, the signal-to-noise ratio of some detections will become quite large (as high as 100 or more), and calibration uncertainty will begin contributing significantly to source parameter estimation uncertainty. With more observing time comes more detections, enabling new tests of general relativity which will be limited by the precision of our detector data. Precision astrophysics demands the best understanding of our calibrated data possible. The methods described in this paper were developed primarily to enable the best science possible from LIGO's gravitational wave detections.

Chapter 5

CORRELATED NOISE

Advanced LIGO operates in DC-readout configuration, with two DC photodetectors (DCPDs) at the antisymmetric port [41, 190]. The signal from differential arm motion (DARM), and therefore gravitational waves (GWs), appears in the sum of these two DCPDs. Noise from the interferometer, known as *correlated noise*, is measured identically in both DCPDs. Sensing noises, like shot noise and photodetector dark noise, are incoherent between each DCPD.

Quantum shot noise tends to dominate the DARM spectrum over most of the GW detection band [22, 47]. It is possible to measure the correlated noise by measuring the cross spectral density (CSD) between the two DCPDs [191, 192]. Correlated noise measurements are useful because they expose noise sources under the quantum shot noise.

5.1 Introduction

There are a number of complications that must be considered for the correlated noise measurement:

1. the DARM loop must be measured and removed,
2. injecting squeezed light correlates the shot noise in the DCPDs,
3. a large number of averages must be taken to integrate away the shot noise to reach the true correlated noise floor, and
4. detector glitches which inject large transients into DARM must be avoided.

The correlated noise measurement requires hours of data to integrate away shot noise. However, large glitches occur on the order of one every fifteen minutes, spoiling the mean-averaged power spectral densities (PSDs) of the DCDPs. Large glitches can be "gated" by removing the glitchy data by hand, but the measurement is still susceptible to small glitches that do not meet our gating criteria.

To mitigate having undetected glitches in the DARM data, it is often useful to use median-averaging for estimating PSDs and CSDs. Median-averaging has an intrinsic bias as opposed to mean-averaging. For PSDs the bias factor is understood to asymptote to $\log(2)$ [57, 193]. For CSDs, the result is covered in Chapter 6. Another useful technique for removing CSDs containing glitches while still using mean-averaging is "PSD rejection", covered in Appendix F.

In this chapter, we first review the DARM control loop, review the method for extracting the correlated noise from the DCPD data, and show results from the correlated noise floor estimation of the LIGO Hanford interferometer during Observing Run 3.

5.2 Method

Here we overview the method of extracting the correlated noise spectrum, assuming there is no squeezing injected into the antisymmetric port. This section follows the procedure in [192].

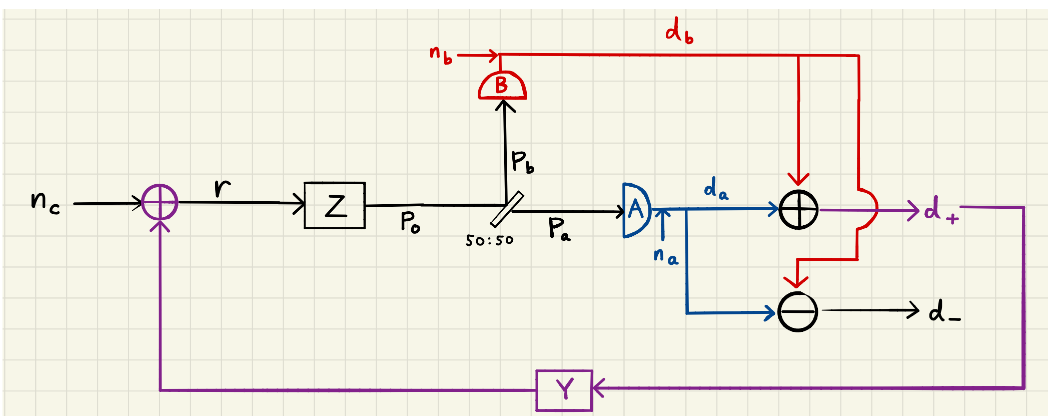

Figure 5.1: Simplified DARM control loop diagram, including the two-photodetector sensing scheme. A and B are the transfer functions of the two DCPDs, plus the photodetector analog electronics and ADCs, with units of $[\mathrm{cts/W}]$. Y is the actuator with units $[\mathrm{m/cts}]$. Z is the interferometer response to differential arm motion, in $[\mathrm{W/m}]$. A 50:50 beamsplitter splits the light from the interferometer onto the two photodetectors. The signals from DCPD A and B, d_a and d_b, are summed to form the DARM error signal, d_+, and subtracted to form the null channel, d_-. The shot noises n_a and n_b are uncorrelated sensing noises added to d_a and d_b, respectively. The correlated noise n_c is here simplified as being entirely displacement noise causing actual motion of the interferometer optics.

5.2.1 Correlated noise without squeezing

Figure 5.1 shows a simplified DARM loop. We would like to measure the correlated noise coming from the interferometer, $\langle n_c, n_c \rangle$. However, shot noises $\langle n_a, n_a \rangle$ and $\langle n_b, n_b \rangle$ drown out the correlated noise for most of the bandwidth. Additionally, the DARM control loop that keeps the interferometer on resonance correlates sensing noise by injecting it into the loop as mirror motion. This is how shot noise on DCPD A can appear on DCPD B, only in the bandwidth of the DARM loop ($< \ 50$ Hz). By cross-correlating the DCPDs $\langle d_a, d_b \rangle$ and removing the DARM loop, the correlated noise can be extracted.

The open loop gain G of the DARM loop is just the product of all the loop transfer functions, while the DARM sensing function C is usually calculated from photon calibrator displacement to DARM error signal in units $[\mathrm{cts/m}]$:

$$G = \frac{1}{2} Y Z (A + B) \tag{5.1}$$

$$C = \frac{1}{2} Z (A + B) \tag{5.2}$$

Solving the loop in Figure 5.1 for the DARM error signal, aka the sum channel d_+, and the null channel d_-:

$$d_+ = \frac{1}{1 - G} (C\, n_c + n_a + n_b) \tag{5.3}$$

$$d_- = n_a - n_b \tag{5.4}$$

Taking the power spectral density of the sum and null channels yields

$$\langle d_+, d_+ \rangle = \frac{1}{|1 - G|^2} \left(|C|^2 \langle n_c, n_c \rangle + \langle n_a, n_a \rangle + \langle n_b, n_b \rangle \right) \tag{5.5}$$

$$\langle d_-, d_- \rangle = \langle n_a, n_a \rangle + \langle n_b, n_b \rangle \tag{5.6}$$

assuming that n_c, n_a, and n_b are all independent, so e.g. $\langle n_a, n_b \rangle = 0$. If we calibrate the DARM error PSD into meters by multiplying $\langle d_+, d_+ \rangle$ by $|1 - G|^2 / |C|^2$, the shot noise terms $\langle n_a, n_a \rangle$ and $\langle n_b, n_b \rangle$ still appear.

Solving the diagram in Figure 5.1 for split DARM error signals d_a and d_b yields

$$d_a = \frac{1}{2(1 - G)} ((2 - G) n_a + G n_b + C n_c) \tag{5.7}$$

$$d_b = \frac{1}{2(1 - G)} (G n_a + (2 - G) n_b + C n_c) \tag{5.8}$$

If we look at the power spectral densities of each individual $\langle d_a, d_a \rangle$ and $\langle d_b, d_b \rangle$

and the cross spectral density $\langle d_a, d_b \rangle$, we get

$$\langle d_a, d_a \rangle = \frac{1}{4|1-G|^2}\big(|2-G|^2\langle n_a, n_a\rangle + |G|^2\langle n_b, n_b\rangle + |C|^2\langle n_c, n_c\rangle\big) \tag{5.9}$$

$$\langle d_b, d_b \rangle = \frac{1}{4|1-G|^2}\big(|G|^2\langle n_a, n_a\rangle + |2-G|^2\langle n_b, n_b\rangle + |C|^2\langle n_c, n_c\rangle\big) \tag{5.10}$$

$$\langle d_a, d_b \rangle = \frac{1}{4|1-G|^2}\big(G(2-G^*)\langle n_a, n_a\rangle + G^*(2-G)\langle n_b, n_b\rangle + |C|^2\langle n_c, n_c\rangle\big). \tag{5.11}$$

Using Eqs. 5.9, 5.10, and 5.11, we can solve for the correlated noise $\langle n_c, n_c \rangle$. Recalling that $\langle d_b, d_a \rangle = \langle d_a, d_b \rangle^*$, the correlated noise is

$$\begin{aligned} |C|^2\langle n_c, n_c\rangle = \Bigg(& |2-G|^2\langle d_a, d_b\rangle + |G|^2\langle d_b, d_a\rangle \\ & - G(2-G^*)\langle d_a, d_a\rangle - G^*(2-G)\langle d_b, d_b\rangle \Bigg) \end{aligned} \tag{5.12}$$

By measuring the individual DCPD signals d_a and d_b and applying the DARM loop gain G and sensing function C, the correlated noise from the interferometer can be directly estimated.

Figure 5.3 shows the O3 correlated noise budget for $\langle n_c, n_c \rangle$.

5.2.2 DC readout with squeezing

Now suppose that squeezed light is injected into the antisymmetric port of the interferometer. This correlates the noise on each DCPD, i.e. $\langle n_a, n_b \rangle \neq 0$ [195]. Here we review the DC readout detection scheme, how squeezing correlates the shot noise appearing on each photodetector, and calculate the squeezed shot noise cross spectral density $\langle n_a, n_b \rangle$ for a DC readout interferometer.

To calculate $\langle n_a, n_b \rangle$ we briefly review shot noise in a DC readout interferometer with split photodetection, as shown in Figure 5.2. This will follow the derivation and notation in [194].

From Figure 5.2, $\vec{L}$ is the local oscillator, $\vec{\ell}$ is the quantum vacuum, $\vec{S}$ is the DARM offset light, and $\vec{s}$ is the gravitational wave signal plus the output squeezed vacuum.

The capital letters refer to the carrier, while the lowercase letters refer to the audio sidebands that beat with the carrier. DC readout operates with no local oscillator

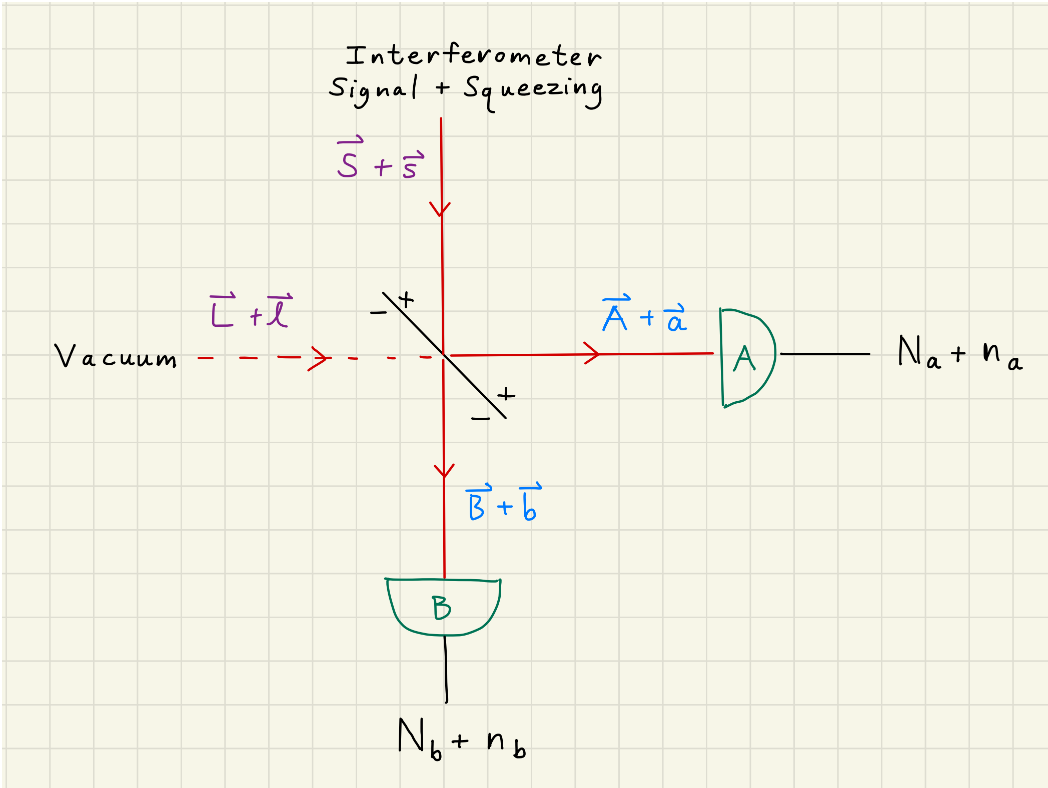

Figure 5.2: Diagram of the DC readout detection scheme with split photodetection [194]. The beamsplitter is 50:50 with the reflection convention designated by the plus and minus signs. The DCPD sum $n_+ = n_a + n_b$ contains the information from the interferometer, including squeezed vacuum and gravitational wave signal. The DCPD null $n_- = n_a - n_b$ contains the information from the quantum vacuum.

$\vec{L} = 0$, and a homodyne angle $\zeta = \pi/2$, which puts the gravitational wave signal entirely in the amplitude quadrature upon exit from the interferometer. $\vec{S}$ is set to some non-zero value to beat against the squeezed vacuum $\vec{s}$. For the shot noise derivation, we assume that the GW signal and interferometer correlated noise is zero.

The light incident on each DCPD $\vec{A} + \vec{a}$ and $\vec{B} + \vec{b}$ is

$$\vec{A} + \vec{a} = \frac{1}{\sqrt{2}}\left(\vec{S} + \vec{s} + \vec{L} + \vec{\ell}\right) \tag{5.13}$$

$$\vec{B} + \vec{b} = \frac{1}{\sqrt{2}}\left(\vec{S} + \vec{s} - \vec{L} - \vec{\ell}\right) \tag{5.14}$$

First, recall that the local oscillator $\vec{L} = 0$. Second, the homodyne angle definition from [22] for the signal $\vec{s}$ is:

$$s_\zeta = s_1 \sin(\zeta) + s_2 \cos(\zeta). \tag{5.15}$$

The Advanced LIGO interferometers are dual-recycled resonant-sideband extraction in DC readout configuration, meaning the signal-recycling cavity is tuned to $\phi = \pi/2$. This puts the GW signal, upon exit from the interferometer, in the *amplitude* quadrature s_1. It also places the static DARM offset light $\vec{S}$, in the amplitude quadrature.

For now, we ignore contrast defect and set the homodyne angle $\zeta = \pi/2$ using Eq. 5.15, so

$$\vec{S} = \begin{pmatrix} S \\ 0 \end{pmatrix} \tag{5.16}$$

picks out the output amplitude quadratures $s_1 = s$ and $\ell_1 = \ell$. The light is converted into current on the DCPDs, represented by $N_a + n_a$ and $N_b + n_b$.

$$N_a + n_a = |\vec{A} + \vec{a}|^2 = \frac{1}{2}S^2 + Ss + S\ell \tag{5.17}$$

$$N_b + n_b = |\vec{B} + \vec{b}|^2 = \frac{1}{2}S^2 + Ss - S\ell \tag{5.18}$$

where terms proportional to s^2, $s\ell$, and ℓ^2 are small enough to be negligible. $N_a = N_b = S^2/2$ represents the DARM offset light being split in half by the beamsplitter, nominally $N_a = N_b \approx 20$ mW in Advanced LIGO.

Removing DC components represented by the capital letters from Eqs. 5.17 and 5.18, we can calculate the shot noise sum n_+ and null n_-:

$$n_+ = 2Ss \tag{5.19}$$

$$n_- = 2S\ell. \tag{5.20}$$

This illustrates how, with the DC readout scheme, the sum signal picks out the squeezed vacuum signal from the interferometer s and the null signal picks out the unsqueezed vacuum ℓ.

The squeeze parameter r is used to quantify how quantum measurement uncertainty increases and decreases between quadratures [196]. The output squeezed vacuum in s is phased such that the maximum squeezing e^{-r} occurs in the amplitude quadrature at the output. The quantum vacuum in ℓ has no squeezing. Calculating the power spectral densities of the sum and null signals:

$$\langle n_+, n_+ \rangle = 4S^2 \langle s, s \rangle = 4S^2 e^{-2r} \tag{5.21}$$

$$\langle n_-, n_- \rangle = 4S^2 \langle \ell, \ell \rangle = 4S^2. \tag{5.22}$$

Here we have assumed that the unsqueezed quantum vacuum shot noise $\langle \ell, \ell \rangle = 1$, which follows from our definition of quadratures in Section A.3. The sum shot noise PSD is reduced by the squeeze factor e^{-2r}, which is often expressed in dB by setting $e^{-2r} = 10^{-\frac{\mathrm{dB_{sqz}}}{10}}$.

The shot noise cross spectral density $\langle n_a, n_b \rangle$ is the important quantity that arises when calculating the correlated noise between DCPDs. From Eqs. 5.17 and 5.18, the cross spectral density is found:

$$\begin{aligned} \langle n_a, n_b \rangle &= S^2 \langle s + \ell, s - \ell \rangle \\ &= S^2 (\langle s, s \rangle - \langle \ell, \ell \rangle + \langle s, \ell \rangle - \langle \ell, s \rangle) \\ \langle n_a, n_b \rangle &= S^2 \left(e^{-2r} - 1 \right) \end{aligned} \tag{5.23}$$

where we have no correlation between our squeezed and unsqueezed vacuum, so $\langle s, \ell \rangle = \langle \ell, s \rangle = 0$, and we have used the definitions from Eqs. 5.21 and 5.22 and written $\langle s, s \rangle = e^{-2r}$ and $\langle \ell, \ell \rangle = 1$.

A key observation here is that, for true squeezing where $r > 0$, $\langle n_a, n_b \rangle$ is *real* and *negative*. This implies that

1. when detecting squeezed light, the power measured on each DCPD is *anti-correlated*,
2. because $\langle n_a, n_b \rangle$ is real, $\langle n_b, n_a \rangle = \langle n_a, n_b \rangle$,
3. correlated noise due to squeezing will have an opposite sign to correlated noise coming from the interferometer $\langle n_c, n_c \rangle$, which must be positive.

Figure 5.5 plots the measured correlated noise with squeezing, illustrating how in the shot noise dominated frequency band the phase is $180°$.

5.2.3 Correlated noise with squeezing

Recalculating the DCPD spectral densities including $\langle n_a, n_b \rangle$ terms yields

$$\langle d_a, d_a \rangle = \frac{1}{4|1-G|^2} \Bigg(|2-G|^2 \langle n_a, n_a \rangle + |G|^2 \langle n_b, n_b \rangle + |C|^2 \langle n_c, n_c \rangle + 2(G + G^* - |G|^2) \langle n_a, n_b \rangle \Bigg) \tag{5.24}$$

$$\langle d_b, d_b \rangle = \frac{1}{4|1-G|^2} \Bigg(|G|^2 \langle n_a, n_a \rangle + |2-G|^2 \langle n_b, n_b \rangle + |C|^2 \langle n_c, n_c \rangle + 2(G + G^* - |G|^2) \langle n_a, n_b \rangle \Bigg) \tag{5.25}$$

$$\langle d_a, d_b \rangle = \frac{1}{4|1-G|^2} \Bigg(G(2-G^*) \langle n_a, n_a \rangle + G^*(2-G) \langle n_b, n_b \rangle + |C|^2 \langle n_c, n_c \rangle + 2(2 - G - G^* + 2|G|^2) \langle n_a, n_b \rangle \Bigg). \tag{5.26}$$

If we compute the correlated noise using the RHS of Eq. 5.12 with Eqs. 5.24, 5.25, and 5.26, we get

$$|C|^2 \langle n_c, n_c \rangle + 4 \langle n_a, n_b \rangle = \Bigg(|2-G|^2 \langle d_a, d_b \rangle + |G|^2 \langle d_b, d_a \rangle - G(2-G^*) \langle d_a, d_a \rangle - G^*(2-G) \langle d_b, d_b \rangle \Bigg) \tag{5.27}$$

Figure 5.4 plots this expression as the correlated noise with squeezing trace.

5.2.4 Squeezing level estimate from correlated noise

If we have already estimated the correlated noise $\langle n_c, n_c \rangle$ during a time without squeezing, then it's possible to estimate the squeezing level. This can be especially useful if the squeezing is frequency dependent, as it was at LIGO Hanford during O3.

We assume here that the correlated noise is the same for both squeezing and non-squeezing times. This is not true where quantum radiation pressure noise (QRPN) is significant, as anti-squeezing will enhance the QRPN contribution to correlated noise [89].

First, we write the new expressions for the DCPD sum and null PSDs including squeezing:

$$\begin{aligned}
\langle d_+, d_+ \rangle &= \frac{1}{|1-G|^2}\big(\langle n_a, n_a\rangle + \langle n_b, n_b\rangle + 2\langle n_a, n_b\rangle + |C|^2\langle n_c, n_c\rangle\big) \\
\langle d_+, d_+ \rangle &= \frac{1}{|1-G|^2}\big(\langle n_+, n_+\rangle + |C|^2\langle n_c, n_c\rangle\big) \\
\langle d_+, d_+ \rangle &= \frac{1}{|1-G|^2}\big(4S^2 e^{-2r} + |C|^2\langle n_c, n_c\rangle\big) \qquad (5.28)\\
\langle d_-, d_- \rangle &= \langle n_a, n_a\rangle + \langle n_b, n_b\rangle - 2\langle n_a, n_b\rangle \\
\langle d_-, d_- \rangle &= \langle n_-, n_-\rangle \\
\langle d_-, d_- \rangle &= 4S^2 \qquad (5.29)
\end{aligned}$$

where we used Eqs. 5.21 and 5.22 to simplify to the final expressions.

Then, solving for the squeeze ratio e^{-2r} using Eqs. 5.28, 5.29, and the correlated noise $\langle n_c, n_c\rangle$ calculated via Eq. 5.12:

$$e^{-2r} = \frac{|1-G|^2\langle d_+, d_+\rangle - |C|^2\langle n_c, n_c\rangle}{\langle d_-, d_-\rangle} \qquad (5.30)$$

Figure 5.6 plots the squeezing levels estimated via Eq. 5.30.

5.3 Results

All spectral densities in this section were taken using median-averaging to avoid the frequent glitches, with phase compensated and mean-to-median biasing corrected according to Chapter 6. They were also verified using the "PSD rejection" technique for removing glitches described in Appendix F.

The unsqueezed correlated spectrum in Figure 5.3 gives a broader picture of the "mystery noise" limiting LIGO sensitivity at 30 Hz and below. Above $3\ \mathrm{kHz}$, the correlated noise is consistent with the laser intensity noise coupling to DARM. Around $300\ \mathrm{Hz}$ the correlated noise approaches the coatings thermal noise limit. Below $300\ \mathrm{Hz}$, conventional "mystery noise" which limits DARM also limits the correlated noise. Around $1\ \mathrm{kHz}$, there is a large gap between the measured correlated noise and the expected sum that is also not understood.

The remaining plots compare correlated noise results from a single fifteen hour lock stretch, where squeezing was injected for six hours, then the squeezer was turned off for nine hours. Figures 5.4 and 5.5 plot the amplitude and phase of the correlated noise as calculated for both the no squeezing time (Eq. 5.12) and the

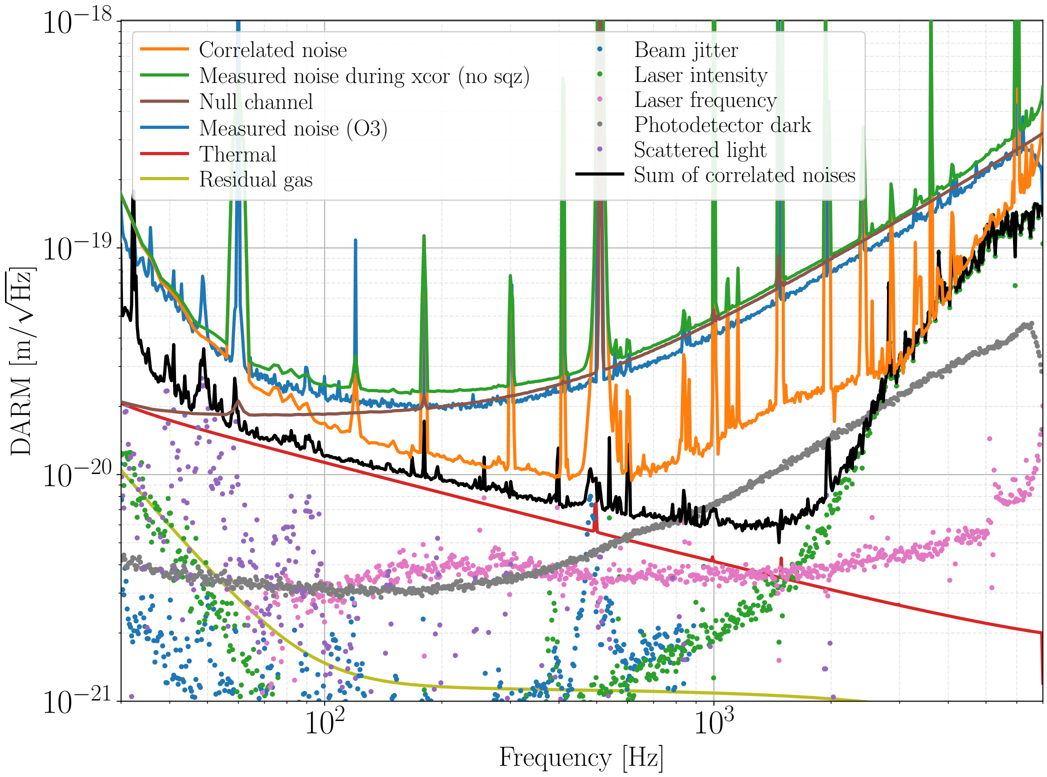

Figure 5.3: LIGO Hanford correlated noise budget during nine hours without squeezing in August 2020. The correlated noise from the interferometer, seen here in orange, is calculated via Eq. 5.12. This can be directly compared to the sum of correlated noise budget traces, seen in black.

squeezing time (Eq. 5.27). Also plotted in Figure 5.4 is the unsqueezed DCPD sum PSD (Eq. 5.5), the squeezed DCPD sum PSD (Eq. 5.28), and the DCPD null (Eq. 5.29).

Eq. 5.27 shows how classical and quantum correlated noise both show up in the final expression. Recall from Eq. 5.23 that the quantum correlated noise is negative. This causes the classical and quantum correlated noise to cancel each other out, leading to classical and quantum correlated noise dominated regimes. In Figure 5.4, the crossing of the squeezed sum PSD in green and the null PSD in brown, corresponds to the dips in the correlated noise, signifying the change from classical- to quantum-dominated correlated noise.

In Figure 5.5, the measured phase of the squeezed correlated noise is shown to be $180°$ in the quantum-dominated regime. The unsqueezed spectrum does not exhibit this phase change.

Because we have the squeezed DCPD sum, null, and the classical correlated noise estimates all from the same lock stretch, we can better estimate the squeezing lev-

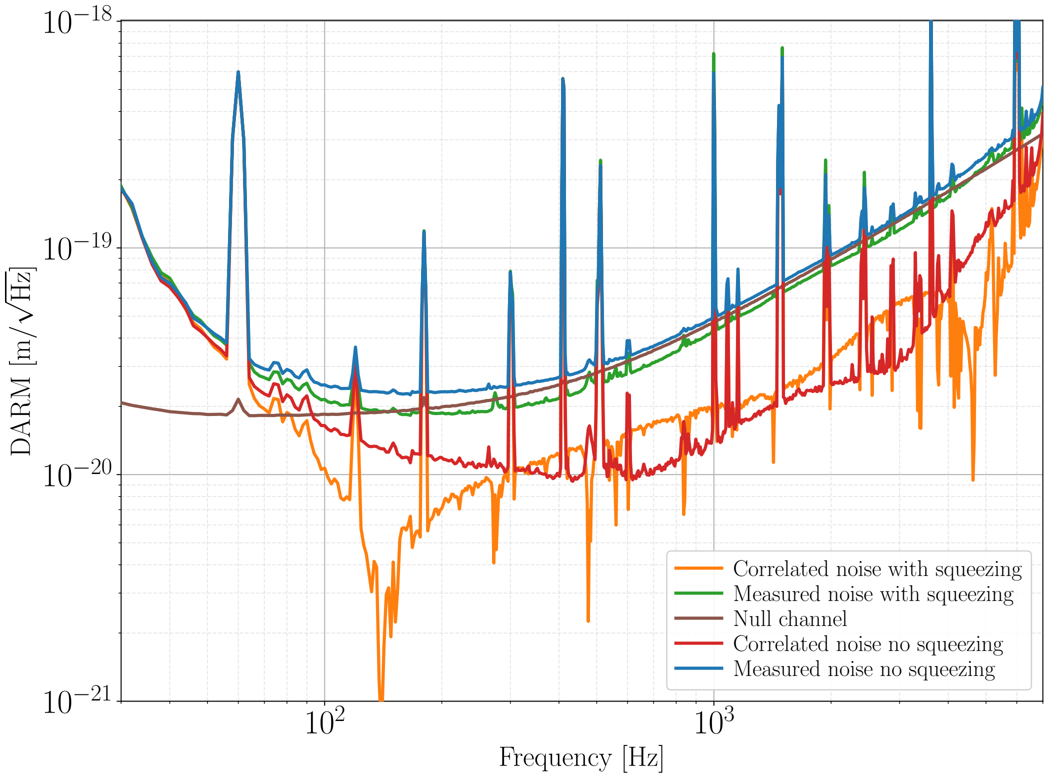

Figure 5.4: LIGO Hanford correlated noise during six hours with squeezing and nine hours without squeezing on September 16, 2020. Squeezed light correlates the shot noise detected on each DCPD, as calculated in Eq. 5.27. The dip at 150 Hz and 5 kHz comes from the interaction of squeezed shot noise and classical noise canceling each other out. Figure 5.5 plots the phase of the correlated noise traces shown here. This plot mirrors a similar study done at LIGO Livingston [197].

els using Eq. 5.30. Expressing the squeeze ratio in terms of dB such that $e^{-2r} = 10^{-\frac{\text{dB}_{\text{sqz}}}{10}}$ yields the estimate shown in Figure 5.6. The squeezing exhibited by the LIGO Hanford detector is frequency-dependent, with the largest squeezing of $\sim$ 2.5 dB in the 100 to 300 Hz region, up to 1 dB above 1 kHz. Hanford's O3 unintentional frequency-dependent squeezing is currently under further investigation.

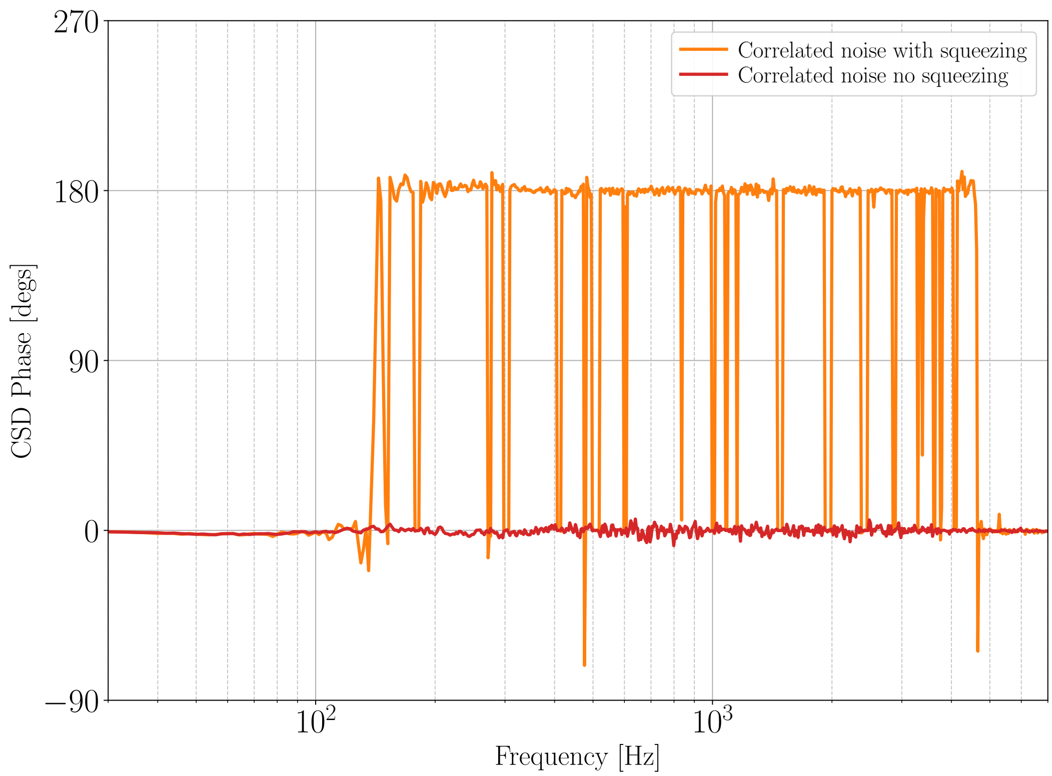

Figure 5.5: Phase of the correlated noise with and without squeezing. The sign flip at at 150 Hz corresponds to the transition from classical correlated noise $\langle n_c, n_c \rangle$ dominating the spectrum to squeezed shot noise $\langle n_a, n_b \rangle$ dominating. Figure 5.4 plots the amplitude of the correlated noise traces shown here.

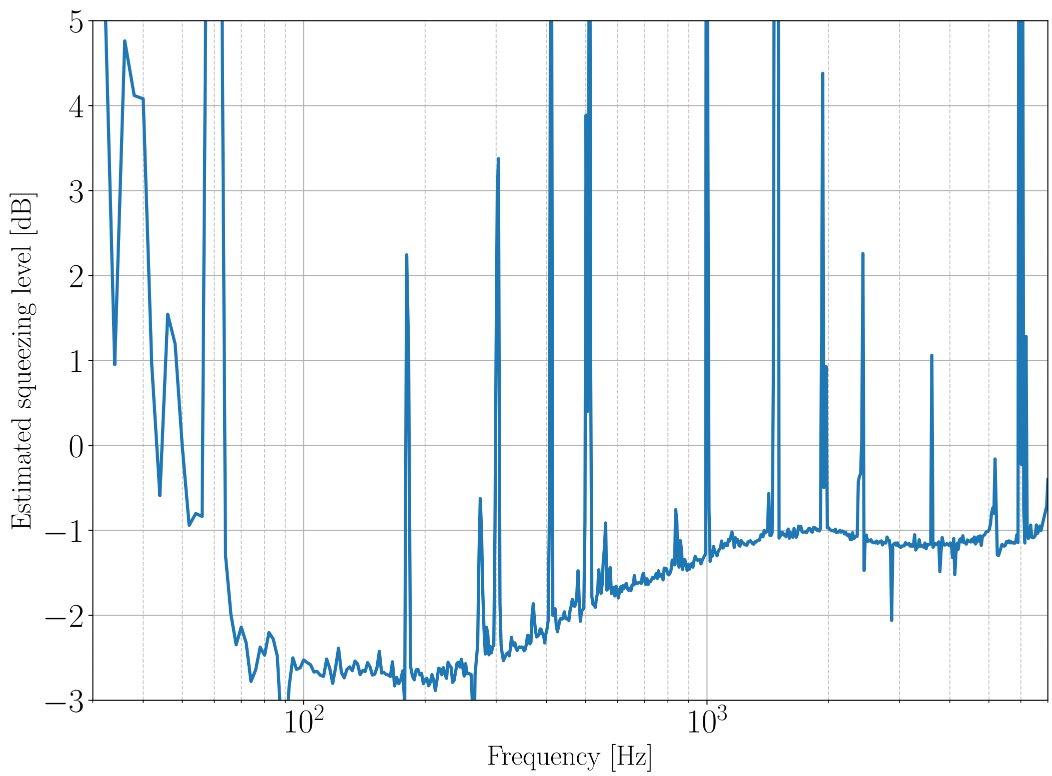

Figure 5.6: Squeezing levels estimated by removing correlated noise from the squeezed DCPD sum, according to Eq. 5.30. This estimate is good in the region where shot noise and correlated noise are about equivalent, or everywhere below ~ 70 Hz.

5.4 Future Work

The correlated noise budget is useful for verifying the DARM noise budget traces, and determining where the classical noise is under the quantum shot noise. The broad range of unexplained noise in the correlated noise, from familiar "mystery noise" at low frequency, to the "correlated mystery noise" between 1 and 3 kHz which is only a factor of 3 below squeezed DARM, means there is much important work remaining to be done understanding what lies below the shot noise. If they can be improved, correlated noise spectra could verify the Advanced LIGO thermal noise floor estimated from the coating Brownian noise for the titania-doped silica tantala optic coatings [32].

The correlated noise could potentially be used to improve sensitivity to continuous gravitational wave signals, such as the stochastic background or continuous waves from spinning neutron stars. The injection of squeezing can confuse such an analysis if the quantum or classical correlated noise is not stable.

Future detectors, including A+, are expected to use balanced homodyne detection, rather than DC readout detection [198]. Balanced homodyne uses a local oscillator to beat with the GW signal rather than light from the interferometer via a DC offset in the DARM loop. The usual two-photodetector detection scheme would not allow for the correlated noise spectrum to be measured, since the amplitude noise on the local oscillator would dominate that spectrum [194]. To recover the correlated noise spectrum a four-photodetector scheme would need to be employed.

Chapter 6

PROBABILITY DISTRIBUTIONS FOR SPECTRAL DENSITIES

Below we review the probability distributions relevant for analyzing noise in the frequency domain. The methods explored below are used in the analysis in Chapter 5.

The goal of this chapter is to derive the probability distributions for power and cross spectral density estimators, and provide convenient formulae and examples for proper statistical treatment of these estimators. Additionally, we discuss the pros and cons of mean-averaging and median-averaging using Welch's method, and explicitly calculate the expected mean-to-median bias for both power and cross spectral densities.

Often, the jitter in an estimate of a power spectral density or transfer function is called "noise", when in fact it is just the manifestation of statistical uncertainty in the estimate. True noise in a gravitational-wave detector is any power measured in the gravitational-wave signal channel which is not gravitational-wave signal. Shot noise, thermal noise, seismic noise are all sources of true noise that obfuscate signal. True noise can never be "averaged away", but statistical uncertainty can be reduced by taking additional averages.

For precision interferometry, it is imperative to understand uncertainty in our estimators. Blackman and Tukey provide a practical understanding of the statistics of power spectra derived from Gaussian noise [199]. Goodman derived the general probability distributions for spectral matrices as complex Wishart distributions, and extended this understanding to the distributions describing multiple coherence functions and transfer functions [200–203]. Bendat and Piersol summarize the statistics of spectral densities and transfer functions, and derive approximations that are good in the limit that the number of averages $n \to \infty$ and coherence $\gamma^2 \to 1$ [174].

In this chapter, we start from Gaussian noise and derive the distributions associated with a single-input single-output linear system. Also provided are numerical verifications of the derived formula.

6.1 Random variables and probability functions

The fundamentals of noise are couched in random variables and their associated probability distributions. A noise process, or random process, is the set of all possible data that could be generated from a random variable $x(t)$ [174]. Although it is impossible to exactly predict an observation from a random variable, the results can characterized via probability density functions (PDFs).

We define $f_{\mathcal{X}}(x)$ to be the probability density function of a stationary, ergodic random variable $\mathcal{X}$ such that $f_{\mathcal{X}}(x) \geq 0 \; \forall x$ and

$$\mathrm{Prob}[a < x < b] = \int_a^b f_{\mathcal{X}}(x) dx. \tag{6.1}$$

In words, the probability that x falls between the values of a and b is the integral of the PDF from a to b. As the name suggests, the PDF can be thought to have units of $[\mathrm{probability/units\,of}\; x]$.

If we let $a \to -\infty$ we arrive at the definition of the cumulative distribution function (CDF) $F_{\mathcal{X}}(x)$:

$$F_{\mathcal{X}}(x) = \int_{-\infty}^{x} f_{\mathcal{X}}(y) dy \tag{6.2}$$

If we let $x \to \infty$ then $F_{\mathcal{X}}(x) \to 1$.

The *expected value*, or mean, μ is a weighted integral over all possible values of x:

$$\langle x \rangle = \mu = \int_{-\infty}^{\infty} x f_{\mathcal{X}}(x) dx. \tag{6.3}$$

The *median* ρ is the value of x such that $F_{\mathcal{X}}(\rho) = 1/2$.

The *variance* σ^2 is the mean square value about the mean of the data:

$$\sigma^2 = \left\langle (x - \langle x \rangle)^2 \right\rangle = \left\langle x^2 \right\rangle - \langle x \rangle^2 = \int_{-\infty}^{\infty} (x - \mu)^2 f_{\mathcal{X}}(x) dx. \tag{6.4}$$

Variance in the time domain will often be synonymous with power in the frequency domain.

The *characteristic function* $\varphi_{\mathcal{X}}(t)$ is the Fourier transform of a probability density function:

$$\varphi_{\mathcal{X}}(t) = \left\langle e^{it\mathcal{X}} \right\rangle = \int_{\mathbb{R}} f_{\mathcal{X}}(x) e^{itx} dx. \tag{6.5}$$

Characteristic functions are another way of describing the random variable $\mathcal{X}$, and are useful tools for deriving relationships between probability density functions, as we will use in Sections 6.10 and 6.11.

6.2 Spectral analysis

The fundamental building block of spectral analysis is the Fourier transform:

$$\mathcal{F}[x(t)] = X(f) = \int_{-\infty}^{\infty} x(t)e^{-i2\pi ft}dt \tag{6.6}$$

A Fourier transform breaks down a signal in the time domain $x(t)$ into its periodic components in the frequency-domain $X(f)$.

In general, the frequency domain extends from negative frequencies to positive frequencies: $f \in (-\infty, \infty)$. If $x(t)$ is real-valued, then its Fourier transform is Hermitian symmetric: $X(f) = X^*(-f)$. Thus all frequency information in $x(t)$ is contained in the positive frequency domain of $X(f)$.

Parseval's theorem for Fourier transforms conserves energy in both the time and frequency domains:

$$\int_{-\infty}^{\infty} |x(t)|^2 dt = \int_{-\infty}^{\infty} |X(f)|^2 df. \tag{6.7}$$

The *autocorrelation* function $R_x(\tau)$ and *cross correlation* function $R_{xy}(\tau)$ are measures of how related a signal $x(t)$ is with itself or another signal $y(t)$ over some time lag τ:

$$R_x(\tau) = \langle x(t)x(t+\tau)\rangle \tag{6.8}$$

$$R_{xy}(\tau) = \langle x(t)y(t+\tau)\rangle \tag{6.9}$$

The *double-sided power spectral density* and *double-sided cross spectral density*, $\Xi_x(f)$ and $\Xi_{xy}(f)$, are the Fourier transforms of the correlation functions $R_x(\tau)$ and $R_{xy}(\tau)$:

$$\Xi_x(f) = \int_{-\infty}^{\infty} R_x(\tau)e^{-i2\pi f\tau}d\tau = \int_{-\infty}^{\infty} \langle x(t)x(t+\tau)\rangle e^{-i2\pi f\tau}d\tau \tag{6.10}$$

$$\Xi_{xy}(f) = \int_{-\infty}^{\infty} R_{xy}(\tau)e^{-i2\pi f\tau}d\tau = \int_{-\infty}^{\infty} \langle x(t)y(t+\tau)\rangle e^{-i2\pi f\tau}d\tau \tag{6.11}$$

on the interval $f \in (-\infty, \infty)$. The spectral densities $\Xi_x(f)$ and $\Xi_{xy}(f)$ can also be expressed as the product of Fourier transforms:

$$\Xi_x(f) = \int_{-\infty}^{\infty} \langle X^*(f)X(g)\rangle e^{-i2\pi(g-f)t}dg \tag{6.12}$$

$$\Xi_{xy}(f) = \int_{-\infty}^{\infty} \langle X^*(f)Y(g)\rangle e^{-i2\pi(g-f)t}dg \tag{6.13}$$

Since we cannot measure negative frequencies in reality, and we measure real-valued signals $x(t)$ and $y(t)$, we define the *single-sided power spectral density* and *single-sided cross spectral density*

$$S_x(f) = \langle x, x \rangle = 2\Xi_x(f) \tag{6.14}$$

$$S_{xy}(f) = \langle x, y \rangle = 2\Xi_{xy}(f) \tag{6.15}$$

on the interval $f \in [0, \infty)$.

Finally, we define the *coherence* of two signals $x(t)$ and $y(t)$

$$\gamma_{xy}^2(f) = \frac{|S_{xy}(f)|^2}{S_x(f)S_y(f)} \tag{6.16}$$

6.3 Estimators

Once the PDF of a random variable is known, its future behavior is as well-known as possible. The above definitions are good if infinite samples are taken over infinite time. The difficulty lies in estimating the PDF from finite data, as well as estimating the confidence of our estimate.

6.3.1 Properties

For this section, we will take Φ to be some statistic we want to know, such as the true mean, and $\hat{\Phi}$ is its estimator, such as the sample mean.

The estimator *bias* b is the expected value of the difference between an estimator and its true value:

$$b_{\hat{\Phi}} = \left\langle \hat{\Phi} - \Phi \right\rangle. \tag{6.17}$$

If $b = 0$, the estimator $\hat{\Phi}$ is unbiased. If $b \neq 0$, the estimator is biased, meaning the expected value of the estimator does not equal the true value.

The quality of an estimator can change with the number of samples n. An estimator is *consistent* if, as the number of samples n increases, the estimator converges to the true value. That is, for any positive real ϵ,

$$\lim_{n \to \infty} \mathrm{Prob}[|\hat{\Phi} - \Phi| < \epsilon] = 1 \tag{6.18}$$

i.e. the probability that the difference in the estimate and true value is less than arbitrary ϵ is one.

We will find that several estimators of important spectral quantities, including the sample coherence, sample median PSD and CSD, are biased but consistent quantities (Sections D.7, D.4, and D.6).

6.3.2 Example estimators

Suppose we have a data set x_i of N samples. The *sample mean* $\hat{\mu}$ is

$$\hat{\mu} = \frac{1}{N}\sum_{i=0}^{N-1} x_i \tag{6.19}$$

The *sample median* $\hat{\rho}$ is the middle value of the sorted x_i. In other words, if x_i are sorted, then $\hat{\rho} = x_{N/2}$. The sample median tends to be more robust to outliers than the sample mean, as we will see in Section 6.14.

The *sample variance* $\hat{\sigma}^2$ is

$$\hat{\sigma}^2 = \frac{1}{N-1}\sum_{i=0}^{N-1} (x_i - \hat{\mu})^2 \tag{6.20}$$

In reality, signals have a finite length T and have an finite sampling frequency f_s. To approximate the Fourier transform for a real signal $x[n] \triangleq x(t = n/f_s)$ with integer $N = Tf_s$ total samples and integer $n \in [0, N-1]$, the discrete Fourier transform is

$$\hat{X}[k] = \sum_{n=0}^{N-1} x[n]e^{-i2\pi kn/N}. \tag{6.21}$$

Here $k \triangleq Nf/f_s$ is a integer proxy for frequency. The lowest non-zero frequency measurable is the *resolution frequency* or *binwidth* and occurs for $k = 1$, or $f_b = f_s/N = 1/T$. The highest frequency measurable is the *Nyquist frequency*, $f_{\mathrm{Nyquist}} = f_s/2$, and occurs when $k = N/2$.

Note that Eq. 6.21 does not preserve power normalization. In section 6.14 we will introduce a prefactor of $\sqrt{1/Nf_s}$, motivated by Parseval's theorem, to conserve power. If $x[n]$ has units of volts V, this will give $X[k]$ units of $\mathrm{V}/\sqrt{\mathrm{Hz}}$. Additionally, if a window is used, the power loss from the windowing function must be accounted for in the prefactor as well.

The *estimated power spectral density* $\hat{S}_x$ and *estimated cross spectral density* $\hat{S}_{xy}$ are

$$\hat{S}_x[k] = \langle x, x\rangle = \frac{2}{Nf_s}|X[k]|^2 \tag{6.22}$$

$$\hat{S}_{xy}[k] = \langle x, y\rangle = \frac{2}{Nf_s}X^*[k]Y[k]. \tag{6.23}$$

$\hat{S}_x[k]$ and $\hat{S}_{xy}[k]$ have units of V^2/Hz. The factor of 2 comes from the fact these are single-sided spectral densities.

The advantage of spectral densities is, no matter the length of the signal, the noise level will remain constant. This is due to the normalization by the sampling frequency f_s, which adjusts for the fact that longer signals will have a finer binwidth f_b, so the same noise power is divided among more bins. The disadvantage is signals are not constant for spectral densities: they grow like the number of samples N. If a signal has infinite fidelity, all of its power will appear in a single bin no matter the binwidth f_b.

These functions include the power normalization constant assuming no windowing. Windows reduce spectral leakage due to aliasing by enforcing periodicity on the signal $x[n]$. If a window $w[n]$ is applied to each data point $x[n]$ before taking the discrete Fourier transform, there will be power loss associated with the window. To preserve power in each bin, replace N in the equations above with $S_2 = \sum_{n=1}^{N-1} w[n]^2$. For a rectangular window where $w[n] = 1$ for all n, $S_2 = N$ and we recover Eqs. 6.22 and 6.23 [204].

6.4 Welch's method

Welch's method is a method of estimating a power spectral density by averaging together many power spectral densities. The variance associated with a single PSD estimate is relatively high, and equals the value of the PSD itself squared, as we will derive in 6.8. Welch's method builds a distribution of M PSDs and finds the mean of each frequency bin k.

Welch's method takes advantage of the data windowing to add more PSD segments by allowing segments to overlap. This method reuses data points, so care must be taken that the overlap is not too high, which falsely correlates the PSD segments.

Welch's method proceeds as follows:

1. Select three of the following four:

   - The total number of samples N
   - The overlap ratio p, i.e. ratio of samples to share between segments.
   - The number of segments, or averages, M
   - The binwidth or frequency resolution f_b

   These are related by the equation $M = \dfrac{N f_b/f_s - 1}{1-p} + 1$, where f_s is the sampling frequency.

2. Split the N data samples into M equal segments, each segment having samples $n = \dfrac{N}{(M-1)(1-p)+1}$.

3. Apply a window to each data segment, if desired.

4. Estimate the PSD $\hat{S}_{x,i}[k]$ of each segment $i \in [1, M]$ using Eq. 6.22 and the desired window function.

5. For each frequency bin k, average all $\hat{S}_{x,i}[k]$ such that $\overline{\hat{S}_x[k]} = \dfrac{1}{M}\sum_{i=1}^{M} \hat{S}_{x,i}[k]$.

Thus $\overline{\hat{S}_x[k]}$ is the mean-averaged PSD estimate.

In the final step, we might take the sample median $\widetilde{\hat{S}_x[k]}$ of $\hat{S}_{x,i}[k]$ rather than the mean. For Gaussian noise, this asymptotes for large M to a factor of $\log(2)$ bias in the median estimate vs the mean estimate: $\overline{\hat{S}_x[k]} = \dfrac{\widetilde{\hat{S}_x[k]}}{\log(2)}$.

6.5 Probability distribution formulae

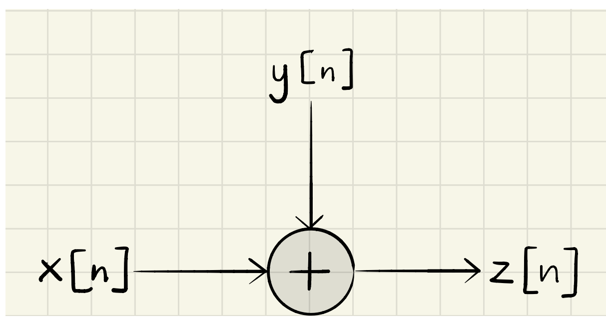

Figure 6.1: Independent Gaussian signals x and y summed into third Gaussian signal z with no delay. These signals will form the basis of the power, amplitude, and cross spectral density probability densities in Sections 6.8, 6.9 and 6.11. Table 6.1 reports the PDFs associated with these signals as derived in this chapter.

In this following sections we will derive the probability distributions of several random variables relevant to signal processing and spectral estimation, introduced in Sections 6.1 and 6.2. Figure 6.1 diagrams the simple Gaussian signals used in this section. Table 6.1 lists the probability distributions investigated here.

We will attempt to chain together the relations of the probability distributions such that the relationship of power in the time and frequency domain is clear. This will motivate why the mean-value of the PSD is a natural estimator of power in the time-domain, so median-value PSD estimators are "biased" and must be corrected back to mean-values.

Through this section we will use the *change of variables formula* for probability densities. Suppose we know the probability distribution $f_{\mathcal{X}}(x)$ of a random variable $\mathcal{X}$, and we have an n-valued function $g(x)$. Then we know the probability density of $\mathcal{G} \sim g(x)$ [174]:

$$f_{\mathcal{G}}(g) = n f_{\mathcal{X}}(x) \left| \left(\frac{dg}{dx} \right)^{-1} \right| \tag{6.24}$$

One important example of the change of variables formula is the scaler formula:

$$g = cx \implies f_{\mathcal{G}}(g) = \frac{1}{c} f_{\mathcal{X}}\left(\frac{g}{c}\right). \tag{6.25}$$

For this example $n = 1$ as g has one-to-one mapping with x. In Section 6.10, we use the two-dimensional verison of the above formula.

The *convolution theorem* states that a Fourier transform of the convolution of two functions $\mathcal{F}(f * g)$ in the time domain is equal to a multiplication in the frequency domain $\mathcal{F}(f)\mathcal{F}(g)$:

$$\mathcal{F}(f * g) = \mathcal{F}(f)\mathcal{F}(g). \tag{6.26}$$

Finally, a sum of random variables $\mathcal{Z} = \mathcal{X} + \mathcal{Y}$ can be expressed as the convolution of PDFs:

$$f_{\mathcal{Z}}(z) = f_{\mathcal{X}}(x) * f_{\mathcal{Y}}(y) = \int_{-\infty}^{\infty} f_{\mathcal{X}}(x) f_{\mathcal{Y}}(z - x) dx. \tag{6.27}$$

The characteristic function of $\mathcal{Z}$, $\varphi_{\mathcal{Z}}$, becomes a product of the characteristic functions of $\mathcal{X}$ and $\mathcal{Y}$:

$$\varphi_{\mathcal{Z}}(t) = \left\langle e^{it\mathcal{Z}} \right\rangle = \left\langle e^{it(\mathcal{X}+\mathcal{Y})} \right\rangle = \left\langle e^{it\mathcal{X}} \right\rangle \left\langle e^{it\mathcal{Y}} \right\rangle = \varphi_{\mathcal{X}}(t)\varphi_{\mathcal{Y}}(t). \tag{6.28}$$

This property of characteristic functions can be extended for the sum of n independent, identically distributed random variables, such as for the sample mean $\hat{\mu} = \frac{1}{n}\sum_{i=1}^{n} x_i$ drawn from random variable $\mathcal{X}$:

$$\varphi_{\hat{\mu}}(t) = \left\langle e^{it\hat{\mu}} \right\rangle = \left\langle e^{\frac{it}{n}\mathcal{X}} \right\rangle^{n} = \varphi_{\mathcal{X}}\left(\frac{t}{n}\right)^{n}. \tag{6.29}$$

This property will come in handy when calculating the sample means of the PSDs and CSDs in Appendix D.

Table 6.1: Summary of probability density functions derived for spectral estimators in this chapter. Power is conserved for every probablility transform. Two independent, white noise time-domain signals $x[n] \sim \mathcal{N}(0, \sigma_x)$ and $y[n] \sim \mathcal{N}(0, \sigma_y)$ with N total samples and sampling frequency f_s form the base signals. Another white noise signal $z[n] = x[n] + y[n]$ gives a signal correlated with $x[n]$, as shown in Figure 6.1. The Fourier transforms of $x[n], y[n]$ are $\hat{X}[k] = \mathcal{A} + i\mathcal{B}$, $\hat{Y}[k] = \mathcal{C} + i\mathcal{D}$. By linearity, $\hat{Z}[k] = (\mathcal{A} + \mathcal{C}) + i(\mathcal{B} + \mathcal{D})$. $\mathcal{A}$ and $\mathcal{B}$ both follow Gaussian distributions $\mathcal{N}(0, \sigma_a)$ where $\sigma_a = \sigma_x\sqrt{N/2}$, as shown by Eq. 6.31. Similarly, $\mathcal{C}, \mathcal{D} \sim \mathcal{N}(0, \sigma_c)$ where $\sigma_c = \sigma_y\sqrt{N/2}$. The PSD prefactor $2/(Nf_s)$ alters the variance of the resulting exponential distribution, as shown in Eq. 6.39. The CSD of correlated signals $\langle x, z\rangle \propto \mathcal{A}^2 + \mathcal{A}\mathcal{C} + \mathcal{B}^2 + \mathcal{B}\mathcal{D} + i(\mathcal{A}\mathcal{D} - \mathcal{B}\mathcal{C})$.

Estimator	Symbol	Probability Distribution	Expression
White time-domain signal	$x[n]$	Gaussian	$\mathcal{N}(0, \sigma_x)$
Discrete Fourier transform	$\mathcal{A}$	Gaussian	$\mathcal{N}(0, \sigma_a)$
Squared Gaussian	$\mathcal{A}^2$	Scaled chi-squared, n = 1	$\chi_1^2(\sigma_a)$
Power spectral density	$\langle x, x\rangle$	Exponential	$\mathrm{Exp}\left(\frac{2\sigma_x^2}{f_s}\right)$
Amplitude spectral density	$\sqrt{\langle x, x\rangle}$	Rayleigh	$\mathrm{Rayleigh}\left(\frac{\sigma_x}{\sqrt{f_s}}\right)$
Product of two Gaussians	$\mathcal{A}\mathcal{C}$	Modified Bessel of the 2nd kind	$\frac{1}{\pi\sigma_a\sigma_c}K_0\left(\frac{\lvert z\rvert}{\sigma_a\sigma_c}\right)$
Dependent Gaussian product	$\mathcal{A}(\mathcal{A} + \mathcal{C})$	Modified Bessel of the 2nd kind	$\frac{e^{s/\sigma_c^2}}{\pi\sigma_a\sigma_c}K_0\left(\frac{\lvert s\rvert\sqrt{\sigma_a^2 + \sigma_c^2}}{\sigma_a\sigma_c^2}\right)$
Cross spectral density (minor)	$\langle x, z\rangle$	Laplace	$\mathrm{Laplace}\left(0, \frac{\sigma_x\sigma_y}{f_s}\right)$
Cross spectral density (major)	$\langle x, z\rangle$	Asymmetric Laplace	$\mathrm{AL}\left(0, \frac{\sigma_x\sigma_y}{f_s}, -\frac{\sigma_x}{\sigma_y}\left(1 - \sqrt{1 + \frac{\sigma_y^2}{\sigma_x^2}}\right)\right)$
Cross spectral density (joint)	$\langle x, z\rangle$	Modified Bessel of the 2nd kind	$\frac{e^{\frac{u\cos(\phi) + v\sin(\phi)}{\sigma_y^2/f_s}}}{2\pi\sigma_x^2\sigma_y^2/f_s^2}K_0\left(\frac{\sqrt{(u^2+v^2)\left(1+\frac{\sigma_y^2}{\sigma_x^2}\right)}}{\sigma_y^2/f_s}\right)$

6.6 Gaussian distribution and the Fourier transform

A random variable $\mathcal{A}$ follows a *Gaussian*, or *normal distribution* $\mathcal{N}(\mu, \sigma)$ if its samples come from

$$\mathcal{N}(\mu, \sigma) = f_{\mathcal{A}}(a|\mu, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(a-\mu)^2}{2\sigma^2}} \tag{6.30}$$

where μ is the mean, σ^2 is the variance, and σ is the standard deviation. Gaussian distributions describe a great many physical processes, including white noise processes like shot noise and Johnson noise. From the central limit theorem, a Gaussian describes the mean of many independent, identically distributed random variables, no matter what the distribution of the random variables is.

One crucial proof is the Fourier transform of a Gaussian is another Gaussian with inverted variance: $\mathcal{F}[\mathcal{N}(0, \sigma)] = \mathcal{N}(0, 1/\sigma)$.

Using this, along with Eqs. 6.25, 6.26, and 6.27, we prove that for Gaussian white noise $x[n] \sim \mathcal{N}(0, \sigma_x)$, the PDF of the real part of the discrete Fourier transform is another Gaussian with variance $\sigma_x^2 N/2$.

If $\mathcal{A} \sim \Re\left(\hat{X}[k]\right) = \sum_{n=0}^{N-1} x[n]\cos\left(\frac{2\pi kn}{N}\right)$, then

$$
\begin{aligned}
f_{\mathcal{A}}(a) &= \text{Prob}\left[\sum_{n=0}^{N-1} x[n]\cos\left(\frac{2\pi kn}{N}\right)\right] \\
f_{\mathcal{A}}(a) &= \text{Prob}\left[x[0]\cos\left(\frac{2\pi k0}{N}\right)\right] * \cdots * \text{Prob}\left[x[n]\cos\left(\frac{2\pi k(N-1)}{N}\right)\right] \\
f_{\mathcal{A}}(a) &\propto \mathcal{N}(0, \sigma_x \cos(2\pi k0/N)) * \cdots * \mathcal{N}(0, \sigma_x \cos(2\pi k(N-1)/N)) \\
\mathcal{F}[f_{\mathcal{A}}] &\propto \prod_{n=0}^{N-1} \mathcal{F}[\mathcal{N}(0, \sigma_x \cos(2\pi k0/N))] \\
\mathcal{F}[f_{\mathcal{A}}] &\propto \prod_{n=0}^{N-1} \mathcal{N}\left(0, \frac{1}{\sigma_x \cos(2\pi kn/N)}\right) \\
\mathcal{F}[f_{\mathcal{A}}](\xi) &\propto \prod_{n=0}^{N-1} \exp\left(-\frac{\xi^2}{2}\sigma_x^2 \cos^2(2\pi kn/N)\right) \\
\mathcal{F}[f_{\mathcal{A}}](\xi) &\propto \exp\left(-\frac{\xi^2}{2}\sigma_x^2 \sum_{n=0}^{N-1} \cos^2(2\pi kn/N)\right) \\
\mathcal{F}[f_{\mathcal{A}}](\xi) &\propto \exp\left(-\frac{\xi^2}{2}\sigma_x^2 \frac{N}{2}\right) \\
f_{\mathcal{A}}(a) &\propto \exp\left(-\frac{a^2}{2}\frac{2}{\sigma_x^2 N}\right) \\
f_{\mathcal{A}}(a) &\propto \mathcal{N}\left(0, \sigma_x\sqrt{\frac{N}{2}}\right)
\end{aligned}
\tag{6.31}
$$

A similar argument with $\sin(2\pi kn/N)$ gives $\Im\left(\hat{X}[k]\right) \sim \mathcal{N}\left(0, \sigma_x\sqrt{\frac{N}{2}}\right)$. Since sin and cos are orthogonal, $\Re\left(\hat{X}[k]\right)$ and $\Im\left(\hat{X}[k]\right)$ are independent. This gives the result reported in Table 6.1.

For the remainder of this section, we define the Fourier transform of $x[n]$ as $\hat{X}[k] = \mathcal{A} + i\mathcal{B}$, and the Fourier transform of $y[n]$ as $\hat{Y}[k] = \mathcal{C} + i\mathcal{D}$. As proven above, $\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{D}$ are all independent Gaussian variables. We define $\sigma_a = \sigma_x\sqrt{N/2}$ and $\sigma_c = \sigma_y\sqrt{N/2}$ so that $\mathcal{A}, \mathcal{B} \sim \mathcal{N}(0, \sigma_a)$ and $\mathcal{C}, \mathcal{D} \sim \mathcal{N}(0, \sigma_c)$.

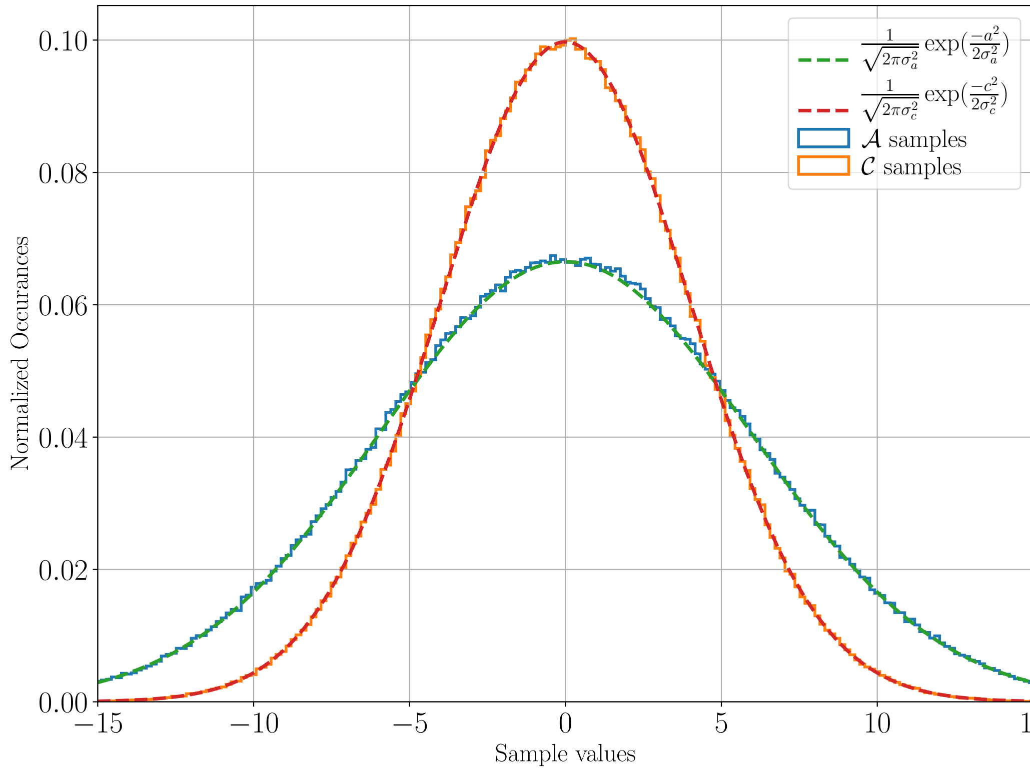

Figure 6.2: Histogram of Gaussian random variables $\mathcal{A}$ and $\mathcal{C}$. Gaussian random variables describe the real and imaginary part of the Fourier transform of Gaussian noise. Equation 6.30 is plotted as the green and red dashed curves. In this example, $\sigma_a = 6$ and $\sigma_c = 4$.

6.7 Chi-squared distribution

A *chi-squared distribution* describes a random variable $\mathcal{Z} = \sum_{i=1}^{n} \mathcal{X}_i^2$ where each $\mathcal{X}_i$ is a standard normal random variable $\mathcal{N}(0, 1)$. The distribution follows

$$\chi_n^2 = f(z|n) = \frac{z^{\frac{n}{2}} e^{-\frac{z}{2}}}{2^{\frac{n}{2}} \Gamma\left(\frac{n}{2}\right)} \qquad z \in [0, \infty) \tag{6.32}$$

where n is the number of degrees of freedom, or random variables in the sum, and Γ is the gamma function.

A chi-squared with one degree of freedom χ_1^2 is equivalent to the product of a standard Gaussian $\mathcal{N}(0, 1)$ with itself: $\mathcal{Z} = \mathcal{N}^2$. For a scaled $\chi_1^2(\sigma)$ coming from a zero-mean, nonstandard Gaussian $\mathcal{N}(0, \sigma)$, we can use change of variables. Let

$z = g(x) = x^2$. From Eq. 6.24, $n = 2$, $\dfrac{dg}{dx} = 2x = 2\sqrt{g}$, and,

$$f_{\mathcal{G}}(g) = 2f_{\mathcal{X}}(\sqrt{g})\left|\frac{1}{2\sqrt{g}}\right| \tag{6.33}$$

$$\chi_1^2(\sigma) = \frac{1}{\sqrt{2\pi\sigma^2 g}}e^{-\frac{g}{2\sigma^2}} \qquad g \in [0, \infty) \tag{6.34}$$

This gives the result reported in Table 6.1.

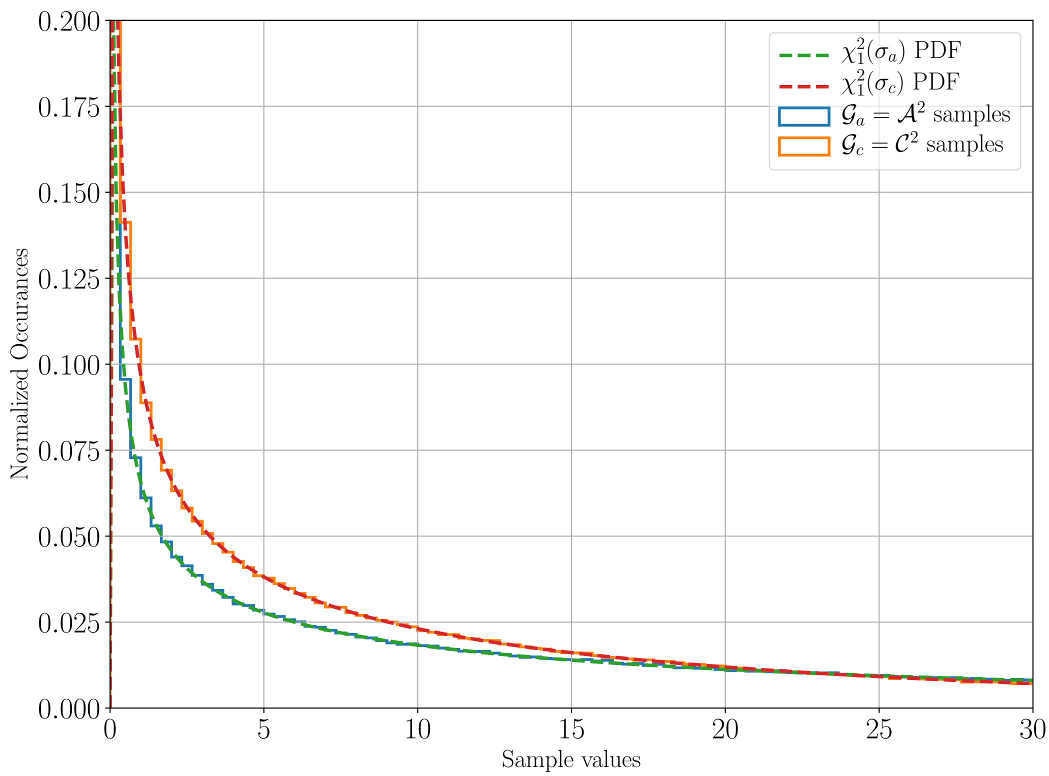

Figure 6.3: Histogram of scaled chi-squared random variables $\mathcal{G}_a$ and $\mathcal{G}_c$. The random variables $\mathcal{G}_a = \mathcal{A}^2$ and $\mathcal{G}_c = \mathcal{C}^2$, where Gaussian random variables $\mathcal{A} \sim \mathcal{N}(0, \sigma_a)$ and $\mathcal{C} \sim \mathcal{N}(0, \sigma_c)$. Equation 6.33 is plotted as the green and red dashed curves. In this example, $\sigma_a = 6$ and $\sigma_c = 4$.

6.8 Exponential distribution and the power spectral density

One special case of the chi-squared distribution is when $n = 2$, we recover the *exponential distribution*:

$$\mathrm{Exp}(\lambda) = f(z|\lambda) = \frac{1}{\lambda}e^{-\frac{z}{\lambda}} \qquad z \in [0, \infty) \tag{6.35}$$

where λ is the mean.

The exponential distribution median m is found by

$$\frac{1}{2} = \int_0^m \frac{1}{\lambda} e^{-\frac{z}{\lambda}} dz = 1 - e^{-\frac{m}{\lambda}}$$
$$m = \lambda \log(2) \tag{6.36}$$

The mean-to-median bias factor b is the ratio of two statistics:

$$b = \frac{m}{\lambda} = \log(2) \tag{6.37}$$

This gives the factor of $\log(2)$ difference in the mean- and median-averaging methods for Welch's method for PSD estimation.

The exponential distribution variance σ^2 is found by

$$\sigma^2 = \left\langle z^2 \right\rangle - \left\langle z \right\rangle^2 = \int_0^\infty z^2 \frac{1}{\lambda} e^{-\frac{z}{\lambda}} dz - \left(\int_0^\infty z \frac{1}{\lambda} e^{-\frac{z}{\lambda}} dz \right)^2 = 2\lambda^2 - \lambda^2$$
$$\sigma^2 = \lambda^2. \tag{6.38}$$

Therefore the variance of a PSD estimate $\mathrm{Var}\left(\hat{S}_x[k]\right)$ is equal to its mean squared.

6.8.1 Power spectral densities

The exponential distribution describes estimated PSDs $\hat{S}_x[k] = \langle x, x \rangle$ of Gaussian random noise with units V^2/Hz. The estimated PSD is the sum of the squares of our Fourier transform real and imaginary components $\mathcal{A}, \mathcal{B}$:

$$\hat{S}_x[k] = \langle x, x \rangle = \frac{2}{N f_s} \left| \hat{X}[k] \right|^2$$
$$\hat{S}_x[k] = \langle x, x \rangle = \frac{2}{N f_s} (\mathcal{A}^2 + \mathcal{B}^2). \tag{6.39}$$

If $\mathcal{A}$ and $\mathcal{B}$ are zero-mean Gaussians with the same distribution $\mathcal{N}(0, \sigma)$, then $\mathcal{A}^2$ and $\mathcal{B}^2$ are both scaled chi-squared random variables as in Eq. 6.33. We can show the random variable $\mathcal{Z} \sim \mathcal{A}^2 + \mathcal{B}^2 \sim \mathrm{Exp}(2\sigma^2)$ via convolution:

$$f_{\mathcal{Z}}(z) = \int_{-\infty}^{\infty} f_{\mathcal{A}^2}(a) f_{\mathcal{B}^2}(z - a) da$$
$$f_{\mathcal{Z}}(z) = \int_0^z \chi_1^2(\sigma) \chi_1^2(\sigma) da$$
$$f_{\mathcal{Z}}(z) = \frac{1}{2\pi\sigma^2} e^{-\frac{z}{2\sigma^2}} \int_0^z \frac{1}{\sqrt{a}\sqrt{z - a}} da$$
$$f_{\mathcal{Z}}(z) = \frac{1}{2\sigma^2} e^{-\frac{z}{2\sigma^2}} = \mathrm{Exp}(2\sigma^2) \tag{6.40}$$

If $\mathcal{A}^2, \mathcal{B}^2 \sim \chi_1^2\left(\sigma_x\sqrt{\frac{N}{2}}\right)$ as in Table 6.1, then $\mathcal{A}^2 + \mathcal{B}^2 \sim \text{Exp}(\sigma_x^2 N)$.

Then, by scaling random variables (Eq. 6.25), we can recover the PSD $\langle x, x\rangle$ distribution in Table 6.1:

$$\frac{2}{N f_s}\text{Exp}(\sigma_x^2 N) \sim \text{Exp}\left(\frac{2\sigma_x^2}{f_s}\right). \tag{6.41}$$

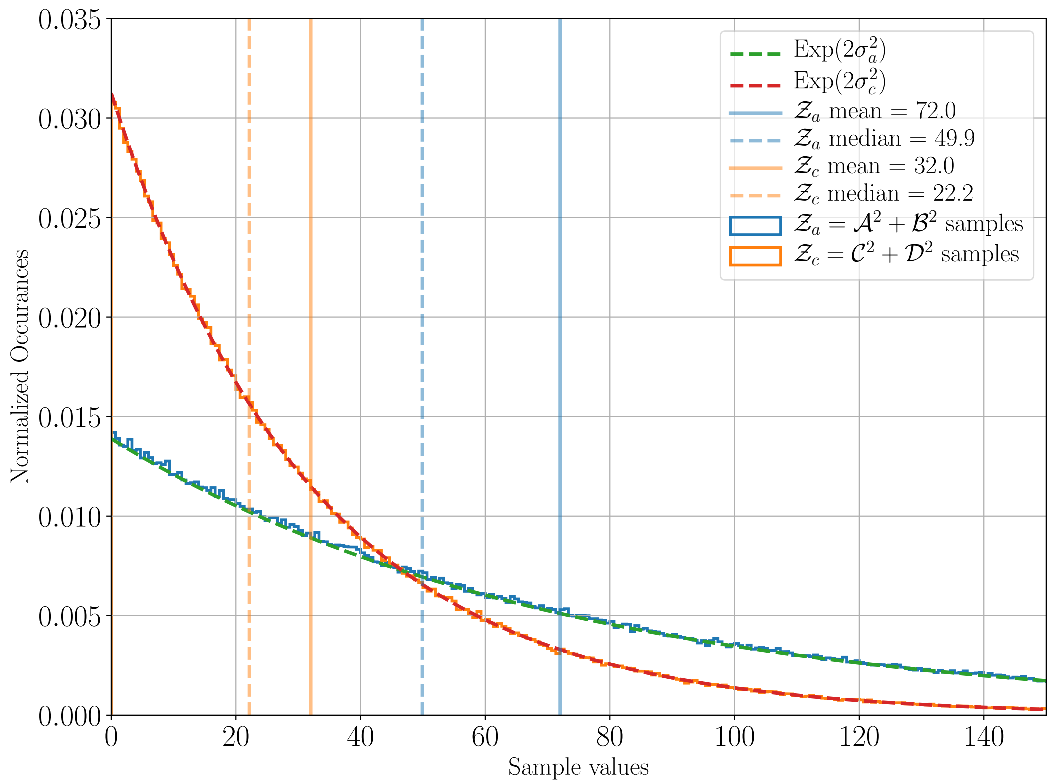

Figure 6.4: Histograms of exponential random variables $\mathcal{Z}_a$ and $\mathcal{Z}_c$. The random variables $\mathcal{Z}_a = \mathcal{A}^2 + \mathcal{B}^2$ and $\mathcal{Z}_c = \mathcal{C}^2 + \mathcal{D}^2$, where Gaussian random variables $\mathcal{A}, \mathcal{B} \sim \mathcal{N}(0, \sigma_a)$ and $\mathcal{C}, \mathcal{D} \sim \mathcal{N}(0, \sigma_c)$. The random variable $\mathcal{Z}_a \propto \langle x, x\rangle$, the PSD estimate of x, and $\mathcal{Z}_c \propto \langle y, y\rangle$. Equation 6.40 is plotted as the green and red dashed curves. In this example, $\sigma_a = 6$ and $\sigma_c = 4$, which gives us the mean of each distribution $2\sigma_a^2 = 72$ and $2\sigma_c^2 = 32$. Both medians are a factor of $\log(2)$ below the mean.

6.9 Rayleigh distribution and the amplitude spectral density

A random variable $\mathcal{Z} = \sqrt{\mathcal{X}}$ which is the square root of a exponential random variable $\mathcal{X}$ follows a *Rayleigh distribution*. A Rayleigh distribution follows

$$\text{Rayleigh}(\upsilon) = f(z|\upsilon) = \frac{z}{\upsilon^2}e^{-\frac{z^2}{2\upsilon^2}} \qquad z \in [0, \infty) \tag{6.42}$$

where υ is the mode of the distribution.

The Rayleigh distribution mean μ is found by

$$\mu = \int_0^\infty z \frac{z}{\upsilon^2} e^{-\frac{z^2}{2\upsilon^2}} dz$$
$$\mu = \upsilon\sqrt{\frac{\pi}{2}} \tag{6.43}$$

The Rayleigh distribution median m is found by

$$\frac{1}{2} = \int_0^m \frac{z}{\upsilon^2} e^{-\frac{z^2}{2\upsilon^2}} dz = 1 - e^{-\frac{m^2}{2\upsilon^2}}$$
$$m = \upsilon\sqrt{2\log(2)} \tag{6.44}$$

The Rayleigh distribution root mean square r is found by

$$r = \sqrt{\langle z^2 \rangle}$$
$$r = \sqrt{\int_0^\infty z^2 \frac{z}{\upsilon^2} e^{-\frac{z^2}{2\upsilon^2}} dz}$$
$$r = \upsilon\sqrt{2} \tag{6.45}$$

The ratio of $m/r = \sqrt{\log(2)}$ is the mean vs median bias factor for amplitude spectral densities. The root mean square is used because Welch's method estimates PSDs, takes their average, then the root is taken to calculate the ASD.

The Rayleigh distribution variance σ^2 is found by combining the mean square r^2 and mean μ above

$$\sigma^2 = \left\langle z^2 \right\rangle - \langle z \rangle^2$$
$$\sigma^2 = \frac{4-\pi}{2}\upsilon^2 \tag{6.46}$$

6.9.1 Amplitude spectral densities

The Rayleigh distribution describes estimated amplitude spectral densities $\sqrt{\hat{S}_x[k]}$ with units $\mathrm{V}/\sqrt{\mathrm{Hz}}$. Amplitude spectral densities are how nearly all of noise spectra in LIGO are displayed. The root mean square r is often what is reported for amplitude spectral density estimates, as Welch's method does mean-averaging with PSDs, and that result is square-rooted.

If $\mathcal{X} \sim \mathrm{Exp}(\lambda)$, then $\mathcal{Z} = \sqrt{\mathcal{X}} \sim \mathrm{Rayleigh}\left(\sqrt{\lambda/2}\right)$ can also be derived via change of variables. Let $z = g(x) = \sqrt{x}$. Then, using Eq. 6.24, we have $n = 1$,

$\frac{dg}{dx} = \frac{1}{2\sqrt{x}} = \frac{1}{2g}$, and

$$f_{\mathcal{G}}(g) = f_{\mathcal{X}}(g^2)|2g| \tag{6.47}$$

$$f_{\mathcal{G}}(g) = \frac{2g}{\lambda} e^{-\frac{g^2}{\lambda}} = \text{Rayleigh}\left(\sqrt{\lambda/2}\right) \qquad g \in [0, \infty). \tag{6.48}$$

If $\mathcal{X} \sim \text{Exp}\left(\frac{2\sigma_x^2}{f_s}\right)$ as in Table 6.1, then

$$\sqrt{\mathcal{X}} \sim \text{Rayleigh}\left(\frac{\sigma_x}{\sqrt{f_s}}\right). \tag{6.49}$$

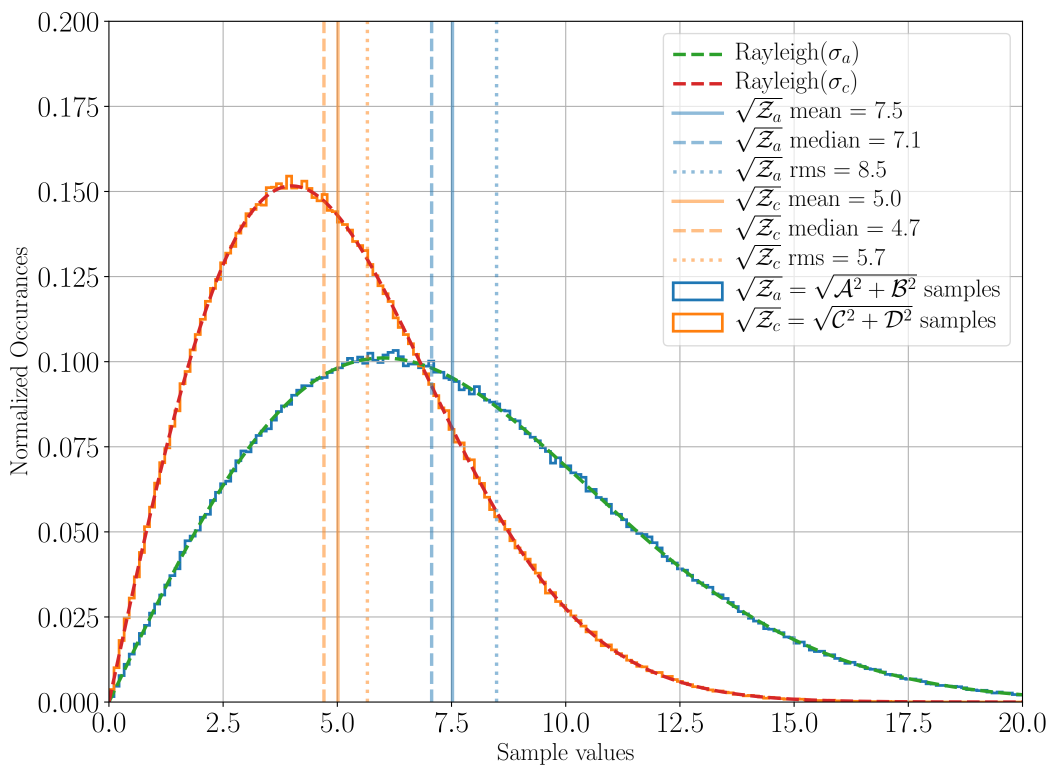

Figure 6.5: Histograms of Rayleigh random variables $\sqrt{\mathcal{Z}_a}$ and $\sqrt{\mathcal{Z}_c}$. The random variables $\sqrt{\mathcal{Z}_a} = \sqrt{\mathcal{A}^2 + \mathcal{B}^2}$ and $\sqrt{\mathcal{Z}_c} = \sqrt{\mathcal{C}^2 + \mathcal{D}^2}$, where Gaussian random variables $\mathcal{A}, \mathcal{B} \sim \mathcal{N}(0, \sigma_a)$ and $\mathcal{C}, \mathcal{D} \sim \mathcal{N}(0, \sigma_c)$. The random variable $\sqrt{\mathcal{Z}_a} \propto \sqrt{\langle x, x \rangle}$, the amplitude spectral density estimate of x, and $\sqrt{\mathcal{Z}_c} \propto \sqrt{\langle y, y \rangle}$. Equation 6.42 is plotted as the green and red dashed curves. In this example, $\sigma_a = 6$ and $\sigma_c = 4$.

6.10 Modified Bessel function of the second kind

We briefly introduce the *zeroth modified Bessel function of the second kind* K_0, as it describes the probability density of product of two different Gaussians $\mathcal{AC}$. First, the definition of the modified Bessel K_ν with $\nu = 0$ from Eq. 10.32.10 of [205] is

$$K_0(z) = \frac{1}{2}\int_0^\infty \exp\left(-t - \frac{z^2}{4t}\right)\frac{dt}{t}. \tag{6.50}$$

which is valid for complex z such that $\arg(z) < \pi/4$.

6.10.1 Product of two Gaussians $\mathcal{AC}$

If we take the convolution of the product of random variables $\mathcal{Z} = \mathcal{AC}$, and let $\mathcal{A} \sim \mathcal{N}(0, \sigma_a)$ and $\mathcal{C} \sim \mathcal{N}(0, \sigma_c)$ be independent, then

$$\begin{aligned} f_\mathcal{Z}(z) &= \int_{-\infty}^{\infty} f_\mathcal{A}(a) f_\mathcal{C}\left(\frac{z}{a}\right)\frac{1}{|a|}da \\ f_\mathcal{Z}(z) &= \frac{1}{2\pi\sigma_a\sigma_c}\int_{-\infty}^{\infty} \exp\left(-\frac{a^2}{2\sigma_a^2}\right)\exp\left(-\frac{z^2}{2a^2\sigma_c^2}\right)\frac{da}{|a|} \\ f_\mathcal{Z}(z) &= \frac{1}{\pi\sigma_a\sigma_c}\int_0^\infty \exp\left(-\frac{a^2}{2\sigma_a^2} - \frac{z^2}{2a^2\sigma_c^2}\right)\frac{da}{a} \end{aligned} \tag{6.51}$$

Let $t = \dfrac{a^2}{2\sigma_a^2}$ and $\dfrac{dt}{da} = \dfrac{a}{\sigma_a^2}$, then $a = \sqrt{2\sigma_a^2 t}$ and $da = \sqrt{\dfrac{\sigma_a^2}{2t}}dt$:

$$\begin{aligned} f_\mathcal{Z}(z) &= \frac{1}{2\pi\sigma_a\sigma_c}\int_0^\infty \exp\left(-t - \frac{z^2}{4t\sigma_a^2\sigma_c^2}\right)\frac{dt}{t} \\ f_\mathcal{Z}(z) &= \frac{1}{\pi\sigma_a\sigma_c}K_0\left(\frac{|z|}{\sigma_a\sigma_c}\right) \end{aligned} \tag{6.52}$$

We introduce the $|z|$ since z is real, $z \in (-\infty, \infty)$, and is symmetric in the integrand.

Setting $\sigma_a = \sigma_x\sqrt{N/2}$ and $\sigma_c = \sigma_y\sqrt{N/2}$ gives the result reported in Table 6.1.

Note that the characteristic function of the PDF described in Eq. 6.52 is [206]

$$\varphi_\mathcal{Z}(t) = \int_{-\infty}^{\infty} f_\mathcal{Z}(z)e^{izt}dz = \frac{1}{\sqrt{\sigma_a^2\sigma_c^2 t^2 + 1}} \tag{6.53}$$

6.10.2 Product of Gaussian with itself and another Gaussian $\mathcal{A}(\mathcal{A}+\mathcal{C})$

In the next section, it will be important that we know the PDF of a random variable $\mathcal{S} = \mathcal{A}(\mathcal{A}+\mathcal{C})$. The difficulty here is the Gaussian $\mathcal{A}$ is used to multiply itself and another Gaussian $\mathcal{C}$, so we must consider the joint PDF.

In the proof we will start with the joint probability distribution of $f_{\mathcal{A}+\mathcal{C}}(a, c)$, which is easy because $\mathcal{A}$ and $\mathcal{C}$ are independent, use change of variables with

$\mathcal{A} \to \mathcal{A}^2$ and $\mathcal{C} \to \mathcal{AC}$ to form the cumulative distribution function $F_{\mathcal{S}}(s)$, then take the derivative of $F_{\mathcal{S}}(s)$ to get the PDF $f_{\mathcal{S}}(s)$.

$$\begin{aligned} f_{\mathcal{A}+\mathcal{C}}(a,c) &= f_{\mathcal{A}}(a) f_{\mathcal{C}}(c) \\ &= \frac{1}{2\pi\sigma_a\sigma_c} \exp\left(-\frac{1}{2}\left(\frac{a}{\sigma_a}\right)^2 - \frac{1}{2}\left(\frac{c}{\sigma_c}\right)^2\right) \end{aligned} \tag{6.54}$$

To do change of variables from $(x_1, x_2) = (a, c)$ to $(y_1, y_2) = (a^2, ac)$, we calculate the determinant of the Jacobian $|J| = \left|\frac{dx_1}{dy_1}\frac{dx_2}{dy_2} - \frac{dx_2}{dy_1}\frac{dx_1}{dy_2}\right|$ to scale the joint distribution:

$$x_1 = \sqrt{y_1} \qquad x_2 = \frac{y_2}{\sqrt{y1}} \tag{6.55}$$

$$\frac{dx_1}{dy_1} = \frac{1}{2\sqrt{y_1}} \qquad \frac{dx_2}{dy_2} = \frac{1}{\sqrt{y_1}} \qquad \frac{dx_2}{dy_1} = -\frac{y_2}{2(y_1)^{3/2}} \qquad \frac{dx_1}{dy_2} = 0 \tag{6.56}$$

$$|J| = \frac{1}{2y_1} \tag{6.57}$$

Now we can write the new joint probability distribution $f_{\mathcal{S}}(y_1, y_2)$ by making substitutions into Eq 6.54. Recall that $n = 2$ for changing variables of a double-valued function $y_1 = x_1^2$.

$$\begin{aligned} f_{\mathcal{S}}(y_1, y_2) &= n|J| f_{\mathcal{A}+\mathcal{C}}(a,c) \\ &= 2\frac{1}{2y_1}\frac{1}{2\pi\sigma_a\sigma_c} \exp\left(-\frac{y_1}{2\sigma_a^2} - \frac{y_2^2}{2y_1\sigma_c^2}\right) \end{aligned} \tag{6.58}$$

The cumulative distribution function $F_{\mathcal{S}}(\mathcal{S} < s) = \text{Prob}(\mathcal{Y}_1 + \mathcal{Y}_2 < s) = \text{Prob}(\mathcal{Y}_2 < s - \mathcal{Y}_1)$ can be written as the double integral over $y_1 \in [0, \infty)$ and $y_2 \in (-\infty, s - y_1)$:

$$F_{\mathcal{S}}(\mathcal{S} < s) = \frac{1}{2\pi\sigma_a\sigma_c} \int_0^\infty dy_1 \int_{-\infty}^{s-y_1} dy_2 \, \frac{1}{y_1} \exp\left(-\frac{y_1}{2\sigma_a^2} - \frac{y_2^2}{2y_1\sigma_c^2}\right) \tag{6.59}$$

Recall that the PDF is the derivative of the CDF:

$$\begin{aligned} f_{\mathcal{S}}(s) &= \frac{dF_{\mathcal{S}}}{ds} \\ &= \frac{1}{2\pi\sigma_a\sigma_c} \int_0^\infty dy_1 \frac{d}{ds}\left(\int_{-\infty}^{s-y_1} dy_2 \, \frac{1}{y_1} \exp\left(-\frac{y_1}{2\sigma_a^2} - \frac{y_2^2}{2y_1\sigma_c^2}\right)\right) \\ &= \frac{1}{2\pi\sigma_a\sigma_c} \int_0^\infty dy_1 \frac{1}{y_1} \exp\left(-\frac{y_1}{2\sigma_a^2} - \frac{(s-y_1)^2}{2y_1\sigma_c^2}\right) \\ &= \frac{1}{2\pi\sigma_a\sigma_c} \exp\left(\frac{s}{\sigma_c^2}\right) \int_0^\infty dy_1 \frac{1}{y_1} \exp\left(\frac{y_1(\sigma_a^2 + \sigma_c^2)}{2\sigma_a^2\sigma_c^2} - \frac{s^2}{2y_1\sigma_c^2}\right) \end{aligned} \tag{6.60}$$

Let $u = \dfrac{y_1(\sigma_a^2 + \sigma_c^2)}{2\sigma_a^2\sigma_c^2}$, then $\dfrac{du}{dy_1} = \dfrac{\sigma_a^2 + \sigma_c^2}{2\sigma_a^2\sigma_c^2}$, and using Eq. 6.50 gives

$$f_{\mathcal{S}}(s) = \frac{1}{2\pi\sigma_a\sigma_c} \exp\left(\frac{s}{\sigma_c^2}\right) \int_0^\infty du \frac{1}{u} \exp\left(-u - \frac{s^2(\sigma_a^2 + \sigma_c^2)}{4\sigma_a^2\sigma_c^4 u}\right)$$

$$f_{\mathcal{S}}(s) = \frac{1}{\pi\sigma_a\sigma_c} \exp\left(\frac{s}{\sigma_c^2}\right) K_0\left(\frac{|s|\sqrt{\sigma_a^2 + \sigma_c^2}}{\sigma_a\sigma_c^2}\right) \tag{6.61}$$

This probability density function is plotted versus numerical samples in Figure 6.6.

Note that the characteristic function of the PDF described in Eq. 6.61 is

$$\varphi_{\mathcal{S}}(t) = \int_{-\infty}^{\infty} f_{\mathcal{S}}(s) e^{-ist} ds = \frac{1}{\sqrt{1 - i2\sigma_a^2 t + \sigma_a^2\sigma_c^2 t^2}} \tag{6.62}$$

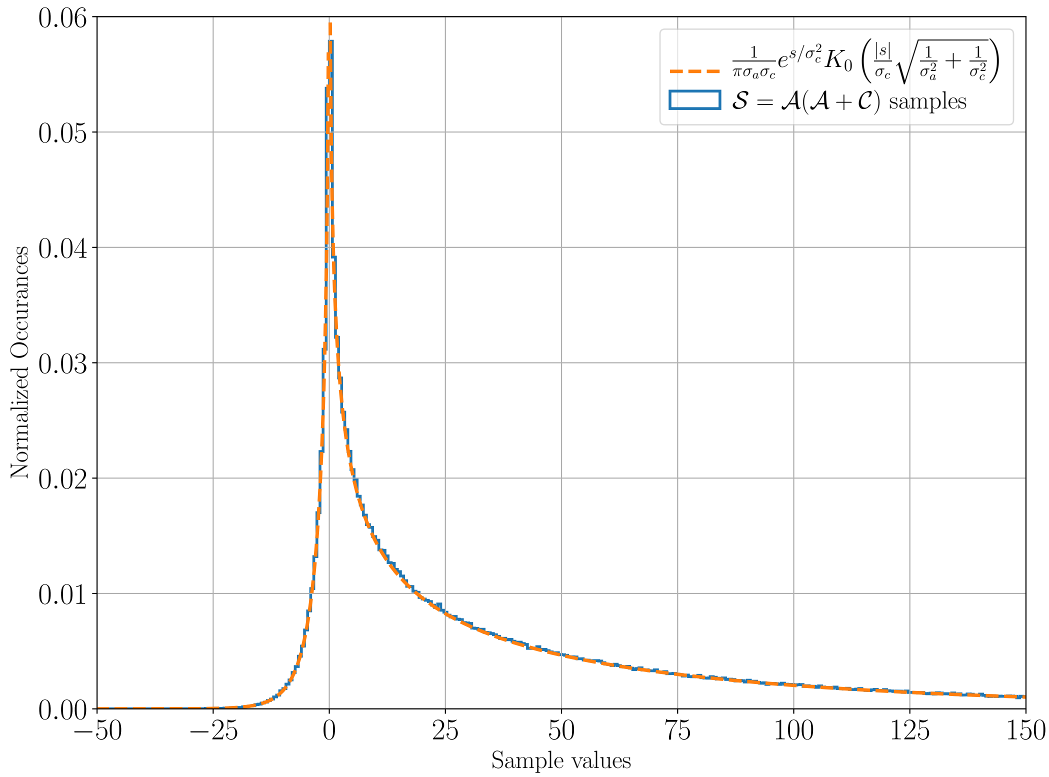

Figure 6.6: Histograms of a random variable $\mathcal{S}$ which follows a distribution described by a modified Bessel function of the second kind. The random variable $\mathcal{S} = \mathcal{A}(\mathcal{A} + \mathcal{C})$, where Gaussian random variables $\mathcal{A} \sim \mathcal{N}(0, \sigma_a)$ and $\mathcal{C} \sim \mathcal{N}(0, \sigma_c)$. Equation 6.61 is plotted as the orange dashed curve. In this example, $\sigma_a = 7$ and $\sigma_c = 3$.

6.11 Laplace, asymmetric Laplace, and the cross spectral density

The *Laplace distribution*, or *double exponential distribution*, is defined as

$$\mathrm{Laplace}(\mu, \gamma) = f(x|\mu, \gamma) = \frac{1}{2\gamma} \exp\left(-\frac{|x-\mu|}{\gamma}\right) \tag{6.63}$$

where μ is the mean and γ is the scale parameter.

The *asymmetric Laplace distribution* is defined as

$$\mathrm{AL}(m, \lambda, \kappa) = f(x|m, \lambda, \kappa) = \frac{1}{\lambda\left(\kappa + \dfrac{1}{\kappa}\right)} \begin{cases} \exp\left(\dfrac{x-m}{\kappa\lambda}\right) & x \leq m \\ \exp\left(-\dfrac{\kappa(x-m)}{\lambda}\right) & x > m \end{cases} \tag{6.64}$$

where m is the location of the peak, λ is the scale parameter, and κ is the asymmetry parameter. When $\kappa = 1$, the asymmetric Laplace becomes the Laplace: $\mathrm{AL}(m, \lambda, 1) = \mathrm{Laplace}(m, \lambda)$.

The asymmetric Laplace distribution mean μ is found by

$$\mu = \int_{-\infty}^{m} x \frac{1}{\lambda\left(\kappa + \frac{1}{\kappa}\right)} \exp\left(\frac{x-m}{\kappa\lambda}\right) dx + \int_{m}^{\infty} x \frac{1}{\lambda\left(\kappa + \frac{1}{\kappa}\right)} \exp\left(-\frac{\kappa(x-m)}{\lambda}\right) dx$$

$$\mu = m + \lambda \frac{1-\kappa^2}{\kappa} \tag{6.65}$$

The asymmetric Laplace distribution median ρ has two solutions depending on whether $\kappa < 1$ or $\kappa > 1$. If $\kappa < 1$, then the distribution is skewed right and the median is greater than m. Otherwise, the distribution is skewed left and the median is less than m. The median is found by

$$\frac{1}{2} = \begin{cases} \displaystyle\int_{\rho}^{\infty} \frac{1}{\lambda\left(\kappa + \frac{1}{\kappa}\right)} \exp\left(-\frac{\kappa(x-m)}{\lambda}\right) dx, & \kappa < 1 \\ \displaystyle\int_{-\infty}^{\rho} \frac{1}{\lambda\left(\kappa + \frac{1}{\kappa}\right)} \exp\left(\frac{x-m}{\kappa\lambda}\right) dx, & \kappa > 1 \end{cases}$$

$$\rho = \begin{cases} m - \dfrac{\lambda}{\kappa} \log\left(\dfrac{1+\kappa^2}{2}\right), & \kappa < 1 \\ m + \kappa\lambda \log\left(\dfrac{1+\kappa^2}{2\kappa^2}\right), & \kappa > 1 \end{cases} \tag{6.66}$$

The asymmetric Laplace distribution variance σ^2 is

$$\sigma^2 = \frac{\lambda^2(1+\kappa^4)}{\kappa^2} \tag{6.67}$$

The mean to median bias b for the asymmetric Laplace distribution has different results depending on κ, as seen from Eq. 6.66. The physically relevant case is where $m = 0$, yielding

$$b = \frac{\rho}{\mu} = \begin{cases} -\dfrac{\log\left(\dfrac{\kappa^2+1}{2}\right)}{1-\kappa^2} & \kappa < 1 \\ \dfrac{\kappa^2 \log\left(\dfrac{\kappa^2+1}{2\kappa^2}\right)}{1-\kappa^2} & \kappa > 1 \end{cases} \tag{6.68}$$

This result will be important for mean- vs median-averaging for CSD estimates.

The characteristic function of the Laplace distribution is

$$\varphi_{\mathrm{Laplace}}(t) = \frac{e^{i\mu t}}{1+\gamma^2 t^2} \tag{6.69}$$

and the characteristic function of the asymmetric Laplace distribution is

$$\varphi_{\mathrm{AL}}(t) = \frac{e^{imt}}{1+\lambda^2 t^2 + it\lambda\left(\kappa - \dfrac{1}{\kappa}\right)} \tag{6.70}$$

6.11.1 Cross spectral densities

The statistics of cross spectral densities of white noise are known to follow a Laplace distribution [201]. The asymmetric Laplace distribution describes estimated CSDs $\hat{S}_{xz}[k] = \langle x, z \rangle$ with units V^2/Hz. Consider the case where $x[n]$ and $y[n]$ are uncorrelated Gaussian noise, and $z[n] = x[n] + y[n]$. Define the Fourier transforms components $\hat{X}[k] = \mathcal{A} + i\mathcal{B}$ and $\hat{Y}[k] = \mathcal{C} + i\mathcal{D}$. As shown in Eq. 6.31, $\mathcal{A}$ and $\mathcal{B}$ are independent and follow the same Gaussian distribution $\mathcal{N}(0, \sigma_a)$. Similarly, $\mathcal{C}, \mathcal{D} \sim \mathcal{N}(0, \sigma_c)$. Then, by linearity in Eq. 6.22,

$$\begin{aligned} \langle x, z \rangle &= \langle x, x \rangle + \langle x, y \rangle \\ &= \frac{2}{N f_s}\left(\left|\hat{X}[k]\right|^2 + \hat{X}^*[k]\hat{Y}[k]\right) \\ &= \frac{2}{N f_s}\left(\mathcal{A}^2 + \mathcal{B}^2 + \mathcal{A}\mathcal{C} + \mathcal{B}\mathcal{D} + i(\mathcal{A}\mathcal{D} - \mathcal{B}\mathcal{C})\right) \\ &= \frac{2}{N f_s}(\mathcal{U} + i\mathcal{V}) \end{aligned} \tag{6.71}$$

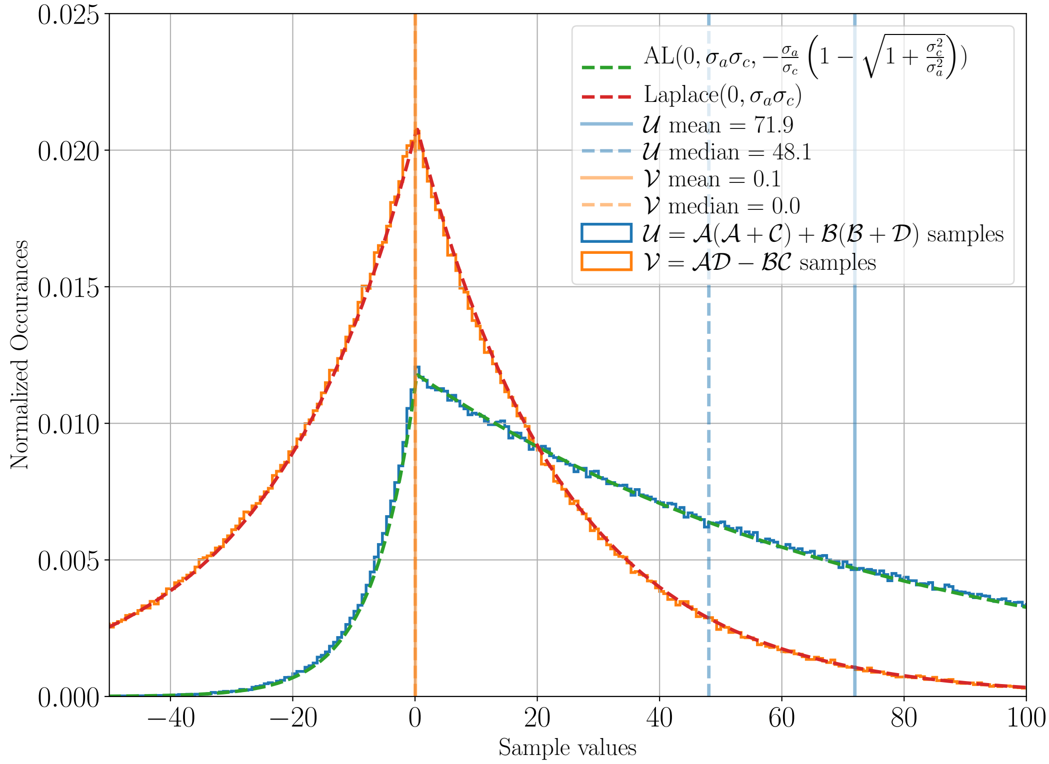

Figure 6.7: Histograms of random variables $\mathcal{U}$ and $\mathcal{V}$ which follow asymmetric Laplace and Laplace distributions, respectively. The random variables $\mathcal{U} = \mathcal{A}(\mathcal{A}+\mathcal{C}) + \mathcal{B}(\mathcal{B}+\mathcal{D})$ and $\mathcal{V} = \mathcal{A}\mathcal{D} - \mathcal{B}\mathcal{C}$, where that $\mathcal{A}, \mathcal{B} \sim \mathcal{N}(0, \sigma_a)$ and $\mathcal{C}, \mathcal{D} \sim \mathcal{N}(0, \sigma_c)$. $\mathcal{U}$ and $\mathcal{V}$ describe the real and imaginary parts of a CSD $\langle x, z\rangle = \langle x, x\rangle + \langle x, y\rangle$ where x and y are uncorrelated Gaussian noise. Equation 6.76 is plotted as the green dashed curve, while equation 6.73 is plotted as the red dashed curve. In this example, $\sigma_a = 6$ and $\sigma_c = 4$, so the Laplace scale factors $\lambda = \gamma = 24$, the asymmetry parameter $\kappa = 0.30$, and the median/mean bias $b = \rho/\mu = 0.67$.

6.11.1.1 Minor axis PDF of the cross spectral density

First, we show that the minor (transverse) axis of Eq. 6.71 with random variable $\mathcal{V} = \mathcal{A}\mathcal{D} - \mathcal{B}\mathcal{C}$ follows a Laplace distribution. In the case from Figure 6.1, the imaginary axis and minor axis are the same.

First, recall from Eq. 6.52 that $\mathcal{A}\mathcal{D}$ and $\mathcal{B}\mathcal{C}$ both independently follow the same modified Bessel distribution. Note that the distribution is symmetric about zero, so $-\mathcal{BC} \sim \mathcal{BC}$. Using Eqs. 6.26, 6.27, 6.53, we can write

$$\begin{aligned}
f_{\mathcal{V}}(v) &= f_{\mathcal{AD}-\mathcal{BC}}(v) \\
&= f_{\mathcal{AD}} * f_{-\mathcal{BC}} \\
\varphi_{\mathcal{V}}(t) &= \mathcal{F}[f_{\mathcal{AD}} * f_{-\mathcal{BC}}] \\
&= \mathcal{F}[f_{\mathcal{AD}}]\mathcal{F}[f_{-\mathcal{BC}}] \\
&= \varphi_{\mathcal{AD}}(t)\varphi_{-\mathcal{BC}}(t) \\
&= \frac{1}{\sqrt{\sigma_a^2\sigma_c^2t^2+1}}\frac{1}{\sqrt{\sigma_a^2\sigma_c^2t^2+1}} \\
\varphi_{\mathcal{V}}(t) &= \frac{1}{\sigma_a^2\sigma_c^2t^2+1} \qquad (6.72) \\
f_{\mathcal{V}}(v) &= \frac{1}{2\sigma_a\sigma_c}\exp\left(-\frac{|v|}{\sigma_a\sigma_c}\right). \qquad (6.73)
\end{aligned}$$

In the last step we observe from Eq. 6.69 the characteristic function of a Laplace distribution with $\mu = 0$ and $\gamma = \sigma_a\sigma_c$. Therefore, the random variable $\mathcal{V}$ that characterizes the imaginary part of the CSD $\langle x, z\rangle$ follows a Laplace distribution [207].

It can similarly be shown that for a completely uncorrelated CSD $\langle x, y\rangle$, the real part $\mathcal{T} = \mathcal{AC} + \mathcal{BD}$ follows the same Laplace distribution as the imaginary part.

If we let $\sigma_a \to \sigma_x\sqrt{N/2}$ and $\sigma_c \to \sigma_y\sqrt{N/2}$, then $\gamma \to N\sigma_x\sigma_y/2$. If we then scale the distribution by $2/(Nf_s)$ from the definition of the CSD Eq. 6.23, then

$$f_{\Im(\langle x,z\rangle)} = \text{Laplace}\left(0, \frac{\sigma_x\sigma_y}{f_s}\right) \qquad (6.74)$$

This gives the result from Table 6.1. The distribution in Eq. 6.77 are plotted on the imaginary axis projection of Figure 6.8.

6.11.1.2 Major axis PDF of the cross spectral density

For cross spectral densities with correlated noise $\langle x, z\rangle$, the major (radial) axis has a random variable $\mathcal{U} = \mathcal{A}^2 + \mathcal{B}^2 + \mathcal{AC} + \mathcal{BD} = \mathcal{A}(\mathcal{A} + \mathcal{C}) + \mathcal{B}(\mathcal{B} + \mathcal{D})$. In the case from Figure 6.1, $z[n] = x[n] + [y]$ has no relative delay between the x and z signals, so the real axis and major axis of $\langle x, z\rangle$ are the same.

By a similar argument as used for Eq. 6.73, just replacing Eq. 6.53 with Eq. 6.62 and using Eq. 6.70 we can write

$$\begin{aligned} f_{\mathcal{U}}(u) &= f_{\mathcal{A}(\mathcal{A}+\mathcal{C})} * f_{\mathcal{B}(\mathcal{B}+\mathcal{D})} \\ \varphi_{\mathcal{U}}(t) &= \varphi_{\mathcal{A}(\mathcal{A}+\mathcal{C})}(t)\varphi_{\mathcal{B}(\mathcal{B}+\mathcal{D})}(t) \\ \varphi_{\mathcal{U}}(t) &= \frac{1}{1 - i2\sigma_a^2 t + \sigma_a^2\sigma_c^2 t^2} \qquad (6.75) \\ f_{\mathcal{U}}(u) &= \frac{1}{2\sigma_a\sqrt{\sigma_a^2+\sigma_c^2}} \begin{cases} \exp\left(\dfrac{u}{\sigma_a\sqrt{\sigma_a^2+\sigma_c^2}-\sigma_a^2}\right) & u \le 0 \\ \exp\left(-\dfrac{u}{\sigma_a\sqrt{\sigma_a^2+\sigma_c^2}+\sigma_a^2}\right) & u > 0 \end{cases} \qquad (6.76) \end{aligned}$$

where, in the last step we recover an asymmetric Laplace distribution with $m = 0$, $\lambda = \sigma_a\sigma_c$, and $\kappa = -\frac{\sigma_a}{\sigma_c}\left(1 - \sqrt{1 + \frac{\sigma_c^2}{\sigma_a^2}}\right)$.

If we let $\sigma_a \to \sigma_x\sqrt{N/2}$ and $\sigma_c \to \sigma_y\sqrt{N/2}$, then $\lambda \to N\sigma_x\sigma_y/2$ and $\kappa = -(\sigma_x/\sigma_y)\left(1 - \sqrt{1+\sigma_y^2/\sigma_x^2}\right)$. If we then scale the distribution by $2/(Nf_s)$ then

$$f_{\Re(\langle x,z\rangle)} = \mathrm{AL}\left(0, \frac{\sigma_x\sigma_y}{f_s}, -\frac{\sigma_x}{\sigma_y}\left(1 - \sqrt{1 + \frac{\sigma_y^2}{\sigma_x^2}}\right)\right) \qquad (6.77)$$

This gives the result from Table 6.1. The distribution in Eq. 6.77 are plotted on the real axis projection of Figure 6.8.

6.11.1.3 Joint probability distribution of the cross spectral density

In general, the CSD is a complex quantity. Here we'll derive the two-dimensional probability distribution function of the CSD, starting from our simple case as defined in Figure 6.1, then generalizing our final result.

We begin with a joint distribution of our four independent Gaussian random variables $\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{D}$:

$$f_{\mathcal{A},\mathcal{B},\mathcal{C},\mathcal{D}}(a,b,c,d) = \frac{1}{4\pi^2\sigma_a^2\sigma_c^2} e^{-\frac{a^2}{2\sigma_a^2}-\frac{b^2}{2\sigma_a^2}-\frac{c^2}{2\sigma_c^2}-\frac{d^2}{2\sigma_c^2}} \qquad (6.78)$$

Here we derive the joint characteristic function of the CSD. We take the Fourier transform of $f_{\mathcal{A},\mathcal{B},\mathcal{C},\mathcal{D}}(a,b,c,d)$ with respect to the random variables $\mathcal{U}$ and $\mathcal{V}$ as defined in Eq. 6.71:

$$\varphi_{\mathcal{U},\mathcal{V}}(s,t) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} da\, db\, dc\, dd\, e^{isu+itv} f_{\mathcal{A},\mathcal{B},\mathcal{C},\mathcal{D}}(a,b,c,d)$$
$$\varphi_{\mathcal{U},\mathcal{V}}(s,t) = \frac{1}{4\pi^2\sigma_a^2\sigma_c^2}\int\int\int\int_{-\infty}^{\infty} da\, db\, dc\, dd\, e^{is(a^2+ac+b^2+bd)+it(ad-bc)} e^{-\frac{a^2}{2\sigma_a^2}-\frac{b^2}{2\sigma_a^2}-\frac{c^2}{2\sigma_c^2}-\frac{d^2}{2\sigma_c^2}}$$
$$\varphi_{\mathcal{U},\mathcal{V}}(s,t) = \frac{1}{1-2i\sigma_a^2 s+\sigma_a^2\sigma_c^2(s^2+t^2)} \tag{6.79}$$

where s and t are the Fourier variables.

Next we Fourier transform back to the joint distribution of the CSD:

$$f_{\mathcal{U},\mathcal{V}}(u,v) = \frac{1}{2\pi}\int\int_{-\infty}^{\infty} ds\, dt\, e^{isu+itv}\varphi_{\mathcal{U},\mathcal{V}}(s,t) \tag{6.80}$$

Let $q = s - i/\sigma_c^2$. Then we can recover circular symmetry to take the double integral easily:

$$f_{\mathcal{U},\mathcal{V}}(u,v) = \frac{1}{2\pi\sigma_a^2\sigma_c^2}\int\int_{-\infty}^{\infty} dq\, dt\, \frac{e^{-iqu-itv+u/\sigma_c^2}}{\frac{1}{\sigma_a^2\sigma_c^2}\left(1+\frac{\sigma_a^2}{\sigma_c^2}\right)+q^2+v^2}$$
$$f_{\mathcal{U},\mathcal{V}}(u,v) = \frac{e^{u/\sigma_c^2}}{2\pi\sigma_a^2\sigma_c^2} K_0\left(\frac{1}{\sigma_c^2}\sqrt{(u^2+v^2)\left(1+\frac{\sigma_c^2}{\sigma_a^2}\right)}\right) \tag{6.81}$$

Finally, we scale the distribution Eq. 6.81 to match the final result of CSDs in practice. First, we generalize the angle of the CSD ϕ, by allowing the variable in the exponent $u \to u\cos(\phi) + v\sin(\phi)$. Second, from Eq. 6.71, let $x = 2u/(Nf_s)$ and $y = 2v/(Nf_s)$ for changing variables, and make the substitutions for $\sigma_a = \sigma_x\sqrt{N/2}$ and $\sigma_c = \sigma_y\sqrt{N/2}$:

$$f_{\langle x,z\rangle}(x,y) = \frac{e^{\frac{x\cos(\phi)+y\sin(\phi)}{\sigma_y^2/f_s}}}{2\pi\sigma_x^2\sigma_y^2/f_s^2} K_0\left(\frac{f_s}{\sigma_y^2}\sqrt{(x^2+y^2)\left(1+\frac{\sigma_y^2}{\sigma_x^2}\right)}\right) \tag{6.82}$$

This gives the result from Table 6.1. The contours of Eq. 6.82 are plotted in Figure 6.8.

6.11.1.4 Mean-to-median bias in the cross spectral density

Here we calculate the bias that results from using median-averaging to estimate a cross spectral density. We focus on the probability distribution describing the major axis of the CSD, Eq. 6.76, since the minor axis of the CSD must always

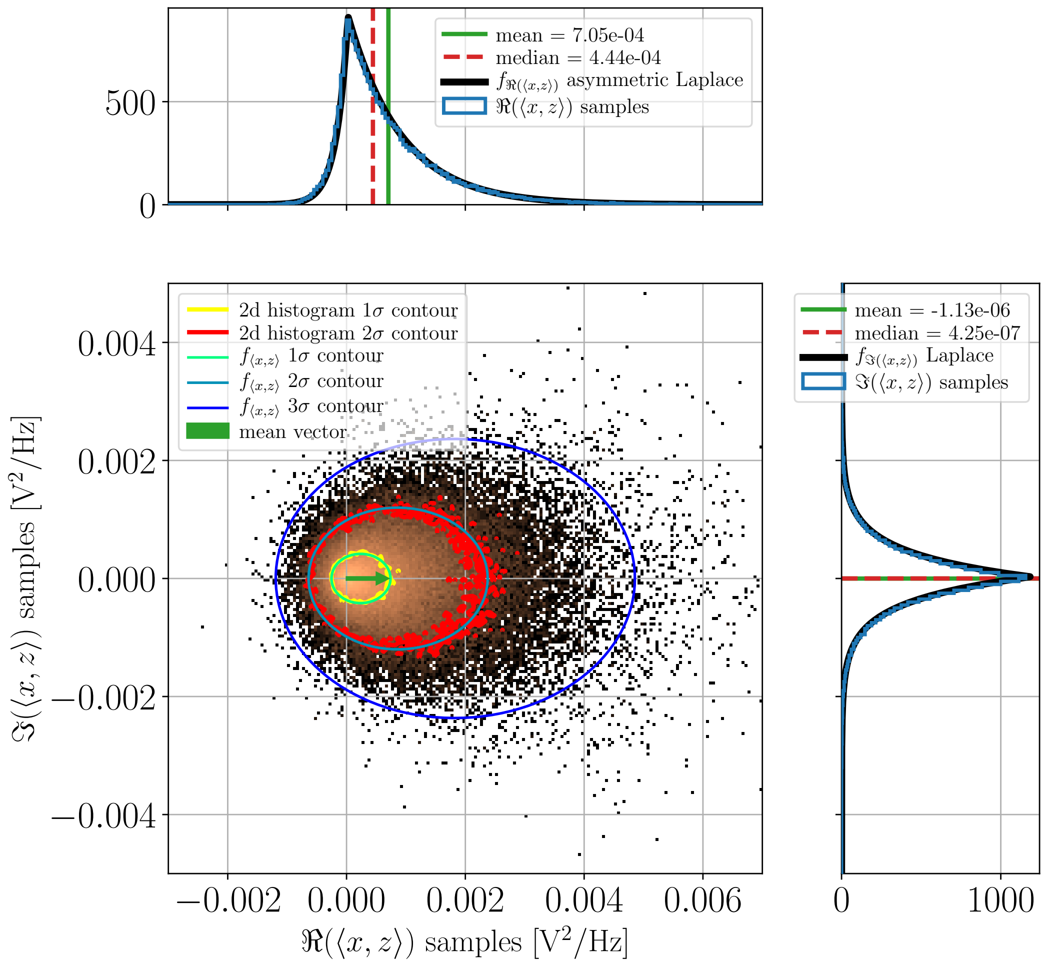

Figure 6.8: Two-dimensional histogram of 100000 samples from a CSD $\langle x, z\rangle$ at a single frequency bin, with contours from the joint distribution $f_{\langle x,z\rangle}(x, y)$ in Eq. 6.82. Above and to the right are plotted one-dimensional histograms of the real and imaginary axes, with the marginal distributions $f_{\Re(\langle x,z\rangle)}(x)$ and $f_{\Im(\langle x,z\rangle)}(y)$ from Eqs. 6.77 and 6.74. In this example, the correlated power $\sigma_x^2/f_s = 3.5 \times 10^{-4}\ \mathrm{V}^2/\mathrm{Hz}$, the uncorrelated power $\sigma_y^2/f_s = 5 \times 10^{-4}\ \mathrm{V}^2/\mathrm{Hz}$, the true coherence $\gamma^2 = 0.412$, the asymmetric Laplace skew parameter for the real axis $\kappa = 0.467$, and the phase $\phi = 0$.

follow a Laplace distribution, which has both mean and median zero and should not contribute to the final CSD significantly. The validity of this assumption is explored in Appendix D.

If we calculate the mean of the asymmetric Laplace from Eq. 6.76 using Eq. 6.65, we recover

$$\mu = 2\sigma_a^2 \tag{6.83}$$

The mean of the CSD depends only on the correlated noise. This is the same as the mean recovered for the exponential distribution in Eqs. 6.40. The invariance of the mean is demonstrated in Figure 6.9.

The median from Eq. 6.76 is more complex, and depends on both the uncorrelated and correlated noise. Using Eq. 6.66:

$$\rho = \begin{cases} -\sigma_a\left(\sqrt{\sigma_a^2+\sigma_c^2}+\sigma_a\right)\log\left(\dfrac{-\sigma_a\sqrt{\sigma_a^2+\sigma_c^2}+\sigma_a^2+\sigma_c^2}{\sigma_c^2}\right), & \kappa < 1 \\ \sigma_a\left(\sqrt{\sigma_a^2+\sigma_c^2}-\sigma_a\right)\log\left(-\dfrac{\sqrt{\sigma_a^2+\sigma_c^2}}{\sigma_a-\sqrt{\sigma_a^2+\sigma_c^2}}\right), & \kappa > 1 \end{cases} \tag{6.84}$$

If we let the correlated noise dominate the uncorrelated noise, $\sigma_a^2 \gg \sigma_c^2$, then the median

$$\rho \to 2\sigma_a^2 \log(2), \qquad \kappa < 1. \tag{6.85}$$

This is the same as the median recovered for the exponential distribution in Eq. 6.36. The median changing with different levels of uncorrelated noise is demonstrated in Figure 6.9.

The bias from Eq. 6.76, using Eq. 6.68, is

$$b = \frac{\rho}{\mu} = \frac{\epsilon}{2\left(1-\sqrt{\epsilon+1}\right)}\log\left(\frac{\epsilon-\sqrt{\epsilon+1}+1}{\epsilon}\right), \quad \kappa < 1 \tag{6.86}$$

where we have defined the uncorrelated power ratio $\epsilon = \sigma_c^2/\sigma_a^2$. Figure 6.10 shows the limits of CSD mean-to-median bias varies between $\log(2)$ for $\epsilon \ll 1$ to $1/2$ for $\epsilon \gg 1$.

In general, the distribution of the CSD $\langle x, z\rangle$ is a two-dimensional asymmetric Laplace describing the real and imaginary components of the CSD. A phase delay in the correlated signal $z[n] = y[n] + x[n-m]$ where m is a time delay, yields a rotation in the 2D asymmetric Laplace. Section 6.14 explores an example with a phase delay.

For the above derivations, the phase is zero. The derivation is true without loss of generality as long as an appropriate phase rotation is applied to the general 2D asymmetric Laplace. The bias will always depend on the power ratio ϵ, but could also depend on the phase of the CSD ϕ, explored in Section 6.13.

In the limit $\kappa \to 0$, the asymmetric Laplace becomes the exponential distribution. For the CSD, this is equivalent to having the correlated noise much greater than the

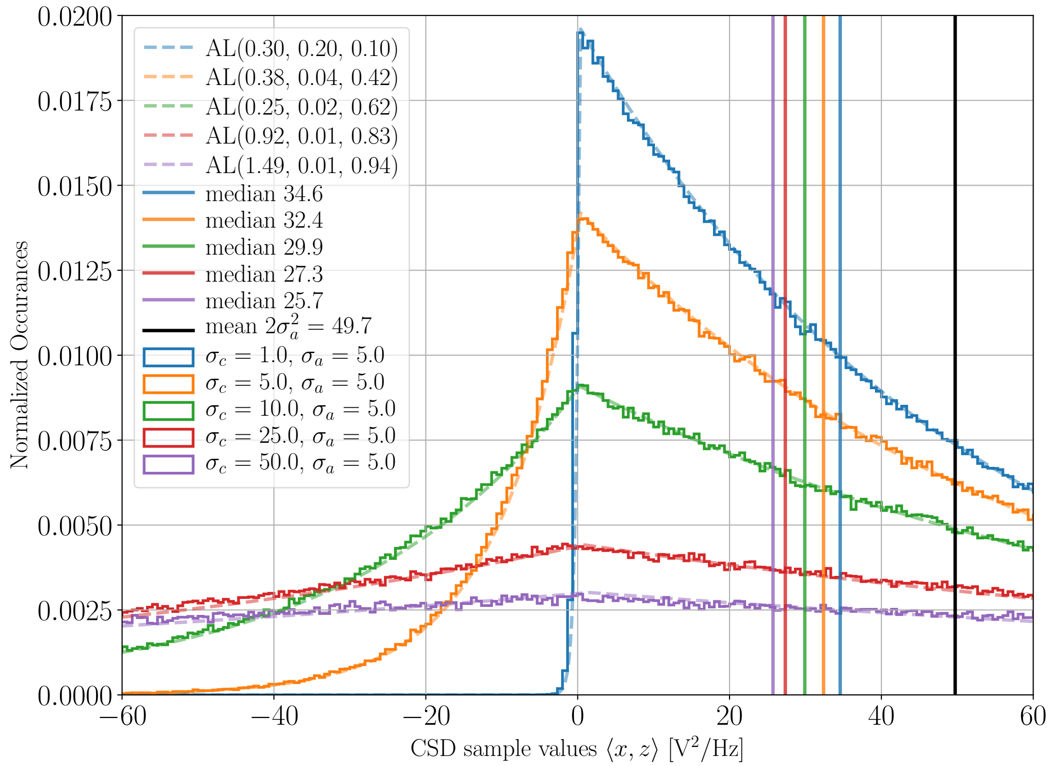

Figure 6.9: Histograms demonstrating the invariance of the mean of an asymmetric Laplace distribution with changing uncorrelated noise σ_c. Equation 6.83 states that the mean of the asymmetric Laplace depends only on the correlated noise σ_a. In this example, correlated noise $\sigma_a = 5$ for all curves. The sample mean of all five curves is shown as the black line, $\hat{\mu} \approx 50$. Equation 6.64 fits are plotted as the dashed curves, while the sample medians are plotted as the solid vertical lines. While the mean is invariant, the medians change based on the level of uncorrelated power in the signal. This means that a CSD $\langle x, z\rangle$ estimated via mean-averaging would yield the same result for all of these curves, but different results if estimated via median-averaging.

uncorrelated noise: $\sigma_a \gg \sigma_c$, so the CSD is approaching the PSD: $\langle x, z\rangle \to \langle x, x\rangle$. The asymmetric Laplace mean to median bias from Eq. 6.68 $\rho/\mu \to \log(2)$, the same bias as the exponential distribution.

We have shown that CSD estimate $\langle x, z\rangle$ follows 2D asymmetric Laplace distributions in general. From Eqs. 6.68 and 6.76, *the bias between mean- and median-averaged CSD estimates depends on the relative power of correlated and uncorrelated noise.*

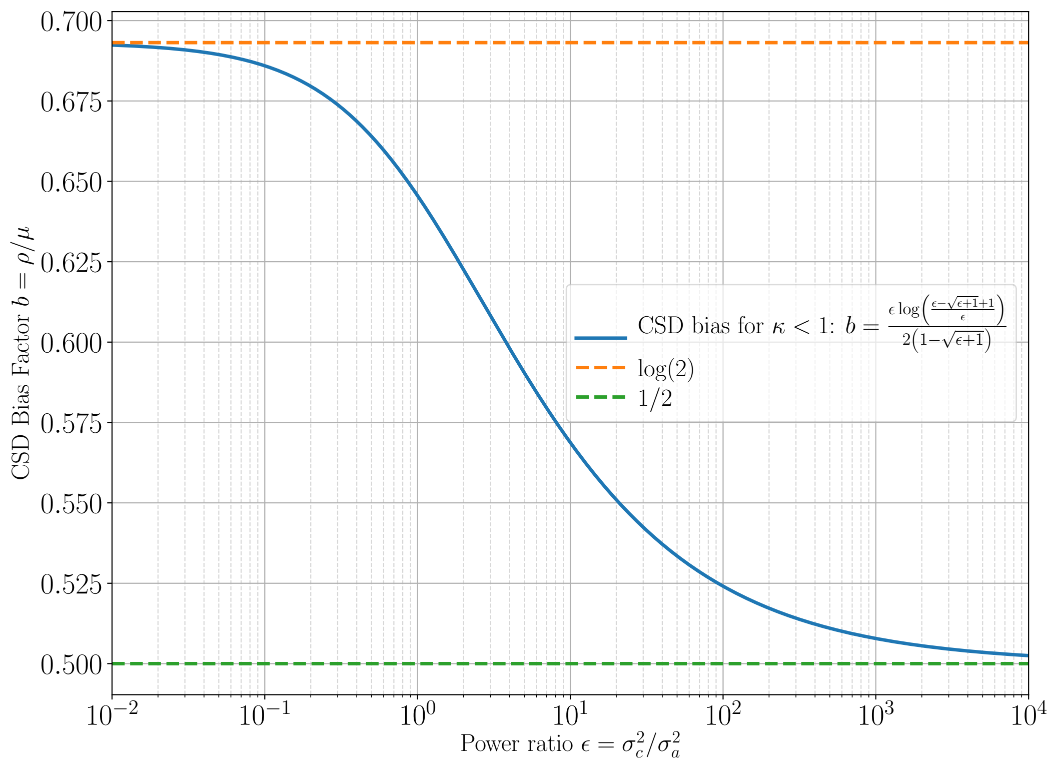

Figure 6.10: Mean-to-median bias factor $b = \rho/\mu$ associated with CSDs as a function of relative power $\epsilon = \sigma_c^2/\sigma_a^2$, where σ_c^2 is the uncorrelated power and σ_a^2 is the correlated power. The bias factor plotted is from Eq. 6.86.

6.12 Coherence

Coherence is a method of computing the correlation of two signals. The definition from Eq. 6.16 is good for infinite signals, but for realistic signals x and $z = x + y$ the mean-averaged PSD and CSD estimates are usually used:

$$\gamma^2 = \frac{|\langle x, z\rangle|^2}{\langle x, x\rangle\langle z, z\rangle} \tag{6.87}$$

Our mean-averaged estimates for the power and cross spectral densities come from Eqs. 6.40 and 6.83, yielding the mean-averaged coherence $\overline{\gamma^2}$:

$$\overline{\langle x, x\rangle} = 2\sigma_a^2 \qquad \overline{\langle z, z\rangle} = 2(\sigma_a^2 + \sigma_c^2) \qquad \overline{|\langle x, z\rangle|} = 2\sigma_a^2 \tag{6.88}$$

$$\overline{\gamma^2} = \frac{\sigma_a^2}{\sigma_a^2 + \sigma_c^2} = \frac{1}{1+\epsilon} \tag{6.89}$$

where $\epsilon = \sigma_c^2/\sigma_a^2$ is the power ratio of uncorrelated over correlated noise. As the signals x and z becomes completely correlated, $\sigma_a^2 \gg \sigma_c^2$ and $\epsilon \to 0$, so the coherence $\overline{\gamma^2} \to 1$. As the signals become uncorrelated, $\sigma_a^2 \ll \sigma_c^2$ and $\epsilon \gg 1$, so the coherence $\overline{\gamma^2} \to 1/\epsilon$.

Using the expression for κ below Eq. 6.76, we express the power ratio $\epsilon = \sigma_c^2/\sigma_a^2$ in terms of κ:

$$\epsilon = \frac{4\kappa^2}{(1-\kappa^2)^2} \tag{6.90}$$

Then the mean-averaged coherence is related to the asymmetric Laplace parameter κ by

$$\overline{\gamma^2} = \left(\frac{1-\kappa^2}{1+\kappa^2}\right)^2 \tag{6.91}$$

This relates the underlying distribution to coherence, which is often used to quantify the quality of the CSD estimate.

The median-averaged PSD and CSD estimates come from Eqs. 6.36 and 6.84 and give the median-averaged coherence $\widetilde{\gamma}^2$:

$$\widetilde{\langle x, x\rangle} = 2\sigma_a^2 \log(2) \qquad \widetilde{\langle z, z\rangle} = 2(\sigma_a^2 + \sigma_c^2)\log(2) \tag{6.92}$$

$$|\widetilde{\langle x, z\rangle}| = -\sigma_a\left(\sqrt{\sigma_a^2+\sigma_c^2} + \sigma_a\right)\log\left(\frac{-\sigma_a\sqrt{\sigma_a^2+\sigma_c^2} + \sigma_a^2 + \sigma_c^2}{\sigma_c^2}\right) \tag{6.93}$$

$$\begin{aligned}
\widetilde{\gamma}^2 &= \frac{\left[-\sigma_a\left(\sqrt{\sigma_a^2+\sigma_c^2} + \sigma_a\right)\log\left(\frac{-\sigma_a\sqrt{\sigma_a^2+\sigma_c^2} + \sigma_a^2 + \sigma_c^2}{\sigma_c^2}\right)\right]^2}{[2\sigma_a^2\log(2)][2(\sigma_a^2+\sigma_c^2)\log(2)]} \\
&= \frac{\sigma_c^4 \log^2\left(\frac{-\sigma_a\sqrt{\sigma_a^2+\sigma_c^2} + \sigma_a^2 + \sigma_c^2}{\sigma_c^2}\right)}{4\log^2(2)(\sigma_a^2+\sigma_c^2)\left(-2\sigma_a\sqrt{\sigma_a^2+\sigma_c^2} + 2\sigma_a^2 + \sigma_c^2\right)} \\
\widetilde{\gamma}^2 &= \frac{\epsilon^2\log^2\left(\frac{1}{\sqrt{\epsilon+1}} + 1\right)}{4(\epsilon+1)\left(\epsilon - 2\sqrt{\epsilon+1} + 2\right)\log^2(2)}
\end{aligned} \tag{6.94}$$

Again, as the signals x and z become completely correlated, $\sigma_a^2 \gg \sigma_c^2$, $\epsilon \to 0$, and $\widetilde{\gamma}^2 \to 1$. As the signals become uncorrelated, $\sigma_a^2 \ll \sigma_c^2$, $\widetilde{\gamma}^2 \to 1/\big(4\epsilon\log^2(2)\big)$. Both the mean- and median-averaged coherences are plotted in Figure 6.11.

Expressing the median-averaged coherence in terms of the mean-averaged coherence yields

$$\widetilde{\gamma}^2 = \frac{(1+\overline{\gamma})^2\log^2(1+\overline{\gamma})}{4\log^2(2)} \tag{6.95}$$

Coherence provides a one-to-one mapping to the uncorrelated over correlated power ratio ϵ. This is key because *we can use coherence to find the mean-to-median*

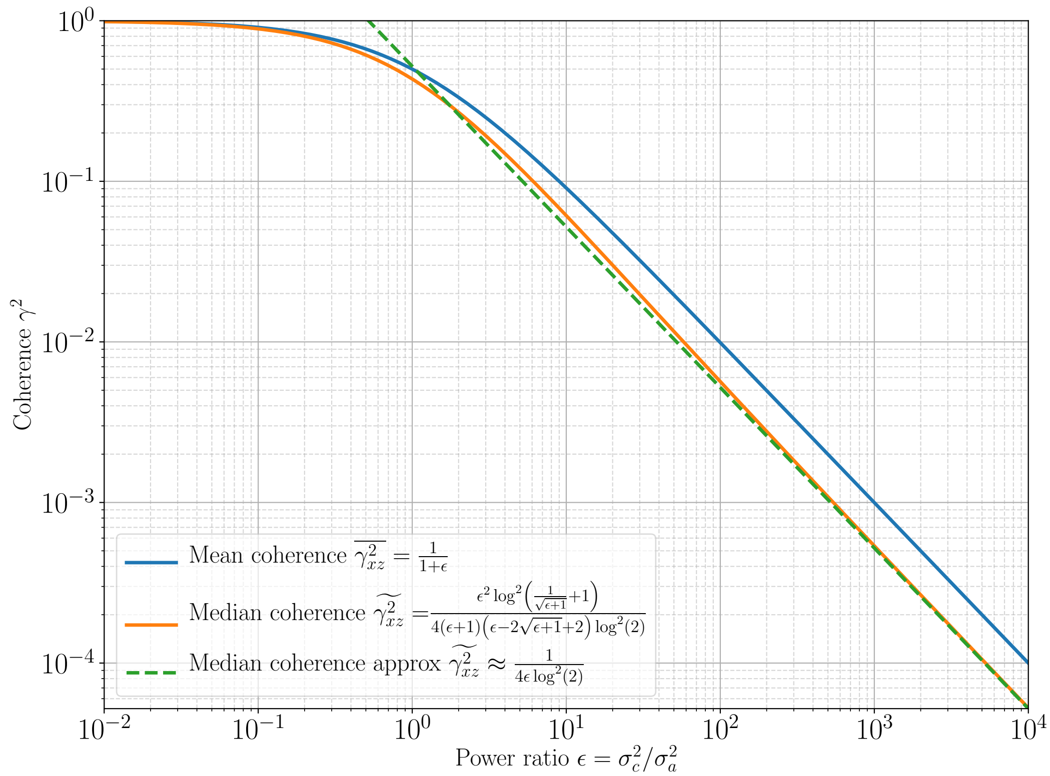

Figure 6.11: Comparison of the coherence estimated with mean-averaging $\overline{\gamma^2}$ and median-averaging $\widetilde{\gamma^2}$.

bias for the CSD. From Eqs. 6.86 and 6.94, we can write the mean-to-median bias in terms of the median-averaged coherence:

$$b = \log(2)\sqrt{(1+\epsilon)\,\widetilde{\gamma^2}} \tag{6.96}$$

The coherence $\widetilde{\gamma^2}$ can be used to solve for ϵ numerically using Eq. 6.94. If $\widetilde{\gamma^2} \approx 1$, then $\epsilon \ll 1$ and the bias $b \approx \log(2)$.

Expressing the CSD mean-to-median bias b in terms of mean-averaged coherence yields

$$b = \frac{(1+\overline{\gamma})\log(1+\overline{\gamma})}{2\overline{\gamma}} \tag{6.97}$$

The CSD mean μ and CSD median ρ in terms of mean-averaged coherence are

$$\mu = \frac{2\lambda\overline{\gamma}}{\sqrt{1-\overline{\gamma^2}}} \tag{6.98}$$

$$\rho = \frac{\lambda(1+\overline{\gamma})\log(1+\overline{\gamma})}{\sqrt{1-\overline{\gamma^2}}} \tag{6.99}$$

The correlated-to-uncorrelated power ratio $\epsilon = \sigma_c^2/\sigma_a^2$ in this derivation refers to the specific case of estimating the CSD $\langle x, z\rangle$ where $z = x + y$. More generally, ϵ can refer to a uncorrelated/correlated power ratio of many Gaussian noises together. In this general case, Eqs. 6.94 and 6.96 are still valid. Section 6.14 explores the case of three Gaussian noises.

6.13 Phase

In the sections above, we have assumed that the phase of the CSD $\phi = 0$. This was because in the CSD $\langle x, z\rangle$ where $z[n] = x[n] + y[n]$, there is no phase difference because x appears in z identically.

In general, there may be a phase difference in our signals. This will rotate our CSD major and minor axes away from the real and imaginary axes.

A phase rotation should not affect the final magnitude of the CSD. Using mean-averaging, this is true. Using median-averaging, because of the logarithm appearing in the median expression Eq. 6.84, the final magnitude can change anywhere from 0% to 4%, depending on the coherence of the signals.

In this section, the relationship of mean and median magnitude and phase is derived, assuming the user is naively taking these statistics along the real and imaginary axes.

6.13.1 General real and imaginary axis distributions

First, we generalize Eq. 6.82 by

$$f_{\langle x,y\rangle}(u,v) = \frac{e^{\frac{\gamma}{\sqrt{1-\gamma^2}\lambda}(u\cos(\phi)+v\sin\phi)}}{2\pi\lambda^2} K_0\left(\frac{1}{\lambda}\sqrt{\frac{u^2+v^2}{1-\gamma^2}}\right) \tag{6.100}$$

where u is the real axis, v is the imaginary axis, $\gamma^2 = 1/(1+\sigma_y^2/\sigma_x^2)$ is the mean-averaged coherence, and $\lambda = \sigma_x\sigma_y/f_s$ is a cross-power scaler. The expressions used here for γ^2 and λ are illustrative for converting from Eq. 6.82, but Eq. 6.100 is good for any γ^2 and λ.

Second, we specify the probability density functions along the real and imaginary axes u and v for any phase ϕ, by marginalizing along each axis:

$$f_{\Re(\langle x,y\rangle)}(u) = \frac{1}{2\lambda\sqrt{\frac{1-\gamma^2\sin^2(\phi)}{1-\gamma^2}}}\exp\left(\frac{u\gamma\cos(\phi)-|u|\sqrt{1-\gamma^2\sin^2(\phi)}}{\lambda\sqrt{1-\gamma^2}}\right) \tag{6.101}$$

$$f_{\Im(\langle x,y\rangle)}(v) = \frac{1}{2\lambda\sqrt{\frac{1-\gamma^2\cos^2(\phi)}{1-\gamma^2}}}\exp\left(\frac{v\gamma\sin(\phi)-|v|\sqrt{1-\gamma^2\cos^2(\phi)}}{\lambda\sqrt{1-\gamma^2}}\right). \tag{6.102}$$

These distributions are both asymmetric Laplace distributions, related by the common parameters γ^2, λ, ϕ.

6.13.2 Mean of the general CSD distributions

If we calculate the mean of Eqs. 6.101 and 6.102, we get

$$\mu_{\Re(\langle x,y\rangle)} = \frac{2\gamma\lambda\cos(\phi)}{\sqrt{1-\gamma^2}} \tag{6.103}$$

$$\mu_{\Im(\langle x,y\rangle)} = \frac{2\gamma\lambda\sin(\phi)}{\sqrt{1-\gamma^2}}. \tag{6.104}$$

Finding the overall mean magnitude $\mu_{\langle x,y\rangle}$,

$$\mu_{\langle x,y\rangle} = \sqrt{\mu^2_{\Re(\langle x,y\rangle)} + \mu^2_{\Im(\langle x,y\rangle)}}$$

$$\mu_{\langle x,y\rangle} = \frac{2\gamma\lambda}{\sqrt{1-\gamma^2}} \tag{6.105}$$

Therefore the magnitude of the CSD mean has circular symmetry, and is not affected in general by the phase ϕ.

6.13.3 Median of the general CSD distributions

Calculating the median of Eqs. 6.101 and 6.102 yields

$$\rho_{\Re(\langle x,y\rangle)} = \frac{\lambda\sqrt{1-\gamma^2}\log\left(\frac{1-\gamma^2+\gamma^2\cos^2(\phi)-\gamma|\cos(\phi)|\sqrt{1-\gamma^2+\gamma^2\cos^2(\phi)}}{1-\gamma^2}\right)}{\gamma\cos(\phi)-\mathrm{sign}(\cos(\phi))\sqrt{1-\gamma^2+\gamma^2\cos^2(\phi)}} \tag{6.106}$$

$$\rho_{\Im(\langle x,y\rangle)} = \frac{\lambda\sqrt{1-\gamma^2}\log\left(\frac{1-\gamma^2+\gamma^2\sin^2(\phi)-\gamma|\sin(\phi)|\sqrt{1-\gamma^2+\gamma^2\sin^2(\phi)}}{1-\gamma^2}\right)}{\gamma\sin(\phi)-\mathrm{sign}(\sin(\phi))\sqrt{1-\gamma^2+\gamma^2\sin^2(\phi)}} \tag{6.107}$$

From Eqs. 6.106 and 6.107, it is clear that the magnitude of the median $\rho_{\langle x,y\rangle} = \sqrt{\rho^2_{\Re(\langle x,y\rangle)} + \rho^2_{\Im(\langle x,y\rangle)}}$ must be affected by the phase.

6.13.4 Phase overview

We've calculated that while phase has no effect on mean-averaged CSD calculations, it does affect the results of median-averaged CSDs, assuming those medians are taken individually along the real and imaginary axes.

Phase also must affect the median-averaged coherence, rendering Eq. 6.94 incorrect in general. The simplest solution is to apply a phase-rotation to every CSDs to set all phases $\phi = 0$ while performing Welch's median-averaging method.

Figure 6.12 illustrates how the CSD median is related to the phase at different levels of coherence. Figure 6.13 shows how the coherence affects the final CSD median result for the phase $\phi = \pi/4$, representing one of the largest possible biases in the median due to phase.

The white noise example in Section 6.14 demonstrates the effect phase can have on an uncorrected median-averaged CSD. Figure 6.15 shows both the uncorrected and "phase-corrected" median-averaged CSDs together.

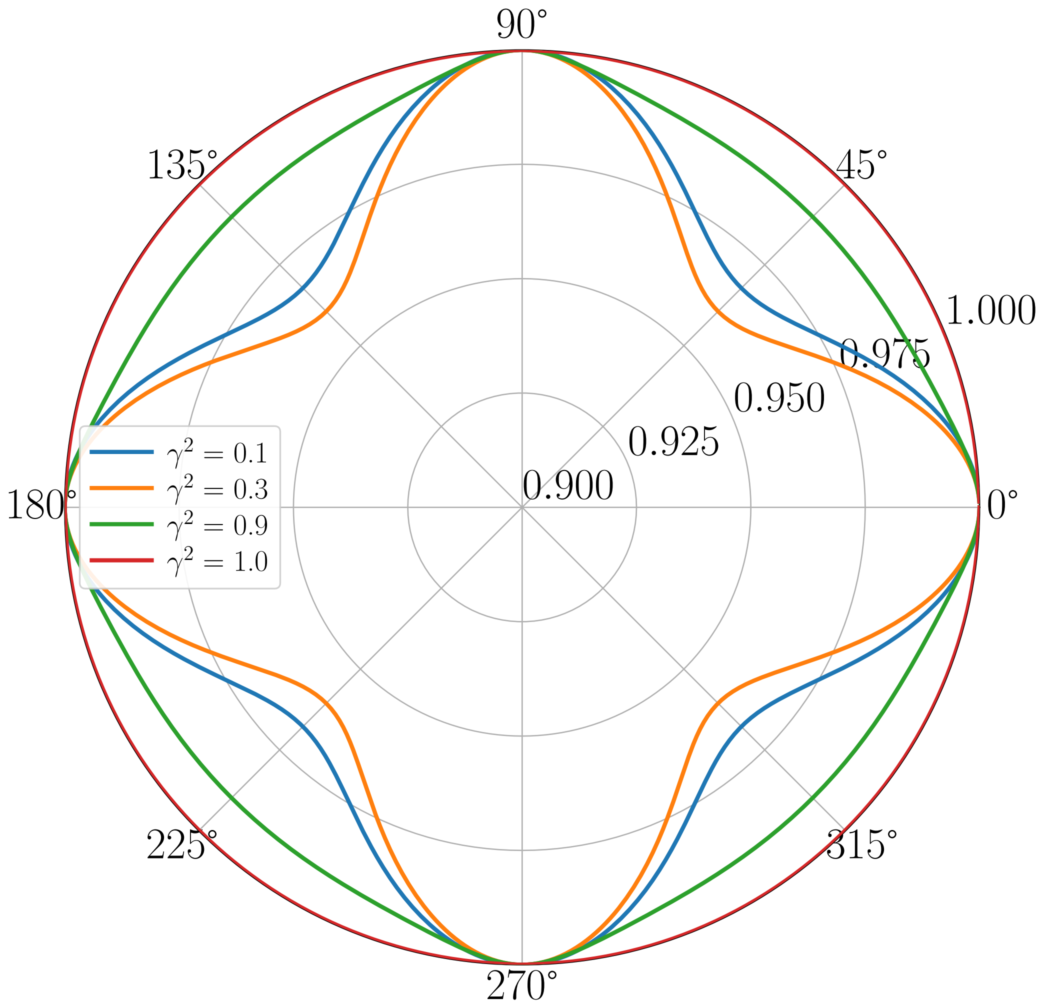

Figure 6.12: Plot of how the phase of a CSD ϕ affects the median-averaged CSD magnitude ρ at different levels of coherence γ^2. The median ρ is normalized against the maximum possible ρ at $\phi = 0$, removing dependance on λ. Ideally, the median would not be affected by the phase ϕ, but this is not true due to the complex expressions for the median (Eqs. 6.106 and 6.107). The median is at a minimum when the phase is $\pi/4$ from the real and imaginary axes. The radial axis of this plot goes from 0.9 to 1.0 to emphasize the small effect of the phase bias (max 4%). Figure 6.13 shows the worst case bias in the median-averaging as a function of coherence.

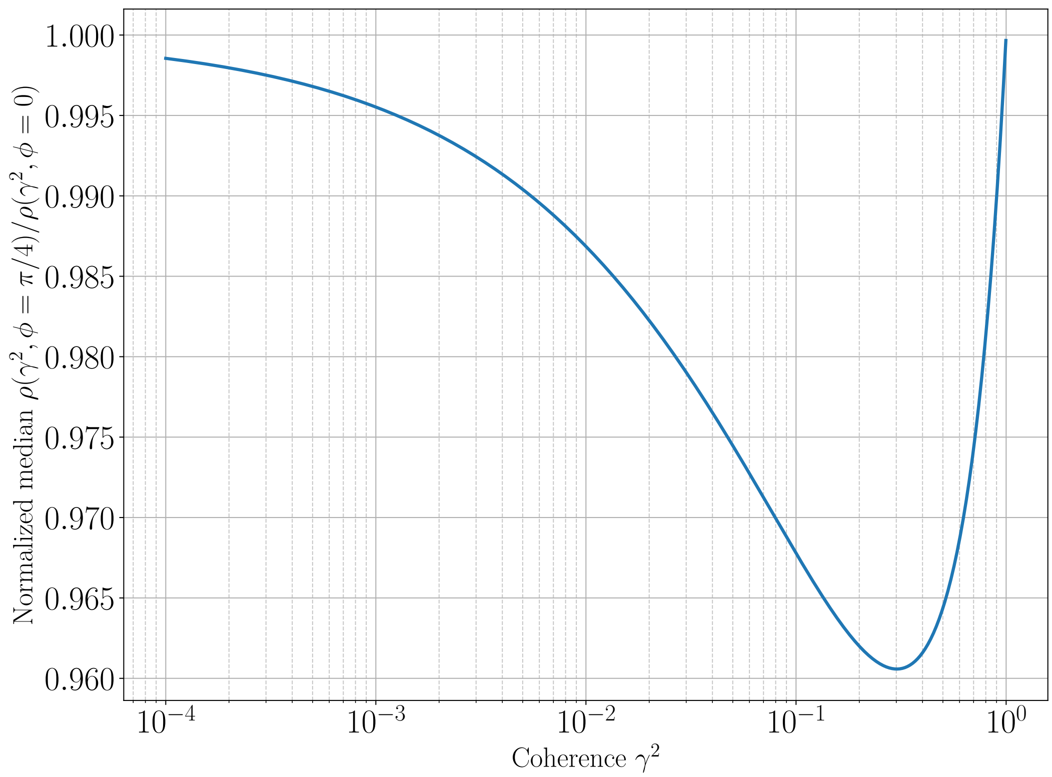

Figure 6.13: Plot of coherence γ^2 vs the median-averaged CSD magnitude ρ, at the minimum possible median at phase $\phi = \pi/4$. Note that the median ρ is normalized against the maximum possible ρ at $\phi = 0$, removing dependance on λ. This plot shows the largest possible phase-based bias we can have from using median-averaging for the CSD, as a function of mean-averaged coherence. The minimum occurs at around $\gamma^2 = 0.3$, shown as the orange line in Figure 6.12.

6.14 White noise example

6.14.1 Problem statement

Suppose we have three uncorrelated white noise signals, a, b, c. We are interested in the noise power in c. Suppose further we only have access to two signals

$$x[n] = a[n] + c[n], \qquad y[n] = b[n] + c[n-2] \tag{6.108}$$

so c is correlated in x and y, with a two-cycle delay in y. The sampling frequency f_s is 10 kHz, so this corresponds to a delay of 0.2 milliseconds.

Unfortunately, the signals x and y are glitchy, i.e. they have short periods of non-Gaussian behavior. Is it possible to estimate the correlated power in x and y?

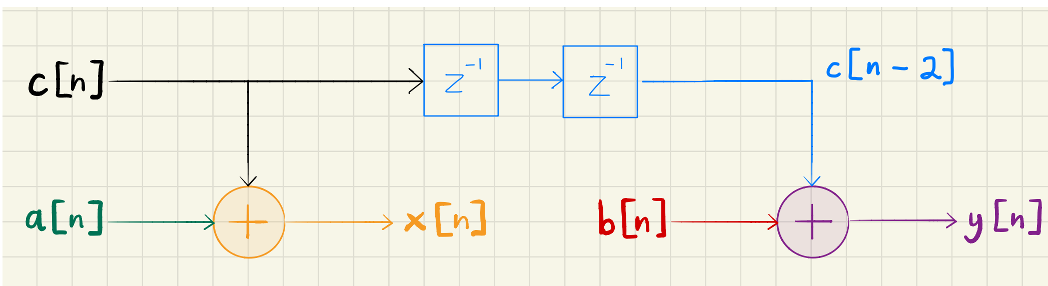

Figure 6.14: Independent Gaussian signals a, b, and c form the signals x and y. These signals will form the basis of the CSD $\langle x, y\rangle$ derivations in this section. c is the correlated signal in both x and y. a is the noise on the input. b is the noise on the output. z^{-1} is a single-cycle delay.

6.14.2 Solution setup

The correlated power can be estimated by taking the CSD $\langle x, y\rangle$, using median-averaging to accommodate glitches, and correcting for the mean-to-median bias, which is estimated from the median-calculated coherence.

The derivation of the probability density function for each frequency bin in $\langle x, y\rangle$ is close to that in Section 6.10. We find the coherence γ^2, power ratio ϵ, and cross-power scaler λ for our noise diagram in Figure 6.14, and relate this back to our derived equations in the previous sections.

Using the definition of mean-averaged coherence, Eq. 6.89:

$$
\begin{aligned}
\gamma^2 &= \frac{|\langle x, y\rangle|^2}{\langle x, x\rangle\langle y, y\rangle} \\
&= \frac{(2\sigma_c^2)^2}{(2(\sigma_a^2+\sigma_c^2))(2(\sigma_b^2+\sigma_c^2))} \\
\gamma^2 &= \frac{1}{1+\frac{\sigma_a^2\sigma_b^2+\sigma_a^2\sigma_c^2+\sigma_b^2\sigma_c^2}{\sigma_c^4}}
\end{aligned} \tag{6.109}
$$

Using the definition for the correlated-to-uncorrelated power ratio ϵ in Eq. 6.89, we get

$$
\epsilon = \frac{\sigma_a^2\sigma_b^2+\sigma_a^2\sigma_c^2+\sigma_b^2\sigma_c^2}{\sigma_c^4}, \tag{6.110}
$$

then the mean-to-median bias b simplifies to Eq. 6.86.

Finally, the cross-power scaler λ from Eq. 6.64 is

$$
\lambda = \sqrt{\sigma_a^2\sigma_b^2+\sigma_a^2\sigma_c^2+\sigma_b^2\sigma_c^2}. \tag{6.111}
$$

With the above definitions, Eq. 6.100 describes the distribution of $\langle x, y\rangle$ for every frequency bin. Figure 6.16 shows a 2d histogram of 100000 individual cross spectral densities at $f = 500$ Hz, along with the contours of Eq. 6.100.

Figure 6.15 shows the results of several stages of the CSD estimation process described in Section 6.15. The first is the mean-averaged CSD, which arrives at the final answer with little difficulty. The second is the naive median-averaged CSD, which is biased away from the final answer by two effects: the usual mean-to-median bias calculated in Eq. 6.86, and the phase-effect from not taking the median of the CSD along its major axis. The third is the "phase-corrected" median-averaged CSD, which simply rotates the axis along which the median-averaging is done, removing the $\sim 4\%$ bias in the median estimate.

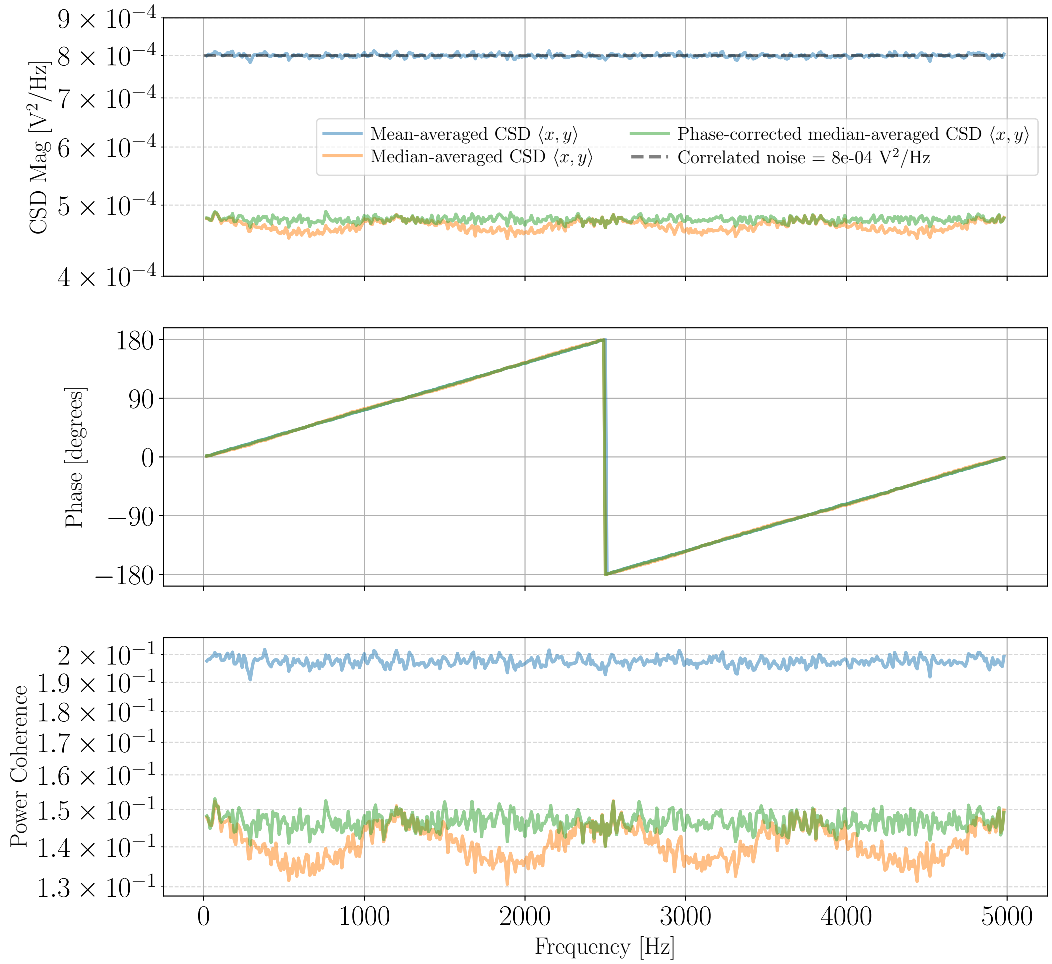

Figure 6.15: Cross spectral densities with 100000 averages for the white noise example shown in Figure 6.14. The effect the phase has on the median-averaged CSD and coherence is seen in the orange lines, also demonstrated in Figure 6.12. The "phase-corrected" median-averaged CSD and coherence are in green. The phase-corrected densities were rotated such that their major axis aligned with the real axis before the median was taken. The power density in the uncorrelated noise $2\sigma_a^2/f_s = 2\sigma_b^2/f_s = 1 \times 10^{-3}$ V^2/Hz, while correlated noise $2\sigma_c^2/f_s = 8 \times 10^{-4}$ V^2/Hz. The maximum bias due to phase in the CSD median seen here is 3.8%. The uncorrelated over correlated power ratio $\epsilon \approx 4.1$. The mean-averaged coherence $\overline{\gamma^2} \approx 0.20$. The median-averaged coherence $\widetilde{\gamma^2} \approx 0.15$. The mean-to-median bias $b \approx 0.60$.

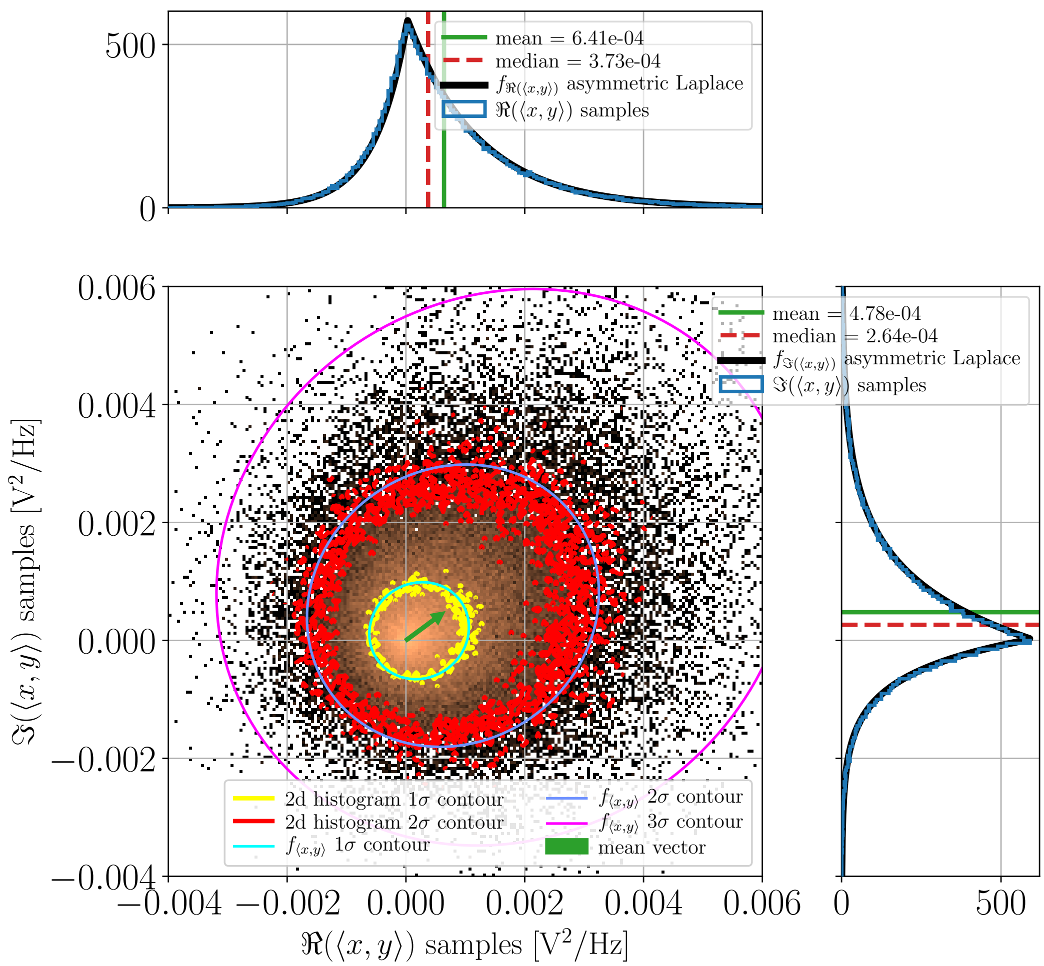

Figure 6.16: 2D histogram of a CSD $\langle x, y\rangle$ at $f = 500$ Hz of white noise processes $x(t)$ and $y(t)$. There is a time delay τ of 0.2 milliseconds between x and y, yielding a phase $\phi = 2\pi f\tau = 0.63$ radians. The correlated noise falls along the major axis, also called the radial axis.

6.15 Discussion

Suppose we have some glitchy data for which mean-averaging is not feasible. The steps of the median-averaged CSD estimate process are:

1. Estimate the PSDs $\widetilde{\langle x, x\rangle}$, $\widetilde{\langle z, z\rangle}$, and CSD $\widetilde{\langle x, z\rangle}$. (For now we assume the densities are well-sampled, i.e. the sample median $\hat{\rho}$ has converged to the actual median ρ for every frequency bin.)
2. Re-estimate the CSD median $\widetilde{\langle x, z\rangle}$ after rotating all of the individual densities by the opposite of the CSD phase $-\phi$. This rotates the CSDs so the median-averaging is done along its major axis, removing the effect phase has on the final magnitude estimate.
3. Estimate the median-averaged coherence $\widetilde{\gamma}^2 = |\widetilde{\langle x, z\rangle}|^2/(\widetilde{\langle x, x\rangle}\widetilde{\langle z, z\rangle})$.
4. Numerically solve for the power ratio ϵ using the coherence $\widetilde{\gamma}^2$ and Eq. 6.94.
5. Calculate the mean-to-median bias from the median-averaged coherence using Eq. 6.96.
6. Apply the bias to the CSD $\langle x, z\rangle$ to recover the mean-averaged CSD.

Here we summarize the main results from the previous sections, and their implications on spectral analysis.

- Power spectral densities $\langle x, x\rangle$ are commonly reported as mean-averaged results from Welch's method in section 6.4. Mean-averaging is useful because the final PSD estimate obeys Parseval's theorem (Eq. 6.7), also known as signal processing conservation of energy. Each frequency bin of a PSD follows an exponential distribution with mean $\lambda = 2\sigma_x^2/f_s$, as reported in Eq. 6.41.
- Median-averaging the PSD can gracefully handle glitchy data, as long as the glitches are infrequent. However, the final PSD estimate does not obey Parseval's theorem. To correct for this, the mean-to-median bias must be applied. For a sufficiently large number of averages, this bias is $\log(2)$, as seen from Eq. 6.37.
- Amplitude spectral densities $\sqrt{\langle x, x\rangle}$ are commonly reported as root-mean-squared averaged results, or just the root of the result of Welch's method.

Each frequency bin of the amplitude spectral density follows a Rayleigh distribution, and from Eq. 6.45 the root mean square of the Rayleigh is equal to the root of the power divided by the Nyquist frequency: $r = \upsilon\sqrt{2} = \sqrt{2\sigma_x^2/f_s}$, as reported in 6.49.

- Cross spectral densities $\langle x, z\rangle$ are more complicated because they have both real and imaginary parts, there are correlated and uncorrelated noises, and the correlated phase is unknown. However, with sufficient averages, mean-averaging "averages away" the uncorrelated noise, yielding only the correlated signal divided by the Nyquist $\mu = 2\sigma_x^2/f_s$. This fulfills Parseval's theorem, recovering only the correlated power in the two signals. The same is not true for median-averaging, where the level of the uncorrelated signal affects the final result, as seen in Eq. 6.84. This is demonstrated in Figure 6.9.
- In the limit of large correlated noise relative to uncorrelated noise $\epsilon \ll 1$, the CSD approaches the PSD of the correlated noise, because the asymmetric Laplace distribution approaches the exponential distribution. This is can be seen by the blue curve in Figure 6.9.
- The mean-to-median bias for PSDs is $b = \dfrac{m}{\lambda} = \log(2)$ for a large number of averages M, as calculated for the exponential distribution in Eq. 6.37.
- The mean-to-median bias for CSDs can be calculated from the median-averaged coherence: $b = \log(2)\sqrt{(1+\epsilon)\widetilde{\gamma^2}}$ (Eq. 6.96). ϵ is the uncorrelated/correlated power ratio, and can be solved numerically from coherence using Eq. 6.94.
- The mean-to-median bias for CSDs ranges from $\log(2)$ for highly correlated CSDs to $1/2$ for totally uncorrelated CSDs, as seen in Figure 6.10.
- The phase of the CSD ϕ can have an small effect on the resulting median-averaged magnitude. The median of the CSD must always be taken along the major axis of the CSD to avoid phase-based biasing in the median calculation.

Chapter 7

FUTURE WORK

This thesis overviewed topics in advanced interferometry and spectral statistics for the purpose of increasing the sensitivity of the Advanced LIGO detectors for observing run three. With the successful introduction of squeezing and high laser power, and the suppression of the technical laser noises that would limit the advantage of those upgrades, Advanced LIGO achieved the lowest quantum noise ever in long-baseline interferometers, and the best sensitivity yet to gravitational waves.

7.1 Noise considerations

The biggest challenge for O4 and beyond is the achievement of design sensitivity. Figure 7.1 shows the difference in DARM sensitivity between sites in O3. Overall, Livingston enjoys a better sensitivity across the entire GW detection band. Livingston's better sensitivity is due to a number of reasons:

1. Better squeezing overall, including less frequency-dependent loss (see Figure 5.6),
2. Higher arm power (see Table 3.3),
3. Much lower length and angular controls noise,
4. Higher detector bandwidth, likely due to lower SRC losses (see Figure 3.24)

The most important frequency region for astrophysical range is the low-frequency region, where binary inspirals spend more time orbiting ($\sim$seconds rather than $\sim$milliseconds at merger). Thus the most critcal region for Hanford to improve sensitivity to GWs is the 30 to 70 Hz region. However, some new astrophysics is being done with high frequency gravitational waves near merger, including higher multipole GW detection [11] and neutron star equation of state measurements [9].

The biggest obstacle to achieving design sensitivity is the "mystery noise" limiting GW sensitivity around 30 Hz. Mystery noise is not coherent with any other witnesses in Advanced LIGO, so mitigation is not straightforward. However, there is

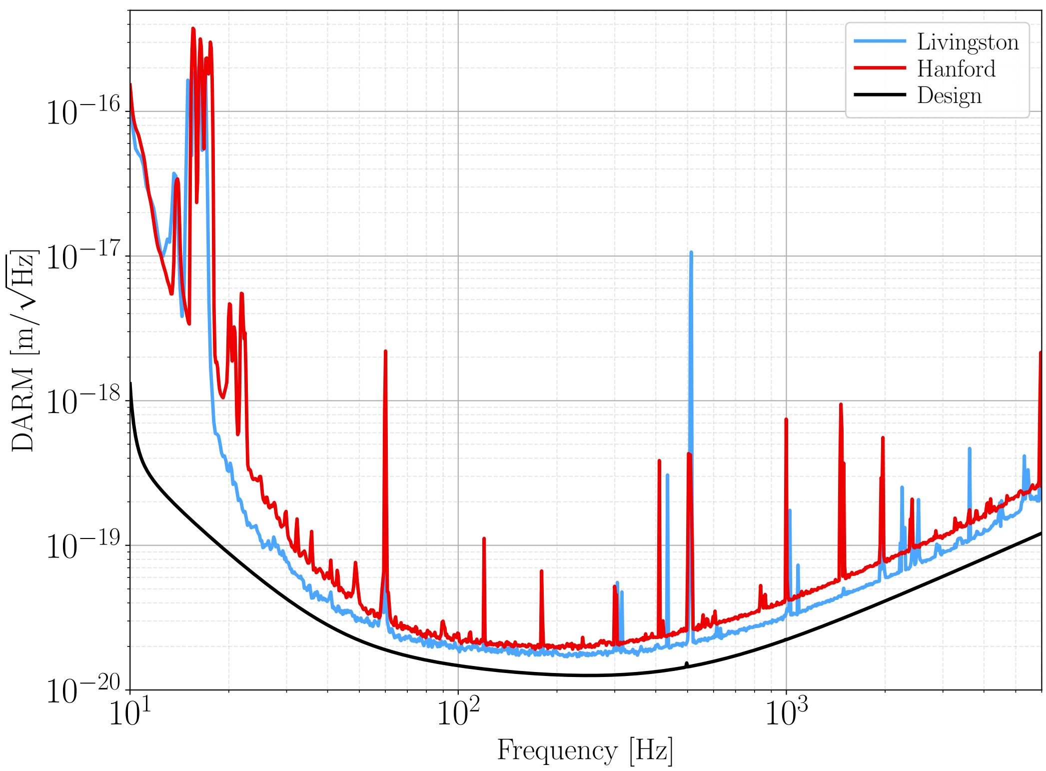

Figure 7.1: DARM sensitivity between Hanford and Livingston during O3a, plus Advanced LIGO design sensitivity (125 W input power, no squeezing).

less mystery noise at Livingston than Hanford, which accounts for most of their increased range compared to Hanford in O3. This suggests some difference between the detectors is causing the excess mystery noise at Hanford. The most likely culprit for mystery noise is excess low frequency motion due to angular and length sensing and control upconverting to higher frequencies, polluting the DARM spectrum.

Angular controls noise is the next biggest obstacle to achieving design sensitivity. This is the dominant noise everywhere below 30 Hz. Angular motion must be suppressed to keep the interferometer locked. Residual angular motion will couple to length degrees of freedom both linearly and nonlinearly, from beam miscentering to spontaneous power fluctuations from poor alignment causing radiation pressure fluctuations.

Sensing the residual angular motion above 10 Hz is limited by the wavefront sensor noise. To filter this noise, aggressive low-pass filters in the angular control loops are engaged. This reduces the angular control gain above 10Hz, and drastically reduces angular control noise coupling to DARM. However, this also reduces the phase margin of the loops to close to few tens of degrees.

With increasing power in the arm cavities, the optical plant of the angular degrees of freedom dynamically change due to increasing radiation pressure torque, which must be compensated. Additionally, the cross-coupling of the loops with length degrees of freedom and other angular degrees of freedom is relatively high with the purposeful beam spot miscentering due to the point absorbers. If cross-coupling is high, then controls noise from one degree of freedom pollutes another degree of freedom, causing additional noise.

Because of the nonlinearity of angular motion coupling to DARM, its very likely controls noise goes hand-in-hand with mystery noise at both sites. Improvements here will either directly mitigate mystery noise, or reveal more of the underlying noise.

Finally, the source of frequency-dependent squeezing must be found and mitigated to achieve the full benefits of injected squeezing. This will be more important with the addition of the filter cavity for O4.

7.2 O4 upgrades

Currently upgrades for observing run four are underway at both Hanford and Livingston.

The biggest upgrade is the addition of the filter cavity, a 300 m long cavity for the injection of frequency-dependent squeezing [208, 209]. This will allow reduction in both quantum shot noise and quantum radiation pressure noise at the DARM readout. Quantum radiation pressure noise has been measured to limit DARM at Livingston [89], and to a lesser extent Hanford. Work on low-frequency controls is required to see the full advantage of the filter cavity.

Next, at Hanford ITMY has been replaced. ITMY had a large point absorber during O3, requiring beam spot moves to mitigate some of its impact. The intentional beam miscentering which adversely affects the angular control loops will no longer be necessary. Additionally, the new ITMY transmission is 1.50%, which matches ITMX and should balance the arm powers and arm poles at Hanford.

7.3 Future projects

Work on building the filter cavity and improving the performance of the squeezer are already underway at the site. Commissioning the filter cavity and achieving higher circulating power in the arms must happen prior to O4.

Other projects may yield significant insight into Hanford's performance. Some future projects for understanding detector noise include

1. Create MIMO control models for interferometer angular and length controls using pytickle [210].
2. Create MIMO control models for angle to length couplings [126].
3. Thoroughly test and model the effects of thermal compensation on DARM and the SRC together.
4. Create a plausible power budget for the interferometer carrier and sidebands.
5. Create plausible RF phase and amplitude noise projections to DARM.
6. Develop infrastructure for long-term median-averaged coherence measurements for "quiescent" transfer functions.
7. Investigation of calibration error requirement for future detectors.
8. Development of sub-percent calibration precision schemes.

7.4 Conclusions

Observing run three by Advanced LIGO and Advanced Virgo was the most successful search for gravitational waves in history, with 39 confirmed detections in only the first half of the run [3]. It was also only the third successful search for gravitational waves. Our sensitivity to these extreme astrophysical phenomena is only in its infancy.

In gravitational wave detector science, a small improvement in DARM sensitivity can have a tremendous impact on astrophysics. Now we push our current detectors to their sensitivity limits, to further expand the horizons of the new field of gravitational wave astronomy.

Appendix A

MODULATION, QUADRATURES, AND ELECTRIC FIELD UNITS

Here we will derive the expected gravitational wave signal from a simple Michelson interferometer, alongside the expected laser frequency and intensity noise couplings, as has been done before [78, 114, 211–213].

In other chapters we explored the frequency and intensity stabilization schemes, as well are the full detector response to GWs for accurate calibration. The purpose of this chapter is to introduce the technical origins of concepts important to this thesis, including the sideband representation, amplitude and phase quadratures, homodyne angle, optical gain, contrast defect, and DC offset.

This appendix helps draw the connection between the theoretical representation and the instrumental results. Often, random scale factors are applied to help data match up with theory. If losses are a parameter we care about, which they are in Advanced LIGO for O3, random scale factors are degenerate with those parameters.

The overview facilitates a direct comparison of analytic results like those in [22, 114], to numerical results like those calculated from Finesse [134, 135], to actual measurements taken at LIGO Hanford.

A.1 Modulation

First we review the physics of modulation in the sideband picture [214]. The modulation picture emphasizes the different frequencies of light that are created by different interactions. GWs are detected as infinitesimal modulations applied to highly stabilized laser light. Noise in the laser light, whether quantum or classical, is modulations not caused by GWs.

A perfectly noiseless electric field E is known as the *carrier*:

$$E = E_0 e^{i\omega_0 t} \tag{A.1}$$

where ω_0 is the carrier frequency, E_0 is the carrier amplitude, and t is time.

A.1.1 Phase modulation

A *phase modulation* of amplitude $\delta\phi$ at frequency ω can be applied to the carrier light:

$$\begin{aligned} E_{\delta\phi} &= E_0 e^{i\phi} \\ E_{\delta\phi} &= E_0 e^{i(\omega_0 t + \delta\phi\cos(\omega t))} \end{aligned} \tag{A.2}$$

This can be thought of as splitting the carrier power, which is always a sine wave at ω_0, off into *sidebands* at frequencies $\omega_0 \pm \omega$. Using the Jacobi-Anger expansion on Eq. A.2 yields:

$$E_{\delta\phi} = E_0 e^{i\omega_0 t} \sum_{n=-\infty}^{\infty} i^n J_n(\delta\phi) \exp(in\omega t) \tag{A.3}$$

where J_n is the n^{th} Bessel function of the first kind.

If we assume $\delta\phi$ is small, then we can ignore the higher-order sidebands $n \geq 2$, and write Eq. A.3 as

$$E_{\delta\phi} = E_0 e^{i\omega_0 t} \left[J_0(\delta\phi) + iJ_1(\delta\phi)e^{i\omega t} - iJ_{-1}(\delta\phi)e^{-i\omega t} \right] \tag{A.4}$$

Finally, using $J_{-1}(\delta\phi) = -J_1(\delta\phi)$, $J_0(\delta\phi) \approx 1$, and $J_1(\delta\phi) \approx \delta\phi/2$, we write the final phase modulation in terms of the carrier ω_0, upper sideband $\omega_0 + \omega$ and lower sideband $\omega_0 - \omega$:

$$E_{\delta\phi} = E_0 e^{i\omega_0 t} \left(1 + i\frac{\delta\phi}{2} E_0 e^{i\omega t} + i\frac{\delta\phi}{2} E_0 e^{-i\omega t} \right) \tag{A.5}$$

The key observation of Eq. A.5 is the relative phase of the sidebands compared with the carrier. The sidebands are aligned with one another when they are orthogonal to the carrier. Calculating the power in the field,

$$\begin{aligned} P_{\delta\phi} &= |E_{\delta\phi}|^2 \\ &= |E_0|^2 \left(1 + \frac{(\delta\phi)^2}{4} \left(e^{i2\omega t} + e^{-i2\omega t} \right) \right) \\ P_{\delta\phi} &\approx |E_0|^2. \end{aligned} \tag{A.6}$$

The sidebands push and pull the phase of the carrier by $\delta\phi$, but to first order do not alter the amplitude. Figure A.1b illustrates the sideband and quadrature picture of phase modulation.

A.1.2 Frequency noise

Frequency noise is mathematically equivalent to a phase modulation. Using the definition of frequency as the time derivative of phase, $d\phi/dt$, and the phase ϕ from Eq. A.2, we calculate the relationship between frequency noise and phase modulation [215]:

$$\begin{aligned}\frac{d\phi}{dt} &= \frac{d}{dt}(\omega_0 t + \delta\phi\cos(\omega t)) \\ \frac{d\phi}{dt} &= \omega_0 - \omega\delta\phi\sin(\omega t)\end{aligned} \tag{A.7}$$

The frequency can be broken down into the carrier term ω_0 and the noise term $\delta\nu$, where

$$\frac{2\pi\delta\nu}{\omega} = \delta\phi \tag{A.8}$$

where $\delta\nu$ is the amplitude of the frequency swing.

We can substitute Eq. A.8 into Eq. A.5 with no change in the final result (except an arbitrary phase advance of $\pi/2$ for both sidebands):

$$E_{\delta\nu} = E_0 e^{i\omega_0 t}\left(1 - \frac{\pi\delta\nu}{\omega}e^{i\omega t} + \frac{\pi\delta\nu}{\omega}e^{-i\omega t}\right) \tag{A.9}$$

Here we recall the distinction between ω_0, ω, and $\delta\omega$. The carrier frequency is ω_0, this is a constant, $\omega_0 = 2\pi c/\lambda = 1.77 \times 10^{15}$ rad/s (The Advanced LIGO laser wavelength $\lambda = 1064$ nm). The modulation frequency itself is ω, this is how fast the carrier frequency changes. The frequency modulation amplitude $\delta\nu$ is how much the carrier frequency changes.

A.1.3 Amplitude modulation

An *amplitude modulation* of amplitude δE at frequency ω can be applied to the carrier light:

$$\begin{aligned}E_{\delta E} &= E_0 e^{i\omega_0 t}\left(1 + \frac{\delta E}{E_0}\cos(\omega t)\right) \\ E_{\delta E} &= E_0 e^{i\omega_0 t}\left(1 + \frac{\delta E}{2E_0}e^{i\omega t} + \frac{\delta E}{2E_0}e^{-i\omega t}\right)\end{aligned} \tag{A.10}$$

Again, the key observation of Eq. A.10 is the relative phase of the sidebands compared with the carrier. This time, the sidebands are aligned with one another when they are also aligned with the carrier. Calculating the power in the field,

$$\begin{aligned} P_{\delta E} &= |E_{\delta E}|^2 \\ &= |E_0|^2\left(1 + \frac{2\delta E}{E_0}\cos(\omega t) + \frac{(\delta E)^2}{E_0^2}\cos(\omega t)^2\right) \\ P_{\delta E} &\approx P_0\left(1 + \frac{2\delta E}{E_0}\cos(\omega t)\right). \end{aligned} \tag{A.11}$$

The sidebands push and pull the amplitude of the carrier by δE, but do not alter the phase. Figure A.1a illustrates the sideband and quadrature picture of amplitude modulation.

A.1.4 Intensity noise

Relative intensity noise (RIN) is mathematically equivalent to amplitude modulation. From Eq. A.11, we can relate relative intensity noise to relative amplitude noise (RAN). Dividing Eq. A.11 by P_0, we define the relative intensity noise in terms of amplitude modulation:

$$\frac{\delta P}{P_0} = \frac{2\delta E}{E_0} \tag{A.12}$$

Going back to the expression for amplitude modulation $E_{\delta E}$ Eq. A.10, we can express the electric field $E_{\delta P}$ and power $P_{\delta E}$ in terms of relative intensity noise $\delta P/P$:

$$E_{\delta P} = E_0 e^{i\omega_0 t}\left(1 + \frac{\delta P}{4P_0}e^{i\omega t} + \frac{\delta P}{4P_0}e^{-i\omega t}\right) \tag{A.13}$$

$$P_{\delta E} = P_0\left(1 + \frac{\delta P}{P_0}\cos(\omega t)\right) \tag{A.14}$$

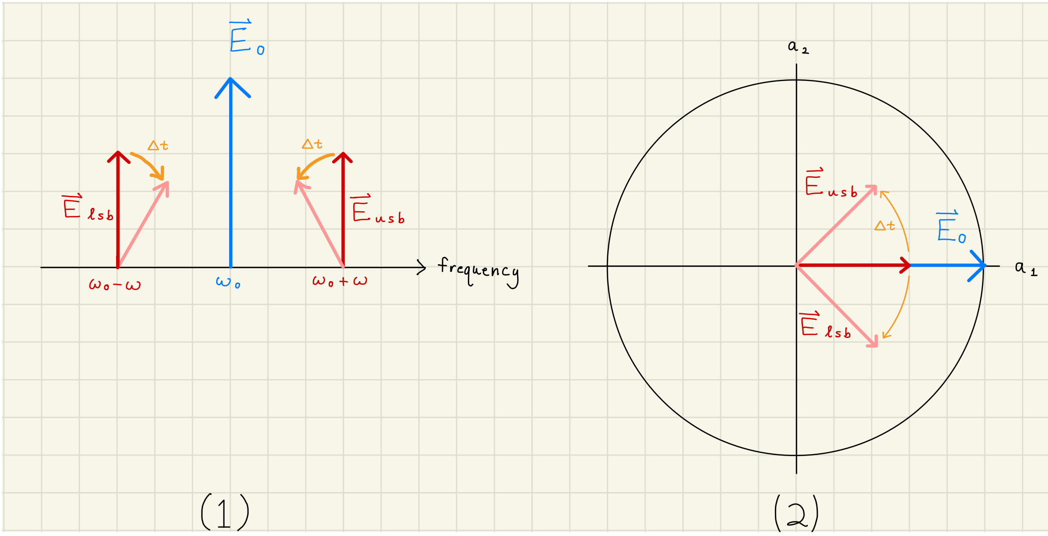

(a) Amplitude modulation

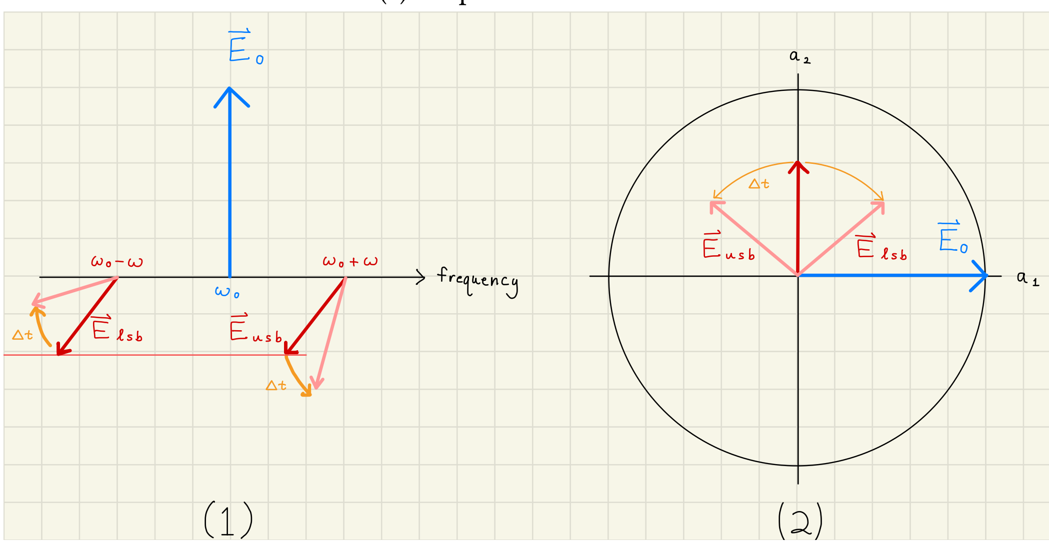

(b) Phase modulation

Figure A.1: Diagram (1) illustrates the sideband picture, with the x-axis representing frequency and the y- and z-axes representing phase. Diagram (2) is a phasor diagram, looking along the frequency axis. The static carrier $\vec{E}_0$ oscillates at the carrier frequency ω_0, the upper sideband $\vec{E}_{usb}$ at the frequency $\omega_0 + \omega$, and the lower sideband $\vec{E}_{lsb}$ at the frequency $\omega_0 - \omega$. $\vec{E}_{usb}$ and $\vec{E}_{lsb}$ rotate relative to the carrier $\vec{E}_0$ in opposite directions.

A.2 General sideband power

The above sections focused on special cases of modulation. Here we derive a general expression for power for some arbitrary modulation at signal frequency ω.

We write the modulated carrier as In the sideband picture, the total electric field can be written

$$\vec{E}_{tot} = E_0 e^{i\omega_0 t}\left[1 + \delta_+ e^{i\omega t} + \delta_- e^{-i\omega t}\right] \tag{A.15}$$
$$= \vec{E}_0 + \vec{E}_+ e^{i\omega t} + \vec{E}_- e^{-i\omega t} \tag{A.16}$$

Looking for the power components at DC, ω, and 2ω:

$$P_tot = |\vec{E}_{tot}|^2 \tag{A.17}$$
$$= |\vec{E}_0 + \vec{E}_+ e^{i\omega t} + \vec{E}_- e^{-i\omega t}|^2 \tag{A.18}$$
$$P_0 = |\vec{E}_0|^2 + |\vec{E}_+|^2 + |\vec{E}_-|^2 \tag{A.19}$$
$$P_\omega = (\vec{E}_0^* \vec{E}_+ + \vec{E}_0 \vec{E}_-^*) e^{i\omega t} + (\vec{E}_0 \vec{E}_+^* + \vec{E}_0^* \vec{E}_-) e^{-i\omega t} \tag{A.20}$$
$$P_{2\omega} = \vec{E}_+ \vec{E}_-^* e^{i2\omega t} + \vec{E}_+^* \vec{E}_- e^{-i2\omega t} \tag{A.21}$$

P_ω carries the linear audio signal we care about. The two components of the sum making up P_ω are complex conjugates of one another, so we can write

$$P_\omega = 2\Re((\vec{E}_0^* \vec{E}_+ + \vec{E}_0 \vec{E}_-^*) e^{i\omega t}) \tag{A.22}$$

or, leaving P_ω as a complex quantity,

$$\tilde{P}_\omega = 2(\vec{E}_0^* \vec{E}_+ + \vec{E}_0 \vec{E}_-^*) \tag{A.23}$$

A.3 Quadratures

Amplitude and phase quadratures offer a convenient way of representing light. Instead of upper and lower sidebands, modulations are broken down into how they affect the carrier.

Using the two-photon formalism of Caves and Schumaker [216, 217], the quantum annihilation and creation operators for photons of frequency ω, a_ω and $a_\omega^\dagger$, are abused to describe the modulation process as annihilating a photon at one frequency and creating another photon at a different, nearby frequency.

Following [22, 114], the upper sideband annihilation operator a_+ and the lower sideband annihilation operator a_- are defined as

$$a_+ = a_{\omega_0+\omega}\sqrt{\frac{\omega_0+\omega}{\omega_0}} \qquad a_- = a_{\omega_0-\omega}\sqrt{\frac{\omega_0-\omega}{\omega_0}}. \tag{A.24}$$

Because the carrier frequency ω_0 is much larger than the signal frequencies we care about ω, we can safely ignore the term in the root.

The amplitude and phase quadrature operators a_1 and a_2 are now defined as

$$a_1 = a_+ + a_-^\dagger \qquad a_2 = i(a_+ - a_-^\dagger). \tag{A.25}$$

a_1 and a_2 do not commute with each other, but do commute with themselves, unlike a_+ and a_-. We use the definition in Eq. A.25 to make shot noise units equal to one (Subsection A.3.1), and make taking the dot product between static and modulated sidebands straightforward (Subsection A.3.2).

The quadrature operators can be written as a vector $\vec{a}$ representing the modulations of the total electric field:

$$\vec{a} = \begin{pmatrix} a_1 \\ a_2 \end{pmatrix}. \tag{A.26}$$

Figure A.1 illustrates the amplitude and phase quadratures on the phasor diagrams.

A.3.1 Quadrature units to electric field units

The quadrature operators a_1 and a_2 are general quantum operators representing field amplitudes oscillating 90° out of phase we each other. a_1 and a_2 do not commute [196]:

$$[a_1, a_2] = 2i \tag{A.27}$$

as long as a_1 and a_2 represent the same signal frequency $\omega_0 \pm \omega$.

The Heisenberg uncertainty for these non-commuting operators is

$$\sigma_{a_1}^2 \sigma_{a_2}^2 \geq \frac{1}{4} |\langle [a_1, a_2] \rangle|^2 \tag{A.28}$$

$$\sigma_{a_1}^2 \sigma_{a_2}^2 = 1 \tag{A.29}$$

For unsqueezed vacuum, we have

$$\langle a_1 \rangle = \langle a_2 \rangle = 0 \tag{A.30}$$

$$\sigma_{a_1}^2 = \sigma_{a_2}^2 = 1. \tag{A.31}$$

We say that a_1 and a_2 have "units of shot noise" such that their variance is equal to one. This is used in Eq. 5.22.

To return to units of electric field, we first note that LIGO literature scales electric fields to have units of $\sqrt{\mathrm{W}}$ so the power in watts can be calculated like

$$P = |E|^2. \tag{A.32}$$

In electricity and magnetism, power in an electric field is expressed like

$$P = AI = \frac{Ac\epsilon_0|E|^2}{2} \tag{A.33}$$

where A is the surface area of the beam, I is the beam intensity in $\mathrm{W/m^2}$, and ϵ_0 is the vacuum permeability. Here, E has units $\mathrm{V/m}$, ϵ_0 has units $\mathrm{N/V^2}$, A has $\mathrm{m^2}$, and c has $\mathrm{m/s}$ Because the entire beam is assumed to be captured by the photodetector, we incorporate the prefactor $Ac\epsilon_0/2$ into our electric field for simplicity.

Finally, we relate the quadrature operators a_1 and a_2 to their equivalent electric fields E_1 and E_2 by

$$a_1 = \sqrt{\hbar\omega_0}E_1 \qquad a_2 = \sqrt{\hbar\omega_0}E_2. \tag{A.34}$$

This converts units of shot noise to the LIGO units of $\sqrt{W}$. (The real units of $\hbar\omega_0$ are joules, but the time average has been incorporated into our LIGO electric field units).

A.3.2 Modulations in quadrature representation

We can express the sideband amplitude and phase modulations from Eqs.A.13 and A.9 in quadrature representation using A.25.

We break up the total electric field $\vec{E}_{tot} = \vec{E} + \vec{e}(\omega)$ into a "static" component $\vec{E}$ centered around the carrier frequency ω_0 and a modulation, or Fourier, component $\vec{e}(\omega)$ at the relative signal frequencies $\omega_0 \pm \omega$. Expressing an electric field modulated in both intensity and frequency $\vec{E}_{tot}$ from Eqs. A.13 and A.9 in quadrature representation yields

$$\vec{E}_{tot} = \vec{E} + \vec{e}(\omega) \tag{A.35}$$

$$\vec{E} = E_0 e^{i\omega_0 t}\begin{pmatrix}1\\0\end{pmatrix} \tag{A.36}$$

$$\vec{e}(\omega) = E_0 e^{i\omega_0 t} e^{i\omega t}\left[\frac{2\pi\delta\nu}{\omega}\begin{pmatrix}0\\1\end{pmatrix} + \frac{\delta P}{2P}\begin{pmatrix}1\\0\end{pmatrix}\right] \tag{A.37}$$

Eq. A.37 is equivalent to Eq. 10 in [114], except for a factor of $\sqrt{2}$ from different quadrature definitions (Eq. A.25 vs Eq. 6 in [114]).

For example, we recover intensity noise from Eq. A.14. Let $\delta\nu \to 0$ for $\vec{e}(\omega)$ in

Eq. A.37, and take the dot product of $\vec{E}_{tot}$ with itself to get power $P_{\delta P}$:

$$\begin{aligned}
P_{\delta P} &= \left|\vec{E} + \vec{e}(\omega)\right|^2 \\
&= \left|\vec{E}\right|^2 + \vec{E}^* \cdot \vec{e}(\omega) + \vec{E} \cdot \vec{e}^*(\omega) + |\vec{e}(\omega)|^2 \\
&\approx E_0^2\left(1 + \frac{\delta P}{2P}e^{i\omega t} + \frac{\delta P}{2P}e^{-i\omega t}\right)\begin{pmatrix}1 & 0\end{pmatrix} \cdot \begin{pmatrix}1 \\ 0\end{pmatrix} \\
P_{\delta P} &= P_0\left(1 + \frac{\delta P}{P}\cos(\omega t)\right) \qquad \text{(A.38)}
\end{aligned}$$

The result from Eq. A.38 matches Eq. A.14.

The same is not true for definitions like Eq. 6 in [114] or Eq. 2.7 in [22] or Eqs. 2.53 and 2.54 in [196]: attention must be paid to the quadrature definition to ensure theory agrees with instrumental results.

A.4 Homodyne angle

The *homodyne angle* ζ is the static angle between the carrier phase and the signal sideband basis. The homodyne angle controls what mixed quadrature a_ζ is measured on the photodetector [47, 196, 218, 219]:

$$a_\zeta = a_1 \cos\zeta + a_2 \sin\zeta \qquad \text{(A.39)}$$

In the sideband picture, the total electric field can be written

$$E_{tot} = E_0 e^{i\omega_0 t}\left[e^{i\zeta} + \frac{\delta_+}{2}e^{i\omega t} + \frac{\delta_-}{2}e^{-i\omega t}\right] \qquad \text{(A.40)}$$

where δ_+, δ_- is some arbitrary complex modulation.

Figure A.2 shows a basic example of homodyne detection, with local oscillator carrier $\vec{L}$ and sidebands $\vec{a}$. The local oscillator phase can be adjusted by the homodyne angle

$$\vec{L} = E_0 e^{i\omega_0 t} e^{i\zeta} \qquad \text{(A.41)}$$

In terms of the total electric field and total power oscillations, using the quadrature

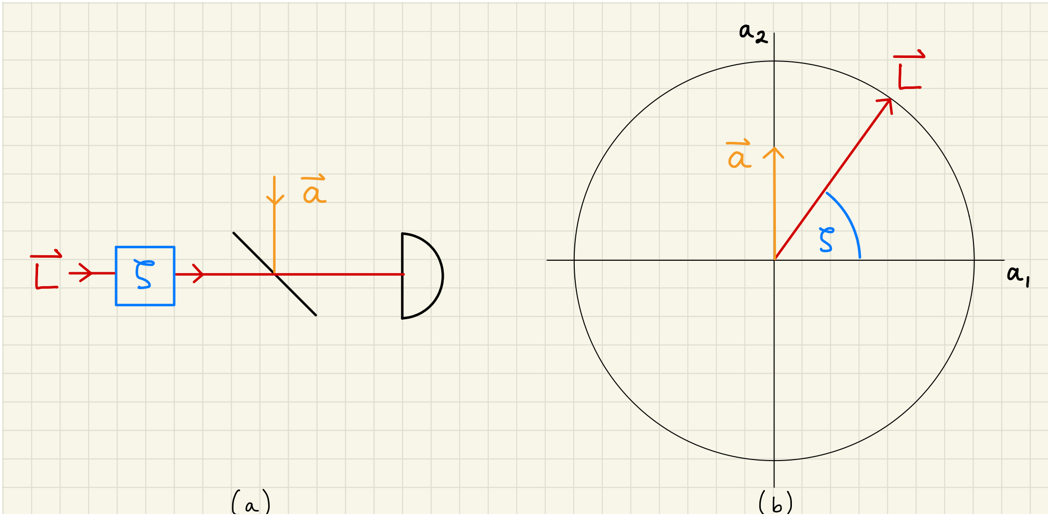

Figure A.2: Diagram (a) illustrates a local oscillator $\vec{L}$ beating with a signal sideband $\vec{a}$ with homodyne angle ζ. Diagram (b) is a phasor diagram illustrating the homodyne angle. In this diagram, $\vec{a}$ is entirely in a_2, so the homodyne angle must be adjusted to $\zeta = \pi/2$ or $\zeta = 3\pi/2$.

representation,

$$\vec{E}_{tot} = \frac{1}{\sqrt{2}}\left(\vec{L} + \vec{a}\right) \tag{A.42}$$

$$P_{tot} = \frac{1}{2}\left|\vec{L} + \vec{a}\right|^2$$

$$P_{tot} = \frac{1}{2}\left|\vec{L}\right|^2 + \vec{L} \cdot \vec{a}$$

$$P_{tot} = \frac{1}{2}\left|\vec{L}\right|^2 + E_0\begin{pmatrix}\cos\zeta & \sin\zeta\end{pmatrix}\begin{pmatrix}a_1 \\ a_2\end{pmatrix}$$

$$P_{tot} = \frac{1}{2}E_0^2 + E_0 a_\zeta$$

$$P_{tot} = \frac{1}{2}P_0 + P_\zeta \tag{A.43}$$

Suppose that $\vec{a}$ is in a_2, as in Figure A.2. Then the detected power oscillations at P_ζ is:

$$P_\zeta = E_0 a_2 \sin\zeta \tag{A.44}$$

and we can detect the phase modulation when the homodyne angle $\zeta = \pi/2$.

A.4.1 Conflicting homodyne angle definition

In Buonanno and Chen's series of landmark Advanced LIGO quantum noise papers [22, 125, 220], the new homodyne angle ζ_{bnc} is defined differently than in Eq. A.39.

For these results, Buonanno and Chen used the homodyne angle definition

$$a_\zeta = a_1 \sin\zeta_{bnc} + a_2 \cos\zeta_{bnc}. \tag{A.45}$$

In this case, $\zeta_{bnc} = 0$ yields a_2, while $\zeta_{bnc} = \pi/2$ yields a_1. The relationship between ζ_{bnc} and ζ is

$$\zeta_{bnc} = \frac{\pi}{2} - \zeta \tag{A.46}$$

Because the results of [22] were so influential, Advanced LIGO literature uses this definition [20, 57, 114, 119, 140, 213]. Advanced LIGO is said to run with a homodyne angle of $\zeta_{bnc} = \pi/2$, as the interferometer is configured such that the signal comes out in the amplitude quadrature b_1 of the input electric field [20, 22].

For consistency with Advanced LIGO literature, this thesis also uses the Buonanno and Chen homodyne angle ζ_{bnc}, (see Chapters 4, 5).

Appendix B

BASIC INTERFEROMETER CONFIGURATIONS

Here we will derive the expected gravitational wave signal from a simple Michelson interferometer, alongside the expected laser frequency and intensity noise couplings, as has been done before [78, 114, 211–213].

In other chapters we explored the frequency and intensity stabilization schemes, as well are the full detector response to GWs for accurate calibration. The purpose of this chapter is to introduce the technical origins of concepts important to this thesis, including optical gain, contrast defect, and DC offset.

This appendix helps draw the connection between the theoretical representation and the instrumental results. Often, random scale factors are applied to help data match up with theory. If losses are a parameter we care about, which they are in Advanced LIGO for O3, random scale factors are degenerate with those parameters.

The overview facilitates a direct comparison of analytic results like those in [22, 114], to numerical results like those calculated from Finesse [134, 135], to actual measurements taken at LIGO Hanford.

B.1 Propagation through space

An electric field propagating through free space only affect its phase. For an input electric field E_{in} propagating a length L to E_{out} through a medium with index of refraction n:

$$\vec{E}_{out} = e^{inkL} \vec{E}_{\text{in}} \tag{B.1}$$

where $k = \omega_0/c$ is the wavenumber. We will exclusively be working with $n = 1$ in this thesis.

In the quadrature representation, it can be convenient to represent spatial propagation as a rotation matrix. This is because if we combine two beams with different quadrature bases, we must first rotate one basis to align with the other before combining the beams:

$$\vec{E}_{out} = \begin{bmatrix} \cos(nkL) & -\sin(nkL) \\ \sin(nkL) & \cos(nkL) \end{bmatrix} \vec{E}_{\text{in}}. \tag{B.2}$$

This is used for combining beams at the beamsplitter in Eq. B.25.

B.2 Reflection from an oscillating mirror

The reflection from oscillating mirror imposes phase modulated sidebands on the incident carrier. Figure B.1 illustrates this basic setup, with a perfectly coherent laser incident on a vibrating mirror.

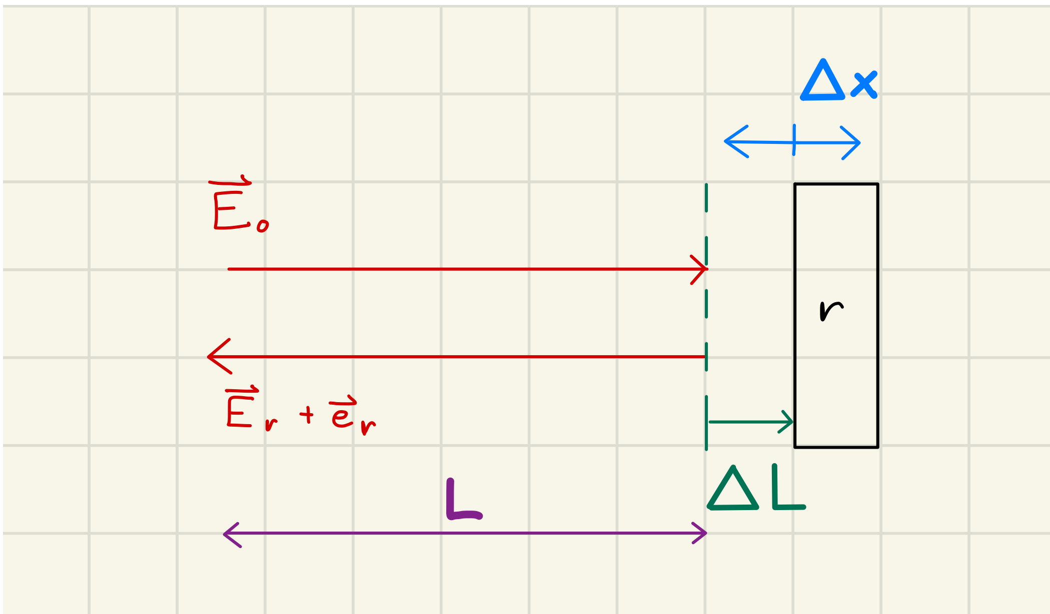

Figure B.1: A noiseless laser $\vec{E}_0$ is incident on a mirror with reflectivity r oscillating at amplitude Δx. The main reflected wave $\vec{E}_r$ carries a phase shift $\phi_D = 2k\Delta L$, while the sideband reflected wave $\vec{e}_r$ holds the information of the oscillation Δx in the phase quadrature.

The incident wave $\vec{E}_0$ travels some distance $L + \Delta L$. L is a distance such that the round-trip phase would be zero:

$$2kL = 2\pi n \tag{B.3}$$

$$L = \frac{\lambda}{2} n \tag{B.4}$$

where $k = 2\pi/\lambda$ is the wavenumber, and n is some integer. ΔL is a static offset less than $\lambda/2$.

The total reflected wave can be broken up as $\vec{E}_r + \vec{e}_r$:

$$\begin{aligned} \vec{E}_r + \vec{e}_r &= E_0 e^{i\omega_0 t} e^{i\phi_D} e^{i2k\Delta x \cos(\omega t)} \\ \vec{E}_r + \vec{e}_r &= E_0 e^{i\omega_0 t} e^{i\phi_D} \left[1 + ik\Delta x e^{i\omega t} + ik\Delta x e^{-i\omega t}\right] \end{aligned} \tag{B.5}$$

or in quadrature representation,

$$\vec{E}_r + \vec{e}_r = E_0 e^{i\omega_0 t} e^{i\phi_D} \left[\begin{pmatrix} 1 \\ 0 \end{pmatrix} + k\Delta x \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right]. \tag{B.6}$$

B.3 Simple Michelson

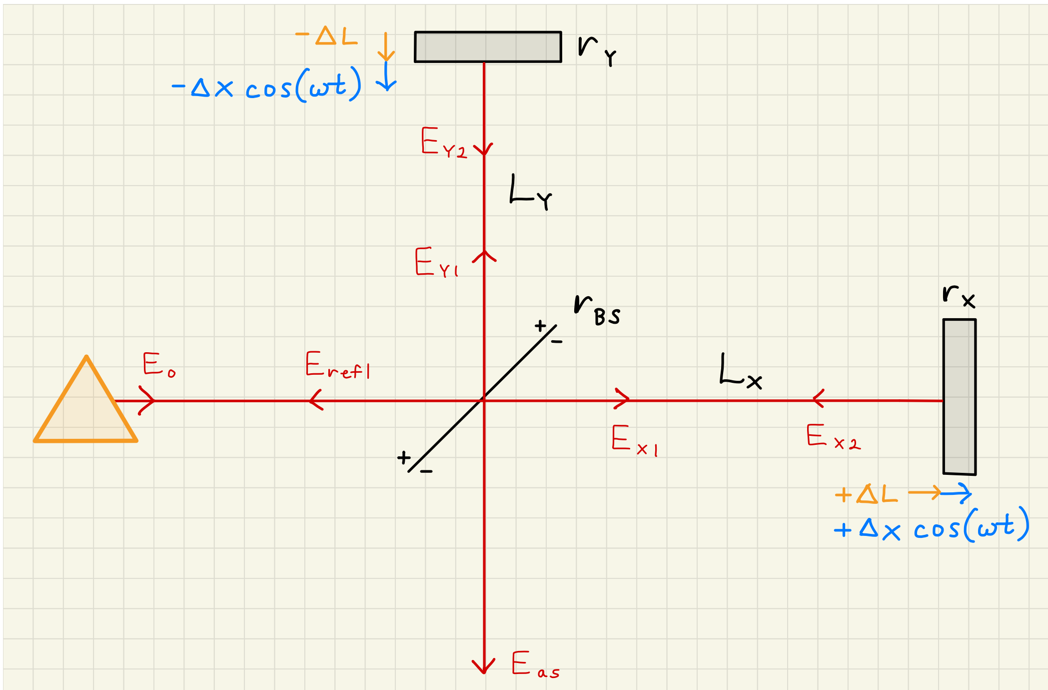

Figure B.2: Diagram of a simple Michelson interferometer with DC readout. A small static offset ΔL allows light to exit the antisymmetric port. A differential audio-band oscillation $\Delta x \cos(\omega t)$ creates a transmitted power oscillation ΔP_{as}.

Figure B.2 shows a simple Michelson with input light $\vec{E}_0$ and light transmitted to the antisymmetric port $\vec{E}_{\text{as}}$. Here we derive the transfer function for light from the input port to the antisymmetric, or transmitted, port. We assume plane waves and thin mirrors with no losses, and use the "+/-" mirror reflection convention as opposed to the "90° transmission" convention $t \rightarrow it$ used in Finesse [134, 215]. For the thin mirrors, we will be working with "amplitude reflectivity" r and "amplitude transmission" t such that $r^2 + t^2 = R + T = 1$, were R and T are the usual reflectivity and transmission of power incident on the mirror.

The input carrier $\vec{E}_0$ is split into the two fields in the arms by the 50:50 beamsplitter, so $r_{bs} = t_{bs} = 1/\sqrt{2}$:

$$\vec{E}_{x1} = \frac{1}{\sqrt{2}} E_0 e^{i\omega_0 t} \tag{B.7}$$

$$\vec{E}_{y1} = \frac{1}{\sqrt{2}} E_0 e^{i\omega_0 t}. \tag{B.8}$$

Fields $\vec{E}_{x1}$ and $\vec{E}_{y1}$ travel the length of the arms and are reflected back to the beamsplitter. Setting the oscillation $\Delta x = 0$ for now, but keeping the differential static offset ΔL yields

$$\vec{E}_{x2} = \frac{r_x}{\sqrt{2}} E_0 e^{i\omega_0 t} e^{i2k(L_x + \Delta L)} \tag{B.9}$$

$$\vec{E}_{y2} = \frac{r_y}{\sqrt{2}} E_0 e^{i\omega_0 t} e^{i2k(L_y - \Delta L)} \tag{B.10}$$

The field "reflected" from the beamsplitter $\vec{E}_{\text{refl}}$ returns to the laser, but the field "transmitted" $\vec{E}_{\text{as}}$ exits the Michelson through the antisymmetric port. The beam from the X-arm gets a sign flip from our reflection convention. Assuming that $L_x = L_y = L$ is some nominal common length, we write the reflected and transmitted fields as

$$\vec{E}_{\text{refl}} = \frac{1}{\sqrt{2}} (E_{y2} + E_{x2}) \tag{B.11}$$

$$\vec{E}_{\text{as}} = \frac{1}{\sqrt{2}} (E_{y2} - E_{x2}) \tag{B.12}$$

$$\vec{E}_{\text{refl}} = \frac{1}{2} E_0 e^{i\omega_0 t} e^{i2kL} \left(r_y e^{-i2k\Delta L} + r_x e^{i2k\Delta L} \right) \tag{B.13}$$

$$\vec{E}_{\text{as}} = \frac{1}{2} E_0 e^{i\omega_0 t} e^{i2kL} \left(r_y e^{-i2k\Delta L} - r_x e^{i2k\Delta L} \right). \tag{B.14}$$

The field $\vec{E}_{\text{as}}$ as written here is correct, but we can write it in a more useful basis by defining

$$r = \frac{r_x + r_y}{2} \qquad \delta r = \frac{r_x - r_y}{2}. \tag{B.15}$$

Here r represents the common, or average, reflectivity of the mirrors, while δr is the *contrast defect* due to reflectivity mismatch.

Using this basis on Eqs. B.13 and B.14 yields

$$\vec{E}_{\text{refl}} = E_0 e^{i\omega_0 t} e^{i2kL} (r \cos\phi_D + i\delta r \sin\phi_D) \tag{B.16}$$

$$\vec{E}_{\text{as}} = -E_0 e^{i\omega_0 t} e^{i2kL} (\delta r \cos\phi_D + ir \sin\phi_D). \tag{B.17}$$

where we have written the static differential phase $\phi_D = 2k\Delta L$. In the quadrature picture,

$$\vec{E}_{\text{refl}} = E_0 e^{i\omega_0 t} e^{i2kL} \begin{pmatrix} r\cos\phi_D \\ \delta r\sin\phi_D \end{pmatrix} \tag{B.18}$$

$$\vec{E}_{\text{as}} = -E_0 e^{i\omega_0 t} e^{i2kL} \begin{pmatrix} \delta r\cos\phi_D \\ r\sin\phi_D \end{pmatrix}. \tag{B.19}$$

Eqs. B.18 and B.19 clearly show what quadrature the carrier light is in. Both experience a phase rotation due to the round-trip length of the arms of $\phi_L = 2kL \mod 2\pi$, which can be ignored in this case as a change of quadrature basis. We set $\phi_L = 0$ for simplicity.

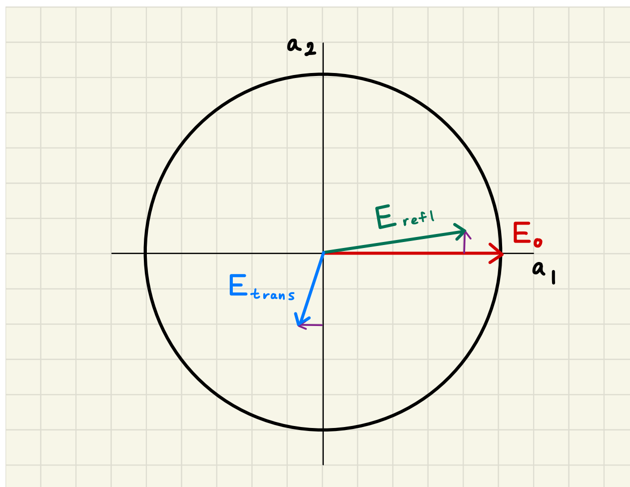

Figure B.3: Simple Michelson phasor diagram for the reflected and transmitted beams from Eqs. B.18 and B.19. The input light $\vec{E}_0$ is along the amplitude quadrature. In an ideal Michelson with no contrast defect and no static offset $\delta r = \Delta L = 0$, $\vec{E}_{\text{refl}} = \vec{E}_0$ and $\vec{E}_{\text{as}} = 0$. However, a small length offset transmits light along the phase quadrature, shortening $\vec{E}_{\text{refl}}$ along the amplitude quadrature. Contrast defect light pushes off the basis vectors, represented by the purple arrows, by transmitting light in the amplitude quadrature and reflecting light in the phase quadrature. Here we have set $r \approx 0.95$ and contrast defect $\delta r \approx 0.05$, with small length offset $\phi_D = 2k\Delta L \approx 0.4$.

For the transmitted light $\vec{E}_{\text{as}}$, the light due to contrast defect δr is in the amplitude quadrature, but the light due to a static length offset ΔL is in the phase quadrature. Figure B.3 illustrates an example phasor diagram with nonzero contrast defect and a static offset.

Calculating the power in the reflected and transmitted fields yields

$$P_{\text{refl}} = |\vec{E}_{\text{refl}}|^2 = P_{\text{in}}\left(r^2\cos(\phi_D)^2 + \delta r^2\sin(\phi_D)^2\right) \tag{B.20}$$

$$P_{\text{as}} = |\vec{E}_{\text{as}}|^2 \;\; = P_{\text{in}}\left(\delta r^2\cos(\phi_D)^2 + r^2\sin(\phi_D)^2\right) \tag{B.21}$$

If we approximate the transmitted power assuming $\Delta L \ll 1$, then $\cos(\phi_D) \to 1$ and $\sin(\phi_D) \to \phi_D$:

$$P_{\text{as}} \approx 4P_{\text{in}}r^2k^2\Delta L^2. \tag{B.22}$$

Directly on the Michelson dark fringe, i.e. when $P_{\text{as}} = 0$, P_{as} is quadratic in ΔL. Figure B.4 plots this approximation alongside the full expressions for reflected and transmitted power.

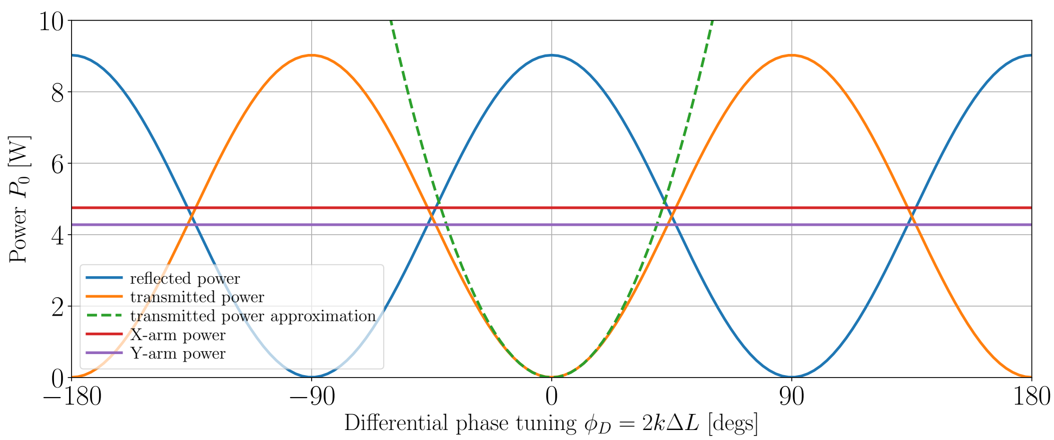

Figure B.4: Simple Michelson differential tuning ϕ_D vs reflected and transmitted power, as derived in Eqs. B.20 and B.21. Here, $P_{\text{in}} = 10\text{W}$, $r = 0.95$, and $\delta r = 0.025$. The power in the arms is unequal because there is significant contrast defect, $r_x = 0.975, r_y = 0.925$. The transmitted power approximation from Eq. B.22 is plotted as the dashed line.

B.3.1 Response to differential motion

Up until now all of the laser light has been at a single carrier frequency ω_0. Now we inject an oscillating differential mirror motion $\pm\Delta x\cos(\omega t)$, which will scatter light from ω_0 to $\omega_0 \pm \omega$, as described in section B.2 Eq. B.6.

Going back to Eqs. B.9 and B.10 and applying the oscillation yields signal sidebands

$\vec{e}_{x2}$ and $\vec{e}_{y2}$:

$$\vec{E}_{x2} + \vec{e}_{x2} = \frac{r_x}{\sqrt{2}} E_0 e^{i\omega_0 t} e^{i\phi_D} \left[\begin{pmatrix} 1 \\ 0 \end{pmatrix} + k\Delta x \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right] \tag{B.23}$$

$$\vec{E}_{y2} + \vec{e}_{y2} = \frac{r_y}{\sqrt{2}} E_0 e^{i\omega_0 t} e^{-i\phi_D} \left[\begin{pmatrix} 1 \\ 0 \end{pmatrix} - k\Delta x \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right]. \tag{B.24}$$

Now $\vec{E}_{\text{as}}$ will also have a signal sideband $\vec{e}_{\text{as}}$. Rotating $\vec{E}_{x2}$ by ϕ_D and $\vec{E}_{y2}$ by $-\phi_D$ using Eq. B.2 so their quadrature bases align, then multiplying yields:

$$\vec{E}_{\text{as}} + \vec{e}_{\text{as}} = E_0 e^{i\omega_0 t} e^{i2kL} \left[\begin{pmatrix} -\delta r \cos\phi_D \\ -r \sin\phi_D \end{pmatrix} + k\Delta x \begin{pmatrix} \delta r \sin\phi_D \\ -r\cos\phi_D \end{pmatrix} \right]. \tag{B.25}$$

Calculating the total power P_{as} yields a static component at DC $P_{trans,0}$ equal to Eq. B.21 and an oscillating component at the signal frequency ω, $P_{\text{as}}(\omega)$. We ignore $(\Delta x)^2$ terms which are extremely small:

$$P_{\text{as}} = |\vec{E}_{\text{as}} + \vec{e}_{\text{as}}|^2 \tag{B.26}$$

$$P_{\text{as}} = P_{trans,0} + P_{trans,\omega}(\omega) \tag{B.27}$$

$$P_{trans,0} = P_{\text{in}}\left(\delta r^2 \cos(\phi_D)^2 + r^2 \sin(\phi_D)^2\right) \tag{B.28}$$

$$P_{\text{as}}(\omega) = P_{\text{in}} k\Delta x (r^2 - \delta r^2) \sin(2\phi_D) \tag{B.29}$$

If we divide the final term Eq. B.29 by the differential mirror motion Δx, we get the *transfer function* from meters to watts for a simple Michelson.

$$\frac{P_{\text{as}}}{\Delta x}(\omega) = P_{\text{in}} k (r^2 - \delta r^2) \sin(2\phi_D) e^{-\frac{i\omega L}{c}}. \tag{B.30}$$

Figure B.5 plots Eq. B.30 for an example 4 km Michelson with different static offsets. We have added a speed-of-light travel delay $\tau = L/c$ for the time it takes the light to travel from the mirror the beamsplitter. This amounts only to a phase shift, and emphasizes that $P_{\text{as}}/\Delta x$ is a complex, frequency-dependent quantity. Subsection B.3.2 derives where this light travel time delay comes from.

Differential arm motion is called *DARM* in Advanced LIGO. The transfer function $P_{\text{as}}/\Delta x$ is the *DARM optical plant*, and in general can be frequency-dependent. The magnitude of the optical plant in units of watts per meters is often called the *optical gain*, although in reality the interferometer is acting as a transducer from mirror motion to power.

We note that high contrast defect δr reduces the optical gain in Eq. B.30. Also, large ΔL such that $\phi_D = \pi/4$ yields the maximum possible optical gain. This is known as *half-fringe* Michelson operation, and while it does increase the signal, it also allows excess laser noise into the antisymmetric port, reducing the SNR overall. We will explore this in the following subsections.

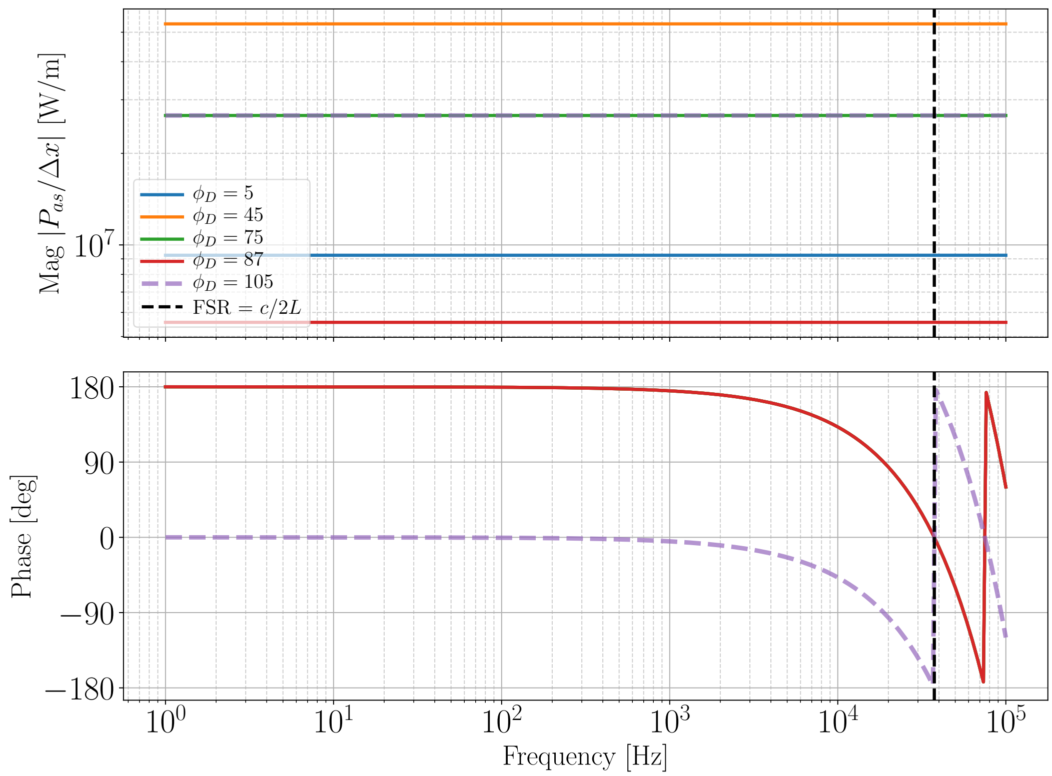

Figure B.5: Simple Michelson DARM transfer functions $P_{\text{as}}/\Delta x$ for several DC offsets ϕ_D, as derived in Eqs. B.30 and plotted via Finesse simulation [134, 135]. The closer to $\phi_D = \pm 90°$, the higher the optical gain of the transfer function. Also, $\phi_D = 105°$ has the same optical gain as $\phi_D = 75°$, but opposite sign. Here, $P_{\text{in}} = 10$ W, $r = 0.95$, $\delta r = 0.025$, and $L = 4$ km.

B.3.2 Gravitational waves to differential motion transfer function

From section B.3.1 we understand how an interferometer responds to DARM motion. However, a gravitational wave modulates the spacetime that the laser is traveling through, not the mirrors' position. The view that the mirrors or any two points in spacetime do not change coordinates when a gravitational wave is incident is a consequence of the *transverse traceless gauge* [221].

Saulson [211] derives the phase shift a perfectly aligned $h_+ = he^{-i\omega t}$ gravitational wave would generate in a simple Michelson by calculating the round-trip time for light in both arms. Recalling with the spacetime interval for light:

$$ds^2 = 0 \tag{B.31}$$

we can break up the spacetime interval into the spatial and temporal components of the light beam traveling in the X-arm, assuming a flat spacetime:

$$0 = -c^2 dt^2 + (1 + he^{-i\omega t})dx^2 \tag{B.32}$$

Moving $c^2 dt^2$ to the LHS, taking the root of both sides, and integrating yields

$$c \int_0^{\tau_{out}} dt = \int_0^L \sqrt{1 + he^{-i\omega t}} dx. \tag{B.33}$$

We've taken the integral over the time it takes for the light wave to propagate out to the X-end mirror τ_{out}, and over the length that the light wave propagates L. Since $h \ll 1$, we can approximate $\sqrt{1+h} \approx 1 + h/2$, substitute $t = x/c$ on the RHS, and solve both integrals for the outbound time

$$\tau_{out} = \frac{L}{c} + h\frac{e^{\frac{-i\omega L}{c}} - 1}{i2\omega} \tag{B.34}$$

For the inbound trip, we substitute in $t = (2L - x)/c$ instead:

$$\tau_{\text{in}} = \frac{L}{c} + h\frac{e^{\frac{-i\omega L}{c}}\left(e^{\frac{-i\omega L}{c}} - 1\right)}{i2\omega} \tag{B.35}$$

The total round-trip time $\tau_{rt} = \tau_{out} + \tau_{\text{in}}$ for the light in the X-arm:

$$\tau_{rt} = \frac{2L}{c} + h\frac{\left(e^{\frac{-i2\omega L}{c}} - 1\right)}{i2\omega} \tag{B.36}$$

Converting this round-trip time to the phase of the X-arm:

$$\phi_x = kc\tau_{rt} \tag{B.37}$$

$$\phi_x = kc\left(\frac{2L}{c} + h\frac{\left(e^{\frac{-i2\omega L}{c}} - 1\right)}{i2\omega}\right) \tag{B.38}$$

The phase of the Y-arm ϕ_y will have nearly the same phase, but $h \to -h$ in the round-trip time from Eq. B.36. Therefore the differential phase $\Delta\phi$ is

$$\Delta\phi = \phi_x - \phi_y \tag{B.39}$$

$$\Delta\phi = h\frac{kc}{i\omega}\left(1 - e^{\frac{-i2\omega L}{c}}\right) \tag{B.40}$$

$$\Delta\phi = 2kLh\,\text{sinc}\left(\frac{\omega L}{c}\right)e^{\frac{-i\omega L}{c}} \tag{B.41}$$

or, in transfer function form,

$$\frac{\Delta\phi}{h}(\omega) = 2kL\,\mathrm{sinc}\left(\frac{\omega L}{c}\right)e^{\frac{-i\omega L}{c}}. \tag{B.42}$$

Finally, we cannot directly measure differential phase $\Delta\phi$, but we can measure antisymmetric power. Relating the transfer functions from Eqs. B.30 and B.42

$$\frac{P_{\mathrm{as}}}{h}(\omega) = \frac{\Delta\phi}{h}(\omega) \times \frac{\Delta x}{\Delta\phi}(\omega) \times \frac{P_{\mathrm{as}}}{\Delta x}(\omega) \tag{B.43}$$

$$\frac{P_{\mathrm{as}}}{h}(\omega) = \left[2kL\,\mathrm{sinc}\left(\frac{\omega L}{c}\right)e^{\frac{-i\omega L}{c}}\right] \times \left[\frac{1}{2k}\right] \times \left[P_{\mathrm{in}}k(r^2 - \delta r^2)\sin(2\phi_D)\right] \tag{B.44}$$

$$\frac{P_{\mathrm{as}}}{h}(\omega) = P_{\mathrm{in}}kL(r^2 - \delta r^2)\sin(2\phi_D)\,\mathrm{sinc}\left(\frac{\omega L}{c}\right)e^{\frac{-i\omega L}{c}} \tag{B.45}$$

Figure B.6 plots the transfer function in Eq. B.45. We note that when the signal frequency $f_{GW} = \omega/2\pi$ equals the *free spectral range* $\mathrm{FSR} = c/2L$, there is a dip to nothing in the response. This is a result of the laser integrating over exactly one period of the GW oscillation, yielding no phase change relative to a Michelson with no GW incident. Beyond the FSR, the response is reduced as the light travel time is expanded and contracted multiple times during one pass through the arm.

Eq. B.45 represents the response of a simple Michelson to a gravitational wave. The largest strain detected by Advanced LIGO is $h \approx 10^{-21}$ at around $f \approx 30$ Hz. For an Advanced LIGO-like simple Michelson operating at the half-fringe, with $\phi_D = \pi/4$, $L = 4$ km, $k = 2\pi/(1064\text{ nm})$, $P = 100$ W, $r = 0.99$, and $\delta r = 0.01$, the simple Michelson has a response $P_{\mathrm{as}}(\omega) \approx 2$ nW. The shot noise limit for the same interferometer is

$$\sqrt{S_P} = \sqrt{2\hbar\omega_0 P_{trans,0}} \tag{B.46}$$

$$\approx 4\frac{\mathrm{nW}}{\sqrt{\mathrm{Hz}}} \tag{B.47}$$

Because the gravitational wave signal only lasts for a few milliseconds, we are going to have to do better to achieve a meaningful detection.

B.3.3 Laser noise

From subsection B.3.2 we understand our detector response to gravitational waves. The signature of gravitational waves on the detected antisymmetric port power P_{as} is called the *signal*.

Noise is anything in the antisymmetric power data which is not signal. There are several sources of noise we could include for our simple Michelson. Shot noise

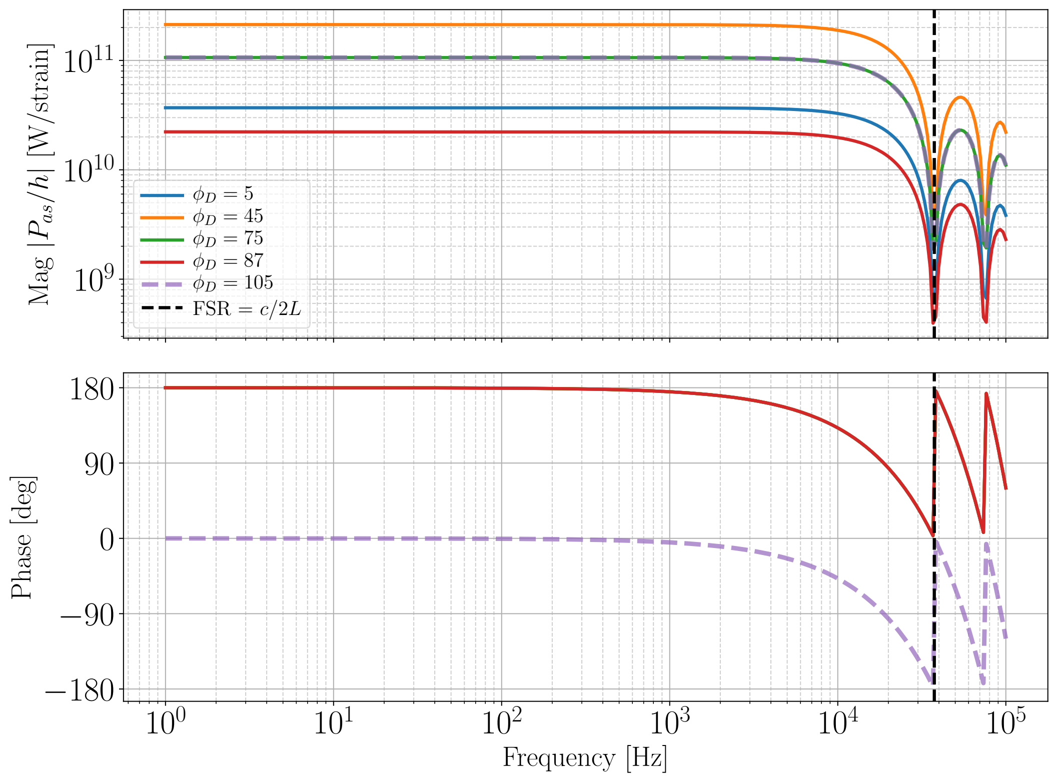

Figure B.6: Simple Michelson strain to antisymmetric power P_{as}/h transfer function for several DC offsets ϕ_D, as derived in Eqs. B.45 and plotted via Finesse simulation. Compared to the DARM transfer function $P_{\text{as}}/\Delta x$ from Eq. B.30 and Figure B.5, P_{as}/h is multiplied by the length of the Michelson L, and the sinc function causing dips at the FSR and a reduced response beyond Here, $P_{\text{in}} = 10$ W, $r = 0.95$, $\delta r = 0.025$, $L = 4$ km, $k = 2\pi/(1064\text{ nm})$.

from the power on the photodetector was explored with Eq. B.46. Seismic noise makes the mirrors move by some $\Delta x_{\text{seismic}}$ will couple to antisymmetric power with the same efficiency as Eq. B.30.

This subsection, and this thesis, will focus on laser noise. This is the natural noise that occurs when a laser emits light at the carrier frequency ω_0. Frequency noise is a measure of how much the carrier frequency moves around its central value, while intensity noise is how much the intensity fluctuates. Appendix A explores how these noises can be expressed as modulated electric fields.

We wish to understand how laser noise couples to the antisymmetric port and obscures the gravitational wave signal. Laser noise is mathematically represented by modulations on the input laser. Recalling Eqs. A.35, we express the input noise

$\vec{e}_0(\omega)$ as

$$\vec{e}_0(\omega) = E_0 e^{i\omega_0 t} e^{i\omega t} \left[\frac{2\pi\delta\nu}{\omega} \begin{pmatrix} 0 \\ 1 \end{pmatrix} + \frac{\delta P}{2P} \begin{pmatrix} 1 \\ 0 \end{pmatrix} \right]. \tag{B.48}$$

We want to estimate the frequency noise transfer function $P_{\text{as}}/\delta\nu$ and the relative intensity noise transfer function $P_{\text{as}}/(\delta P/P_{\text{in}})$.

Following the example in [114], we first define the transfer function from the bright port to the dark port $\mathcal{T}(\omega)$ using Eq. B.17. We also fully express the wavenumber $k = (\omega_0 + \omega)/c$ to capture the suppressed frequency-dependence, write the differential DC offset explicitly in terms of the carrier $\Delta L = c\phi_D/(2\omega_0)$, and divide out the delay factor e^{i2kL} for simplicity:

$$\mathcal{T}(\omega) = \frac{\vec{E}_{\text{as}}}{\vec{E}_0} = -\delta r \cos\left(\frac{\phi_D(\omega + \omega_0)}{\omega_0}\right) - ir \sin\left(\frac{\phi_D(\omega + \omega_0)}{\omega_0}\right) \tag{B.49}$$

Then, using Appendix A Eq. (A9) of [220], the quadrature transfer matrix $\boldsymbol{\mathcal{T}}(\omega)$ is

$$\boldsymbol{\mathcal{T}}(\omega) = \frac{1}{2} \begin{bmatrix} \mathcal{T}(\omega) + \mathcal{T}^*(-\omega) & i(\mathcal{T}(\omega) - \mathcal{T}^*(-\omega)) \\ -i(\mathcal{T}(\omega) - \mathcal{T}^*(-\omega)) & \mathcal{T}(\omega) + \mathcal{T}^*(-\omega) \end{bmatrix} \tag{B.50}$$

Substituting in Eq. B.49:

$$\boldsymbol{\mathcal{T}}(\omega) = \begin{bmatrix} -\cos(\phi_D)\left(\delta r \cos\left(\frac{\phi_D\omega}{\omega_0}\right) + ir \sin\left(\frac{\phi_D\omega}{\omega_0}\right)\right) & \sin(\phi_D)\left(-ir \cos\left(\frac{\phi_D\omega}{\omega_0}\right) + \delta r \sin\left(\frac{\phi_D\omega}{\omega_0}\right)\right) \\ -\sin(\phi_D)\left(-ir \cos\left(\frac{\phi_D\omega}{\omega_0}\right) + \delta r \sin\left(\frac{\phi_D\omega}{\omega_0}\right)\right) & -\cos(\phi_D)\left(\delta r \cos\left(\frac{\phi_D\omega}{\omega_0}\right) + ir \sin\left(\frac{\phi_D\omega}{\omega_0}\right)\right) \end{bmatrix} \tag{B.51}$$

The transfer matrix $\boldsymbol{\mathcal{T}}(\omega)$, when applied to the input electric field sidebands $\vec{e}_0(\omega)$, converts them into the transmitted sidebands $\vec{e}_{\text{as}}(\omega)$ we would expect at the anti-symmetric port:

$$\vec{e}_{\text{as}}(\omega) = \boldsymbol{\mathcal{T}}(\omega)\vec{e}_0(\omega) \tag{B.52}$$

$$\vec{e}_{\text{as}}(\omega) =$$

$$\begin{pmatrix} \frac{2\pi\delta\nu}{\omega} E_0 \sin(\phi_D)\left(r \cos\left(\frac{\omega\phi_D}{\omega_0}\right) + i\delta r \sin\left(\frac{\omega\phi_D}{\omega_0}\right)\right) - \frac{\delta P}{2P_{\text{in}}} E_0 \cos(\phi_D)\left(\delta r \cos\left(\frac{\omega\phi_D}{\omega_0}\right) + ir \sin\left(\frac{\omega\phi_D}{\omega_0}\right)\right) \\ -\frac{\delta P}{2P_{\text{in}}} E_0 \sin(\phi_D)\left(r \cos\left(\frac{\omega\phi_D}{\omega_0}\right) + i\delta r \sin\left(\frac{\omega\phi_D}{\omega_0}\right)\right) - \frac{2\pi\delta\nu}{\omega} E_0 \cos(\phi_D)\left(\delta r \cos\left(\frac{\omega\phi_D}{\omega_0}\right) + ir \sin\left(\frac{\omega\phi_D}{\omega_0}\right)\right) \end{pmatrix} \tag{B.53}$$

We also need the static carrier at the antisymmetric port $\vec{E}_{\text{as}}$ from Eq. B.19. This time we set $k = \omega_0/c$ since we are dealing on with carrier light, and again we cut the phase factors in front for simplicity:

$$\vec{E}_{\text{as}} = -E_0 \begin{pmatrix} \delta r \cos\phi_D \\ r \sin\phi_D \end{pmatrix}. \tag{B.54}$$

Calculating the power fluctuations at the antisymmetric port due to laser noise:

$$P_{\text{as}}(\omega) = 2\vec{E}_{\text{as}} \cdot \vec{e}_{\text{as}}(\omega) \tag{B.55}$$

$$= \frac{\delta P}{P_{\text{in}}} E_0^2 \left(\cos\left(\frac{\omega\phi_D}{\omega_0}\right) \left(r^2 \sin^2(\phi_D) + \delta r^2 \cos^2(\phi_D)\right) + ir\delta r \sin\left(\frac{\omega\phi_D}{\omega_0}\right)\right) \tag{B.56}$$

$$+ i\frac{2\pi\delta\nu}{\omega} E_0^2 (r^2 - \delta r^2) \sin(2\phi_D) \sin\left(\frac{\omega\phi_D}{\omega_0}\right) \tag{B.57}$$

Writing Eq. B.56 in terms of transfer functions from laser frequency and laser intensity to antisymmetric power, and multiplying back in the round-trip phase delay $e^{-i2\omega L/c}$:

$$\frac{P_{\text{as}}}{\delta\nu}(\omega) = i\frac{2\pi E_0^2}{\omega}(r^2 - \delta r^2)\sin(2\phi_D)\sin\left(\frac{\omega\phi_D}{\omega_0}\right) e^{-\frac{i2\omega L}{c}} \tag{B.58}$$

$$\frac{P_{\text{as}}}{\delta P/P_{\text{in}}}(\omega) = E_0^2 \left(\cos\left(\frac{\omega\phi_D}{\omega_0}\right) \left(r^2 \sin^2(\phi_D) + \delta r^2 \cos^2(\phi_D)\right) + ir\delta r \sin\left(\frac{\omega\phi_D}{\omega_0}\right)\right) e^{-\frac{i2\omega L}{c}} \tag{B.59}$$

Using the small angle approximation on Eqs. B.58 and B.59 by letting $\sin(\omega\phi_D/\omega_0) \rightarrow \omega\phi_D/\omega_0$ and $\cos(\omega\phi_D/\omega_0) \rightarrow 1$, and dropping the small imaginary part of the intensity transfer function gives

$$\frac{P_{\text{as}}}{\delta\nu}(\omega) \approx i\frac{2\pi\phi_D P_{\text{in}}}{\omega_0}(r^2 - \delta r^2)\sin(2\phi_D) e^{-\frac{i2\omega L}{c}} \tag{B.60}$$

$$\frac{P_{\text{as}}}{\delta P/P_{\text{in}}}(\omega) \approx P_{\text{in}}\left(r^2 \sin^2(\phi_D) + \delta r^2 \cos^2(\phi_D)\right) e^{-\frac{i2\omega L}{c}} \tag{B.61}$$

Eqs. B.60 and B.61 are the full expressions for frequency and intensity transfer functions to antisymmetric power with a non-negligible differential DC offset ϕ_D. Both are flat in frequency-dependence, except for the delay due to the round-trip of the audio signal. Figures B.7 and B.8 plot the transfer functions for frequency noise and intensity noise appearing in the antisymmetric port.

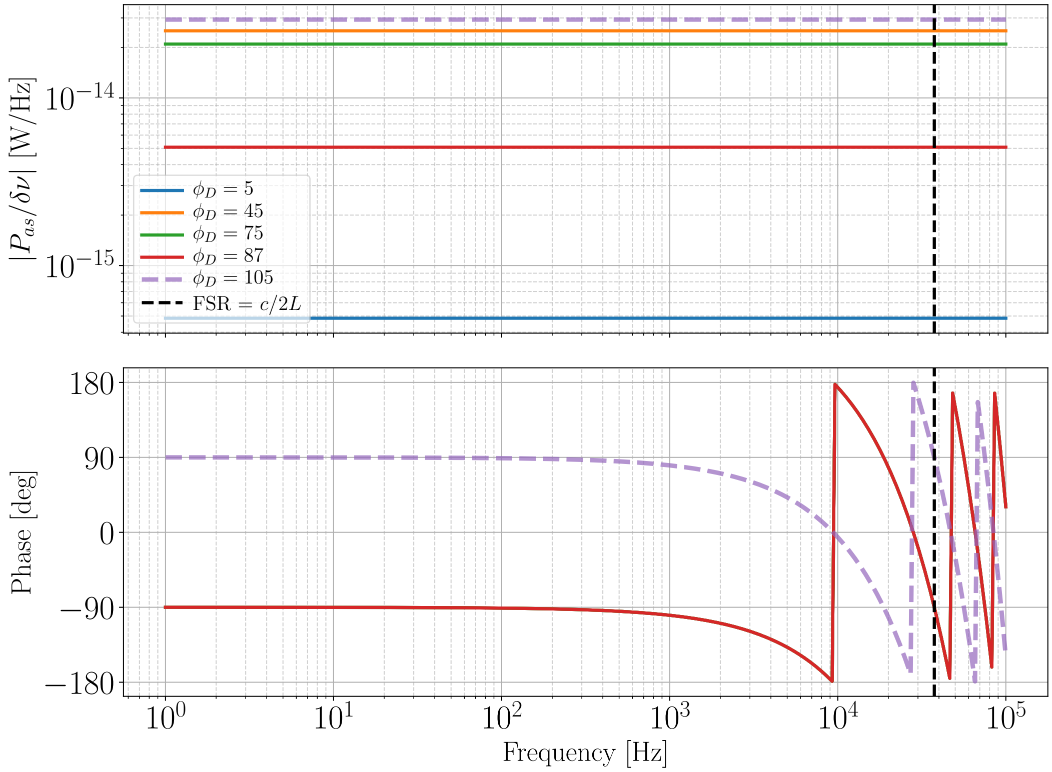

Figure B.7: Simple Michelson frequency noise to antisymmetric power transfer function $P_{\text{as}}/\delta\nu$ for several DC offsets ϕ_D, as derived in Eqs. B.60 and plotted via Finesse simulation. Here, $P_{\text{in}} = 10$ W, $r = 0.95$, $\delta r = 0.025$, $L = 4$ km.

Examining the frequency noise transfer function Eq. B.60, we see primary coupling mechanism is through the static DC offset ϕ_D, the same as the gravitational wave transfer function from Eq. B.42. Frequency noise is especially diabolical because of this fact: excess frequency noise will directly mask a gravitational wave signal because it propagates in the same quadrature as the GW signal.

The primary coupling for the relative intensity transfer function Eq. B.61 depends on the value of ϕ_D. If ϕ_D is not close to zero, then intensity noise couples through the large DC offset power directly. This can be circumvented by limiting the power leaking to the antisymmetric port.

If we let $\phi_D \ll 1$, then Eqs. B.60 and B.61 simplify further:

$$\frac{P_{\text{as}}}{\delta\nu}(\omega) \approx i\frac{4\pi\phi_D^2 P_{\text{in}}}{\omega_0}(r^2 - \delta r^2)e^{-\frac{i2\omega L}{c}} \tag{B.62}$$

$$\frac{P_{\text{as}}}{\delta P/P_{\text{in}}}(\omega) \approx P_{\text{in}}\left(r^2\phi_D^2 + \delta r^2\right)e^{-\frac{i2\omega L}{c}} \tag{B.63}$$

Now both frequency and intensity coupling to the antisymmetric port can be made

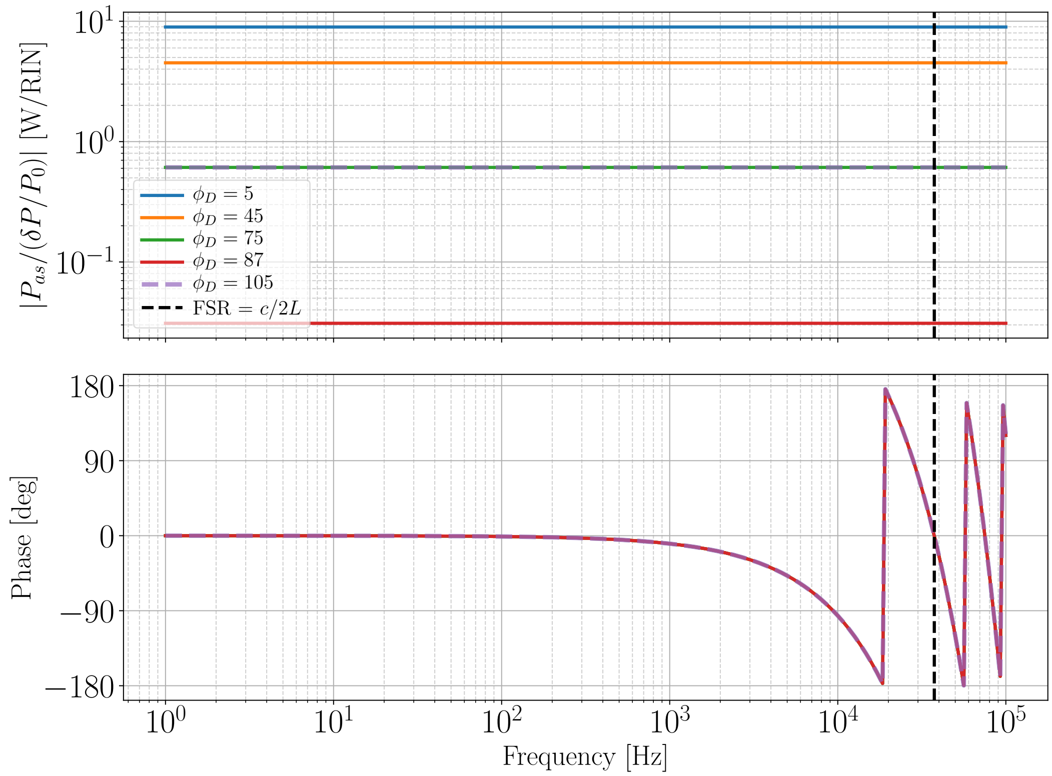

Figure B.8: Simple Michelson relative intensity noise to antisymmetric power transfer function $P_{\text{as}}/(\delta P/P_{\text{in}})$ for several DC offsets ϕ_D, as derived in Eqs. B.61 and plotted via Finesse simulation. Here, $P_{\text{in}} = 10$ W, $r = 0.95$, $\delta r = 0.025$, $L = 4$ km.

small, up to the contrast defect limit from δr^2 for intensity coupling. Using radio-frequency (RF) Pound-Drever-Hall locking, the simple Michelson can be operated exactly on the dark port $\phi_D = 0$, at the expense of added RF sidebands in the dark port.

A useful expression is the *signal-referred transfer function*. The signal-referred TF is the relationship between the transfer function from a gravitational wave to some signal, and the transfer function from some noise source to the same signal. The GW signal-referred frequency noise at the antisymmetric port is calculated:

$$\frac{\frac{P_{\text{as}}}{\delta\nu}(\omega)}{\frac{P_{\text{as}}}{h}(\omega)} = \frac{i\frac{2\pi\phi_D P_{\text{in}}}{\omega_0}(r^2 - \delta r^2)\sin(2\phi_D)e^{-\frac{i2\omega L}{c}}}{P_{\text{in}}kL(r^2 - \delta r^2)\sin(2\phi_D)\,\text{sinc}\left(\frac{\omega L}{c}\right)e^{-\frac{i\omega L}{c}}} \tag{B.64}$$

$$\frac{h}{\delta\nu}(\omega) = \frac{i2\pi c e^{-\frac{i\omega L}{c}}}{L\omega_0^2\,\text{sinc}(\omega L/c)} \tag{B.65}$$

where we have let $k = \omega_0/c$ in the second expression. Eq. B.64 quantifies how frequency noise $\delta\nu$ will look like GW strain h when we measure antisymmetric power P_{as}. Eq. B.64 emphasizes that the only parameters we can change to improve the simple Michelson's resilience to frequency noise are the carrier frequency of the laser ω_0 and the length of the interferometer L.

B.4 Fabry-Perot cavity

In this section we will very briefly overview the Fabry-Perot cavity interferometer, also known as a two-mirror resonator. Then we will review basics of cavity geometry, Gouy phase, and higher order modes [222–224]. These discussion will be relevant for the frequency and intensity noise discussions in Sections 3.4.5 and 3.5.3.

The beam trace math shown here is used in Section C.2 for modeling the arm and SRC cavity beam profiles.

B.4.1 Basics

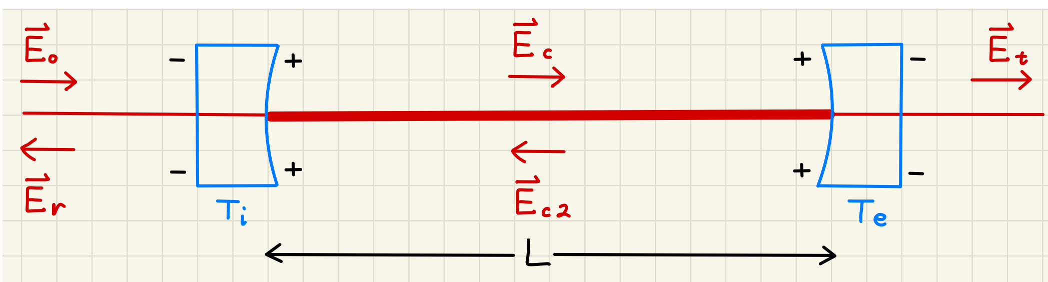

Figure B.9: A laser $\vec{E}_0$ is incident on the input mirror with transmission T_i. Part of the laser is promptly reflected $\vec{E}_r$, and part is transmitted into the cavity $\vec{E}_c$. The cavity field propagates to the end mirror with transmission T_e, where part is transmitted $\vec{E}_t$ and part is reflected in the cavity again $\vec{E}_{c2}$.

The Fabry Perot cavity is formed by placing two mirrors facing one another. Assuming we use lossless thin mirrors and use the plane-wave approximation for electric fields, and we know the input light carrier frequency ω_0, the Fabry-Perot cavity field dynamics are completely solved with only three parameters:

1. the input mirror transmission T_i,
2. the end mirror transmission T_e,
3. the length of the cavity L.

Figure B.9 shows the Fabry Perot cavity parameters and fields.

The Fabry Perot works by storing light of the correct frequency in the cavity for many reflections. The "correct frequency" here means the light is constructively interfering with itself inside the cavity, building up to higher levels than were input. The higher the mirror reflectivities, the more light builds up in the cavity. The buildup of light power is known as the *cavity gain*, or *finesse* of the cavity.

To quantify these properties, we will solve the Fabry-Perot transfer functions. From Figure B.9 we can set up the following system of equations:

$$E_r = -r_i E_0 + t_i E_{c2} \tag{B.66}$$

$$E_c = t_i e^{ikL} E_0 + r_i e^{ikL} E_{c2} \tag{B.67}$$

$$E_{c2} = r_e e^{ikL} E_c \tag{B.68}$$

$$E_t = t_e E_{c2} \tag{B.69}$$

$$\tag{B.70}$$

where e^{ikL} represents the space propagation inside the cavity, $k = \omega_0/c$ is the wave number, and r_i, r_e, t_i, and t_e are the amplitude reflectivity and transmission of the input and end mirrors. Solving for E_c, E_r, E_t yields the transfer functions

$$\frac{E_c}{E_0}(f) = \frac{t_i e^{ikL}}{1 - r_i r_e e^{i2kL}} \tag{B.71}$$

$$\frac{E_r}{E_0}(f) = \frac{-r_i + r_e(r_i^2 + t_i^2)e^{i2kL}}{1 - r_i r_e e^{i2kL}} \tag{B.72}$$

$$\frac{E_t}{E_0}(f) = \frac{t_i t_e e^{ikL}}{1 - r_i r_e e^{i2kL}}. \tag{B.73}$$

The frequency dependence of these transfer functions is hidden in k.

First, we set the resonance condition to be when there is a maximum amount of light in the cavity. This is when the denominator of E_c/E_0 is as small as possible, or equivalently when $e^{i2kL} = 1$. This is achieved when

$$\phi_{rt} = 2kL = \frac{4\pi L}{\lambda} = 2\pi n \tag{B.74}$$

$$L = \frac{\lambda}{2} n \tag{B.75}$$

where n is some huge integer.

The *free spectral range*, or FSR, is the frequency spacing between cavity resonances. This is associated with each consecutive n from Eq. B.74 if we changed λ

instead of L, and can be written

$$L = \frac{c}{2\nu_1} n, \qquad L = \frac{c}{2\nu_2}(n+1) \tag{B.76}$$

$$\mathrm{FSR} = \nu_2 - \nu_1 = \frac{c}{2L}. \tag{B.77}$$

One-way cavities use the more general expression $\mathrm{FSR} = c/L_{rt}$ for round-trip length.

The *cavity amplitude gain* g is the level of enhanced electric field in the resonant cavity, and the *cavity power gain* $G = g^2$ is the enhanced power levels:

$$g = \frac{t_i}{1 - r_i r_e} \tag{B.78}$$

A cavity can be *undercoupled*, *critically-coupled*, or *overcoupled* based on the input and end mirror reflectivities:

$$\text{Undercoupled}: \quad r_i < r_e \tag{B.79}$$

$$\text{Critically coupled}: \quad r_i = r_e \tag{B.80}$$

$$\text{Overcoupled}: \quad r_i > r_e \tag{B.81}$$

A resonant cavity that is critically-coupled will transmit all incident light, as can be seen from Eq. B.73. Most Advanced LIGO cavities are overcoupled: the end mirrors are made as reflective as possible to avoid transmission losses. We note here that for beams reflected from a cavity (Eq. B.72), if the beam is *not* resonant in that cavity, it will experience a $180°$ phase flip. If the beam *is* resonant, it will experience no phase flip.

Now we know the location of the resonances, but the cavity dynamics around resonance determine how sensitive the cavity is to changes in length or frequency. A good proxy for the cavity linewidth is the *cavity pole*. Setting the denominator to zero, and setting the wavenumber $k = 2\pi f_p/c$ and solving for f_p yields:

$$0 = 1 - r_i r_e e^{\frac{i2(2\pi f_p)L}{c}} \tag{B.82}$$

$$\rightarrow i f_p = -\frac{c}{4\pi L}\log(r_i r_e) \tag{B.83}$$

The *finesse* of a cavity is the ratio of the cavity linewidth to its free spectral range. The cavity linewidth is the full-width half-maximum power from a frequency sweep given from the cavity power gain G. In the high-finesse limit, twice the cavity pole is approximately equal to the cavity linewidth:

$$\mathcal{F} = \frac{\mathrm{FSR}}{f_{\mathrm{FWHM}}} \tag{B.84}$$
$$\approx \frac{\mathrm{FSR}}{2 f_p} \tag{B.85}$$
$$= \frac{-\pi}{\log(r_i r_e)} \tag{B.86}$$

This approximation is very good above $\mathcal{F} \approx 10$, which describes all cavities in LIGO.

B.4.2 Cavity geometry

The frequency of a laser incident on a cavity is not the only aspect that governs how well the laser will resonate. The spatial distribution, or *transverse modes*, of the propagating laser is also important.

The *cavity geometry* refers to the Gaussian beam eigenmodes defined by a resonator's parameters. A two-mirror resonator with mirrors of radii of curvature R_1 and R_2, spaced a length L apart, fully defines the eigenmodes of a Fabry-Perot cavity. Incident light in the transverse beam eigenmodes will be resonantly enhanced inside the cavity. Incident light not in the cavity modes will attenuate and be preferentially reflected, and is sometimes referred to as "junk" light

First, we discuss the basics of ABCD beam propagation, as covered in depth by Siegman [223], overviewed by Kogelnik and Li [222], and employed for Advanced LIGO cavities by Arai [42, 225]. Next, we discuss cavity stability and the calculation of transverse eigenmodes of a general cavity. Finally, we discuss Gouy phase, transverse mode spacing, and the capabilities and limitations of a cavity to act as a frequency and spatial mode sieve.

B.4.2.1 Beam spatial eigenmodes

A laser beam's intensity can be decomposed into spatial modes transverse to their direction of propagation. One usual decomposition employed is the Hermite-Gauss (HG) modes, corresponding to a Cartesian coordinate eigenmode representation (see e.g. Kogelnik and Li Section 3.3a [222]). The other is the cylindrical coordinate eigenmodes, Laguerre-Gauss (LG) modes (see Kogelnik and Li Section 3.3b).

The lowest order intensity mode is the Gaussian mode (Chapter 17, Eq. 6 of [223]):

$$I(r,z) = \frac{2P}{\pi w(z)^2} \exp\left(-2\frac{r^2}{w(z)^2}\right) \tag{B.87}$$

where I is the intensity profile in $\mathrm{W/m^2}$, z is the cavity axis defined by the point of maximum beam intensity, $r = \sqrt{x^2 + y^2}$ is the radial distance from the cavity axis z, P is the total power in watts, and $w(z)$ is the beam radius defined at any point along the cavity axis.

In Figures B.10 and C.3, the lowest order TEM_{00} beam waist profiles are represented by the red lines. Quantitatively, the beam waist represents the radial point where the beam intensity falls off by e^{-2}, or equivalently, where the laser amplitude falls by e^{-1}, or where $r = w(z)$ in Eq. B.87. This gives an impression of the changing intensity profile of the beam as it propagates along the cavity axis. A real beam is not perfectly Gaussian, and can be *astigmatic*, meaning the beam has different profiles for the x and y plane intensity projections.

Higher order modes (HOMs) correspond to all other modes, e.g. TEM_{nm} where either $n \neq 0$ or $m \neq 0$. HOMs have a different intensity profile and faster Gouy phase accumulation. This faster phase velocity of HOMs leads to different resonance frequencies for different modes incident on the same cavity (Sections 3.3 and 3.5 [222]). This is the fundamental principle behind mode cleaners, which LIGO employs both in input and output to remove HOMs from the main beam.

B.4.2.2 Beam propagation and ABCD matrices

A laser beam $\vec{w}$ can be defined by its beam radius w and ray angle with respect to the cavity axis θ. A laser beam propagating through a series of mirrors, lenses, and spaces defines a ray transfer matrix, also known as an *ABCD matrix*, which transforms an input beam $\vec{w}_1$ to an output beam $\vec{w}_2$. as shown in Figure B.10:

$$\begin{bmatrix} w_2 \\ \theta_2 \end{bmatrix} = \begin{bmatrix} A & B \\ C & D \end{bmatrix} \begin{bmatrix} w_1 \\ \theta_1 \end{bmatrix}. \tag{B.88}$$

An ABCD matrix must be unitary:

$$AD - BC = 1 \tag{B.89}$$

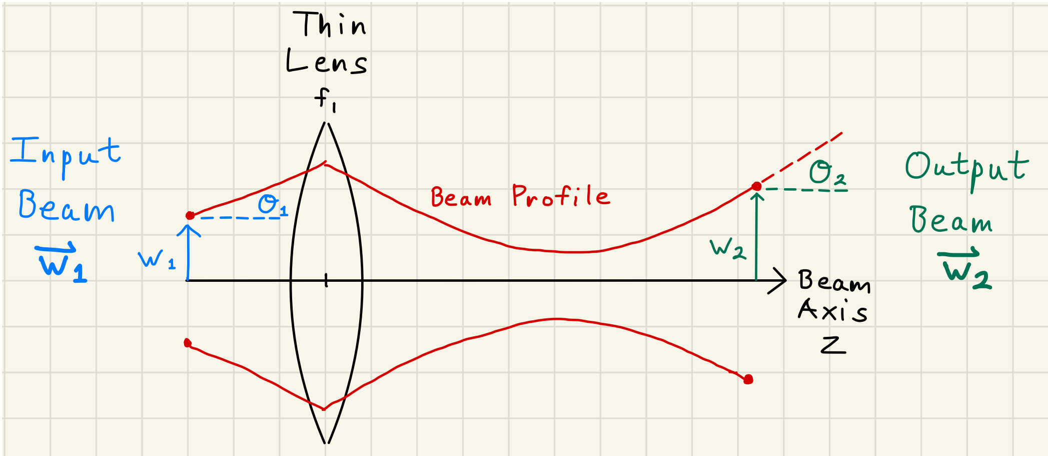

Figure B.10: An example input and output beam profile with a thin lens and space to propagate. The main beam propagates along the cavity axis z, with the point of maximum intensity always directly on the z-axis. The red line represents the beam waist, which is usually defined as the point in the x-y plane with e^{-2} of the maximum intensity.

A final ABCD matrix is made up of several individual matrices multiplied in order. Some component ABCD matrices are

$$\text{Space propagation of length } L : \mathcal{S}(L) = \begin{bmatrix} 1 & L \\ 0 & 1 \end{bmatrix} \tag{B.90}$$

$$\text{Thin lens with focal length } f : \mathcal{L}(f) = \begin{bmatrix} 1 & 0 \\ -\frac{1}{f} & 1 \end{bmatrix} \tag{B.91}$$

$$\text{Spherical mirror with radius of curvature } R : \mathcal{M}(R) = \begin{bmatrix} 1 & 0 \\ -\frac{2}{R} & 1 \end{bmatrix}. \tag{B.92}$$

For propagation through a medium with index of refraction n, replace L with L/n in Eq. B.90.

Returning to the fundamental Gaussian beam profile from Eq. B.87, we define the *complex beam parameter* q in terms of two real parameters, beam radius $w(z)$ and the wavefront radius of curvature $R(z)$:

$$\frac{1}{q} = \frac{1}{R(z)} - i\frac{\lambda}{\pi w(z)^2} \tag{B.93}$$

where λ is the wavelength of the laser. We can see immediately that the real and imaginary components of the inverse complex beam parameter $1/q$ separately represent the radius of curvature and beam radius.

Finally, we write the relationship between the q parameter and ABCD matrices, sometimes called the *ABCD law* (Chapter 20 [223]):

$$q_2 = \frac{Aq_1 + B}{Cq_1 + D} \tag{B.94}$$

where q_2 is the output beam parameter, and q_1 is an input beam parameter.

From Eq. B.94 and Eqs. B.90, B.91, and B.92, we can take any input beam q_1 and a set of spherical optics in its path, and easily model the output beam we expect.

B.4.2.3 Cavity resonance

Now we move on to understanding resonance in a cavity for transverse beam modes. We will focus on the fundamental Gaussian beam to begin.

A cavity can be resonant for any transverse mode if a set of optics allows for the input beam q_1 to equal the output beam q_2 at the same point along the cavity axis z. By combining Eqs. B.89 and B.94, it is possible to show a cavity is stable if its round-trip ABCD matrix obeys

$$-1 < \frac{A + D}{2} < 1 \tag{B.95}$$

$$0 < \frac{A + D + 2}{4} < 1 \tag{B.96}$$

For example, if we calculate a simple two-mirror Fabry-Perot cavity's ABCD matrix with length between mirrors L and mirror radii of curvature R_1, R_2, we get

$$\mathcal{M}_{FP}(L, R_1, R_2) = \mathcal{M}(R_1)\, \mathcal{S}(L)\, \mathcal{M}(R_2)\, \mathcal{S}(L) \tag{B.97}$$

$$= \begin{bmatrix} 1 - \dfrac{2L}{R_2} & 2L\left(1 - \dfrac{L}{R_2}\right) \\ -\dfrac{2}{R_1} - \dfrac{2}{R_2} + \dfrac{4L}{R_1 R_2} & 1 - \dfrac{2L}{R_2} - \dfrac{4L}{R_1} + \dfrac{4L^2}{R_1 R_2} \end{bmatrix} \tag{B.98}$$

then for the Fabry-Perot to be geometrically stable we must satisfy Eq. B.96:

$$0 < \left(1 - \frac{L}{R_1}\right)\left(1 - \frac{L}{R_2}\right) < 1 \tag{B.99}$$

$$0 < g_1 g_2 < 1. \tag{B.100}$$

Here in Eq. B.100 we have introduced the g-factor: $g = 1 - L/R$. The g-factor is a convenient parameter for characterizing the stability of a Fabry-Perot.

From Eq. B.100, both g-factors must have the same sign for a stable cavity. If both g-factors are positive, then $R > L$ and the cavity is said to be a *planar* resonator. If both g-factors are negative, then $R < L$ and the cavity is said to be a *concentric* resonator. If one mirror is flat, then $R_1 \to \infty$ and $g_1 = 1$ implies that $0 < g_2 < 1$, and the cavity is *half-symmetric*. If one or both mirrors have $R = L$, then $g = 0$ and the cavity is *confocal*. If one mirror is convex, then $R_1 < 0$ and $g_1 > 1$ implies that $0 < g_2 < 1/g_1$, and the cavity is *convex-concave*. Figure B.11 plots example beam waists of each type of cavity.

B.4.2.4 Gouy phase

This section follows Siegman Chapter 19.3 [223] and Arai [225].

Gouy phase is the additional phase $\varphi(z)$ acquired by a Gaussian laser beam, as opposed to a plane wave, traveling along its cavity axis from some reference location z_0 to z:

$$\varphi(z) - \varphi(z_0) = \arctan\left(\frac{\lambda z}{\pi w_0^2}\right) \tag{B.101}$$

$$\varphi(z) - \varphi(z_0) = \arctan\left(\frac{z}{z_R}\right) \tag{B.102}$$

where z is the cavity axis, z_0 is a reference beam location on the cavity axis, λ is the laser wavelength, w_0 is the beam waist, i.e. the minimum beam radius, and $z_R = \pi w_0^2/\lambda$ is the Rayleigh range. We will suppress z_0 in remaining equations, but fundamentally the Gouy phase is the difference in accumulated phase between two locations by the Gaussian beam.

Higher order modes acquire Gouy phase faster than the fundamental. For Hermite-Gauss transverse modes TEM_{nm}:

$$\varphi(z|n,m) = (n+m+1)\arctan\left(\frac{\lambda z}{\pi w_0^2}\right) \tag{B.103}$$

The value defined in Eq. B.101 is sometimes called the *Gouy phase shift*, or *accumulated Gouy phase*, and represents the difference in phase accumulated between the fundamental mode TEM_{00} and the first order modes TEM_{01} or TEM_{10}.

The round-trip Gouy phase accumulation inside a cavity of length L can be written in multiple ways. Siegman and Kogelnik and Li write it in terms of the g-factors:

$$\varphi(2L) = 2\arccos(\pm\sqrt{g_1 g_2}) \tag{B.104}$$

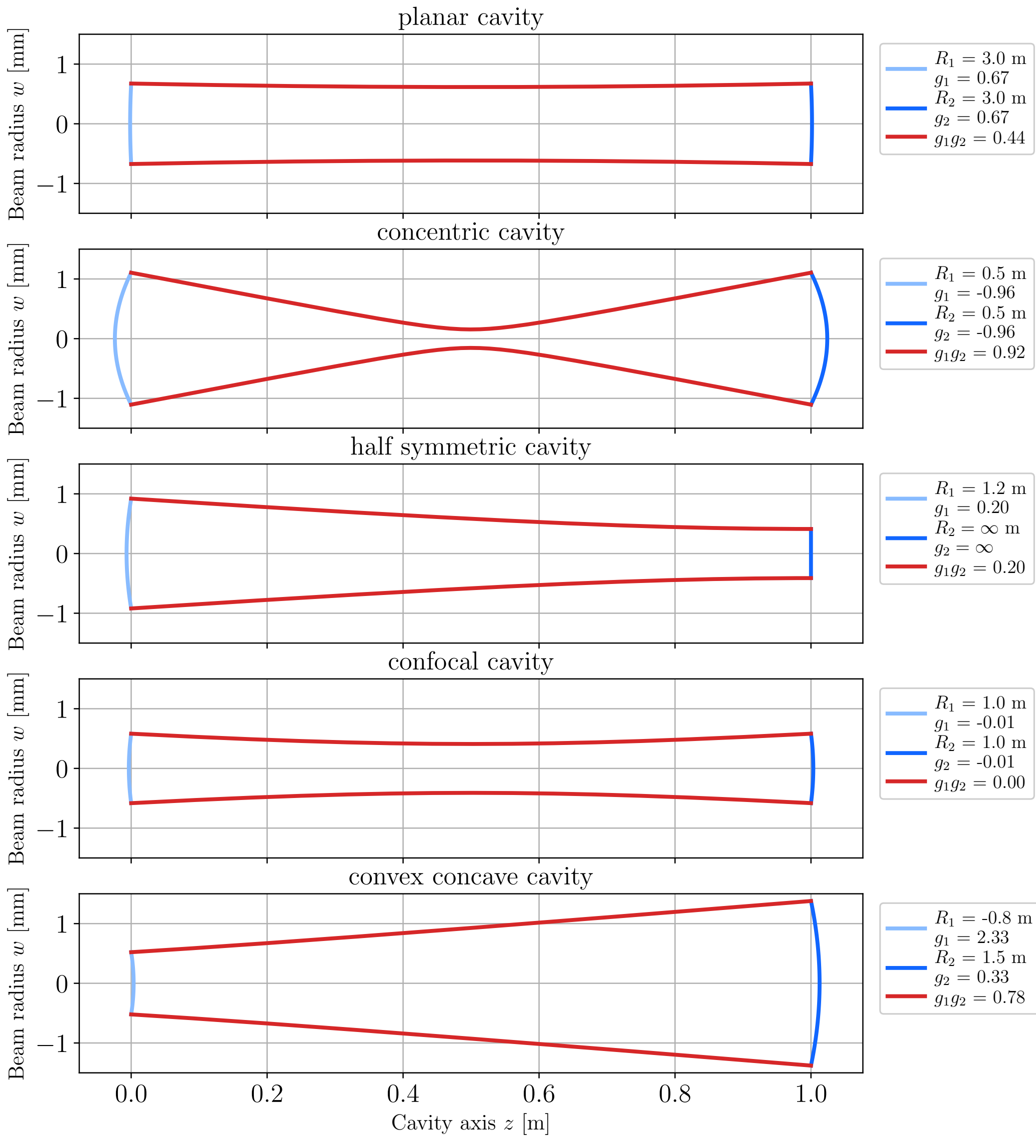

Figure B.11: Beam profiles exemplary of some possible resonant cavities. The confocal cavity has the smallest possible average beam over the entire cavity length, including at the mirrors. The concentric cavity has a sharply focused beam, and a very small beam waist. This plot was made using the python3-based `beamtrace` library.

where the plus sign applies for positive g-factors $g_1 > 0, g_2 > 0$, and the minus sign for negative g-factors. The $z = 2L$ argument to the Gouy phase emphasizes the round-trip, but for a resonant cavity $\varphi(2L) = 2\varphi(L)$.

Arai expresses the cavity round-trip Gouy phase accumulation more generally, in terms of the ABCD matrix:

$$\varphi(2L) = \operatorname{sign}(B)\arccos\left(\frac{A+D}{2}\right) \tag{B.105}$$

$$= 2\arccos\left(\operatorname{sign}(B)\sqrt{\frac{A+D+2}{4}}\right) \tag{B.106}$$

Equation B.106 expression works for any resonating series of mirrors and lens.

Figure B.12 plots the accumulated Gouy phase for a single pass through a resonant cavity.

B.4.2.5 Transverse mode spacing

For a resonant cavity, the total round-trip phase shift must be equal to 2π times some (large) number q. We have the usual round-trip plane wave phase contribution ϕ from the laser frequency ν and the length of the cavity L, $\phi = 2kL$. But we also must remove the additional round-trip accumulated Gouy phase $2\varphi(L|n,m)$ for our TEM_{nm} mode.

From these factors we can calculate the *transverse mode spacing* ν_{TMS} of the cavity. Recall that the free spectral range $\nu_{\mathrm{FSR}} = c/(2L)$. The transverse mode spacing is the frequency spacing between higher order modes resonant in the cavity:

$$2\pi q = \frac{4\pi L\nu}{c} - (n+m+1)\varphi(2L) \tag{B.107}$$

$$= \frac{2\pi\nu}{\nu_{\mathrm{FSR}}} - (n+m+1)\varphi(2L) \tag{B.108}$$

$$\nu_{nm} = \nu_{\mathrm{FSR}}\left[q + (n+m+1)\frac{\varphi(2L)}{2\pi}\right] \tag{B.109}$$

$$\Rightarrow \nu_{\mathrm{TMS}} = \nu_{\mathrm{FSR}}\frac{\varphi(2L)}{2\pi}. \tag{B.110}$$

q is some (large) integer corresponding to the number of wavelengths a plane wave would take to make a round-trip in the cavity. ν_{nm} are the frequencies at which the TEM_{nm} mode resonates.

Figure B.13 illustrates the resonant frequencies for some example transverse modes. One key observation from Eq. B.110 is the spacing of the transverse mode resonances is determined by their round-trip accumulated Gouy phase. This means its possible to control what modes are resonant, or near resonance, by managing the cavity geometry.

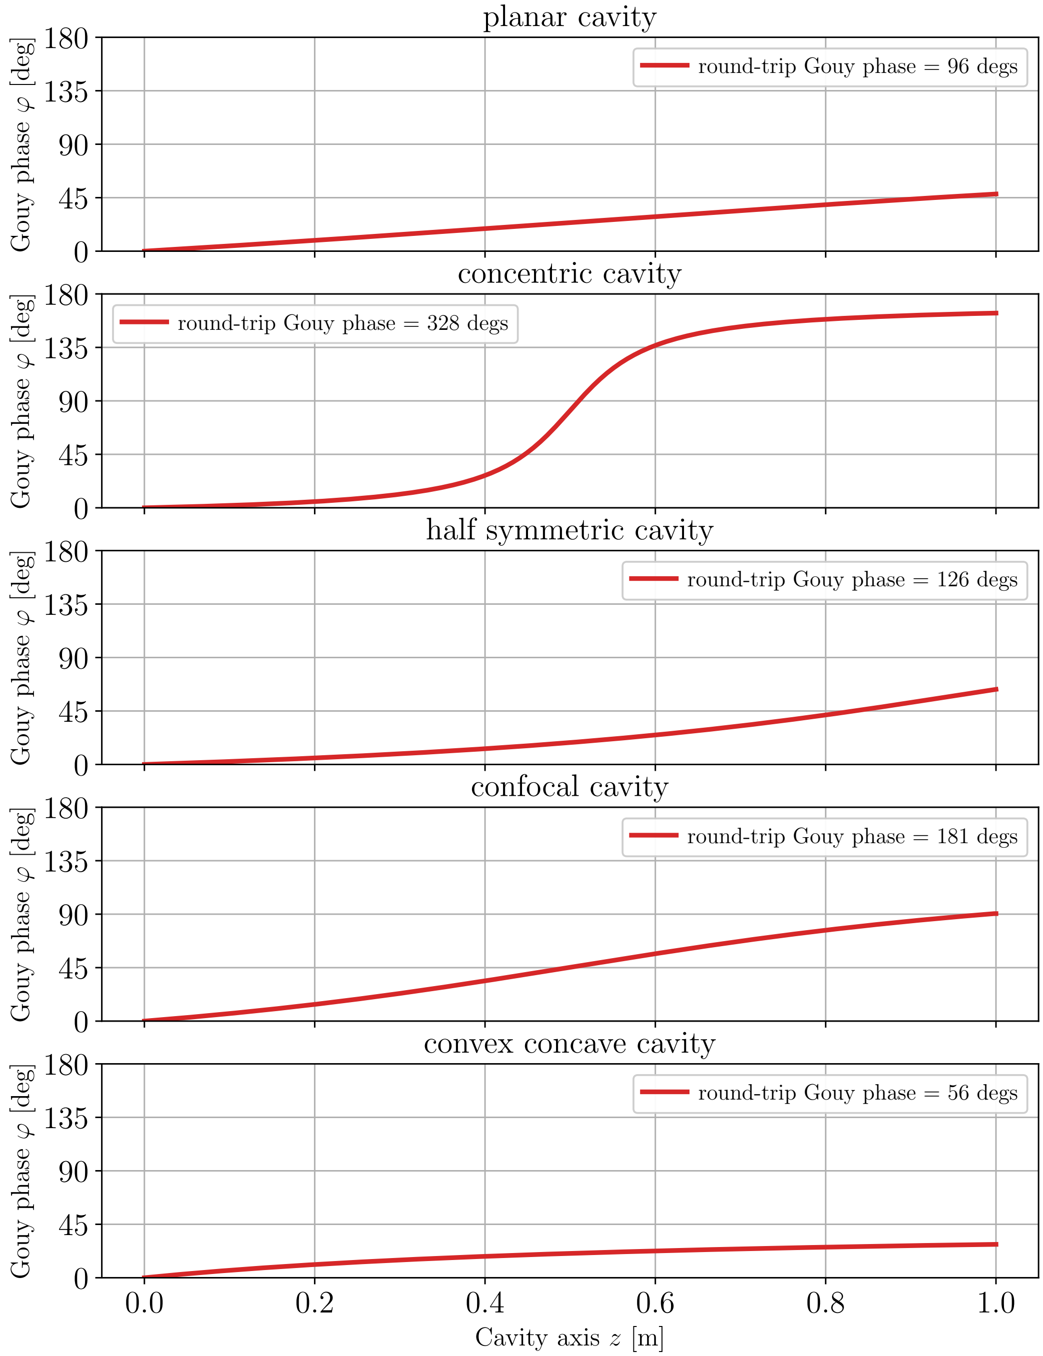

Figure B.12: Single-pass accumulated Gouy phase (half of Eq. B.106) for some possible resonant cavities. The confocal cavity has a single-pass Gouy phase equal to exactly 90 degrees, which will be important for its transverse mode spacing. The concentric cavity quickly accumulates Gouy phase when passing through its shape beam waist. This plot was made using the python3-based `beamtrace` library.

In Advanced LIGO, mode cleaner cavity geometry is managed such that the early higher order modes ($n+m < 10$) do not come close to resonance with the TEM_{00}. However, uncertainty in the cavity length and mirror radii of curvature due to thermal effects must be considered so small perturbations do not hurt the cavity performance.

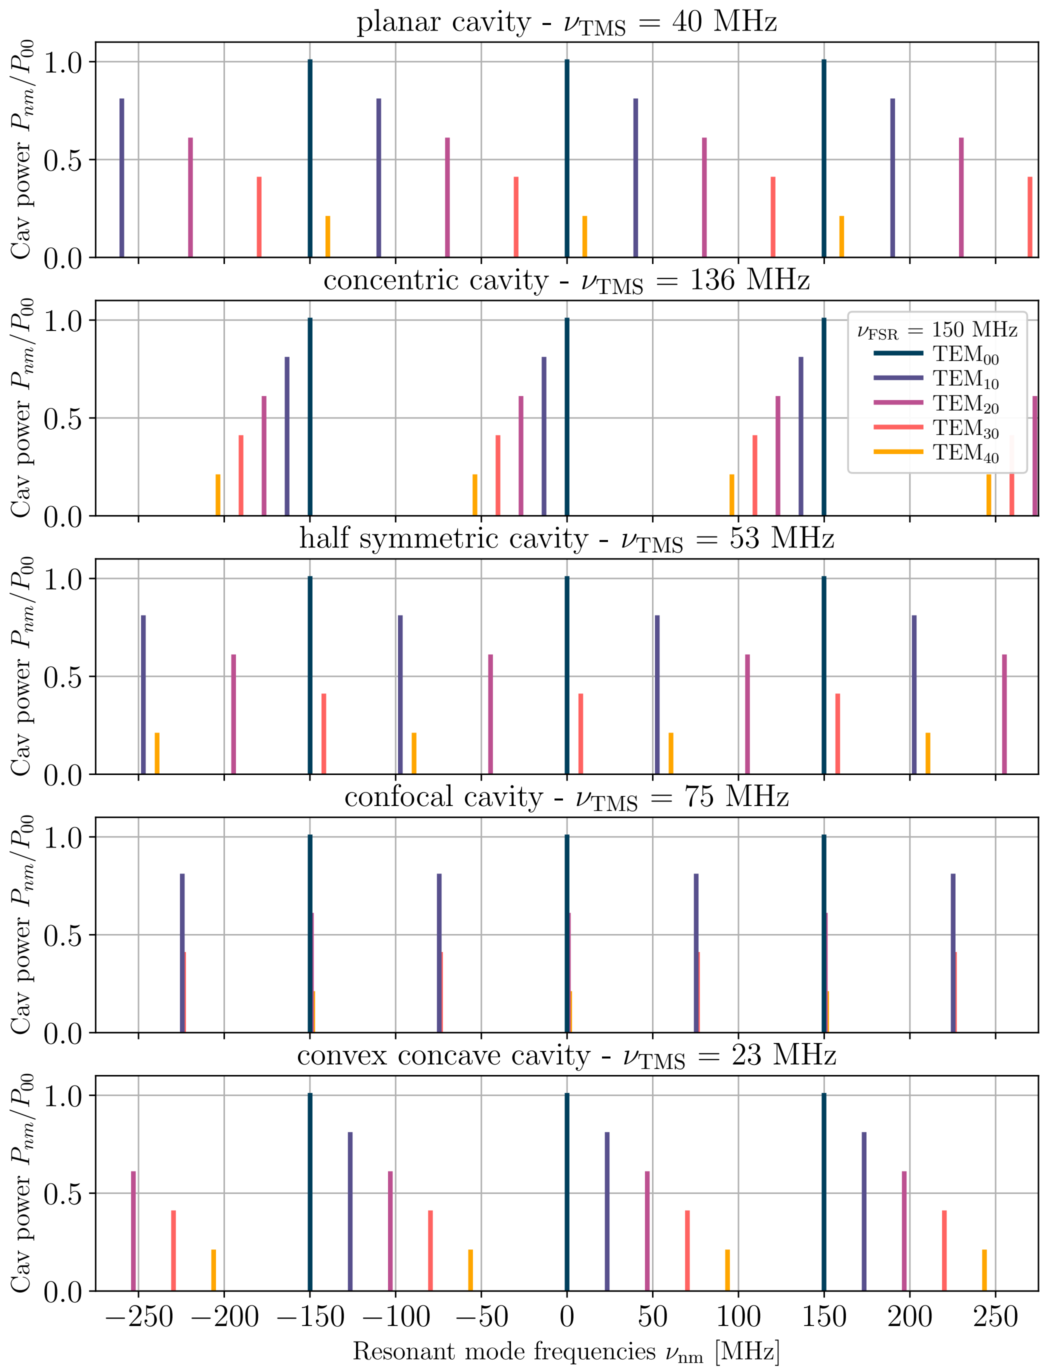

Figure B.13: Resonant frequencies (Eq. B.109) for some possible resonant cavities, around three free spectral ranges for TEM_{00}. The free spectral range for each one meter cavity $\nu_{\mathrm{FSR}} = 150$ MHz. The transverse mode spacing ν_{TMS} from Eq. B.110 is shown in the title of each plot. The convex concave cavity and the planar cavity have small transverse mode spacings, so the main TEM_{00} mode and TEM_{10} mode resonate at nearby frequencies. The confocal cavity has perfect mode overlap for even modes, due to its 180 degree round-trip Gouy phase. This means TEM_{00}, TEM_{20}, $\mathrm{TEM}_{40}, \ldots$, all resonate at the same frequency in a confocal cavity. The concentric cavity has a high transverse modes spacing relative to the free spectral range, so the q TEM_{10} resonant mode is near to the $q+1$ TEM_{00} mode.

Appendix C

TOPICS IN ADVANCED INTERFEROMETRY

This appendix will overview assorted small studies done prior to O3. They are collected here because they represent relatively small, technically detailed techniques required for advanced interferometry.

C.1 Calibration of radio-frequency photodetector response

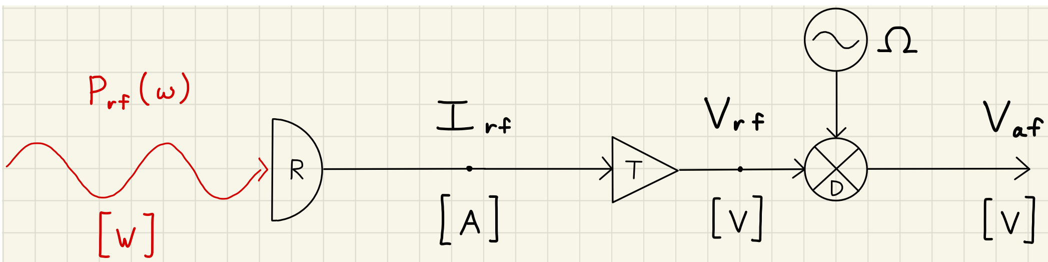

Figure C.1: Radio-frequency photodetector signal chain.

Understanding the signal-chain response of the radio-frequency photodetectors (RFPDs) is important to an accurate calibration of the length error signals in Advanced LIGO. The CARM, PRCL, SRCL, and MICH lengths all rely on Pound-Drever-Hall (PDH) locking, which detects the radio-frequency beatnote between the carrier and RF sidebands. Good overviews of PDH locking are found in [226] and [57]. The results here were used in the CARM calibration in Section 3.4.3.

C.1.1 RF signal chain diagram

Figure C.1 shows a simplified schematic of the radio-frequency sensing chain electronics. The full schematic is shown in [227].

The incoming light power has both an RF P_{rf} and DC P_{dc} component. The RFPD senses and records both in two separate signal paths.

The power is converted to current by the photodiode, in this case wide-area InGaAs diodes [111]. The responsivity R of the photodiode to light is

$$R = \frac{e\lambda}{hc} \tag{C.1}$$

where e is the electron charge, λ is the laser wavelength, and h is Planck's constant. This assumes that each incident photon with energy hc/λ excites a single electron in the biased semiconducting surface of the photodiode. In reality, not all power incident on the photodiode is perfectly absorbed. This is captured by the quantum efficiency η, so that the DC current I_{dc} produced by the incident power P_{dc} is

$$I_{dc} = \eta \frac{e\lambda}{hc} P_{dc} \tag{C.2}$$

Next, the RF current is converted to voltage early by a transimpedance circuit. The gain of this circuit is referred to as a the transimpedance T.

Finally, the output is sent to an demodulation electronics board to bring the voltage signal down from RF to audio frequency [228]. Usually demodulation has a factor of $1/2$ associated with it as half of the signal goes to DC and the other half goes to twice the RF frequency 2Ω. However, the demod board has other gain factors associated with it, yielding a total demod gain of $D = 5.4$.

This yields the final audio frequency voltage V_{af} carrying the signal at radio frequency Ω. The signal chain be written as the product of all of the above components.

C.1.2 RFPD shot noise

Shot noise can be used to measure the sensing chain. First, we must think about the how the shot noise will appear on for the RFPD [229]. Shot noise is white over the full bandwidth of the detector. For a DC current I_{dc}, the shot noise power spectral density is

$$S_I(\omega) = 2eI_{dc} \left[\frac{\mathrm{A}^2}{\mathrm{Hz}}\right] \tag{C.3}$$

If the current is dominated by RF sidebands, then we have *cyclostationary shot noise* [230–232], where shot noise is increased in one quadrature readout due to the cyclic nature of the RF power on the photodetector. We will assume that the power on the RFPDs are roughly constant, i.e. that DC carrier "junk light" dominates the light incident on the detector.

The same level of shot noise from Eq. C.3 appears at both audio sidebands on the RF sideband, i.e $S_I \Omega - \omega)$ and $S_I(\Omega + \omega)$. When we demodulate at Ω, we "fold" the noise from negative frequency $-\omega$ on top of the noise at ω, leading to an additional factor of two of shot noise:

$$S_I^{\Omega}(\omega) = 4eI_{dc} \left[\frac{\mathrm{A}^2}{\mathrm{Hz}}\right] \tag{C.4}$$

Now, propagating the shot noise ASD $\sqrt{S_V^{shot,\Omega}(\omega)}$ through the sensing chain block diagram in Figure C.1 to the output audio frequency, and then substituting in DC power from Eq. C.2 yields

$$\sqrt{S_V^{\Omega}(\omega)} = 2TD\sqrt{eI_{dc}} \tag{C.5}$$

$$= 2TD\sqrt{\eta\frac{e^2\lambda}{hc}P_{dc}}\left[\frac{\mathrm{V}}{\sqrt{\mathrm{Hz}}}\right] \tag{C.6}$$

C.1.3 RFPD dark noise

The electronics that make up the photodetector and demod board have fundamental thermal noise associated their components as well, like Johnson noise. This is known as "dark noise", because it is noise intrinsic to the photodetector even with no like on the photodiode.

In our case, the most likely source of noise is the current induced on the photodiode due to thermal noise. This noise will also be flat in frequency, but will not depend on the incident power P_{dc}, only the temperature and electronics of the photodiode.

Because the dark noise is always there, it is equivalent to some constant level of light always incident on the photodiode producing shot noise. We define this light level as P_{dark}, and add it to our RFPD shot noise equation:

$$\sqrt{S_V^{\Omega}(\omega)} = 2TD\sqrt{\eta\frac{e^2\lambda}{hc}(P_{dc} + P_{dark})}\left[\frac{\mathrm{V}}{\sqrt{\mathrm{Hz}}}\right] \tag{C.7}$$

C.1.4 Shot noise calibration

We can use shot noise from Eq. C.7 to calibrate the RFPD sensing chain. First, we measure the dark noise voltage output of an RFPD by shutting off the laser. Then, we increase the laser power on the RFPD until the noise starts increasing, and measure the flat ASD at several light levels P_{dc}. Finally, we fit Eq. C.7 to our measured ASD levels.

For the fit, we assume we know the demod gain $D = 5.4$, and collect all unknowns in the effective transimpedance T. Also fit is the P_{dark} parameter, which is entirely determined by the dark noise measurement.

Figure C.2 shows the shot noise calibration results for the 9 MHz REFL A and B PDs, and the 24 MHz IMC REFL PD. These numbers were important for the calibration of CARM, see Section 3.4.3.

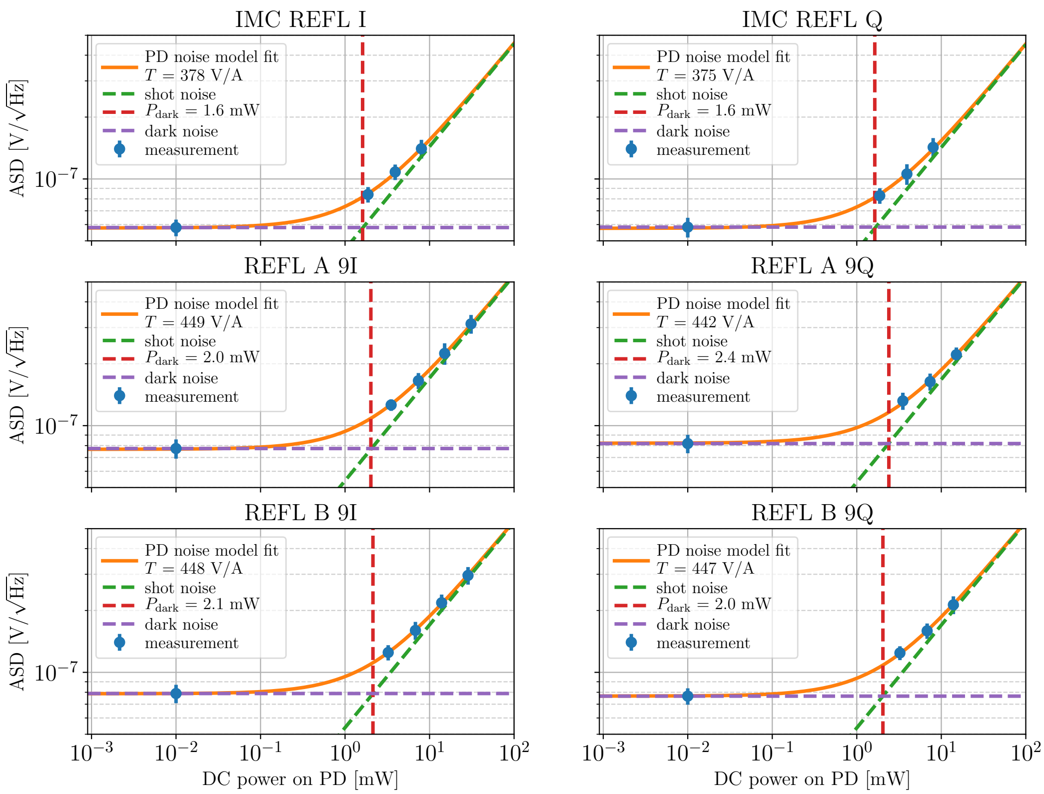

Figure C.2: Measured Hanford RFPD responses for dark noise and shot noise [112]. The dark noise equivalent power P_{dark} is the amount of power on the PD required such that shot noise and dark noise contributions are equal.

IMC REFL previously had only $1\ \mathrm{mW}$ of light on it in full lock, which was less than the equivalent dark noise power, meaning the IMC sensing noise was dominated by dark noise. This was increased to around $9\ \mathrm{mW}$, putting IMC REFL squarely in the shot-noise limited regime.

C.2 DARM cavity beam profile

In O3, the DARM optical plant exhibited an unexpectedly low DARM cavity pole (411 Hz) compared to design (457 Hz), see Figure 3.24. In O1, the SRC exhibited some higher order mode-hopping when locking the dual-recycled Michelson interferometer [121]. Also in O1, SRCL detuning was first measured via the DARM optical spring, and has been a problem in all runs since [5, 57, 94].

A low DARM pole is indicative of excessive SRC losses. Excess SRC losses could be caused by mode mismatch between the arms cavity beams and the SRC beam.

The DARM cavity is made up of two coupled cavities, the 4 km arm cavity and the 56 m signal recycling cavity. As the arm cavities achieve high power, thermal effects cause the cavity geometry to change. In a case with high point absorbers on the ITMs, the cavity geometry can change significantly.

Below are the modeling efforts showing the high sensitivity of the SRC Gouy phase accumulation to the input beam parameter. It quantifies how the SRC is particularly vulnerable to position and mirror radius of curvature changes.

A python3 based beam profile library, `beamtrace`, was used for these simulations. A general overview of beam profile tracing is provided in Section B.4.2.

C.2.1 Arm cavity geometric parameters

The design of the Advanced LIGO arm cavity geometry is a balancing act of several considerations.

The biggest concern for the arm cavity geometry design is the mirror size. LIGO mirrors are 34 cm in diameter, so the beam radius at the mirror is set to be 6 cm to avoid excessive scatter losses. The beam must be large to mitigate excessively high power density and absorption, and to minimize thermal noise as well. However, a large beam radius w also dramatically increases the hard mode radiation-pressure based torque on the test masses [233].

The second concern for arm cavity geometry design is angular controls. In initial LIGO, the arm cavity geometry was planar, with $R_{\mathrm{ITM}} = 7.4$ km and $R_{\mathrm{ETM}} = 14.6$ km. In Advanced LIGO, the arm cavity geometry is concentric, with $R_{\mathrm{ITM}} = 1.940$ km, and the ETM $R_{\mathrm{ETM}} = 2.247$ km. This change was largely due to help handle the higher expected radiation-pressure based torque from the higher circulating arm power. Although an statically unstable torsional mode is always present in the interferometer at high resonating power, there is a choice of which mode is unstable via the cavity geometry. For Advanced LIGO, the hard mode degrees of freedom were chosen to be stable due to its much larger torsional constant κ, while the soft mode degrees of freedom were chosen to be unstable [233].

When a mirror in a concentric resonator is misaligned, the cavity axis will mostly tilt, i.e. displace in the hard mode. When a mirror in a planar resonator is misaligned, the cavity axis will most shift, i.e. displace in the soft mode. Because soft modes are statically unstable for high arm power, it is preferable to have concentric cavities where misalignment causes self-restoring hard mode motion.

Table C.1 shows the arm cavity geometric parameters. Figure C.3 plots the expected fundamental Gaussian beam profile inside the arm cavity.

Table C.1: Table of the Advanced LIGO arm cavity geometric parameters. The arm cavity beam profile is plotted in Figure C.3.

Parameters	Variable	Value	Units
Arm length	L	3.9944	km
ITM radius of curvature	R_{ITM}	1.940	km
ETM radius of curvature	R_{ETM}	2.247	km
ITM g-factor	g_1	-1.06	-
ETM g-factor	g_2	-0.78	-
Cavity g-factor	$g_1 g_2$	0.82	-
Beam waist location	z_0	1.833	km
Beam waist	w_0	1.22	cm
ITM beam radius	w_1	5.25	cm
ETM beam radius	w_2	6.12	cm
Round-trip Gouy phase	φ	310	deg
Free spectral range	ν_{FSR}	37.5	kHz
Transverse mode spacing	ν_{TMS}	32.3	kHz
ABCD matrix	$\begin{bmatrix} A & B \\ C & D \end{bmatrix}$	$\begin{bmatrix} -2.56 & -6.21 \times 10^3 \\ 1.74 \times 10^{-3} & 3.85 \end{bmatrix}$	-

C.2.2 Signal recycling cavity geometric parameters

The signal recycling cavity (SRC) is a folded convex-concave cavity, with three convex mirrors (ITMY, SR2, SRM), and one strongly concave mirror (SR3).

The SRC is folded to bring the large arm cavity beam down to a manageable size (2 mm) at the antisymmetric port. The designed SRC Gouy phase accumulation is 19 degrees [234]. The model defined here gives a 18 degree Gouy phase accumulation. Measurements at LIGO Hanford suggest a much higher SRC Gouy phase, 25 degrees [131, 132].

The SRC is designed to accept the beam from the arm cavities. If the cavity Gouy phase is badly off, its possible the arm to SRC mode mismatch was high.

The Gouy phase discrepancy between measurement and model is likely due to small perturbations in the as-built cavity geometry versus the design. Figure C.4

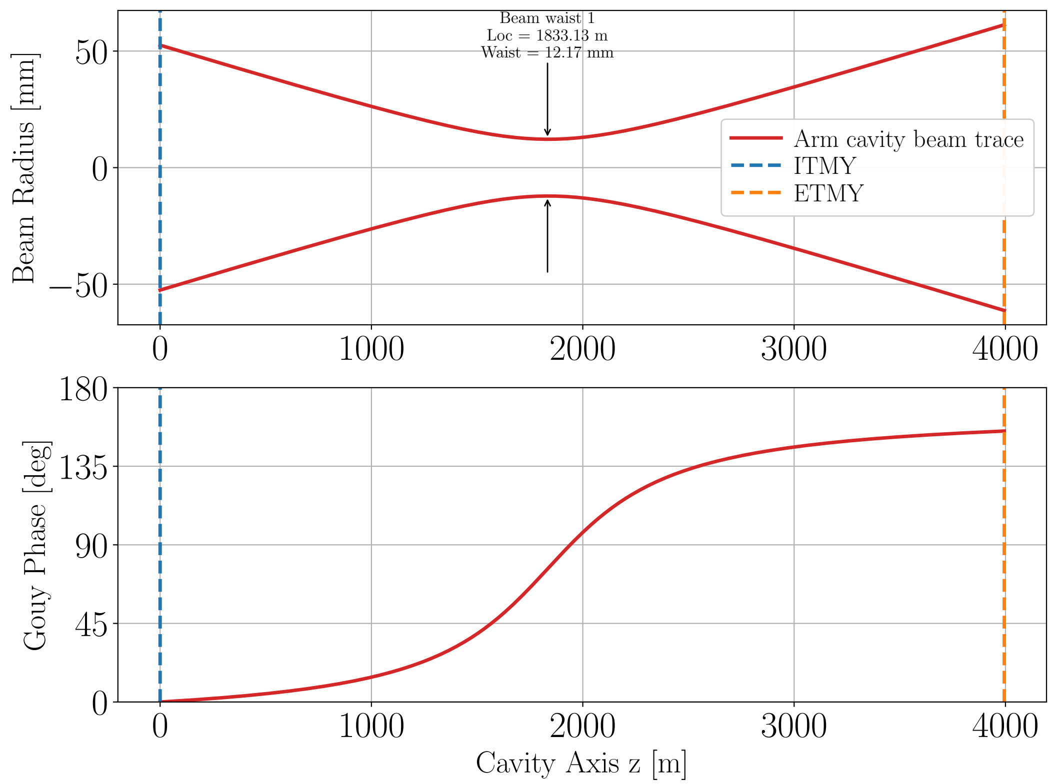

Figure C.3: The Advanced LIGO arm cavity transverse beam profile. The top plot shows the beam waist at any given point along the cavity axis z. The bottom plot shows the single-pass accumulated Gouy phase along the cavity axis. The arrows denote the beam waist, the region of highest beam intensity. Table C.1 shows the arm cavity geometric parameters. This plot was made using the python-based `beamtrace` library.

plots the expected fundamental Gaussian beam profile inside the signal recycling cavity. The Gouy phase is very dynamic at the end of the single-pass of the beam. This is due to the fact that the cavity eigenmode beam waist occurs very near the SRM reflective face. Gouy phase accumulates very fast near the beam waist, leading to the high uncertainty in overall SRC Gouy phase.

Table C.2 reports the signal recycling cavity geometric parameters.

Figures C.5 and C.6 show how the SRC Gouy phase model changes with small changes to the nominal parameters defined in Table C.2. Also plotted are the actual measurements of the SRC Gouy phase from [131, 132].

Most likely, the cavity lengths are built very slightly different than designed. Figure C.5 illustrates the effect of optic moves on the Gouy phase. Displacement of SR2 and SR3 by around 7 mm backward along the cavity axis would approximately

Table C.2: Table of the Advanced LIGO signal recycling cavity geometric parameters. The path from ITMY to SRM is used because it doesn't require two transmissions through the beamsplitter substrate. The beam waist occurs very near the SRM surface. This leads to to a highly variable Gouy phase under cavity parameter uncertainty. The signal recycling cavity beam profile is plotted in Figure C.4.

Parameters	Variable	Value	Units
Optic index of refraction	n	1.45	-
Lengths			
ITMY thickness	t_{ITMY}	0.2	m
ITMYAR to BS	$L_{\mathrm{ITMYAR-BS}}$	5.013	m
BS substrate	t_{BS}	0.069	m
BSAR to SR3	$L_{\mathrm{BSAR-SR3}}$	19.37	m
SR3 to SR2	$L_{\mathrm{SR3-SR2}}$	15.46	m
SR2 to SRM	$L_{\mathrm{SR2-SRM}}$	15.74	m
Radii of curvature			
ITM	R_{ITM}	-1940	m
BS	R_{BS}	∞	m
SR3	R_{SR3}	36.013	m
SR2	R_{SR2}	-6.424	m
SRM	R_{SRM}	-5.678	m
ITMY lens	f_{ITMY}	34.5	km
Beam waist	w_0	2.0	mm
ITM beam radius	w_1	51.75	mm
SRM beam radius	w_2	2.0	mm
Round-trip Gouy phase	φ	40.4	deg
Free spectral range	ν_{FSR}	2.683	MHz
Transverse mode spacing	ν_{TMS}	0.301	MHz
ABCD matrix	$\begin{bmatrix} A & B \\ C & D \end{bmatrix}$	$\begin{bmatrix} -3.07 & 5.13 \times 10^{3} \\ -2.95 \times 10^{-3} & 4.59 \end{bmatrix}$	-

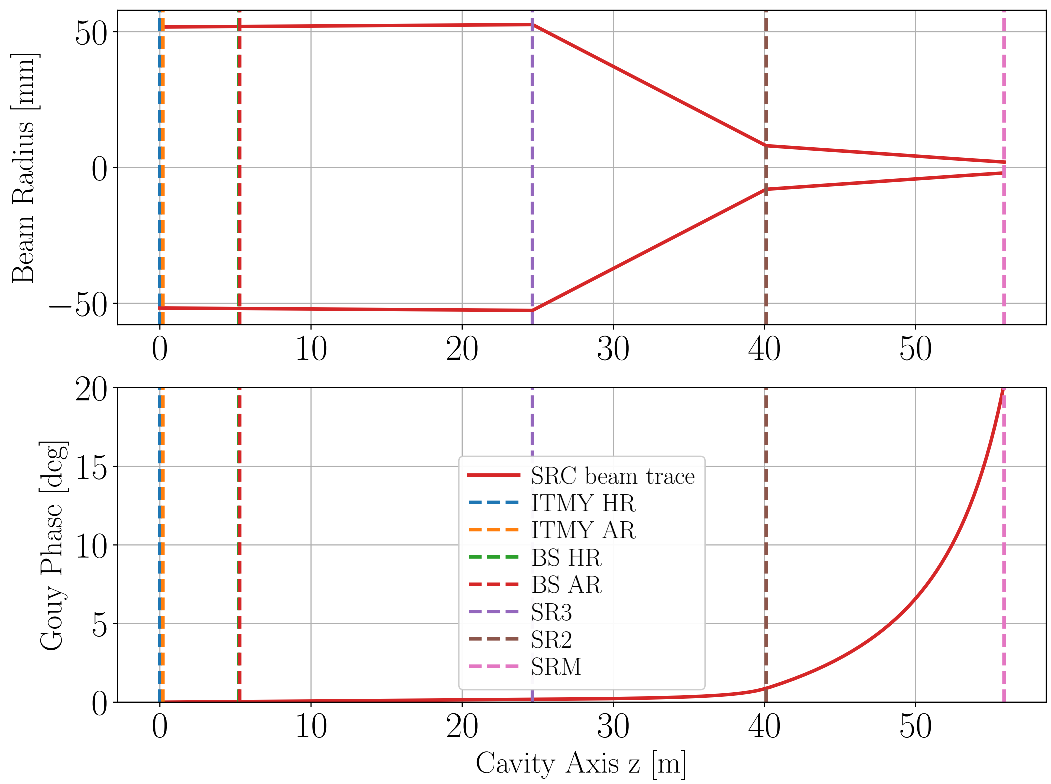

Figure C.4: The Advanced LIGO signal recycling cavity transverse beam profile. The top plot shows the beam waist at any given point along the cavity axis z. The bottom plot shows the single-pass accumulated Gouy phase along the cavity axis. The beam waist occurs very near the SRM face. Therefore, the Gouy phase accumulation rate there is high. This makes the SRC Gouy phase very sensitive to cavity parameter uncertainty. This plot was made using the python-based `beamtrace` library.

match the measured Gouy phase from [132]. Displacement of ITMY or SRM does not significantly impact the beam parameter.

There are two measurements of the SRC Gouy phase, one with the SR3 disk heater off (25.3 degs) and one with it on at 4 W (29.0 degs). The SR3 disk heater is designed to heat the SR3 mirror to adjust the beam parameter resonant in the SRC, such that it is better matched with the beam parameter coming in from the arms. The SR3 disk heater is has been measured to move the SR3 radius of curvature by -3.05 ± 0.15 mm/Watt [130]. Thus, the SR3 radius of curvature should move by around -12 mm when on at 4 W. These measurements are consistent with the analysis in Figure C.6.

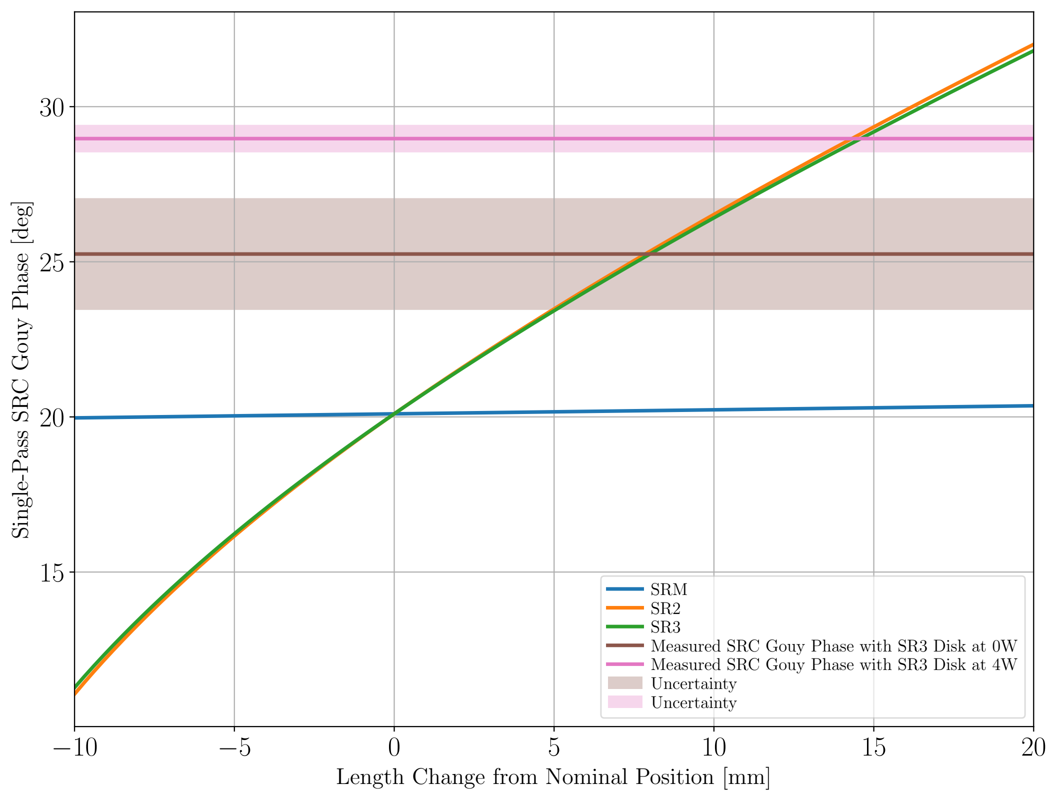

Figure C.5: Single-pass SRC Gouy phase φ when displacing SRC optics along the cavity axis z.

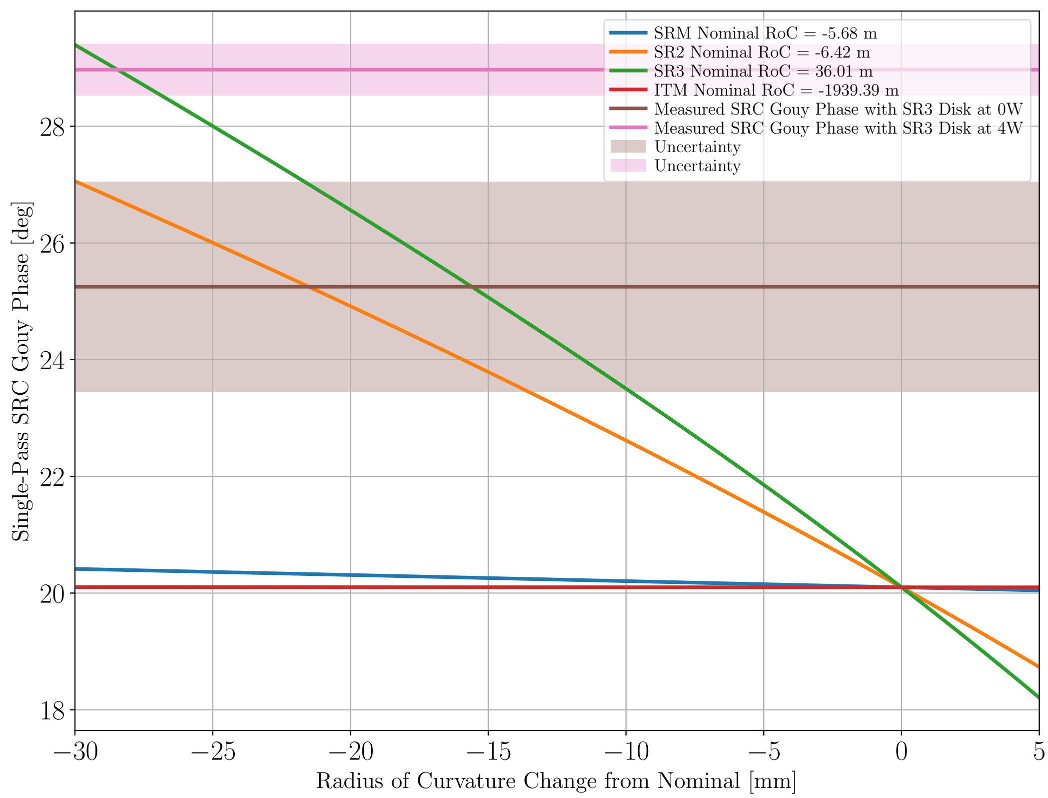

Figure C.6: Single-pass SRC Gouy phase φ when changing the radii of curvature of SRC optics.

Appendix D

TRANSFER FUNCTION ESTIMATES

In cases where data is limited, it can be difficult to properly estimate spectral densities and transfer functions. Transfer functions are best estimated with a large, known excitation applied to a system, to increase coherence between the excitation input and the output. However, in some systems such as the LIGO angular controls or seismic isolation, large excitations are not possible, either because they are physically infeasible (seismic waves cannot be duplicated), or will cause lock-losses.

When calibrating the interferometer, precision and accuracy in the transfer function measurement is key to a good understanding of the astrophysical strain data. However, at frequencies below 10 Hz where seismic noise is high, and frequencies above 1 kHz where the photon calibrators have low actuation authority over the shot noise, it can be difficult to achieve high coherence with many averages. As the SNR of gravitational-wave detections approaches infinity, the accuracy and precision of the calibrated data will be the dominant source of uncertainty in astrophysics done with the data, including parameter estimates and tests of general relativity.

In these systems where low frequency information is valuable, only a few number of averages can be achieved for Welch's method. If true coherence is low, then cross spectral density, transfer function, and coherence estimates become dominated by noise and are rendered useless without sufficient sampling. Additionally, if there is significant noise on the input, then transfer function estimates can become biased. Under these circumstances, understanding the biases and uncertainty in estimates of spectral densities, coherence, and transfer functions is crucial to making accurate measurements. The full expression of the probability distribution functions gives us the best understanding possible of the estimates of spectral quantities.

The first part of this chapter derives the Cramér-Rao bound for cross spectral density estimates, and Bayes factor expression for determining at what number of samples we have "resolved" the cross spectral density, depending on the coherence.

The later part of this chapter will reproduce some important spectral estimators' probability distributions, first derived by N. R. Goodman [200, 202]. Goodman describes a joint distribution the spectral matrix of two zero-mean Gaussian signals as a complex Wishart distribution. From this joint distribution, Goodman changes variables and marginalizes over others to yield exact probability distribution functions of spectral densities, sample coherence, and transfer function gain and phase.

Approximations are useful in certain limits, and their origin lies in the full probability distribution converging with a large number of samples n. Bendat and Piersol have derived the variance in the transfer function estimate $\hat{H} = |\hat{H}|e^{i\hat{\phi}}$ of the true transfer function $H = |H|e^{i\phi}$ [174]:

$$\sigma^2_{|\hat{H}|/|H|} = \sigma^2_{\hat{\phi}} = \frac{1-\gamma^2}{2n\gamma^2} \tag{D.1}$$

where γ^2 is the true coherence and n is the number of averages. This result is valid in the regime $n \to \infty$, coherence $\gamma^2 \to 1$, and noise on the input is negligible so the transfer function bias goes to zero. Equation D.1 will be compared to the full distributions.

D.1 Maximum likelihood estimators of asymmetric Laplace parameters

The asymmetric Laplace distribution describes the distribution of samples from a cross spectral density, as derived in section 6.11. Maximum likelihood estimates for the asymmetric Laplace are known [235, 236]. The covariance matrix is asymptotically normal for large samples n. We focus on the case where the peak location parameter m is known to be zero, as it is for cross-spectral density estimates. Equation (17) from [235] and Table 1 Case 5 from [236] are reproduced here. I assume the peak location parameter $m = 0$, as this is the relevant case for cross spectral densities.

First, the likelihood function of the asymmetric Laplace for n independent samples $[x_1, \ldots, x_n]$ is

$$L(\lambda, \kappa | x_i) = \frac{1}{\lambda^n \left(\kappa + \frac{1}{\kappa}\right)^n} \exp\left(-\frac{n}{\lambda}\left(\kappa\alpha + \frac{\beta}{\kappa}\right)\right) \tag{D.2}$$

where $\alpha = \frac{1}{n}\sum_{i=0}^{n} x_i^+$ and $\beta = \frac{1}{n}\sum_{i=0}^{n} x_i^-$, and $x_i^+ = \max(0, x_i)$ and $x_i^- = \max(0, -x_i)$. α and β are the averages of the positive and negative samples x_i, respectively. Note that, as defined here, β is a positive number.

The maximum likelihood estimators $\hat{\lambda}$ and $\hat{\kappa}$ come from setting the derivative of the log likelihood $\frac{\partial \log(L)}{\partial \lambda}$ and $\frac{\partial \log(L)}{\partial \kappa}$ to zero and solving. This yields the following estimators:

$$\hat{\lambda} = (\alpha\beta)^{1/4}\left(\alpha^{1/2} + \beta^{1/2}\right) \tag{D.3}$$

$$\hat{\kappa} = \left(\frac{\beta}{\alpha}\right)^{1/4} \tag{D.4}$$

We define the mean estimator vector $\vec{\theta} = [\hat{\lambda}, \hat{\kappa}]^T$.

The Fisher information matrix $\mathcal{I}(\lambda, \kappa)$ is the expected value of the product of derivatives of the log likelihood $\left\langle \frac{\partial \log(L)}{\partial \theta_i} \frac{\partial \log(L)}{\partial \theta_j} \right\rangle$. The Fisher information matrix for the asymmetric Laplace with known m is

$$\mathcal{I}(\lambda, \kappa) = \begin{bmatrix} \dfrac{1}{\lambda^2} & \dfrac{\kappa^2 - 1}{(\kappa^3 + \kappa)\lambda} \\ \dfrac{\kappa^2 - 1}{(\kappa^3 + \kappa)\lambda} & \dfrac{\kappa^4 + 6\kappa^2 + 1}{(\kappa^3 + \kappa)^2} \end{bmatrix} \tag{D.5}$$

Inverting the Fisher information matrix yields the Cramér-Rao bound, the covariance matrix $\Sigma(\lambda, \kappa)$:

$$\Sigma(\lambda, \kappa) = \frac{1}{\sqrt{n}} \begin{bmatrix} \dfrac{(\kappa^4 + 6\kappa^2 + 1)\lambda^2}{8\kappa^2} & \dfrac{\lambda - \kappa^4\lambda}{8\kappa} \\ \dfrac{\lambda - \kappa^4\lambda}{8\kappa} & \dfrac{1}{8}(\kappa^2 + 1)^2 \end{bmatrix} \tag{D.6}$$

The maximum likelihood estimators in this section rely on partial sample means α and β, and so are susceptible to glitches. If sample means are not useable, partial sample medians could be substituted, with bias corrections if needed (biases will divide out for $\hat{\kappa}$, but not for $\hat{\lambda}$). Fitting routines robust to outliers can also be used.

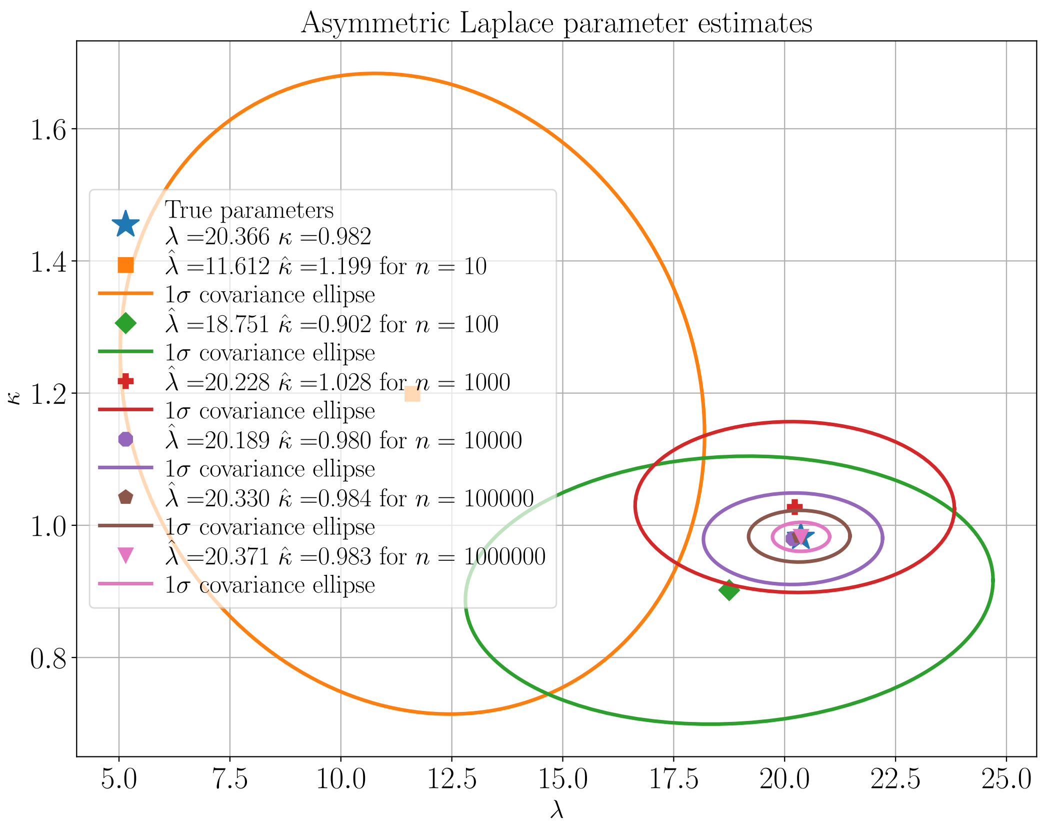

Figure D.1: Maximum likelihood estimates and associated covariance ellipses. Plotted are examples of Eqs. D.3, D.4, and D.6.

D.2 Bayes factor for cross spectral density model comparison

We have explored the usefulness of median-averaging for cross spectral density estimation in this chapter. However, the results above depend on the underlying asymmetric Laplace distribution being *well-sampled*. *Well-sampled* means that sufficient samples n have been taken of our distribution such that its parameters have been accurately determined.

Practically, we would like to know how many cross spectral density averages n we must take until we achieve convergence upon the correlated noise. If the correlated noise is far below the uncorrelated noise, many samples are required. It can be difficult to acquire enough samples, as samples require time and stationary data.

The cross spectral density at some frequency bin is *well-sampled* if, for some number of samples n, the log Bayes factor comparing the symmetric Laplace to asymmetric Laplace is greater than $\log(100)$.

A symmetric Laplace distribution is simply an asymmetric Laplace where $\kappa =$

1. If the signals in the cross spectral density were completely uncorrelated, the symmetric Laplace distribution would describe the samples of each frequency bin, as proven in Section 6.11. By determining the cross spectral density distribution is asymmetric Laplace, we also determine that that the two signals are correlated.

The Bayes factor BF is a ratio of the evidence for the two models given the sample data. The larger the Bayes factor, the more support for the model in the numerator, in this case the asymmetric Laplace. The asymmetric Laplace has two parameters, λ and κ, while the symmetric Laplace has only λ. We assume flat priors across all of the parameter space for both models: $P(\lambda) = 1 \;\; \forall\, \lambda \in [0, \infty)$ and $P(\kappa) = 1 \;\; \forall\, \kappa \in [0, \infty)$. The Bayes factor marginalizes over the entire extra parameter space κ of the asymmetric Laplace, preferring the model with fewer parameters if possible, the manifestation of Occam's razor in Bayes factor calculations.

The calculation of the Bayes factor model comparison can be done analytically, for n samples $\vec{x}$:

$$
\begin{aligned}
BF &= \frac{\displaystyle\int P(\vec{x}|\vec{\theta}, \text{AsymmetricLaplace}) P(\vec{\theta}) d\vec{\theta}}{\displaystyle\int P(\vec{x}|\vec{\theta}, \text{Laplace}) P(\vec{\theta}) d\vec{\theta}} \\
&= \frac{\displaystyle\int_0^\infty \int_0^\infty L(\lambda, \kappa|\vec{x}) d\lambda d\kappa}{\displaystyle\int_0^\infty L(\lambda, \kappa = 1|\vec{x}) d\lambda} \\
&= \frac{\displaystyle\int_0^\infty \int_0^\infty \frac{1}{\lambda^n \left(\kappa + \frac{1}{\kappa}\right)^n} e^{-\frac{n}{\lambda}\left(\kappa\alpha + \frac{1}{\kappa}\beta\right)} d\lambda d\kappa}{\displaystyle\int_0^\infty \frac{1}{(2\lambda)^n} e^{\frac{-n(\alpha+\beta)}{\lambda}} d\lambda} \\
&= \frac{\frac{1}{2}\alpha(\alpha n)^{-n}\Gamma(n-1)^2\Gamma(n+1)\, {}_2F_1\left(n-1, n-1; 2n-1; 1-\frac{\beta}{\alpha}\right)}{2^{-n}\Gamma(n-1)(n(\alpha+\beta))^{1-n}}
\end{aligned}
$$

$$
BF = 2^{n-1}\Gamma(n-1)\Gamma(n)\left(1 + \frac{\beta}{\alpha}\right)^{n-1} {}_2\tilde{F}_1\left(n-1, n-1; 2n-1; 1-\frac{\beta}{\alpha}\right) \quad \text{(D.7)}
$$

where ${}_2\tilde{F}_1$ is the regularized hypergeometric function and Γ is the gamma function. We've used Eq. D.2 for the likelihood $L(\lambda, \kappa|\vec{x})$, and α and β are the positive and negative averages, defined below D.2.

The log Bayes factor is plotted in Figure D.2. The negative and positive mean ratio β/α determines at what n the Bayes factor will swing toward supporting the asymmetric Laplace model. The further β/α is from 1, the fewer samples n

Table D.1: Numbers of samples required to confidently resolve a cross spectral density with levels of coherence $\overline{\gamma^2}$. Also listed are the four equivalent ways of quantifying the skewness of a cross spectral density estimate: mean-averaged coherence and power ratio from Eq. 6.89, asymmetric Laplace parameter κ from Eqs. 6.90 and 6.91, and the negative over positive mean ratio from Eq. D.4. These results are plotted in Figure D.2. The number of required samples approximately equals $8(1/\overline{\gamma^2} - 1)$, as seen in Figure D.3.

Samples n	Coherence $\overline{\gamma^2}$	Power ratio ϵ	κ	Mean ratio β/α
23	0.333	2	0.518	0.072
69	0.100	9	0.721	0.270
212	0.033	29	0.831	0.478
749	0.010	99	0.905	0.669
2400	0.003	299	0.944	0.794
8600	0.001	999	0.969	0.881

required to achieve a decisive Bayes factor of 100. If $\beta/\alpha = 1$, the distribution is symmetric, and the Bayes factor will only support the symmetric Laplace more and more with additional samples. Setting $\beta/\alpha = 1$ yields

$$BF = \frac{\sqrt{\pi}\,\Gamma(n-1)}{\Gamma\left(n - \frac{1}{2}\right)} \tag{D.8}$$

$$BF \approx \sqrt{\frac{\pi}{n}} \tag{D.9}$$

The log of D.9 is plotted as the dashed line in Figure D.2.

A table of required samples n to resolve a mean-averaged coherence $\overline{\gamma^2}$ is provided in Table D.1. The number of required samples to reach the decisive log Bayes factor of $\log(100)$ approximately equals $8(1/\overline{\gamma^2} - 1)$, as shown in Figure D.3. This approximation is only good for a log Bayes factor of $\log(100)$, the line may move up or down if we raise or lower the decisive log Bayes factor level.

Blackman and Tukey found that for a true coherence $\gamma^2 = 0$, the mean-averaged sample coherence $\overline{\hat{\gamma}^2} = 1/n$ [199]. This is plotted in Figure D.3 to illustrate the zero-coherence case. This study is also important to the sample coherence in Section D.7.

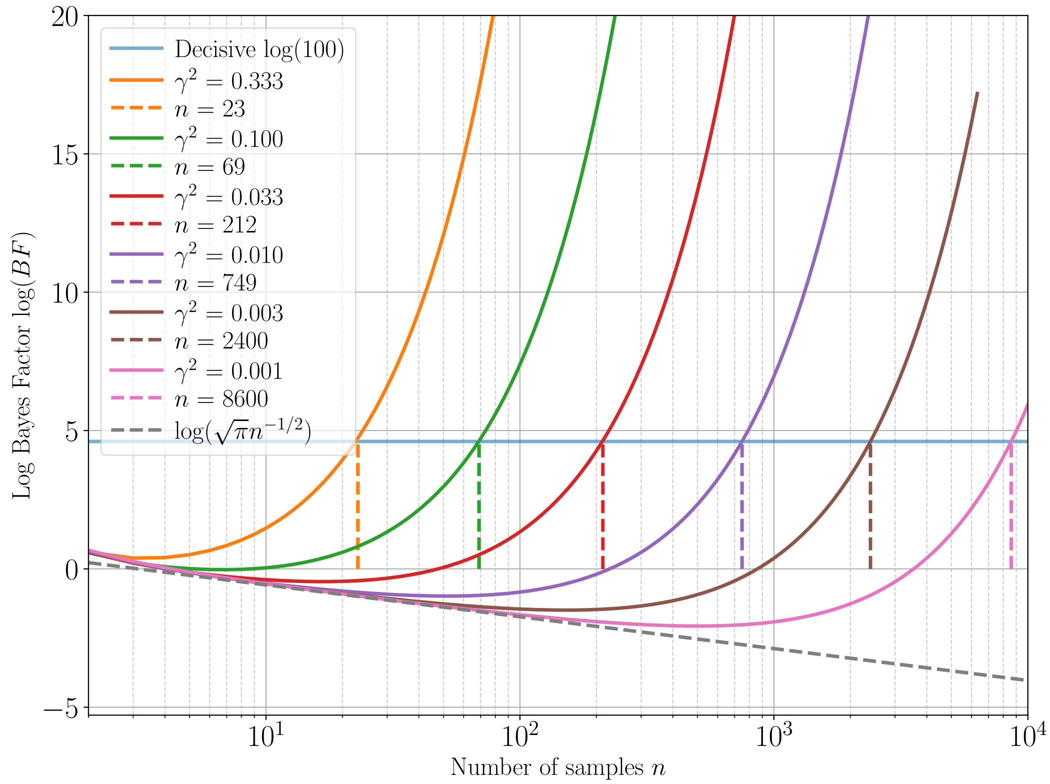

Figure D.2: Log Bayes factor vs number of samples n. The decision point in favor of the asymmetric Laplace model is $BF \geq 100$, the blue line. This is the number of samples n when a cross spectral density measurement can be said to be converged to the correlated noise level. Plotted are log versions of Eq. D.7 for different coherences γ^2. The dashed grey line is Eq. D.9, the maximum support possible for the symmetric Laplace model. Table D.1 records the minimum samples required to resolve a given coherence.

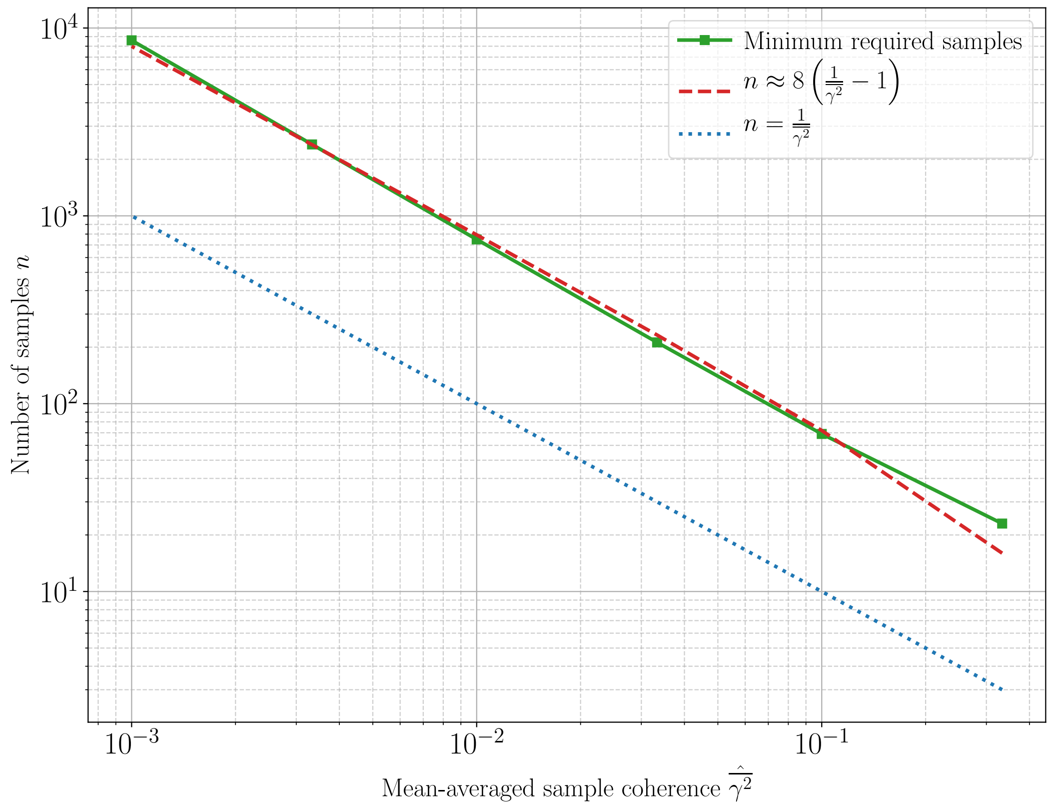

Figure D.3: Mean-averaged sample coherence $\overline{\hat{\gamma^2}}$ vs number of samples required to confidently resolve coherence level. The green dots are calculated directly from the log Bayes factor Eq. D.7. The dashed red line is a convenient approximation for determining the number of samples required up to the decisive log Bayes factor of $\log(100)$. The dotted blue line is the sample coherence expected from two signals with true coherence $\gamma^2 = 0$ after n averages, according to Tukey [199].

D.3 Sample mean for power spectral densities

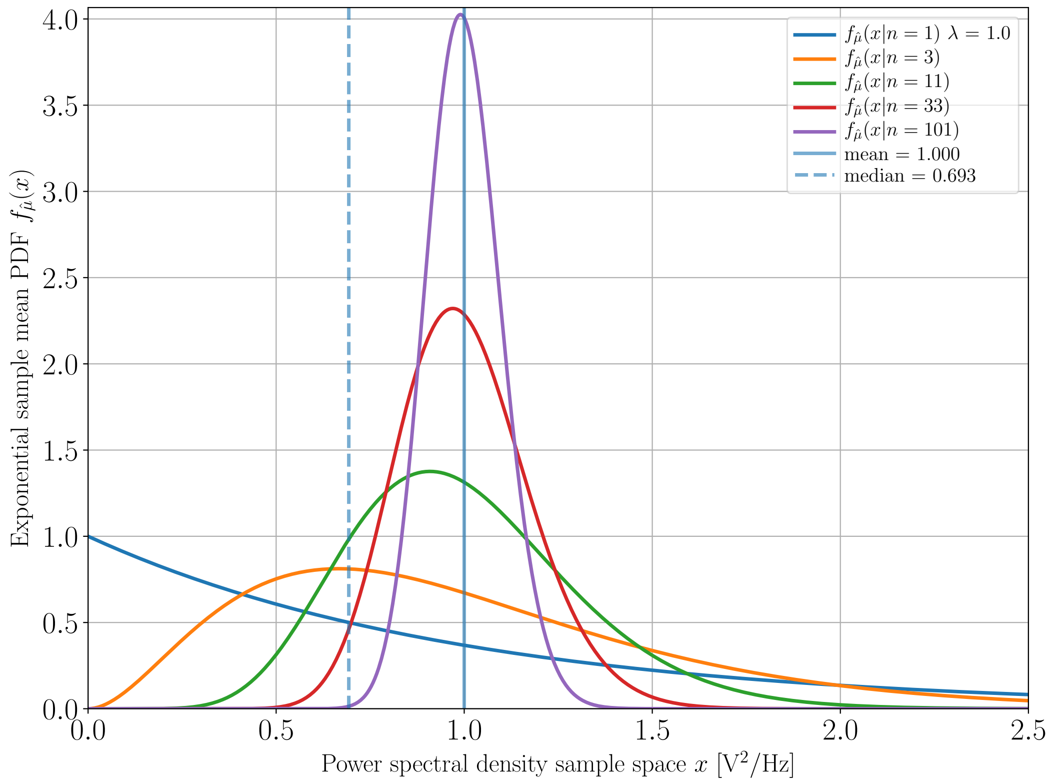

Figure D.4: Probability distributions of the sample mean for power spectral densities $f_{\hat{\mu}}$ for different numbers of samples n. Plotted from Eq. D.10.

The probability density of the PSD sample mean $\hat{\mu} = \frac{1}{n}\sum_{i=0}^{n-1} x_i$ of the exponential distribution is reported here. Each x_i is a power spectral density estimate at some frequency bin used in Welch's method.

The PSD sample mean distribution is derived by taking a convolution of the exponential with itself n times. The probability distribution function of the PSD sample mean of the exponential distribution $f_{\hat{\mu}}$ on the power spectral density sample space x is

$$f_{\hat{\mu}}(x) = n\frac{(nx)^{n-1}}{\lambda^n\Gamma(n)}e^{-\frac{nx}{\lambda}} \tag{D.10}$$

where n is the number of samples, λ is the mean of the exponential, and Γ is the gamma function. This is equivalent to a chi-squared distribution with $2n$ degrees of freedom, scaled by $2n/\lambda$: $f_{\hat{\mu}}(x|n) \sim \chi^2_{2n}(2nx/\lambda)$.

Using Eq. D.10, the expected value of the sample PSD

$$\langle\hat{\mu}\rangle = \lambda \tag{D.11}$$

making $\hat{\mu}$ an unbiased estimator.

The variance of Eq. D.10 is

$$\sigma_{\hat{\mu}}^2 = \frac{\lambda^2}{n}, \tag{D.12}$$

meaning the variance in a sample PSD is equal to itself squared divided by the number of samples.

The probability density in Eq. D.10 and their mean are plotted in Figure D.4.

D.4 Sample median for power spectral densities

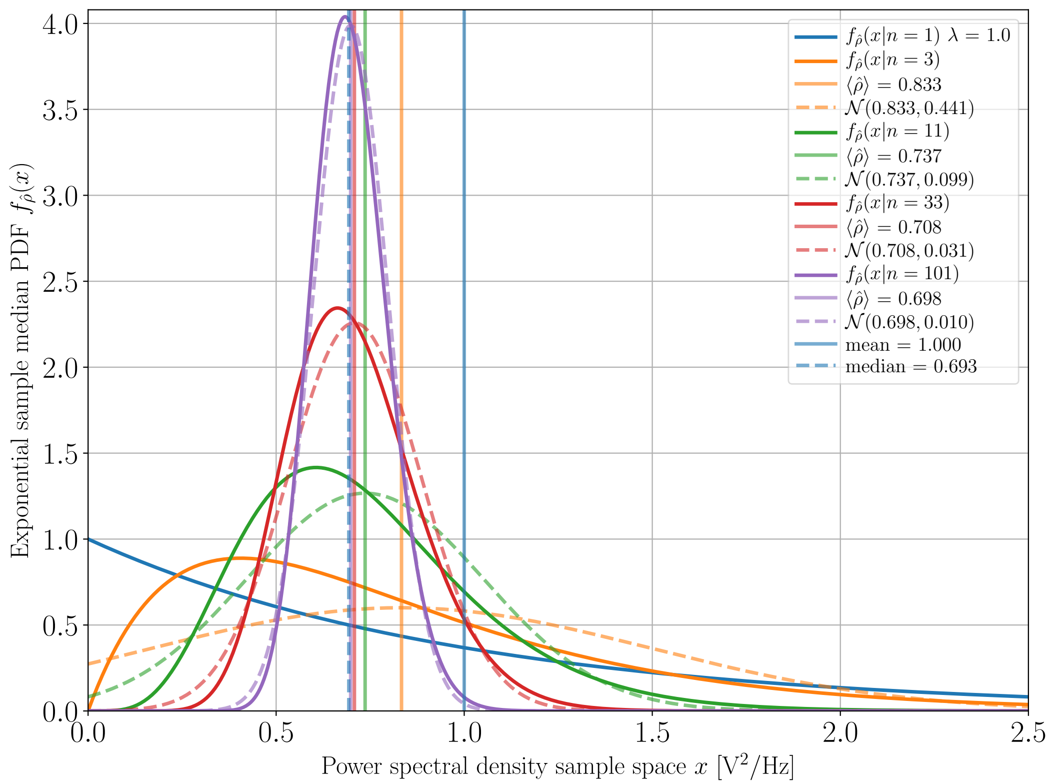

Figure D.5: Probability distributions of the sample median for power spectral densities $f_{\hat{\rho}}$ for different numbers of samples n (Eq. D.14). The mean of each sample median distribution $\langle\hat{\rho}\rangle$ is plotted as a solid vertical line (Eq. D.15). Gaussian approximations to each distribution are plotted as dashed lines, and converge as n increases.

The probability density of the PSD sample median $\hat{\rho} = x_{1/2}$ of the exponential distribution is reported here. These distributions describe a power spectral density at some frequency bin estimated via Welch's method using median-averaging.

This derivation closely follows that of Appendix B in [193], which derives the bias in the sample median estimator for low numbers of samples for power spectral densities. Methods from section 28.5 of [237] are employed here. A similar calculation for the Laplace distribution has also been performed [238].

First, we must know the cumulative distribution function $F(x)$ of the exponential. This is found from Eq. 6.2 to be

$$F(x) = 1 - e^{-\frac{x}{\lambda}} \tag{D.13}$$

The probability distribution of the sample median $f_{\hat{\rho}}(x)$ comes from calculating the probability that half of all samples will be less than the median times the probability that half will be greater. The probability that half of all samples will be less than the median is $F(x)^{(n-1)/2}$, while the probability of half of all samples greater is $(1 - F(x))^{(n-1)/2}$. This yields the sample median PDF $f_{\hat{\rho}}(x)$:

$$\begin{aligned} f_{\hat{\rho}}(x|n) &= F(x)^{\frac{n-1}{2}}(1 - F(x))^{\frac{n-1}{2}} f(x) \\ f_{\hat{\rho}}(x|n) &= \frac{\left(e^{-\frac{x}{\lambda}}\right)^{\frac{n+1}{2}}\left(1 - e^{-\frac{x}{\lambda}}\right)^{\frac{n-1}{2}}}{\lambda B\left(\frac{n+1}{2}, \frac{n+1}{2}\right)} \end{aligned} \tag{D.14}$$

where $B(x, y)$ is the Beta function and $f(x)$ is the parent exponential distribution from which all samples are drawn.

Unlike the sample mean, the sample median $\hat{\rho}$ is a biased estimator, as $\langle\hat{\rho}\rangle \neq \rho$. It is a *consistent* estimator, however. The expected value of the sample median $\langle\hat{\rho}\rangle$ for a small number of samples n, assuming that n is odd, is [193]

$$\langle\hat{\rho}\rangle = \lambda \sum_{k=1}^{n} \frac{(-1)^{k+1}}{k} \tag{D.15}$$

As $n \to \infty$, $\langle\hat{\rho}\rangle \to \rho$, and the sample median variance $\sigma^2_{\hat{\rho}} \to 1/(4nf(\rho)^2)$ [237]. Figure D.5 plots the normal approximations to the sample median probability density functions with mean $\langle\hat{\rho}\rangle$ and variance $1/(4nf(\rho)^2)$ as a dashed line.

D.5 Sample mean for cross spectral densities

The joint probability density of the sample mean $\hat{\mu}$ for cross spectral densities is reported here.

The sample mean distribution can be found by using the characteristic function theorem for summing n independent, identically distributed random variables, Eq. 6.29. The characteristic function of the CSD joint distribution is found in

Eq. 6.79. Taken as the base characteristic function in Eq. 6.29, we find the characteristic function of the CSD sample mean $\varphi_{\hat{\mu}}$:

$$\begin{aligned}\varphi_{\hat{\mu}}(s,t) &= \varphi_{\mathcal{U},\mathcal{V}}\left(\frac{s}{n},\frac{t}{n}\right)^n \\ \varphi_{\hat{\mu}}(s,t) &= \left(\frac{n^2}{n^2-2i\sigma_a^2 ns+\sigma_a^2\sigma_c^2(s^2+t^2)}\right)^n\end{aligned}\tag{D.16}$$

Then the joint distribution $f_{\hat{\mu}}$ of the sample mean for n samples is

$$\begin{aligned}f_{\hat{\mu}}(u,v) &= \frac{1}{2\pi}\int\int_{-\infty}^{\infty} ds\,dt\,e^{isu+itv}\varphi_{\hat{\mu}}(s,t) \\ f_{\hat{\mu}}(u,v) &= \frac{n^{n+1}((1-\gamma^2)(u^2+v^2))^{\frac{n-1}{2}}e^{\frac{n\gamma}{\lambda\sqrt{1-\gamma^2}}(u\cos(\phi)+v\sin(\phi))}}{\pi 2^n\lambda^{n+1}\Gamma(n)}K_{n-1}\left(\frac{n}{\lambda}\sqrt{\frac{u^2+v^2}{1-\gamma^2}}\right)\end{aligned}\tag{D.17}$$

Using Eq. D.17, we find the expected value of the sample mean

$$\langle\hat{\mu}\rangle = \frac{2\gamma\lambda}{\sqrt{1-\gamma^2}}\tag{D.18}$$

for all n, making $\hat{\mu}$ an unbiased estimator.

The probability density in Eq. D.17 and the mean in Eq. D.18 are plotted in Figure D.6.

Now we marginalize the joint sample mean distribution to recover the sample distribution along the major and minor axes. From Chapter 6 section 6.11, we know the major axis distribution is an asymmetric Laplace, while the minor axis has a Laplace distribution.

The sample mean probability distribution along the major axis u and minor axis v are

$$f_{\hat{\mu}}(u) = \frac{n^{n+\frac{1}{2}}(1-\gamma^2)^{\frac{1}{2}(n-\frac{1}{2})}|u|^{n-\frac{1}{2}}e^{\frac{n\gamma u}{\lambda\sqrt{1-\gamma^2}}}}{2^{n-\frac{1}{2}}\lambda^{n+\frac{1}{2}}\sqrt{\pi}\,\Gamma(n)}K_{n-\frac{1}{2}}\left(\frac{n}{\lambda}\frac{|u|}{\sqrt{1-\gamma^2}}\right)\tag{D.19}$$

$$f_{\hat{\mu}}(v) = \frac{n^{n+\frac{1}{2}}|v|^{n-\frac{1}{2}}}{2^{n-\frac{1}{2}}\lambda^{n+\frac{1}{2}}\sqrt{\pi}\,\Gamma(n)}K_{n-\frac{1}{2}}\left(\frac{n}{\lambda}|v|\right)\tag{D.20}$$

The expected value of the major axis sample mean (Eq. D.19) is equal to Eq. D.18. The expected value of the minor axis sample mean (Eq. D.20) is equal to zero.

The variance of the major axis sample mean $\sigma_{\hat{\mu}}^2 = \langle \hat{\mu}^2 \rangle - \langle \hat{\mu} \rangle^2$ is equal to

$$\sigma_{\hat{\mu}}^2 = \frac{2\lambda^2(1+\gamma^2)}{n(1-\gamma^2)} \tag{D.21}$$

The variance of the minor axis sample mean is equal to

$$\sigma_{\hat{\mu}}^2 = \frac{2\lambda^2}{n} \tag{D.22}$$

The probability density in Eq. D.19 and the mean in Eq. D.18 are plotted in Figure D.7.

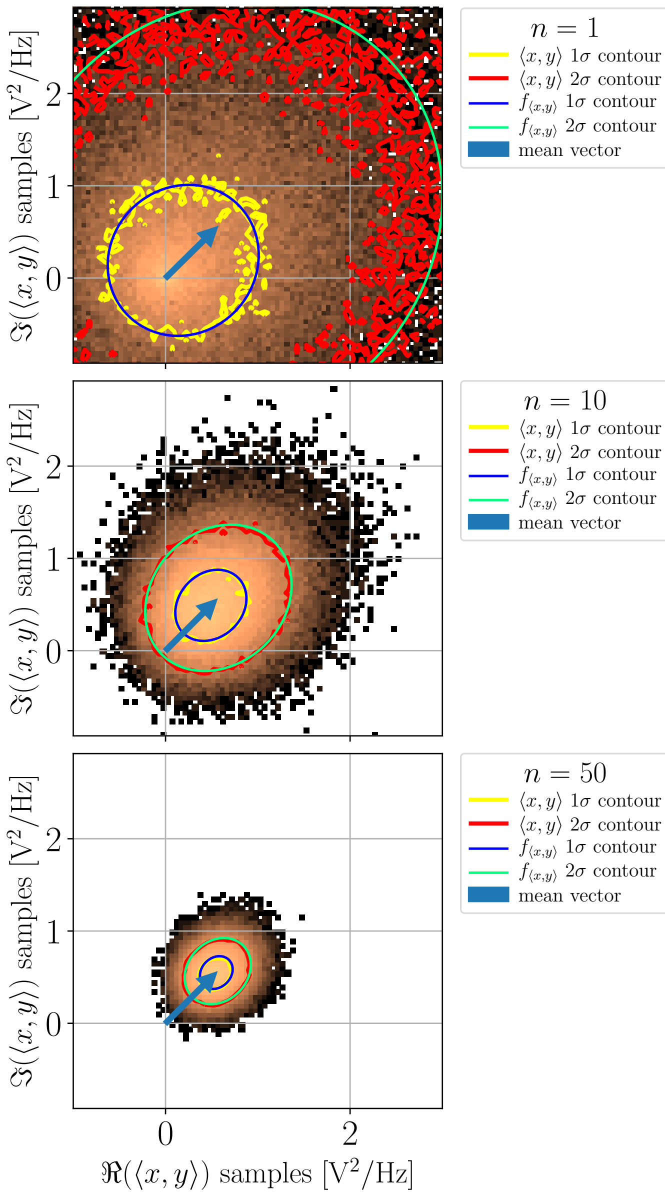

Figure D.6: Cross spectral density sample mean joint probability distributions, 2d histograms, and data contours for different numbers of averages n. Each data point of the histogram represents a CSD average. For instance, for the center plot with $n = 10$, each data point represents an average of ten CSDs. Each plot features 100000 averages. Every point in the histogram represents a potential result for a measurement of a CSD. This plot illustrates how the mean vector, represented by the blue arrow, always remains the same: $\mu = 0.8$ V^2/Hz (cf. Eq. D.18). It also shows how quickly the results converge around the mean, with the 1σ and 2σ data contours decreasing rapidly with increasing averages. Finally, the data contours are matched to the contours drawn from the joint probability distribution Eq. D.17. In this example, the coherence $\gamma^2 = 0.198$, the cross power scaler $\lambda = 0.806$ V^2/Hz, and the phase $\phi = 45$ degs.

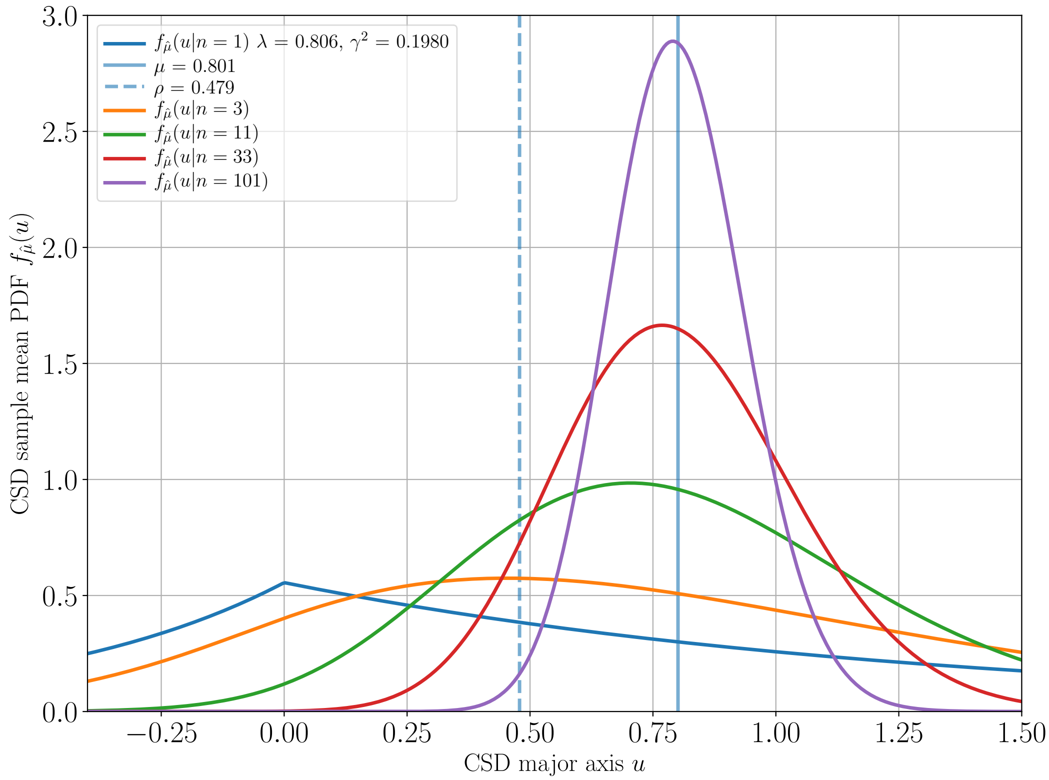

Figure D.7: Probability distributions of the sample mean along the major axis of a cross spectral density for different numbers of samples n. The major axis sample mean distribution $f_{\hat{\mu}}$ is plotted from Eq. D.19. This plot illustrates how the range of potential CSD sample means converge to the true mean with increasing samples n. The mean of all distributions is always equal to $\mu = 0.8$ V^2/Hz, making the sample mean $\hat{\mu}$ an unbiased estimator. In this example, the coherence $\gamma^2 = 0.198$ and the cross power scaler $\lambda = 0.806$ V^2/Hz.

D.6 Sample median for cross spectral densities

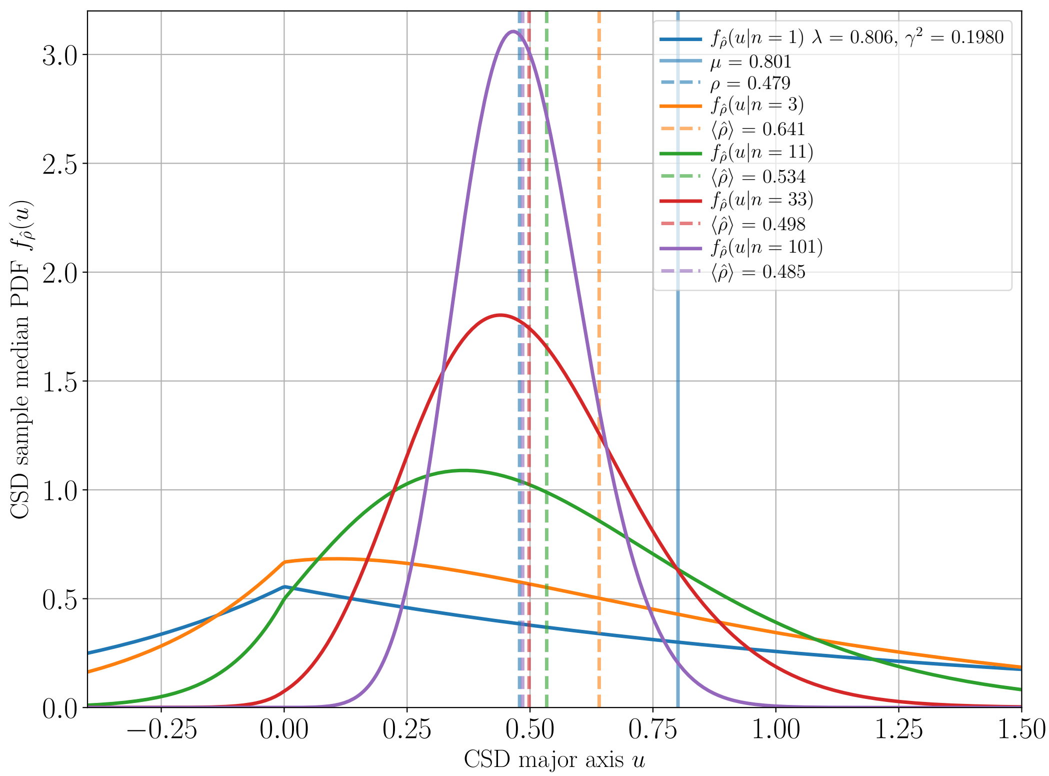

Figure D.8: Probability distributions of the sample median of cross spectral densities for different numbers of samples n. The major axis sample median distribution $f_{\hat{\rho}}$ is plotted from Eq. D.24. This plot illustrates how the range of potential CSD sample medians converge to the true median with increasing samples n. The sample median $\hat{\rho}$ is a biased, but consistent, estimator, as illustrated by the asymptotic convergence of the expected value of the sample median to the true median, $\rho = 0.478$ V^2/Hz. In this example, the coherence $\gamma^2 = 0.198$ and the cross power scaler $\lambda = 0.806$ V^2/Hz.

The probability density of the sample median $\hat{\rho}$ for cross spectral densities is reported here.

The full CSD probability distribution is a 2d joint distribution. The median of a joint distribution is ill-defined. Instead, this section will focus on the median of the 1d CSD distributions.

The cumulative distribution function $F(u)$ of the asymmetric Laplace is

$$F(u) = \begin{cases} \dfrac{1-\gamma}{2} e^{\frac{u(1+\gamma)}{\lambda\sqrt{1-\gamma^2}}} & u \leq 0 \\ 1 - \dfrac{1+\gamma}{2} e^{\frac{-u(1-\gamma)}{\lambda\sqrt{1-\gamma^2}}} & u > 0 \end{cases} \tag{D.23}$$

The sample median distribution is calculated similarly to Section D.4. The sample median probability distribution along the major axis of the CSD $f_{\hat{\rho}}$ is:

$$f_{\hat{\rho}}(u) = F(u)^{\frac{n-1}{2}}(1-F(u))^{\frac{n-1}{2}} f(u)$$

$$f_{\hat{\rho}}(u) = \begin{cases} \dfrac{\sqrt{1-\gamma^2}e^{\frac{u}{\lambda}\sqrt{\frac{1+\gamma}{1-\gamma}}}}{2^n \lambda\, B\left(\frac{n+1}{2},\frac{n+1}{2}\right)}\left((1-\gamma)e^{\frac{u}{\lambda}\sqrt{\frac{1+\gamma}{1-\gamma}}}\left(2-(1-\gamma)e^{\frac{u}{\lambda}\sqrt{\frac{1+\gamma}{1-\gamma}}}\right)\right)^{\frac{n-1}{2}} & u \leq 0 \\ \dfrac{\sqrt{1-\gamma^2}e^{-\frac{u}{\lambda}\sqrt{\frac{1-\gamma}{1+\gamma}}}}{2^n \lambda\, B\left(\frac{n+1}{2},\frac{n+1}{2}\right)}\left((1+\gamma)e^{-\frac{u}{\lambda}\sqrt{\frac{1-\gamma}{1+\gamma}}}\left(2-(1+\gamma)e^{-\frac{u}{\lambda}\sqrt{\frac{1-\gamma}{1+\gamma}}}\right)\right)^{\frac{n-1}{2}} & u > 0 \end{cases} \tag{D.24}$$

where $B(x,y)$ is the Beta function and $f(u)$ is the parent asymmetric Laplace distribution from which all samples are drawn.

Unlike the sample mean, the sample median $\hat{\rho}$ is a biased estimator, as $\langle\hat{\rho}\rangle \neq \rho$. It is a *consistent* estimator, however. The expected value of the sample median for a small number of samples n is

$$\begin{aligned}\langle\hat{\rho}\rangle = \frac{\lambda(\gamma+1)^{\frac{n}{2}}\Gamma(n+1)}{2^{\frac{n+1}{2}}\sqrt{1-\gamma}}\Bigg((1+\gamma)\,{}_3\tilde{F}_2\left(\frac{1-n}{2},\frac{n+1}{2},\frac{n+1}{2};\frac{n+3}{2},\frac{n+3}{2};\frac{1+\gamma}{2}\right) \\ -(1-\gamma)\left(\frac{1-\gamma}{1+\gamma}\right)^{\frac{n+1}{2}}{}_3\tilde{F}_2\left(\frac{1-n}{2},\frac{n+1}{2},\frac{n+1}{2};\frac{n+3}{2},\frac{n+3}{2};\frac{1-\gamma}{2}\right)\Bigg)\end{aligned} \tag{D.25}$$

where ${}_3\tilde{F}_2$ is the regularized hypergeometric function and Γ is the gamma function.

As $n \to \infty$, $\langle\hat{\rho}\rangle \to \rho = \lambda(1+\gamma)\log(1+\gamma)/\sqrt{1-\gamma^2}$, where ρ is the true mean of the asymmetric Laplace describing the major axis of the CSD, as seen from Eq 6.99.

As $n \to \infty$, $\langle\hat{\rho}\rangle \to \rho$, and the sample median variance $\sigma^2_{\hat{\rho}} \to 1/(4nf(\rho)^2)$ [237]. Figure D.8 plots the sample median distribution Eq. D.24 alongside each PDF's associated mean from Eq. D.25. This plot shows the importance of having a sufficiently sampled CSD before correcting for any mean-to-median biasing: an insufficiently sampled CSD will not have converged to the true median. A good heuristic for sufficient sampling is whether the coherence γ^2 is greater than $1/n$, this is explored in depth in Section D.2.

D.7 Sample coherence distribution

The probability density of the sample coherence $\hat{\gamma^2}$ is reported here. Estimating the coherence well, and understanding its variance, is crucial for accurate transfer function and cross spectral density estimation.

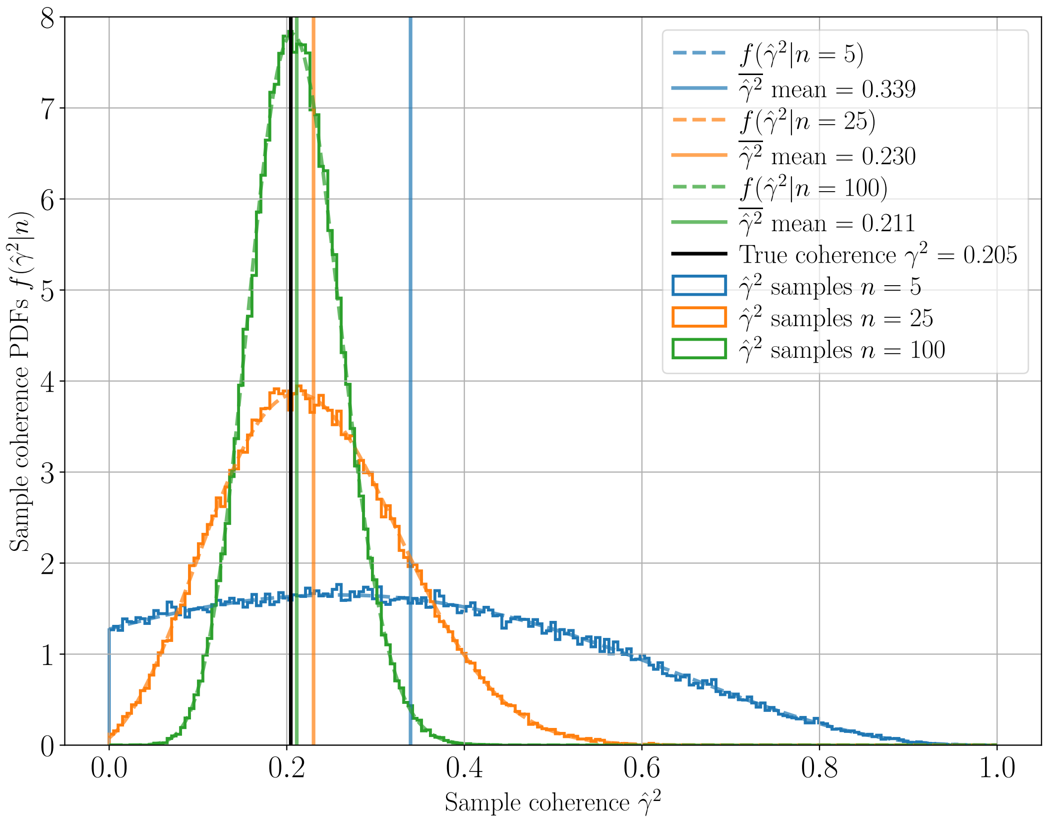

Figure D.9: Probability distributions of the sample coherence for different numbers of samples n. This plot illustrates how the range of potential sample coherences converge to the true coherence with increasing samples n. The sample coherence $\hat{\rho}$ is a biased, but consistent, estimator, as illustrated by the asymptotic convergence of the expected value of the sample coherence to the true coherence. In this example, the true coherence $\gamma^2 = 0.205$ and is plotted as the black line.

When estimating the true coherence γ^2 between two signals, the sample coherence starts equal to exactly one for only one sample, as the relationship between the two signals is unknown. As more samples are taken, the sample coherence falls approximately like $1/n$ to its true value, making the sample coherence $\hat{\gamma}^2$ a biased, consistent estimator.

The sample coherence probability density is derived from a complex Wishart distribution by Goodman [200, 202]. Setting the variable $z = \hat{\gamma}^2$ for ease of notation, the sample coherence probability density is

$$f(z) = (n-1)(1-\gamma^2)^n (1-z)^{n-2} \, {}_2F_1\big(n, n; 1; z\gamma^2\big) \tag{D.26}$$

where ${}_2F_1$ is the hypergeometric function and n is the number of samples.

The mean of the sample coherence $\left\langle \hat{\gamma^2} \right\rangle$ is found from Eq. D.26:

$$\left\langle \hat{\gamma^2} \right\rangle = (1-\gamma^2)^n \Gamma(n) \, {}_3\tilde{F}_2\big(2, n, n; 1, n+1; \gamma^2\big) \tag{D.27}$$

where ${}_3\tilde{F}_2$ is the regularized hypergeometric function, and Γ is the gamma function. When $n \ll 1/\gamma^2$, the sample coherence mean $\left\langle \hat{\gamma^2} \right\rangle \approx 1/n$. When $n \gtrsim 1/\gamma^2$, $\left\langle \hat{\gamma^2} \right\rangle$ asymptotes to γ^2.

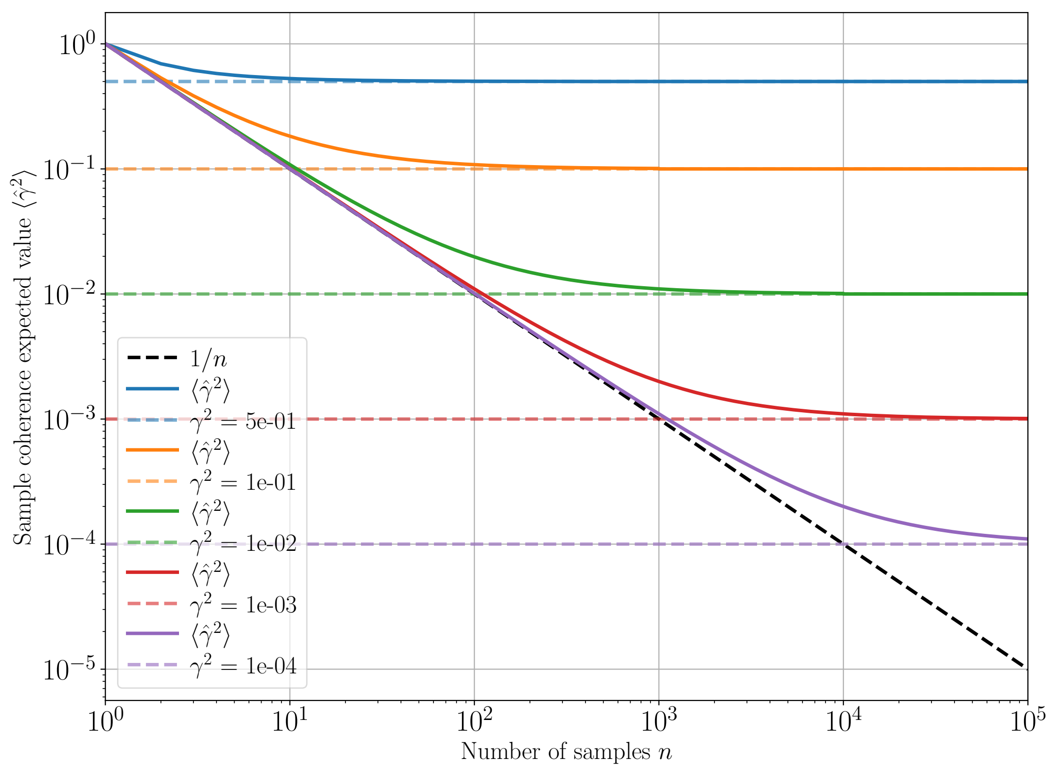

Figure D.10: Expected values of the sample coherence $\left\langle \hat{\gamma^2} \right\rangle$ vs number of samples n, for several levels of true coherence γ^2 (Eq. D.27). This plot illustrates how the sample coherence falls like $1/n$ before converging to the true coherence. Accurate coherence estimates are necessary for accurate estimates of transfer functions and cross spectral densities.

The variance of the sample coherence $\sigma^2_{\hat{\gamma^2}}$ is found from Eq. D.26:

$$\sigma^2_{\hat{\gamma^2}} = 2\big(1-\gamma^2\big)^n \Gamma(n) \Bigg({}_3\tilde{F}_2\big(3, n, n; 1, n+2; \gamma^2\big) \\ - \frac{1}{2}\big(1-\gamma^2\big)^n \Gamma(n) \, {}_3\tilde{F}_2\big(2, n, n; 1, n+1; \gamma^2\big)^2 \Bigg) \tag{D.28}$$

Here, when $n \gtrsim 1/\gamma^2$, $\sigma^2_{\hat{\gamma^2}} \approx 2\gamma^2(1-\gamma^2)^2/n$.

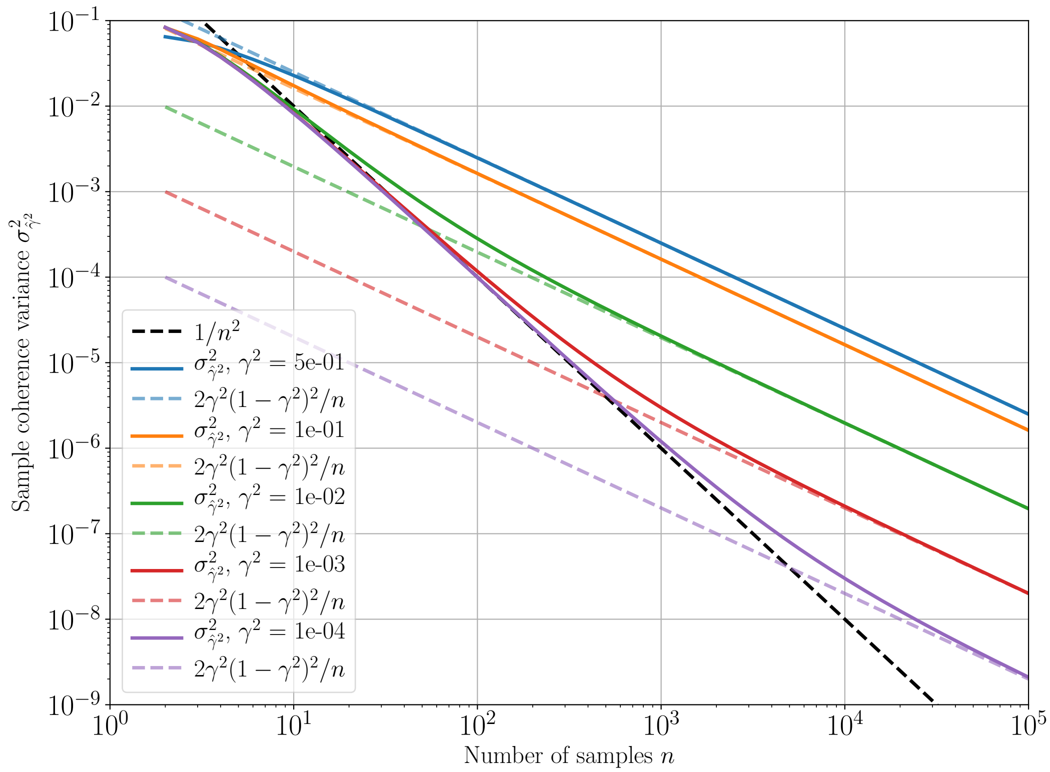

Figure D.11: Expected values of the sample coherence $\sigma^2_{\hat{\gamma}^2}$ vs number of samples n, for several levels of true coherence γ^2 (Eq. D.28). This plot illustrates how the sample coherence variance falls like $1/n^2$ at first, then once the sample coherence expected value approaches the true coherence $\left\langle \hat{\gamma}^2 \right\rangle \to \gamma^2$, the variance approaches $2\gamma^2(1-\gamma^2)^2/n$.

D.8 Sample transfer function

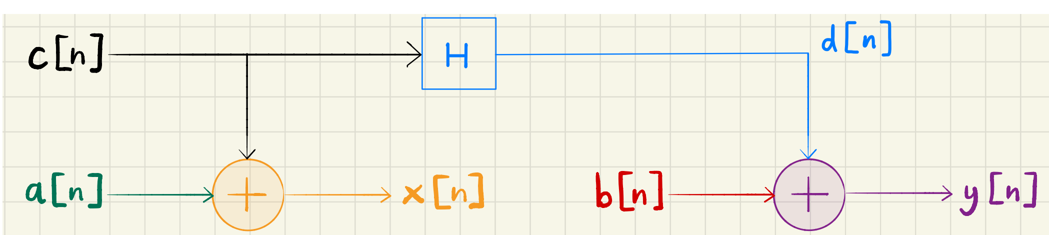

Figure D.12: Independent Gaussian noises a, b, c, and d sum to form the measurable signals x and y. A single-input single-output linear system H transforms the signal c. These signals will form the basis of the derivations in this section.

In this section we explore the sample transfer function (TF) $\hat{H}$, an estimator for the true transfer function H, reporting the explicit probability density functions, common approximations for the mean and variance, and sources of potential bias in TF estimates.

Transfer functions, or frequency response functions, measure how an input signal at some frequency f is linearly transformed by a system H. The diagram in Figure D.12 illustrates the input signal $c[n]$ being transformed by the system H into the output signal $d[n]$, with the measured input $x[n] = a[n] + c[n]$ and the measured output $y[n] = b[n] + d[n]$.

The sample TF $\hat{H}$ is an estimate of the true TF based on the sample power and cross spectral densities of our measurable signals x and y:

$$\hat{H} = |\hat{H}|e^{i\phi} = \frac{\langle x, y\rangle}{\langle x, x\rangle} \tag{D.29}$$

where $|\hat{H}|$ is the magnitude of the sample TF, ϕ is the phase of the sample TF, $\langle x, x\rangle$ is the mean PSD of x, and $\langle x, y\rangle$ is the mean CSD between x and y.

The true transfer function H represents how a sinusoid input at frequency f into a linear system produces a sinusoid output at the same frequency f. From Figure D.12, the true TF is

$$H = |H|e^{i\phi_0} = \frac{\langle c, d\rangle}{\langle c, c\rangle} \tag{D.30}$$

where $|H|$ is the magnitude of the sample TF, and ϕ_0 is the phase of the sample TF.

Expressing the sample TF in Eq. D.29 in terms of a, b, c, and d illustrates where bias creeps into the measurement. Assuming that only signals c and d are correlated with each other, and all others are independent such that e.g. $\langle a, b\rangle \to 0$ as the number of samples $n \to \infty$, we can write the sample TF as

$$\begin{aligned}
\hat{H} &= \frac{\langle a+c, b+d\rangle}{\langle a+c, a+c\rangle} \\
\hat{H} &= \frac{\langle a, b\rangle + \langle a, d\rangle + \langle c, b\rangle + \langle c, d\rangle}{\langle a, a\rangle + \langle a, c\rangle + \langle c, a\rangle + \langle c, c\rangle} \\
\hat{H} &= \frac{\langle c, d\rangle}{\langle a, a\rangle + \langle c, c\rangle} \\
\hat{H} &= H\left(\frac{1}{1 + \frac{\langle a,a\rangle}{\langle c,c\rangle}}\right).
\end{aligned} \tag{D.31}$$

The term in parentheses in Eq. D.31 is the bias term in TF estimation that arises when there is noise a on the measured input x. For the rest of this section, we will assume noise on the input is negligible, that is, $\langle a, a\rangle \ll \langle c, c\rangle$ so $\hat{H} \to H$. This is a reasonable assumption when large excitation inputs are possible for TF measurement. Figure D.13 illustrates a case where significant bias is present due to high noise on the input.

The joint probability density function for the sample TF to true TF ratio $\hat{H}/H = re^{i(\phi-\phi_0)}$ is reported from Goodman [200], Eq. 4.81:

$$f(r,\phi) = \frac{n\gamma(1-\gamma^2)^n r}{\pi(1-2\gamma^2 r\cos(\phi-\phi_0)+\gamma^2 r^2)^{n+1}} \tag{D.32}$$

The probability function for the sample TF gain ratio $r = |\hat{H}/H|$ is found by marginalizing over the sample phase ϕ (Goodman Eq. 4.97):

$$f(r) = \frac{2n\gamma^2\Gamma\left(n+\frac{1}{2}\right)(1-\gamma^2)^n r}{\sqrt{\pi}\,\Gamma(n+1)(1+2\gamma^2 r+\gamma^2 r^2)(1-2\gamma^2 r+\gamma^2 r^2)^{n+\frac{1}{2}}} \tag{D.33}$$

$$\times\ {}_2F_1\left(\frac{1}{2}, -n; \frac{1}{2}-n; \frac{1-2\gamma^2 r+\gamma^2 r^2}{1+2\gamma^2 r+\gamma^2 r^2}\right) \tag{D.34}$$

The probability function for the sample TF phase ϕ is found by marginalizing over the sample phase ϕ (Goodman Eq. 4.104):

$$f(\phi) = \frac{(1-\gamma^2)^n}{\pi}\left(1+\frac{n\gamma\cos(\phi-\phi_0)}{(1-\gamma^2\cos^2(\phi-\phi_0))^{n+\frac{1}{2}}}\right. \tag{D.35}$$

$$\left.\times\left(\frac{\sqrt{\pi}\,\Gamma\left(n+\frac{1}{2}\right)}{\Gamma(n+1)} + B_{1-\gamma^2\cos^2(\phi-\phi_0)}\left(n+\frac{1}{2}, \frac{1}{2}\right)\right)\right) \tag{D.36}$$

where B is the Incomplete Beta function and $|\phi-\phi_0| < \pi/2$.

The expected values and variances of Eqs. D.33 and D.35 are not calculable analytically for low sample numbers n, but in certain limits they can be derived from the expected values and variances of the PSD and CSD from which the sample TF is built. These estimates are valid in the regime where the sample coherence has converged to the true coherence, as explored in Section D.7 and shown in Figure D.10. In this regime, the sample PSD, CSD and TF estimates are all "normalized" via the central limit theorem.

Assuming no biases exist and the sample coherence has converged to the true coherence, the expected values for the sample TF gain ratio and phase are

$$\left\langle \left| \frac{\hat{H}}{H} \right| \right\rangle = 1 \tag{D.37}$$

$$\langle \phi \rangle = \phi_0. \tag{D.38}$$

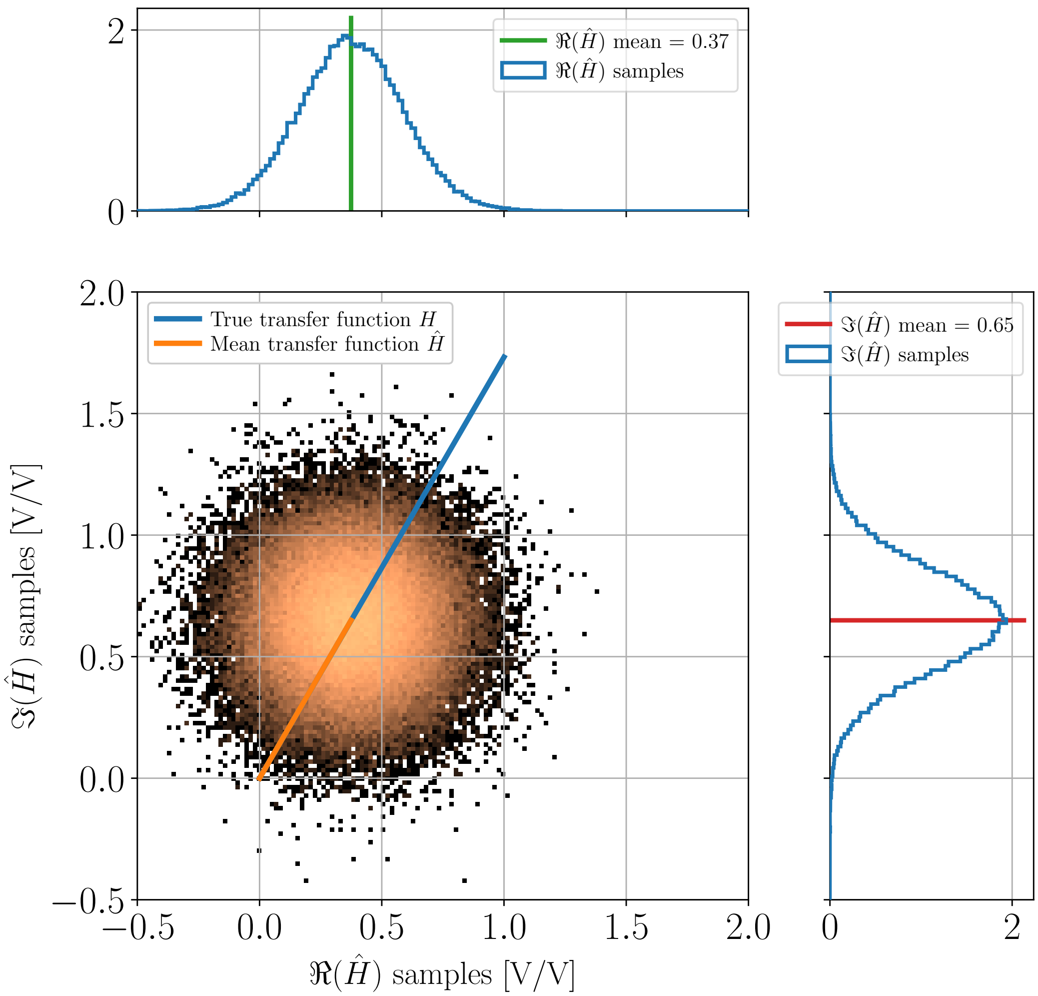

Figure D.13: Histogram of the joint sample transfer function $\hat{H}$, with projections along the real and imaginary axis. This plot illustrates a case where significant uncorrelated noise exists on the measured input x, biasing the sample TF away from its true value H. Equation D.32 describes the contours of the sample TF shown. Referring to Figure D.12, in this example, the true TF $H = 2e^{i\pi/3}$, the coherence $\gamma^2 = 0.205$, the input noise PSD $\langle a, a \rangle = 1.0$ V^2/Hz, the input signal PSD $\langle c, c \rangle = 0.6$ V^2/Hz, making the sample TF bias 0.375. The output noise PSD $\langle b, b \rangle = 2.0$ V^2/Hz.

The variances for the sample TF gain ratio and phase, using Eqs. D.11, D.12, D.18

and D.21, are

$$\sigma^2_{|\hat{H}/H|} = \sigma^2_{\phi} = \frac{\sigma^2_{\langle x,y\rangle}}{|\langle x,y\rangle|^2} - \frac{\sigma^2_{\langle x,x\rangle}}{|\langle x,x\rangle|^2}$$
$$= \frac{1+\gamma^2}{2n\gamma^2} - \frac{1}{n}$$
$$\sigma^2_{|\hat{H}/H|} = \sigma^2_{\phi} = \frac{1-\gamma^2}{2n\gamma^2} \tag{D.39}$$

Eq. D.39 is the oft-cited result for sample TF uncertainty found in Bendat and Piersol, Eq. 9.89 [174]. Figure D.14 illustrates the convergence of the sample TF to the true TF

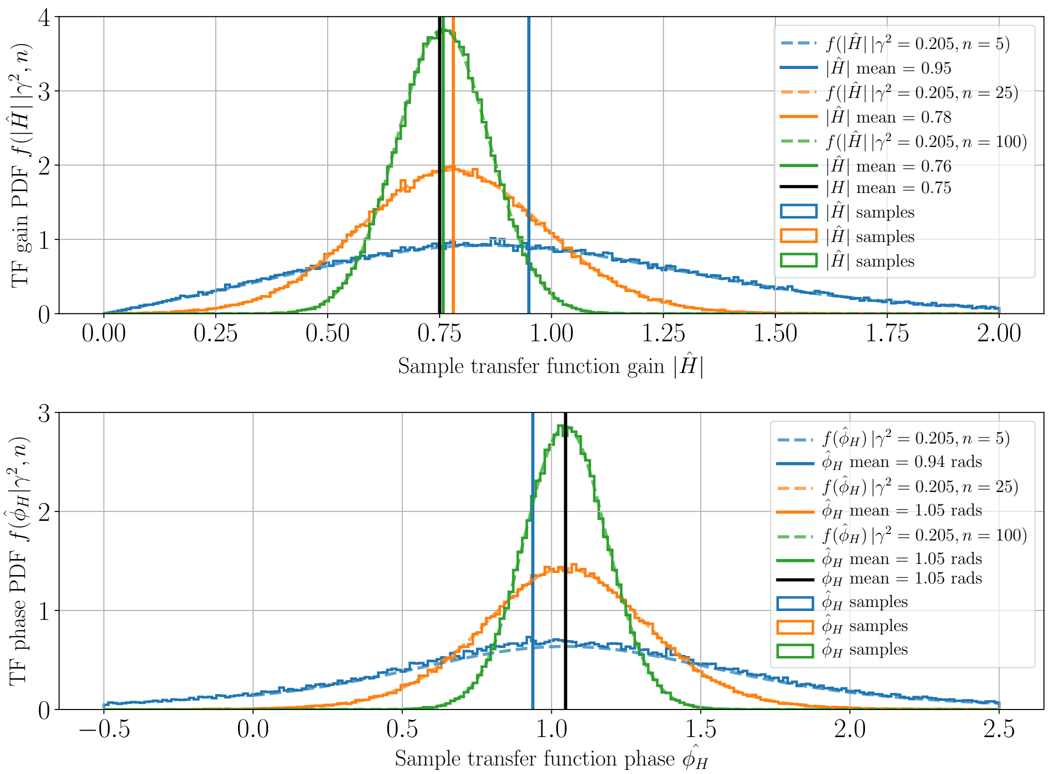

Figure D.14: Probability distributions of the sample transfer function gain $|\hat{H}|$ and phase $\hat{\phi}_H$ for different numbers of samples n. When the number of samples for the TF estimate is too low, in this case $n = 5$, the sample TF expected value $\langle \hat{H} \rangle$ is overestimated and the variance is large. When the number of samples is sufficient, in this case $n = 100$, the sample TF distributions resemble normal distributions with expected values and variances described in Eqs. D.37 and D.38. Equations D.33 and D.35 are the dashed lines for the sample TF shown. In this example, the coherence $\gamma^2 = 0.205$.

Appendix E

LINE INJECTION UNCERTAINTY

Calibration lines are used in LIGO to monitor the DARM response over time, and swept sine transfer functions employ a series of lines to measure the relation between two signals. There are some circumstances where a line lies near or below the the spectral density we want to calibrate, like driving the photon calibrator above DARM at high frequencies $f > 1$ kHz. The SoCal technique may require low lines, and we want to achieve record low uncertainty and bias with those lines.

This appendix will briefly derive and present the Rice distribution which describes a continuous sine wave injection in the presence of Gaussian noise. Relevant results, including the intrinsic uncertainty and bias of an injection, are presented. This appendix will rely on derivations from the distributions in Chapter 6. The results here are used for the line uncertainty and bias calculated for the SoCal calibration lines in Section 4.9.1.

E.1 Basics

When injecting a calibration line into a spectral density, there are multiple important considerations that must be made for the accuracy and precision of the measurement. The first is the injection power. We want the line to be large enough that it quickly dominates the noise of the spectral density and achieves an acceptable SNR, but not so large that the response to the line becomes nonlinear, the line starts exhibiting significant spectral leakage, or the actuation range runs out and the interferometer loses lock. Therefore, some middle ground in the injection power must be found such that the uncertainty and bias levels required are achieved in a timely manner with as low an injection as possible.

First, an important quantity is the *signal-to-noise ratio*, or SNR, of the line with power P_{cal} in the power spectral density $S_P(f)$:

$$\mathrm{SNR} = \frac{P_{signal}}{P_{noise}} = \frac{P_{cal}}{f_b S_P(f)} \tag{E.1}$$

where f_b is the frequency binwidth, or frequency resolution. f_b is the minimum measurable frequency with the data segment given, and for a single average, is equal to

$$f_b = \frac{1}{T} \tag{E.2}$$

where T is the total time of the signal.

The SNR in Eq. E.1 can be increased in three ways:

1. increase the line power P_{cal},
2. lower the noise $S_P(f)$,
3. lower the frequency binwidth f_b by increasing the measurement time T

A power spectral density of Gaussian noise like $S_P(f)$ has units V^2/Hz, and remains constant no matter the measurement time length. A power spectrum $f_b S_P(f)$, on the other hand, has units V^2, and the measurement noise decreases as the same noise power is divided amongst more bins. The calibration line, however, has infinite frequency precision, or at least is assumed to in Eq. E.1. So as the power spectrum $f_b S_P(f)$ of noise decreases with more time T, the calibration line power P_{cal} at exactly the calibration line frequency f_{cal} will remain constant.

For real spectral densities, a window must be used to enforce periodicity of our data. The scale factor between spectrum and spectral density depends on this window, because the window throws away some of the power in the signal. The real scale factor is the *equivalent noise bandwidth*, which depends on the frequency binwidth f_b. For the derivation in Section E.4, we assume a boxcar window function, which is no window, so the equivalent noise bandwidth $= f_b$. Heinzel Chapter 8 [204] discusses windows and the equivalent noise bandwidth.

E.2 Problem setup

We would like to know the uncertainty and bias associated with calibration lines.

When we measured a calibration line in a spectral density, the Gaussian noise power of the density sums with the power in the line:

$$P_{meas} = P_{cal} + f_b S_P(f_{cal}). \tag{E.3}$$

If the SNR is high, then $P_{meas} \approx P_{cal}$ and the effect of noise is negligible. If the SNR is low, then noise contributes significantly to the calibration line, and produces a bias. Figure E.1 illustrates both the high and low SNR case.

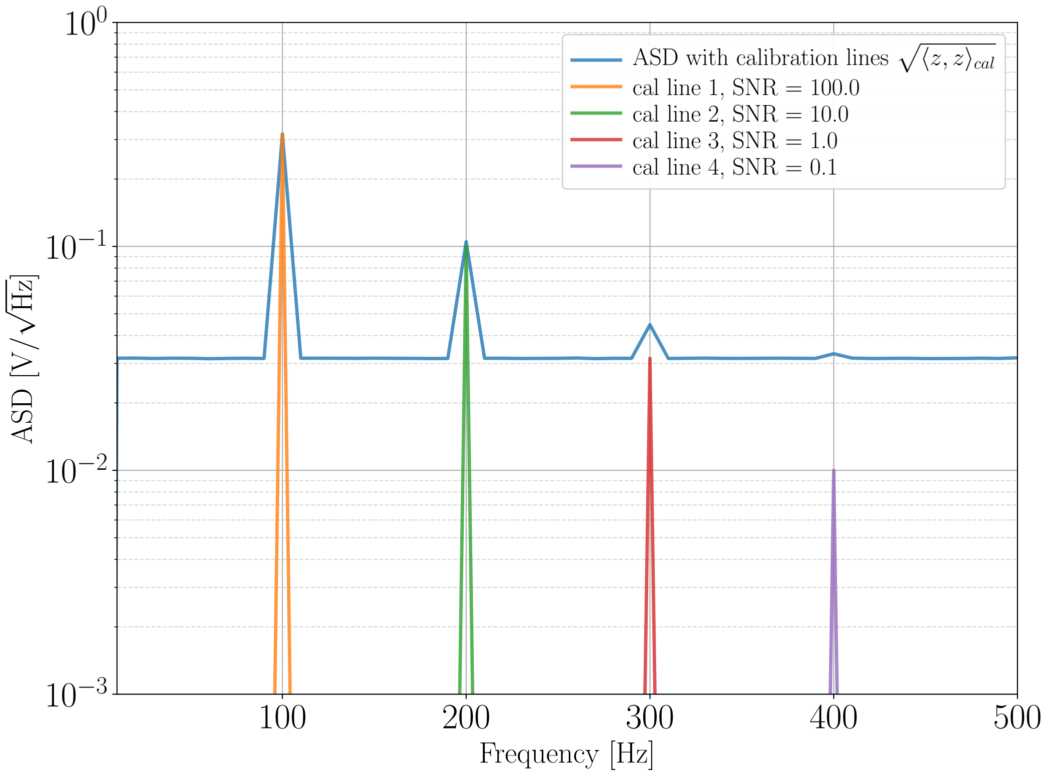

Figure E.1: Four calibration lines f_1, f_2, f_3, f_4 summed into an amplitude spectral density $\sqrt{S_P(f)}$. The high SNR line f_1 shows in the sum spectral density with negligible noise. The SNR ~ 1 line f_3 shows weakly in the sum spectrum, with a significant bias between the line peak and sum peak and noise causing fluctuations on the peak height. To produce this plot, 100000 ASDs were taken using a boxcar window function and averaged together. The frequency binwidth $f_b = 10$ Hz.

E.3 Bias

The bias between the measured peak height and the actual calibration line peak falls out of Eq. E.3:

$$\frac{P_{meas}}{P_{cal}} = 1 + \frac{f_b S_P(f_{cal})}{P_{cal}} = 1 + \frac{1}{\text{SNR}} \tag{E.4}$$

$$\rightarrow b = \frac{1}{\text{SNR}} \tag{E.5}$$

So even with an SNR of 100, you will expect a 1% bias in the calibration line.

The best way to accurately calculate the SNR and bias is to take a long spectral density measurement with no line to measure $S_P(f_{cal})$ very accurately (see Section D.3 and D.4). Then turn on the line, and measure for a time $T = 1/f_b$ seconds such that the minimum acceptable bias is achieved. Note that if you take several averages n, you will have to measure the line for nT seconds.

E.4 Distribution

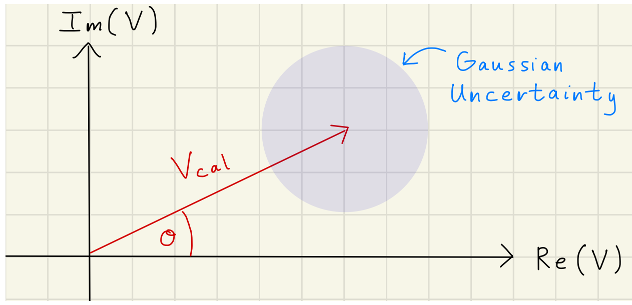

Figure E.2: Diagram of a moderate SNR calibration line V_{cal} with Gaussian noise in the Fourier domain. The noise along the vector will have more of an effect than the noise orthogonal to the line. The Gaussian noise is the same as that derived in Section 6.6.

The uncertainty in the line height depends on the ratio of Gaussian noise relative to the line amplitude, i.e. the SNR. Figure E.2 shows the Fourier domain for the calibration line amplitude V_{cal} with Gaussian noise. We will start the derivation with Figure E.2, and rotate our frame by $-\theta$ so the vector is entirely along with the real axis.

First, we recall our independent Gaussian random variables $\mathcal{A}, \mathcal{B}$ from Section 6.6 to describe the noise:

$$\mathcal{A}, \mathcal{B} \sim \mathcal{N}(0, \sigma) \tag{E.6}$$

where σ is the standard deviation.

Next, we write the distribution of $\mathcal{A} + V_{cal}$. We assume that the line itself is noiseless, yielding

$$\mathcal{A} + V_{cal} \sim \mathcal{N}(\mu, \sigma) \tag{E.7}$$

where μ is the amplitude of V_{cal}.

We care about the measured power P_{meas} calculated from these signals, so we define a random variable $\mathcal{Z}$ such that

$$\mathcal{Z} \sim (\mathcal{A} + V_{cal})^2 + \mathcal{B}^2 \tag{E.8}$$

We recall the zero-mean Gaussian squared $\mathcal{B}^2$ is a general chi-squared distribution with one degree of freedom (Section 6.7):

$$\mathcal{B}^2 \sim \chi_1^2 = \frac{1}{\sqrt{2\pi\sigma^2 b}} e^{-\frac{b}{2\sigma^2}} \qquad b \in [0, \infty) \tag{E.9}$$

The nonzero-mean Gaussian squared follows a noncentral chi-squared distribution with one degree of freedom:

$$(\mathcal{A} + V_{cal})^2 \sim \chi_1'^2 = \frac{1}{\sqrt{2\pi\sigma^2 a}} \cosh\left(\frac{\mu\sqrt{a}}{\sigma^2}\right) e^{-\frac{(a+\mu^2)}{2\sigma^2}} \qquad a \in [0, \infty) \tag{E.10}$$

E.4.1 Power spectral density with line

Now we find the distribution for the power spectral density with a line $\langle z, z \rangle_{cal}$, which follows $\mathcal{Z}$. To find the distribution of $\mathcal{Z}$, we convolve Eqs. E.9 and E.10:

$$\langle z, z \rangle_{cal} \sim p(z) = \int_{-\infty}^{\infty} \chi_1'^2(x) \chi_1^2(z - x) dx \tag{E.11}$$

$$= \frac{1}{2\pi\sigma^2} e^{-\frac{(z+\mu^2)}{2\sigma^2}} \int_0^z \frac{1}{\sqrt{x(z-x)}} \cosh\left(\frac{\mu\sqrt{x}}{\sigma^2}\right) dx \tag{E.12}$$

$$p(z) = \frac{1}{2\sigma^2} e^{-\frac{(z+\mu^2)}{2\sigma^2}} I_0\left(\frac{\mu\sqrt{z}}{\sigma^2}\right) \tag{E.13}$$

where I_0 is the modified Bessel function of the first kind. Eq. E.13 is the probability distribution for how a line with mean power $P_{cal} = \mu^2 f_b$ appears in a power spectral density bin with mean noise power density $2\sigma^2$. If we let $\mu \to 0$, we recover the usual exponential distribution of the power spectral density from Eq. 6.40.

Note that the units here are still V^2/Hz. We recall here that $\sigma^2 = \sigma_x^2 / f_s$ where f_s is the sampling frequency and σ_x^2 is the measured power in V^2 of a white noise time-domain signal. In Eq. E.13, $\mu^2 = P_{mean}/f_b$ is the line power divided by the equivalent noise bandwidth, which in our case is just f_b.

The SNR from Eq. E.1 can be expressed here as

$$\mathrm{SNR} = \frac{\mu^2}{2\sigma^2}. \tag{E.14}$$

E.4.2 Amplitude spectral density with line

Finally, we also derive the distribution for the amplitude spectral density with a line $\sqrt{\langle z, z \rangle_{cal}}$, which follows $\sqrt{\mathcal{Z}}$. This can be done with change of variables $r = \sqrt{z}$, similar to Section 6.9.1, and yields the distribution

$$\sqrt{\langle z, z \rangle_{cal}} \sim p(r) = \frac{r}{\sigma^2} e^{-\frac{(r^2+\mu^2)}{2\sigma^2}} I_0\left(\frac{\mu r}{\sigma^2}\right) \tag{E.15}$$

Eq.E.15 is the Rice distribution. The units here are $\mathrm{V}/\sqrt{\mathrm{Hz}}$.

If we let $\mu \to 0$, we recover the usual Rayleigh distribution amplitude spectral density from Eq. 6.49. If we let $\mathrm{SNR} > 5$, the Rice distribution is well-approximated

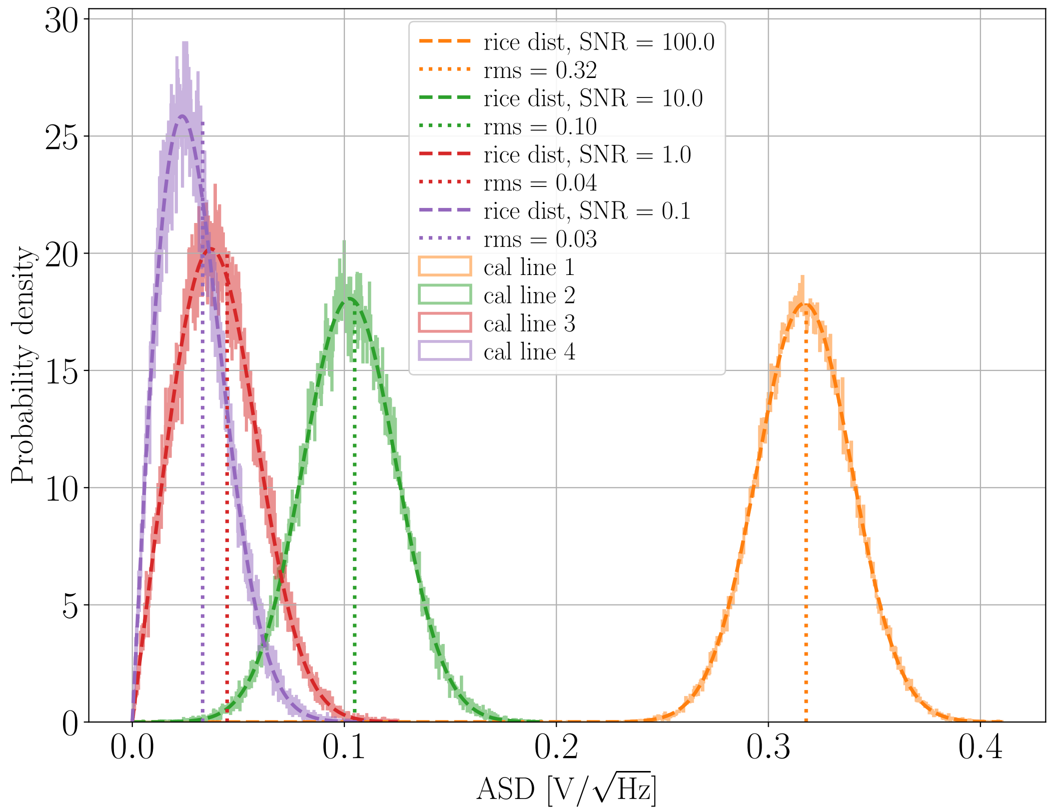

Figure E.3: Histograms of amplitude spectral densities at calibration line frequencies vs estimated rice distributions. 100000 ASDs of the calibration lines plus Gaussian noise from Figure E.1 are plotted as the histograms. The Rice distributions are calculated via Eq. E.15.

by a Gaussian distribution with mean μ and variance σ^2. Figure E.3 shows the Rice distributions of the calibration lines from Figure E.1.

The mean of the Rice distribution in Eq.E.15 is

$$\lambda = \sqrt{\frac{\pi}{2}}\frac{1}{2\sigma}e^{-\frac{\mu^2}{4\sigma^2}}\left[(\mu^2+2\sigma^2)I_0\left(\frac{\mu^2}{4\sigma^2}\right)+\mu^2 I_1\left(\frac{\mu^2}{4\sigma^2}\right)\right] \tag{E.16}$$

The variance is

$$s^2 = \mu^2 + 2\sigma^2 - \lambda^2 \tag{E.17}$$

where λ is the mean from E.16.

The root mean square is

$$\psi = \sqrt{\mu^2 + 2\sigma^2} \tag{E.18}$$

The noncentral fourth-order moment is

$$\mu_4 = \int_0^\infty r^4 p(r)dr = \mu^4 + 8\mu^2\sigma^2 + 8\sigma^4 \tag{E.19}$$

E.4.3 Line height estimation and uncertainty

There are a couple of ways to estimate the true calibration line height $\sqrt{f_b}\mu$, including the method of moments and numerical maximum likelihood estimation. If the $\mathrm{SNR} > 5$, the Rice distribution approximates to a Gaussian, and the method of moments works well to quickly estimate the line height and its uncertainty from the data.

However, some of these techniques fall apart in the regime of low SNR, where the uncertainty would allow unphysical estimates like negative line amplitudes. In those low SNR cases, it will be best to implement a model estimate which allows priors to be incorporated.

E.4.3.1 Estimation via mean squared error

The first and most straightforward way to estimate the line amplitude μ is to take the measured line height and subtract the measured noise.

Quantitatively, we find the sample noncentral second order moment $\hat{\psi}^2$, also known as the mean squared error, from the ASDs. The root mean square $\hat{\psi}$ is the value naturally returned by Welch's method of ASD estimation. Then combine $\hat{\psi}^2$ with a "quiescent" measurement taken with the line off to estimate the average noise $\hat{S}_P(f) = 2\hat{\sigma}^2$:

$$\hat{\mu} = \sqrt{\hat{\psi}^2 - 2\hat{\sigma}^2}. \tag{E.20}$$

To find the uncertainty in the estimate in Eq. E.20, we combine the uncertainty in the mean squared error estimate with the uncertainty in the noise estimate. The variance in the noncentral second order sample moment $\hat{\psi}^2$ estimate is

$$\mathrm{Var}\left[\hat{\psi}^2\right] = \frac{1}{n}\left(\mu_4 - \psi^4\right) \tag{E.21}$$

where n is the number of ASDs taken, and μ_4 is the noncentral fourth-order sample moment. Eq. E.21 is similar to the expression for variance on the sample variance $\mathrm{Var}[\hat{s}^2]$ [239].

The variance in the noise comes from the pure noise measurement with the line off. The variance in the *power* spectral density estimate $2\hat{\sigma}^2$ comes from Eq. D.12:

$$\mathrm{Var}\left[\hat{S}_P(f)\right] = \mathrm{Var}\left[2\hat{\sigma}^2\right] = \frac{4\sigma^4}{n} \tag{E.22}$$

Propagating these uncertainties forward to μ via the usual Taylor series expansion method:

$$\begin{aligned}\mathrm{Var}[\hat{\mu}] &= \frac{\sigma^2(\mu^2 + 2\sigma^2)}{\mu^2 n} &\text{(E.23)}\\ &= \frac{\sigma^2}{n}\left(1 + \frac{1}{\mathrm{SNR}}\right) &\text{(E.24)}\end{aligned}$$

Expressing Eq. E.24 in terms of relative uncertainty using Eqs. E.14:

$$\begin{aligned}\frac{\mathrm{Var}[\hat{\mu}]}{\hat{\mu}^2} &= \frac{\frac{\sigma^2}{n}\left(1 + \frac{1}{\mathrm{SNR}}\right)}{2\sigma^2\mathrm{SNR}}\\ &= \frac{1}{2n\mathrm{SNR}}\left(1 + \frac{1}{\mathrm{SNR}}\right) &\text{(E.25)}\end{aligned}$$

The variance of our estimate of $\hat{\mu}$ decreases linearly with the number of averages n. The absolute variance in Eq. E.24 also flattens out starting at around $\mathrm{SNR} = 1$: absolute uncertainty does not improve much above $\mathrm{SNR} = 5$, but gets much worse for $\mathrm{SNR} < 1$. This suggests that above an SNR of around 5, the experimenter may be better off using data to achieve more averages n rather than increasing the SNR by integrating over more time to reduce the frequency binwidth f_b. Figure E.4 plots the variance of the line estimate as a function of SNR.

Briefly, we overview the regime of unphysical uncertainty. An unphysical result is when the uncertainty of the line height is greater than the line height itself, potentially yielding negative line heights. We can always take more ASDs n to reduce uncertainty to a physical regime, but increasing the SNR helps faster. This can be described in terms of SNR using Eqs. E.14 and E.23

$$\begin{aligned}0 &< \hat{\mu} - \sqrt{\mathrm{Var}[\hat{\mu}]} &\text{(E.26)}\\ 0 &< \sqrt{2\sigma^2\mathrm{SNR}} - \sqrt{\frac{\sigma^2}{n}\left(1 + \frac{1}{\mathrm{SNR}}\right)} &\text{(E.27)}\\ \rightarrow n &> \frac{\mathrm{SNR} + 1}{2\mathrm{SNR}^2} &\text{(E.28)}\end{aligned}$$

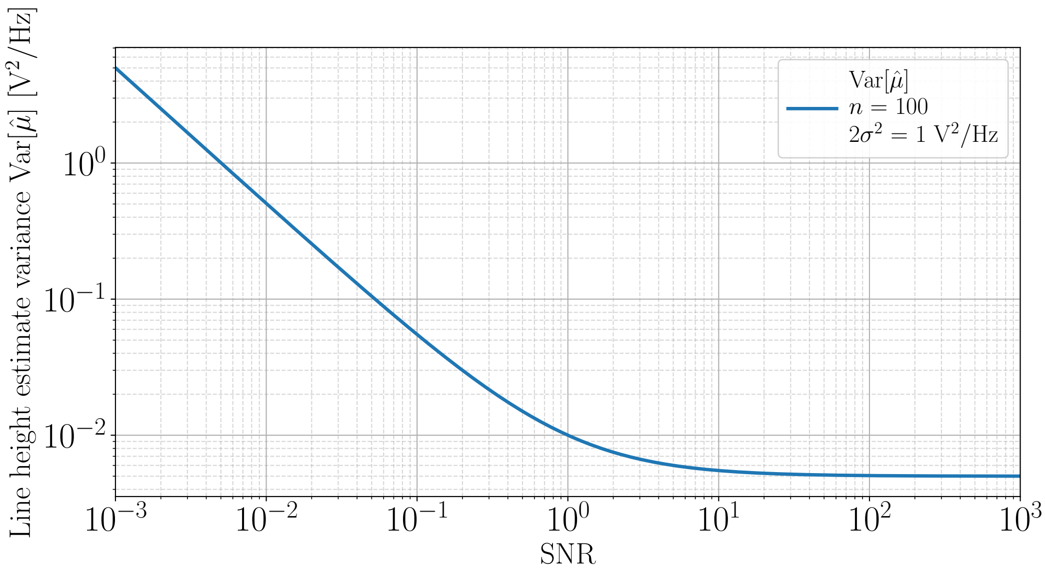

Figure E.4: Variance of the calibration line height estimate $\mathrm{Var}[\hat{\mu}]$ via the method of moments in Eq E.24. The variance in the line height estimate flattens out at high SNR as more line power does not help resolve versus the Gaussian noise further. In this plot, $n = 100$ and $S_P(f) = 2\sigma^2 = 1$ V^2/Hz.

If $\mathrm{SNR} \geq 1$, then $n > 1$ is sufficient for a physical 1σ uncertainty. If $\mathrm{SNR} < 1$, then $n > 1/(2\mathrm{SNR}^2)$ is required.

This method is simple, works well with high SNR lines, and the uncertainties easily calculated and propagated from the noncentral sample moments. However, this method requires two measurements, one each for the line on and off. The state of the interferometer cannot be changing much during this time. Also, the uncertainties are not guaranteed to yield physical results in the very low SNR regime. More samples may be taken to reduce uncertainty in the line height to be physical. The sample moment estimates are not robust to glitches and non-Gaussian noise.

E.4.3.2 Estimation via fourth order moment

If one does not want to perform a "noise-only" measurement with the calibration line off, the noncentral fourth-order sample moment $\hat{\mu}_4$ may be used to estimate the line amplitude $\hat{\mu}$:

$$\hat{\mu} = \left(\hat{\mu}_4 - 2\hat{\psi}^4\right)^{\frac{1}{4}} \tag{E.29}$$

The uncertainties in Eq. E.29 can be propagated similarly to the previous section. The variances associated with Eq. E.29 depend on the noncentral eighth-order sample moments. Analysis suggests the final variance in $\hat{\mu}$ from Eq. E.29 is $\sim 100\times$ higher than Eq. E.23. Again, the uncertainties are not guaranteed to yield physical results in the very low SNR regime, and the sample moment estimates are still susceptible to glitches.

E.5 Future work

The above estimation of calibration line height and uncertainty in the presence of Gaussian noise is useful for quantifying the levels SNR and number of averages required for reasonable uncertainty. For situations where high SNR lines may not be feasible, an MCMC which incorporates priors about the Rice distribution parameters may be necessary.

Future work in this area include examining the uncertainty statistics using medians rather than mean moments. Another useful tool would be code to perform rigorous Bayesian model selection between Rayleigh and Rice distributions to decide whether a line is present or not given the data.

Appendix F

POWER SPECTRAL DENSITY REJECTION FOR GLITCH GATING

Loud, transient glitches in Advanced LIGO data is a problem for spectral analysis. The worst effect glitches can have is when they occur during gravitational wave events, spoiling astrophysical data, although the effects can be partially mitigated [240, 241]. Other issues glitches include causing locklosses, spoiling detector backgrounds [242, 243], interfering with search pipelines and astrophysical parameter estimation [244–246], and spoiling spectral density estimates [247]. This chapter will focus on improving spectral density and transfer function estimates by throwing out glitchy data using the statistics of PSDs.

F.1 Method

Here, we will briefly examine "PSD rejection" which avoids the difficulties associated with median-averaging while still removing glitches from the PSD estimation process. The procedure is

1. Compute all PSDs as normal for Welch's method, say we have n samples
2. Choose a frequency bin f_0 that is sensitive to glitches (in LIGO this is the 20-40 Hz range),
3. Calculate the median ρ of all the PSDs in frequency bin f_0,
4. Choose a rejection threshold r,
5. Reject the PSDs with values higher than the threshold,
6. Use mean-averaging on the rest.

When we calculate the median, we completely characterize the Gaussian noise distribution we care about. We know that a PSD follows an exponential distribution, with a median $\rho = \log(2)\lambda$. By using the cumulative distribution function of the exponential

$$F(x) = 1 - e^{-\frac{x}{\lambda}}, \tag{F.1}$$

we can known exactly how many PSDs we might expect to reject given our rejection threshold r.

For instance, suppose we have $n = 1000$ PSDs. If we choose $r = 10\rho$, then $F(r) = 0.9990$, so we might expect to reject $n(1 - F(r)) \approx 1$ PSD, if there were no glitchy PSDs in the sample. If we reject much more than one, we can safely assume at least some of those PSDs were taken during a glitch.

F.2 Discussion

One might ask, "if I've already characterized my Gaussian noise using the the median, why bother with the rejection?" Indeed, if the PSD is the sole result one cares about, the analysis can stop there, simply correcting from the mean-to-median PSD bias. The above process is most useful for CSD and TF rejection. The individual signals comprising the CSDs may have their PSDs estimated first, and the CSDs associated with the glitchy PSDs can be removed from the analysis before computing the final mean-averaged CSD.

This method works very well for the short-duration, loud, and frequent glitches seen in LIGO data. It avoids the difficulties of median-averaging for CSDs and TFs, and avoids the difficulties of time-domain glitch gating, including data storage costs and spectral leakage due to gating.

This method requires the median to be a decent characterization of the PSD distribution, and so needs a large amount of data, and can miss some quiet glitches that do not exceed the rejection threshold in the bin we've chosen, and can bias the mean-averaged result down if the rejection threshold is too low, cutting off the upper tail of the distribution.

F.3 Example

PSD rejection was used to simplify the analysis of correlated noise in the interferometer, as explored in Chapter 5. For the correlated noise measurement, very long stretches of glitchy data is required to resolve the correlated noise. In the example from Chapter 5, more than 9 hours of data was taken, yielding 67156 sample PSDs and CSDs. At frequency $f_0 = 40$ Hz, I chose a rejection threshold $r = 14\rho$, which should have rejected $67156(1 - F(r)) \approx 4$ CSDs. In fact, 53 CSDs were rejected, meaning about 49 glitches were removed from the analysis.

BIBLIOGRAPHY

[1] B. P. Abbott et al. "Observation of Gravitational Waves from a Binary Black Hole Merger". *Phys. Rev. Lett.* 116 (6 2016), p. 061102.

[2] A. Buikema, **Cahillane, C.**, et al. "Sensitivity and performance of the Advanced LIGO detectors in the third observing run". *Physical Review D* (2020). doi: 10.1103/PhysRevD.102.062003
C.C. was a member of the four-person paper writing team for this Advanced LIGO colloboration-reviewed paper. C.C. produced several of the results presented, including the intensity, frequency, and H1 correlated noise budgets, as well as the arm power measurement. This work is reproduced in Chapter III of this thesis. ISSN: 24700029.

[3] R. Abbott et al. *GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run.* 2020.

[4] B. P. Abbott et al. "GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs". *Phys. Rev. X* 9 (3 2019), p. 031040.

[5] C. Cahillane, J. Betzwieser, et al. "Calibration uncertainty for Advanced LIGO's first and second observing runs". *Phys. Rev. D* 96 (10 2017), p. 102001.

[6] B. P. Abbott et al. "Properties of the Binary Black Hole Merger GW150914". *Phys. Rev. Lett.* 116 (24 2016), p. 241102.

[7] B. P. Abbott et al. "Tests of General Relativity with GW150914". *Phys. Rev. Lett.* 116 (22 2016), p. 221101.

[8] B. P. Abbott et. al. (LIGO Scientific Collaboration). "GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral". *Phys. Rev. Lett.* 119 (16 2017), p. 161101.

[9] B. P. Abbott et al. "GW170817: Measurements of Neutron Star Radii and Equation of State". *Physical Review Letters* (2018). ISSN: 10797114.

[10] B. P. Abbott et al. "A gravitational-wave standard siren measurement of the Hubble constant". *Nature* (2017). ISSN: 14764687.

[11] R. Abbott et al. "GW190814: Gravitational Waves from the Coalescence of a 23 Solar Mass Black Hole with a 2.6 Solar Mass Compact Object". *The Astrophysical Journal* 896.2 (2020), p. L44.

[12] *Gravitational Wave Interferometer Noise Calculator (GWINC).* gwinc git repo.

[13] B. Caron, A. Dominjon, et al. "The VIRGO interferometer for gravitational wave detection". *Nuclear Physics B - Proceedings Supplements* (1997). ISSN: 09205632.

[14] H. Grote, B. Allen, et al. "The status of GEO 600". *Classical and Quantum Gravity* 22.10 (2005), S193–S198.

[15] B. P. Abbott, R. Abbott, et al. "LIGO: the Laser Interferometer Gravitational-Wave Observatory". *Reports on Progress in Physics* 72.7 (2009), p. 076901.

[16] H. G. and. "The GEO 600 status". *Classical and Quantum Gravity* 27.8 (2010), p. 084003.

[17] T. Accadia, F. Acernese, et al. "Virgo: a laser interferometer to detect gravitational waves". *Journal of Instrumentation* 7.03 (2012), P03012–P03012.

[18] J. Aasi, J. Abadie, et al. "Characterization of the LIGO detectors during their sixth science run". *Classical and Quantum Gravity* 32.11 (2015), p. 115012.

[19] J. Aasi, B. P. Abbott, et al. "Advanced LIGO". *Classical and Quantum Gravity* (2015). ISSN: 13616382.

[20] R. Abbott, R. Adhikari, et al. "Advanced LIGO Length Sensing and Control Final Design". *Tech. rep. LIGO-T1000298-v2* (2010).

[21] B. J. Meers. "Recycling in laser-interferometric gravitational-wave detectors". *Physical Review D* (1988). ISSN: 05562821.

[22] A. Buonanno and Y. Chen. "Quantum noise in second generation, signal-recycled laser interferometric gravitational-wave detectors". *Physical Review D* 64.4 (2001).

[23] M. A. Arain and G. Mueller. "Design of the Advanced LIGO recycling cavities". *Optics Express* (2008). ISSN: 1094-4087.

[24] N. A. Robertson et al. "Quadruple suspension design for Advanced LIGO". *Classical and Quantum Gravity* 19.15 (2002), p. 4043.

[25] F. Matichard et al. "Advanced LIGO two-stage twelve-axis vibration isolation and positioning platform. Part 1: Design and production overview". *Precision Engineering* 40 (2015), pp. 273–286.

[26] L. Carbone et al. "Sensors and actuators for the Advanced LIGO mirror suspensions". *Classical and Quantum Gravity* 29.11 (2012), p. 115005.

[27] S. M. Aston et al. "Update on quadruple suspension design for Advanced LIGO". *Classical and Quantum Gravity* 29.23 (2012), p. 235004.

[28] Y. Levin. "Internal thermal noise in the LIGO test masses: A direct approach". *Phys. Rev. D* 57 (2 1998), pp. 659–663.

[29] V. Braginsky and S. Vyatchanin. "Thermodynamical fluctuations in optical mirror coatings". *Physics Letters A* 312.3 (2003), pp. 244–255. ISSN: 0375-9601.

[30] W. Yam, S. Gras, et al. "Multimaterial coatings with reduced thermal noise". *Phys. Rev. D* 91 (4 2015), p. 042002.

[31] S. Gras, H. Yu, et al. “Audio-band coating thermal noise measurement for Advanced LIGO with a multimode optical resonator”. *Physical Review D* (2017). ISSN: 24700029.

[32] S. Gras and M. Evans. “Direct measurement of coating thermal noise in optical resonators”. *Physical Review D* 98.12 (2018).

[33] P. Fritschel, R. Bork, et al. “Readout and control of a power-recycled interferometric gravitational-wave antenna”. *Appl. Opt.* 40.28 (2001), pp. 4988–4998.

[34] F. Seifert, P. Kwee, et al. “Laser power stabilization for second-generation gravitational wave detectors”. *Optics Letters* (2006). ISSN: 0146-9592.

[35] P. Kwee, B. Willke, et al. “Shot-noise-limited laser power stabilization with a high-power photodiode array”. *Optics Letters* (2009). ISSN: 0146-9592.

[36] P. Kwee, C. Bogan, et al. “Stabilized high-power laser system for the gravitational wave detector advanced LIGO”. *Opt. Express* 20.10 (2012), pp. 10617–10634.

[37] C. L. Mueller, M. A. Arain, et al. “The advanced LIGO input optics”. *Review of Scientific Instruments* (2016). ISSN: 10897623.

[38] A. J. Mullavey, B. J. J. Slagmolen, et al. “Arm-length stabilisation for interferometric gravitational-wave detectors using frequency-doubled auxiliary lasers”. *Optics Express* (2012). ISSN: 1094-4087.

[39] K. Izumi, K. Arai, et al. “Multicolor cavity metrology”. *Journal of the Optical Society of America A* (2012). ISSN: 1084-7529.

[40] A. Staley, D. Martynov, et al. “Achieving resonance in the Advanced LIGO gravitational-wave interferometer”. *Classical and Quantum Gravity* (2014). ISSN: 13616382.

[41] T. T. Fricke, N. D. Smith-Lefebvre, et al. “DC readout experiment in Enhanced LIGO”. *Classical and Quantum Gravity* 29.6 (2012), p. 065005.

[42] K. Arai, S. Barnum, et al. *Output Mode Cleaner Design.* Tech. rep. 2013, T1000276.

[43] D. V. Martynov, E. D. Hall, et al. “Sensitivity of the Advanced LIGO detectors at the beginning of gravitational wave astronomy”. *Physical Review D* (2016). ISSN: 24700029.

[44] M. Tse, H. Yu, et al. “Quantum-Enhanced Advanced LIGO Detectors in the Era of Gravitational-Wave Astronomy”. *Physical Review Letters* 123.23 (2019).

[45] S. Hild, M. Abernathy, et al. “Sensitivity studies for third-generation gravitational wave observatories”. *Classical and Quantum Gravity* 28.9 (2011), p. 094013.

[46] K. Danzmann and the LISA study team. "LISA: laser interferometer space antenna for gravitational wave measurements". *Classical and Quantum Gravity* 13.11A (1996), A247–A250.

[47] H. J. Kimble, Y. Levin, et al. "Conversion of conventional gravitational-wave interferometers into quantum nondemolition interferometers by modifying their input and/or output optics". *Phys. Rev. D* 65 (2 2001), p. 022002.

[48] C. M. Caves. "Quantum-Mechanical Radiation-Pressure Fluctuations in an Interferometer". *Phys. Rev. Lett.* 45 (2 1980), pp. 75–79.

[49] C. M. Caves. "Quantum-mechanical noise in an interferometer". *Phys. Rev. D* 23 (8 1981), pp. 1693–1708.

[50] V. B. Braginsky, F. Y. Khalili, et al. *Quantum Measurement.* Cambridge University Press, 1992.

[51] T. Hong, H. Yang, et al. "Brownian thermal noise in multilayer coated mirrors". *Phys. Rev. D* 87 (8 2013), p. 082001.

[52] H. Callen and R. Greene. "On a Theorem of Irreversible Thermodynamics". *Phys. Rev.* 86.5 (1952), pp. 702–710. ISSN: 0031-899X.

[53] R. Kubo. "The fluctuation-dissipation theorem". *Reports on Progress in Physics* 29.1 (1966), pp. 255–284.

[54] P. R. Saulson. "Thermal noise in mechanical experiments". *Phys. Rev. D* 42 (8 1990), pp. 2437–2445.

[55] J. B. Johnson. "Thermal Agitation of Electricity in Conductors". *Phys. Rev.* 32 (1 1928), pp. 97–109.

[56] H. Nyquist. "Thermal Agitation of Electric Charge in Conductors". *Phys. Rev.* 32 (1 1928), pp. 110–113.

[57] E. D. Hall. PhD thesis. California Institute of Technology, 2017.

[58] N. Nakagawa, A. M. Gretarsson, et al. "Thermal noise in half-infinite mirrors with nonuniform loss: A slab of excess loss in a half-infinite mirror". *Phys. Rev. D* 65 (10 2002), p. 102001.

[59] T. Chalermsongsak, F. Seifert, et al. "Broadband measurement of coating thermal noise in rigid Fabry–Pérot cavities". *Metrologia* 52.1 (2014), pp. 17–30.

[60] S. Biscans, S. Gras, et al. "Suppressing parametric instabilities in LIGO using low-noise acoustic mode dampers". *Phys. Rev. D* 100 (12 2019), p. 122003.

[61] G. M. Harry, M. R. Abernathy, et al. "Titania-doped tantala/silica coatings for gravitational-wave detection". *Classical and Quantum Gravity* 24.2 (2006), pp. 405–415.

[62] M. Granata, A. Amato, et al. "Amorphous optical coatings of present gravitational-wave interferometers". *Classical and Quantum Gravity* 37.9 (2020), p. 095004.

[63] J. Steinlechner. “Development of mirror coatings for gravitational-wave detectors”. *Phil. Trans. R. Soc. A.* (376 2018), p. 20170282.

[64] F. Matichard, B. Lantz, et al. “Seismic isolation of Advanced LIGO: Review of strategy, instrumentation and performance”. *Classical and Quantum Gravity* 32.18 (2015), p. 185003.

[65] S. Wen, R. Mittleman, et al. “Hydraulic external pre-isolator system for LIGO”. *Classical and Quantum Gravity* 31.23 (2014), p. 235001.

[66] S. Biscans, J. Warner, et al. “Control strategy to limit duty cycle impact of earthquakes on the LIGO gravitational-wave detectors”. *Classical and Quantum Gravity* 35.5 (2018), p. 055004.

[67] J. Harms. “Terrestrial Gravity Fluctuations”. *Living Reviews in Relativity* 18.1 (2015), p. 3. ISSN: 1433-8351.

[68] J. C. Driggers, J. Harms, et al. “Subtraction of Newtonian noise using optimized sensor arrays”. *Phys. Rev. D* 86 (10 2012), p. 102001.

[69] M. Coughlin, N. Mukund, et al. “Towards a first design of a Newtonian-noise cancellation system for Advanced LIGO”. *Classical and Quantum Gravity* 33.24 (2016), p. 244001.

[70] M. W. Coughlin, J. Harms, et al. “Implications of Dedicated Seismometer Measurements on Newtonian-Noise Cancellation for Advanced LIGO”. *Phys. Rev. Lett.* 121 (22 2018), p. 221104.

[71] P. R. Saulson. “Terrestrial gravitational noise on a gravitational wave antenna”. *Phys. Rev. D* 30 (4 1984), pp. 732–736.

[72] S. A. Hughes and K. S. Thorne. “Seismic gravity-gradient noise in interferometric gravitational-wave detectors”. *Phys. Rev. D* 58 (12 1998), p. 122002.

[73] J. Harms, E. L. Bonilla, et al. “Observation of a potential future sensitivity limitation from ground motion at LIGO Hanford”. *Phys. Rev. D* 101 (10 2020), p. 102002.

[74] R. W. Drever, J. L. Hall, et al. “Laser phase and frequency stabilization using an optical resonator”. *Applied Physics B Photophysics and Laser Chemistry* (1983). ISSN: 14320649.

[75] K. A. Kuns. “Future Networks of Gravitational Wave Detectors: Quantum Noise and Space Detectors”. English. Copyright - Database copyright ProQuest LLC; ProQuest does not claim copyright in the individual underlying works; Last updated - 2020-12-01. PhD thesis. 2019, p. 158. ISBN: 978-1-392-10049-3.

[76] K. Izumi and D. Sigg. *Frequency Response of the aLIGO Interferometer: part1.* Tech. rep. 2016, T1500325.

[77] D. V. Martynov. “Lock Acquisition and Sensitivity Analysis of Advanced LIGO Interferometers”. en. PhD thesis. 2015.

[78] D. Sigg. “Frequency Response of the LIGO Interferometer”. *Tech. rep. LIGO-T970084-x0* (1997).

[79] J. Driggers. *Length feedforward calculations.* Tech. rep. 2016, T1600504.

[80] A. N. Staley. PhD thesis. Columbia University, 2015.

[81] K. Arai, M. Ando, et al. “New signal extraction scheme with harmonic demodulation for power-recycled Fabry–Perot–Michelson interferometers”. *Physics Letters A* 273.1 (2000), pp. 15–24. ISSN: 0375-9601.

[82] J. G. Rollins. “Distributed state machine supervision for long-baseline gravitational-wave detectors”. *Review of Scientific Instruments* 87.9 (2016), p. 094502.

[83] T. Hardwick. “High Power and Optomechanics in Advanced LIGO Detectors”. PhD thesis. Louisiana State University, 2019.

[84] J. Aasi et al. “Enhanced sensitivity of the LIGO gravitational wave detector by using squeezed states of light”. *Nature Photonics* 7.8 (2013), pp. 613–619. ISSN: 1749-4893.

[85] J. Abadie et al. “A gravitational wave observatory operating beyond the quantum shot-noise limit”. *Nature Physics* 7.12 (2011), pp. 962–965. ISSN: 1745-2481.

[86] H. Grote, K. Danzmann, et al. “First Long-Term Application of Squeezed States of Light in a Gravitational-Wave Observatory”. *Phys. Rev. Lett.* 110 (18 2013), p. 181101.

[87] Á. Fernández-Galiana, L. McCuller, et al. “A compact actively damped vibration isolation platform for optical experiments in ultra-high vacuum”. 2019.

[88] E. Oelker, G. Mansell, et al. “Ultra-low phase noise squeezed vacuum source for gravitational wave detectors”. *Optica* 3.7 (2016), pp. 682–685.

[89] H. Yu, L. McCuller, et al. “Quantum correlations between light and the kilogram-mass mirrors of LIGO”. *Nature* 583.7814 (2020), pp. 43–47.

[90] R. Dolesi, M. Hueller, et al. “Brownian force noise from molecular collisions and the sensitivity of advanced gravitational wave observatories”. *Phys. Rev. D* 84 (6 2011), p. 063007.

[91] M. Evans, P. K. Fritschel, et al. *Gas Damping Monte Carlo.* Tech. rep. LIGO-T0900582. 2011.

[92] J. C. Driggers, S. Vitale, et al. “Improving astrophysical parameter estimation via offline noise subtraction for Advanced LIGO”. *Physical Review D* (2019). ISSN: 24700029.

[93] P. B. Covas, A. Effler, et al. “Identification and mitigation of narrow spectral artifacts that degrade searches for persistent gravitational waves in the first two observing runs of Advanced LIGO”. *Physical Review D* (2018). ISSN: 24700029.

[94] L. Sun, E. Goetz, et al. "Characterization of systematic error in Advanced LIGO calibration". *Classical and Quantum Gravity* (2020). ISSN: 13616382.

[95] L. S. Finn and D. F. Chernoff. "Observing binary inspiral in gravitational radiation: One interferometer". *Phys. Rev. D* 47 (1993), pp. 2198–2219.

[96] L. S. Finn. "Binary inspiral, gravitational radiation, and cosmology". *Phys. Rev. D* 53 (6 1996), pp. 2878–2894.

[97] H.-Y. Chen, D. E. Holz, et al. "Distance measures in gravitational-wave astrophysics and cosmology". 2017.

[98] *GraceDB—–O3 Gravitational-Wave Candidate Event Database.*

[99] *LIGO Scientific Collaboration and VIRGO Collaboration event information.*

[100] C. Cahillane. *SRCL Dither to Arm Power Measurement.* Tech. rep. 2019, T1900529.

[101] K. Izumi, D. Sigg, et al. *Frequency Response of the aLIGO Interferometer: part3.* Tech. rep. 2016, T1500559.

[102] G. Mansell, C. Cahillane, et al. *ALS COMM locked - IMC length locked to x-arm.* LHO alog 42959.

[103] D. Sigg. *New ALS Parameters.* LHO alog 42384.

[104] S. Dwyer, A. Staley, et al. *measurement of IR noise with COMM locked.* LHO alog 10511.

[105] C. Cahillane and D. Sigg. *Out of loop ALS COMM frequency measurement via IR.* LHO alog 43119.

[106] C. Cahillane, S. Dwyer, et al. *Out of Loop ALS COMM frequency measurement - Take Two.* LHO alog 43214.

[107] B. Willke, S. Brozek, et al. "Frequency stabilization of a monolithic Nd:YAG ring laser by controlling the power of the laser-diode pump source". *Optics Letters* (2000). ISSN: 0146-9592.

[108] K. Arai. *OLD/NEW EOM configuration comparison.* LHO alog 41435.

[109] D. Sigg, C. Vorvick, et al. *Changing the power on the IMC REFL PD.* LHO alog 49067.

[110] S. Dwyer, E. Hall, et al. *Another look at frequency noise.* LHO alog 30549.

[111] *Large Area InGaAs PIN Photodiodes.* C30619, C30641, C30642, C30665 and C30723 Series. Excelitas Technologies. 2011.

[112] C. Cahillane and K. Arai. *IMC REFL and REFL A/B shot and dark noise levels.* LHO alog 46483.

[113] C. Whittle and J. Kissel. *IMC PDH daughter board functionality.* LHO alog 29735.

[114] K. Somiya, Y. Chen, et al. "Frequency noise and intensity noise of next-generation gravitational-wave detectors with RF/DC readout schemes". *Phys. Rev. D* 73 (12 2006), p. 122005.

[115] G. Mansell, C. Cahillane, et al. *Demodulating injected lines during ring heater steps so far*. LHO alog 47968.

[116] C. Cahillane. *Frequency Noise Injection while Changing RF Sideband Amplitudes*. LHO alog 46243.

[117] B. Willke, P. King, et al. *Pre-Stabilized Laser Design Requirements*. Tech. rep. 2011, T050036.

[118] P. King, G. Vajente, et al. *ISS measurements*. LHO alog 43838.

[119] R. L. Ward. PhD thesis. California Institute of Technology, 2010.

[120] B. P. Abbott et. al. (LIGO Scientific Collaboration). "Calibration of the Advanced LIGO detectors for the discovery of the binary black-hole merger GW150914". *Phys. Rev. D* 95 (6 2017).
doi: 10.1103/PhysRevD.95.062003
C.C. produced the calibration uncertainty budgets for Advanced LIGO's first detection of gravitational waves, GW150914.

[121] P. Fulda, D. Brown, et al. "Simulation investigation of mode hopping in the H1 signal recycling cavity". *Tech. rep. LIGO-T1500230-v3* (2015).

[122] A. Effler, D. Brown, et al. *36W lock*. LHO alog 47606.

[123] C. Cahillane. *DARM plant detuning changes sign with SR3 heater engaged.* LHO alog 47604.

[124] C. Cahillane and R. Adhikari. *Contrast Defect Measurement at 2 watts input.* LHO alog 46142.

[125] A. Buonanno and Y. Chen. "Signal recycled laser-interferometer gravitational-wave detectors as optical springs". *Phys. Rev. D* 65 (4 2002), p. 042001.

[126] J. Kissel, S. Dwyer, et al. *Calibration Measurements: Analysis of Measuring Sensing Function Under A2L (ETMY Spot Position) and ASC (DHARD) Changes*. LHO alog 50498.

[127] H. Yu and S. Dwyer. *Squeezer optical loss budget*. LHO squeezing gitlab.

[128] G. Mansell, C. Cahillane, et al. *SR3 heater steps down*. LHO alog 48406.

[129] C. Cahillane. *DARM calibration time constant*. LHO alog 52139.

[130] A. Brooks. *SR3 heater test - 2W into SR3. dR = -3.05mm per Watt*. LLO alog 27262.

[131] S. Ballmer and V. Srivastava. *SRC Gouy Phase measurement (without even locking SRMI)*. LHO alog 52639.

[132] S. Ballmer and V. Srivastava. *SRC Gouy phase with SR3 ring heater off: 25.25deg +-0.06deg (statistical) +-1.8deg (systematic)*. LHO alog 52658.

[133] C. Cahillane. *SRC gouy phase modeling*. LHO alog 52731.

[134] D. D. Brown and A. Freise. *Finesse*. The software and source code is available at `http://www.gwoptics.org/finesse`. 2014.

[135] D. D. Brown, P. Jones, et al. "Pykat: Python package for modelling precision optical interferometers". *SoftwareX* 12 (2020), p. 100613.

[136] B. P. Abbott et al. "Calibration of the Advanced LIGO detectors for the discovery of the binary black-hole merger GW150914". *Phys. Rev. D* 95 (6 2017), p. 062003.

[137] B. Allen. *LIGO Calibration Accuracy*. Tech. rep. 1996, T960189.

[138] L. Lindblom. "Optimal calibration accuracy for gravitational-wave detectors". *Physical Review D* 80.4 (2009).

[139] J. Veitch et al. "Parameter estimation for compact binaries with ground-based gravitational-wave observations using the LALInference software library". *Physical Review D - Particles, Fields, Gravitation and Cosmology* (2015). ISSN: 15502368.

[140] E. D. Hall, **Craig Cahillane**, et al. "Systematic calibration error requirements for gravitational-wave detectors via the Cramér–Rao bound". *Classical and Quantum Gravity* 36.20 (2019).
doi: 10.1088/1361-6382/ab368c
C.C. verified the Jacobian numerical results of this paper analytically.

[141] S. Vitale, W. Del Pozzo, et al. "Effect of calibration errors on Bayesian parameter estimation for gravitational wave signals from inspiral binary systems in the advanced detectors era". *Physical Review D - Particles, Fields, Gravitation and Cosmology* (2012). ISSN: 15507998.

[142] S. Vitale, W. D. Pozzo, et al. "Effect of calibration errors on Bayesian parameter estimation for gravitational wave signals from inspiral binary systems in the advanced detectors era: Further investigations". *Journal of Physics: Conference Series* 484 (2014), p. 012026.

[143] S. Vitale, C.-J. Haster, et al. *physiCal: A physical approach to the marginalization of LIGO calibration uncertainties*. 2020.

[144] Y. Huang, C.-J. Haster, et al. *Statistical and systematic uncertainties in extracting the source properties of neutron star - black hole binaries with gravitational waves*. 2020.

[145] S. Vitale. "Three observational differences for binary black holes detections with second- and third-generation gravitational-wave detectors". *Physical Review D* (2016). ISSN: 24700029.

[146] B. P. A. et. al. "Exploring the sensitivity of next generation gravitational wave detectors". *Classical and Quantum Gravity* 34.4 (2017), p. 044001.

[147] V. Morozova, D. Radice, et al. "The Gravitational Wave Signal from Core-collapse Supernovae". *The Astrophysical Journal* 861.1 (2018), p. 10.

[148] B. P. Abbott et al. "The rate of binary black hole mergers inferred from advanced LIGO observations surrounding GW150914". *Astrophys. J. Lett.* 833.1 (2016), p. L1.

[149] B. P. Abbott, R. Abbott, et al. "SUPPLEMENT: The rate of binary black hole mergers inferred from advanced LIGO observations surrounding GW150914". *The Astrophysical Journal Supplement Series* 227.2 (2016), p. 14.

[150] A. B. P. et al. "Binary Black Hole Population Properties Inferred from the First and Second Observing Runs of Advanced LIGO and Advanced Virgo". *The Astrophysical Journal* 882.2 (2019), p. L24.

[151] B. P. Abbott et al. "GW170104: Observation of a 50-Solar-Mass Binary Black Hole Coalescence at Redshift 0.2". *Phys. Rev. Lett.* 118 (22 2017), p. 221101.

[152] B. P. Abbott et al. "Tests of general relativity with the binary black hole signals from the LIGO-Virgo catalog GWTC-1". *Physical Review D* (2019). ISSN: 24700029.

[153] M. Isi, K. Chatziioannou, et al. "Hierarchical Test of General Relativity with Gravitational Waves". *Physical Review Letters* (2019). ISSN: 10797114.

[154] N. Yunes, K. Yagi, et al. "Theoretical physics implications of the binary black-hole mergers GW150914 and GW151226". *Phys. Rev. D* 94 (8 2016), p. 084002.

[155] K. Chamberlain and N. Yunes. "Theoretical physics implications of gravitational wave observation with future detectors". *Phys. Rev. D* 96 (8 2017), p. 084039.

[156] B. P. Abbott et al. "First search for gravitational waves from known pulsars with Advanced LIGO". *Astrophys. J.* 839.1 (2017), p. 12.

[157] B. P. Abbott et al. "First low frequency Einstein at Home all-sky search for continuous gravitational waves in advanced LIGO data". *arxiv:1707.02669* ().

[158] B. P. Abbott et al. "All-sky search for continuous gravitational waves from isolated neutron stars using Advanced LIGO O2 data". *Physical Review D* (2019). ISSN: 24700029.

[159] B. P. Abbott et al. "Upper Limits on the Stochastic Gravitational-Wave Background from Advanced LIGO's First Observing Run". *Phys. Rev. Lett.* 118 (12 2017), p. 121101.

[160] B. P. Abbott et al. "Directional Limits on Persistent Gravitational Waves from Advanced LIGO's First Observing Run". *Phys. Rev. Lett.* 118 (12 2017), p. 121102.

[161] T. Callister, M. Fishbach, et al. "Shouts and Murmurs: Combining Individual Gravitational-wave Sources with the Stochastic Background to Measure the History of Binary Black Hole Mergers". *The Astrophysical Journal* 896.2 (2020), p. L32.

[162] D. F. Chernoff and L. S. Finn. "Gravitational radiation, inspiraling binaries, and cosmology". *Astrophys. J.* 411 (1993), pp. L5–L8.

[163] H. Y. Chen, M. Fishbach, et al. "A two per cent Hubble constant measurement from standard sirens within five years". *Nature* (2018). ISSN: 14764687.

[164] N. Jackson. *The hubble constant.* 2015.

[165] A. G. Riess, S. Casertano, et al. *Cosmic Distances Calibrated to 1% Precision with Gaia EDR3 Parallaxes and Hubble Space Telescope Photometry of 75 Milky Way Cepheids Confirm Tension with LambdaCDM.* 2020.

[166] N. Aghanim et al. "Planck 2018 results: VI. Cosmological parameters". *Astronomy and Astrophysics* (2020). ISSN: 14320746.

[167] M. Soares-Santos et al. "First Measurement of the Hubble Constant from a Dark Standard Siren using the Dark Energy Survey Galaxies and the LIGO Virgo Binary Black-hole Merger GW170814". *The Astrophysical Journal* 876.1 (2019), p. L7.

[168] S. Vitale and H. Y. Chen. "Measuring the Hubble Constant with Neutron Star Black Hole Mergers". *Physical Review Letters* (2018). ISSN: 10797114.

[169] S. Khan, S. Husa, et al. "Frequency-domain gravitational waves from non-precessing black-hole binaries. II. A phenomenological model for the advanced detector era". *Physical Review D* (2016). ISSN: 24700029.

[170] C. M. Biwer, C. D. Capano, et al. "PyCBC inference: a python-based parameter estimation toolkit for compact binary coalescence signals". *Publications of the Astronomical Society of the Pacific* (2019). ISSN: 00046280.

[171] B. F. Schutz. "Determining the Hubble constant from gravitational wave observations". *Nature* (1986). ISSN: 00280836.

[172] D. E. Holz, S. A. Hughes, et al. "Measuring cosmic distances with standard sirens". *Physics Today* 71.12 (2018), pp. 34–40.

[173] V. Varma, S. E. Field, et al. "Surrogate models for precessing binary black hole simulations with unequal masses". *Phys. Rev. Research* 1 (3 2019), p. 033015.

[174] J. S. Bendat and A. G. Piersol. *Random Data: Analysis and Measurement Procedures.* 4th. Vol. 729. John Wiley & Sons, New York, 1966.

[175] B. S. Sheard, M. B. Gray, et al. "Observation and characterization of an optical spring". *Phys. Rev. A* 69 (5 2004), p. 051801.

[176] M. Rakhmanov, R. Savage, et al. "Dynamic resonance of light in Fabry-Perot cavities". *Phys. Lett. A* 305.5 (2002), pp. 239–244.

[177] M. Rakhmanov. "Response of test masses to gravitational waves in the local Lorentz gauge". *Physical Review D* 71.8 (2005).

[178] M. Rakhmanov, J. D. Romano, et al. "High-frequency corrections to the detector response and their effect on searches for gravitational waves". *Classical and Quantum Gravity* 25.18 (2008), p. 184017.

[179] D. Tuyenbayev et al. "Improving LIGO calibration accuracy by tracking and compensating for slow temporal variations". *Classical and Quantum Gravity* 34.1 (2017), p. 015002.

[180] S. Karki et al. "The Advanced LIGO photon calibrators". *Rev. Sci. Instrum.* 87.11 (2016), pp. 114503–114503.

[181] D. Bhattacharjee, Y. Lecoeuche, et al. "Fiducial displacements with improved accuracy for the global network of gravitational wave detectors". *Classical and Quantum Gravity* 38.1 (2020), p. 015009.

[182] J. D. Hunter. "Matplotlib: A 2D graphics environment". *Comput. Sci. Eng.* 9.3 (2007), pp. 90–95.

[183] D. Foreman-Mackey et al. "emcee: the MCMC hammer". *Publications of the Astronomical Society of the Pacific* 125.925 (2013), p. 306.

[184] D. Foreman-Mackey. "corner.py: Scatterplot matrices in Python". *The Journal of Open Source Software* 24 (2016).

[185] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. Vol. 1. MIT Press, 2006.

[186] F. Pedregosa et al. "Scikit-learn: Machine Learning in Python". *Journal of Machine Learning Research* 12.Oct (2011), pp. 2825–2830.

[187] B. P. Abbott et al. "Binary Black Hole Mergers in the First Advanced LIGO Observing Run". *Phys. Rev. X* 6 (4 2016), p. 041015.

[188] D. Estevez, B. Lieunard, et al. "First tests of a Newtonian calibrator on an interferometric gravitational wave detector". *Classical and Quantum Gravity* 35.23 (2018), p. 235009.

[189] D. Estevez, B. Mours, et al. *Newtonian calibrator tests during the Virgo O3 data taking*. 2020.

[190] S. Hild, H. Grote, et al. "DC-readout of a signal-recycled gravitational wave detector". *Classical and Quantum Gravity* 26.5 (2009), p. 055012.

[191] D. V. Martynov, V. V. Frolov, et al. "Quantum correlation measurements in interferometric gravitational-wave detectors". *Phys. Rev. A* 95 (4 2017), p. 043831.

[192] K. Izumi. *Time domain implementation of DCPD cross correlation.* Tech. rep. 2017, T1700131.

[193] B. Allen, W. G. Anderson, et al. "FINDCHIRP: An algorithm for detection of gravitational waves from inspiraling compact binaries". *Phys. Rev. D* 85 (12 2012), p. 122006.

[194] E. Hall and D. Sigg. *Balanced homodyne detection with a null channel.* Tech. rep. https://dcc.ligo.org/LIGO-P1800012. 2018, P1800012.

[195] L. McCuller. *Squeezing effect on OMC PD Cross Correlation.* Tech. rep. https://dcc.ligo.org/LIGO-T1800110. 2018, T1800110.

[196] C. Gerry and P. Knight. *Introductory Quantum Optics.* Cambridge University Press, 2004.

[197] V. Frolov, A. Buikema, et al. *OMC PD cross spectrum with squeezing.* LLO alog 44882.

[198] P. Fritschel, M. Evans, et al. "Balanced homodyne readout for quantum limited gravitational wave detectors". *Opt. Express* 22.4 (2014), pp. 4224–4234.

[199] R. B. Blackman and J. W. Tukey. *The Measurement of Power Spectra from the Point of View of Communications Engineering.* 1958.

[200] N. R. Goodman. *On the joint estimation of the spectra, cospectrum and quadrature spectrum of a two-dimensional stationary Gaussian process.* English. Ed. by New York University and Princeton University. Scientific paper (New York University. Engineering Statistics Laboratory) ; no. 10. New York: Research Division, College of Engineering, New York University, 1957.

[201] N. R. Goodman. "Statistical Analysis Based on a Certain Multivariate Complex Gaussian Distribution (An Introduction)". *The Annals of Mathematical Statistics* 34.1 (1963), pp. 152–177.

[202] N. R. Goodman. *Measurement of matrix frequency response functions and multiple coherence functions.* Tech. rep. 1965.

[203] L. D. Enochson and N. R. Goodman. *GAUSSIAN APPROXIMATIONS TO THE DISTRIBUTION OF SAMPLE COHERENCE.* Tech. rep. 1965.

[204] G. Heinzel, A. Rüdiger, et al. "Spectrum and spectral density estimation by the Discrete Fourier transform (DFT), including a comprehensive list of window functions and some new flat-top windows". *Max Plank Inst* 12 (2002).

[205] F. W. J. Olver, A. B. O. Daalhuis, et al. *NIST Digital Library of Mathematical Functions.* 2020.

[206] F. McNolty. "Some Probability Density Functions and Their Characteristic Functions". *Mathematics of Computation* 27 (123 1973).

[207] D. Huppenkothen and M. Bachetti. "On the Statistical Properties of Cospectra". *The Astrophysical Journal Supplement Series* 236.1 (2018), p. 13.

[208] L. McCuller, C. Whittle, et al. "Frequency-Dependent Squeezing for Advanced LIGO". *Physical Review Letters* (2020). ISSN: 10797114.

[209] L. McCuller and L. Barsotti. *Design Requirement Document of the A+ filter cavity and relay optics for frequency dependent squeezing*. Tech. rep. 2020, T1800447–v7.

[210] *pytickle*. pytickle git repo.

[211] P. R. Saulson. *Fundamentals of Interferometric Gravitational Wave Detectors*. WORLD SCIENTIFIC, 1994.

[212] J. B. Camp, H. Yamamoto, et al. "Analysis of light noise sources in a recycled Michelson interferometer with Fabry–Perot arms". *Journal of the Optical Society of America A* 17.1 (2000), p. 120.

[213] K. Izumi and D. Sigg. "Advanced LIGO: length sensing and control in a dual recycled interferometric gravitational wave antenna". *Classical and Quantum Gravity* 34.1 (2017), p. 015001.

[214] M. W. Regehr. "Signal extraction and control for an interferometric gravitational wave detector". en. PhD thesis. 1995.

[215] G. Heinzel. "Advanced optical techniques for laser-interferometric gravitational-wave detectors". PhD thesis. Hannover: Universität, 1999.

[216] C. M. Caves and B. L. Schumaker. "New formalism for two-photon quantum optics. I. Quadrature phases and squeezed states". *Phys. Rev. A* 31 (5 1985), pp. 3068–3092.

[217] B. L. Schumaker and C. M. Caves. "New formalism for two-photon quantum optics. II. Mathematical foundation and compact notation". *Phys. Rev. A* 31 (5 1985), pp. 3093–3111.

[218] J. Harms, Y. Chen, et al. "Squeezed-input, optical-spring, signal-recycled gravitational-wave detectors". *Physical Review D* (2003). ISSN: 05562821.

[219] Y. Enomoto, K. Nagano, et al. "Standard quantum limit of angular motion of a suspended mirror and homodyne detection of a ponderomotively squeezed vacuum field". *Physical Review A* (2016). ISSN: 24699934.

[220] A. Buonanno and Y. Chen. "Scaling law in signal recycled laser-interferometer gravitational-wave detectors". *Physical Review D* 67.6 (2003).

[221] C. W. Misner, K. S. Thorne, et al. "MTW-Gravitation". *Gravitation*. 2018. ISBN: 9780691177793.

[222] H. Kogelnik and T. Li. "Laser Beams and Resonators". *Applied Optics* (1966). ISSN: 0003-6935.

[223] A. Siegman. *Lasers.* University Science Books, 1986. ISBN: 9780935702118.

[224] D. Z. Anderson. "Alignment of resonant optical cavities". *Applied Optics* (1984). ISSN: 0003-6935.

[225] K. Arai. *On the accumulated round-trip Gouy phase shift for a general optical cavity.* Tech. rep. LIGO-T1300189. 2013.

[226] J. C. Driggers. "Noise Cancellation for Gravitational Wave Detectors". en. PhD thesis. 2015.

[227] R. Abbott. "aLIGO LSC RF Photodetector Schematic". *D1101124-v8* (2020).

[228] R. Abbott. "aLIGO LSC IQ Demodulator Board". *D0902745-v5* (2017).

[229] M. Ando. "Power recycling for an interferometric gravitational wave detector". PhD thesis. University of Tokyo, 1998.

[230] T. M. Niebauer, R. Schilling, et al. "Nonstationary shot noise and its effect on the sensitivity of interferometers". *Phys. Rev. A* 43 (9 1991), pp. 5022–5029.

[231] M. B. Gray, A. J. Stevenson, et al. "Harmonic demodulation of nonstationary shot noise". *Opt. Lett.* 18.10 (1993), pp. 759–761.

[232] T. Fricke. "Homodyne detection for laser-interferometric gravitational wave detectors". PhD thesis. Louisiana State University, 2011.

[233] J. A. Sidles and D. Sigg. "Optical torques in suspended Fabry–Perot interferometers". *Physics Letters A* 354.3 (2006), pp. 167–172. ISSN: 0375-9601.

[234] G. M. Muzammil A. Arain Chris Mueller. *Optical Layout and Parameters for the Advanced LIGO Cavities.* Tech. rep. LIGO-T0900043. 2012.

[235] T. Kozubowski and K. Podgorski. "A Class Of Asymmetric Distributions". *ARCH* 1 (1999), pp. 113–134.

[236] S. Kotz, T. J. Kozubowski, et al. *Annals of the Institute of Statistical Mathematics* 54.4 (2002), pp. 816–826.

[237] H. Cramer. *Mathematical Methods of Statistics (PMS-9).* Princeton University Press, 1999. ISBN: 9780691005478.

[238] J. Lawrence. "Distribution of the Median in Samples from the Laplace Distribution". *Open Journal of Statistics* 03.06 (2013), pp. 422–426.

[239] E. Cho, M. J. Cho, et al. "The variance of the variance of samples from a finite population". *Proceedings of the JSM, Section on Survey Research Methods* (2003).

[240] C. Pankow, K. Chatziioannou, et al. "Mitigation of the instrumental noise transient in gravitational-wave data surrounding GW170817". *Phys. Rev. D* 98 (8 2018), p. 084016.

[241] J. Powell. "Parameter estimation and model selection of gravitational wave signals contaminated by transient detector noise glitches". *Classical and Quantum Gravity* 35.15 (2018), p. 155017.

[242] B. Zackay, T. Venumadhav, et al. *Detecting Gravitational Waves in Data with Non-Gaussian Noise*. 2019.

[243] P. Godwin, R. Essick, et al. *Incorporation of Statistical Data Quality Information into the GstLAL Search Analysis*. 2020.

[244] N. J. Cornish and T. B. Littenberg. "Bayeswave: Bayesian inference for gravitational wave bursts and instrument glitches". *Classical and Quantum Gravity* 32.13 (2015), p. 135012.

[245] A. H. Nitz. "Distinguishing short duration noise transients in LIGO data to improve the PyCBC search for gravitational waves from high mass binary black hole mergers". *Classical and Quantum Gravity* 35.3 (2018), p. 035016.

[246] D. Davis, L. V. White, et al. "Utilizing aLIGO glitch classifications to validate gravitational-wave candidates". *Classical and Quantum Gravity* 37.14 (2020), p. 145001.

[247] T. B. Littenberg and N. J. Cornish. "Bayesian inference for spectral estimation of gravitational wave detector noise". *Physical Review D* 91.8 (2015).