

LABORATORY STUDIES OF ARM-LOCKING USING THE LASER
INTERFEROMETRY SPACE ANTENNA SIMULATOR AT THE UNIVERSITY OF
FLORIDA

By

JAMES IRA THORPE

A DISSERTATION PRESENTED TO THE GRADUATE SCHOOL
OF THE UNIVERSITY OF FLORIDA IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

UNIVERSITY OF FLORIDA

2006

Copyright 2006

by

James Ira Thorpe

For my parents, who provided everything and demanded nothing

ACKNOWLEDGMENTS

Many of my colleagues here at UF and elsewhere were truly instrumental in the production of this work: Shawn Mytrik, my right-hand man with the electronics; Volker Quetschke, the resident computer wizard; Rodrigo Delgadillo, Michael Hartman, and Gabriel Boothe, for their hours of faithful service in the lab; Daniel Shaddock, who provided insight into arm-locking, phasemeters, and other aspects of LISA; Rachel Cruz, who built the optics side of the LISA simulator; and especially my advisor, Guido Mueller, for guiding me to this point. Most importantly, I would like to thank my wife, Suzanne, who put up with nearly three years of separation while I completed this degree.

TABLE OF CONTENTS

	<u>page</u>
ACKNOWLEDGMENTS	4
LIST OF TABLES	8
LIST OF FIGURES	9
KEY TO ABBREVIATIONS	13
KEY TO SYMBOLS	15
ABSTRACT	18
CHAPTER	
1 INTRODUCTION	20
1.1 Motivation for Gravitational Wave Astronomy	20
1.2 Gravitational Wave Detectors	21
1.3 LISA at the University of Florida	22
2 GRAVITATIONAL WAVES	24
2.1 Overview	24
2.1.1 Relativity	24
2.1.2 Weak-field GR and Gravitational Waves	27
2.1.3 Properties of Gravitational Waves	28
2.1.4 Interaction with Matter	30
2.1.5 Generation of Gravitational Waves	32
2.1.6 Energy Carried by Gravitational Waves	34
2.2 Sources of Gravitational Waves	35
2.3 Detection of Gravitational Waves	41
2.3.1 Indirect Detection	41
2.3.2 Direct Detection	42
2.3.2.1 Doppler-tracking of spacecraft	43
2.3.2.2 Pulsar timing	44
2.3.2.3 Resonant mass detectors	45
2.3.2.4 Interferometric detectors	48
3 THE LASER INTERFEROMETER SPACE ANTENNA	54
3.1 Introduction	54
3.2 Sources	54
3.3 Mission Design	55
3.4 The Disturbance Reduction System (DRS)	56
3.5 The Interferometric Measurement System (IMS)	57
3.5.1 IMS Overview	58

3.5.2	The Optical Bench	60
3.5.3	Time Delay Interferometry	62
3.5.3.1	Visualizing TDI	64
3.5.3.2	Extensions to TDI	65
3.5.3.3	The zero-signal TDI variable	66
3.5.3.4	Limitations and noise sources	66
3.5.4	Arm-locking	67
3.5.4.1	Closed-loop system dynamics	68
3.5.4.2	Steady-state arm-locking performance	70
3.5.4.3	Transient response	75
3.5.4.4	Alternative arm-locking schemes	77
3.5.4.5	GW signals	82
3.5.4.6	Interaction with pre-stabilization system	82
4	THE UF LISA INTERFEROMETRY SIMULATOR	86
4.1	Background	86
4.2	The EPD Concept	86
4.3	Optical Components	90
4.3.1	Layout	90
4.3.2	Pre-stabilization	91
4.4	Electronic Components	92
4.5	Phasemeters	95
4.5.1	Overview	95
4.5.2	Phasemeters for LISA-like signals	98
4.5.3	An IQ phasemeter with a tracking LO	99
4.5.4	A Software Phasemeter	103
4.5.4.1	Design	103
4.5.4.2	Results	105
4.5.5	A Real-time Hardware Phasemeter	108
4.5.5.1	Front-end design	108
4.5.5.2	Back-end design	113
4.5.5.3	Single-signal PM test with a VCO	115
4.5.5.4	Single-signal PM test with optical signals	117
4.5.5.5	Entangled-phase PM test with VCO	120
4.5.5.6	Entangled-phase PM test with optical signals	123
4.5.5.7	Performance limitations	127
4.6	EPD Unit	129
4.6.1	Second-generation EPD unit	130
4.6.2	Third-generation EPD unit	132
5	ARM-LOCKING IN THE UF LISA INTERFEROMETRY SIMULATOR	136
5.1	Introduction	136
5.2	Electronic Model	137
5.2.1	Method	137

5.2.2	Results	138
5.2.3	Discussion	139
5.3	Initial Optical Model	140
5.3.1	Method	140
5.3.2	Results	143
5.3.2.1	Frequency counter measurements	144
5.3.2.2	Phasemeter data	147
5.3.2.3	Error-point noise	149
5.3.3	Discussion	150
5.4	Improved Optical Model	151
5.4.1	System Characterization	152
5.4.2	Filter Design	155
5.4.3	Results	158
6	CONCLUSION	159
6.1	Phasemeters and EPD Units	159
6.2	Arm-Locking	160
APPENDIX		
A	DIGITAL SIGNAL PROCESSING	162
A.1	Introduction	162
A.2	Sampling	162
A.2.1	Aliasing	164
A.2.2	Upconversion	165
A.3	Digital Signals	168
A.3.1	Binary Fractions	169
A.3.2	Multiplication and Other Operations	169
A.3.3	Floating-point Representations	171
A.3.4	Digitization Noise	172
A.4	Digital Filtering	174
A.4.1	Time-domain Response	175
A.4.2	Frequency Response	175
A.4.3	Design Methods	178
A.4.3.1	FIR Filters - windowed impulse response method	178
A.4.3.2	IIR Filters - bilinear transform method	180
A.4.4	Realization and Practicalities	181
A.4.4.1	Filter structures	181
A.4.4.2	Latency	182
A.4.5	CIC Filters	183
A.4.6	Fractional-Delay Filters	184
	REFERENCES	186
	BIOGRAPHICAL SKETCH	191

LIST OF TABLES

<u>Table</u>	<u>page</u>
2-1 Suggested frequency bands for GWs	40
2-2 Operational GW bar detectors	47
2-3 Major ground-based GW interferometers	51
4-1 Major LISA IMS components/signals and their EPD equivalents	90
4-2 Reconstruction algorithms for the hardware PM	114
4-3 Beat note frequencies and amplitudes for optical entangled-phase measurement .	123
4-4 Progression of EPD units	130
A-1 IEEE standard floating point representations	171

LIST OF FIGURES

<u>Figure</u>	<u>page</u>
2-1 Tidal distortion of an initially-circular ring of freely-falling test particles	32
2-2 A hypothetical laboratory generator of GWs	35
2-3 A binary star system as a generator of GWs	37
2-4 Observed shift of periastron for PSR 1913+16	42
2-5 Concept for Doppler-tracking detection of GWs	44
2-6 Sensitivity of ALLEGRO bar detector	47
2-7 A Michelson interferometer as a detector of GWs	48
2-8 Aerial photograph of the LIGO GW detector	52
2-9 Sensitivity curves for the LIGO observatory	53
3-1 Sources in the LISA observational window	55
3-2 Orbital configuration of the LISA constellation	56
3-3 Diagram of the LISA IMS	58
3-4 Diagram of a LISA optical bench	61
3-5 “rabbit-ear” diagram for the first-generation TDI X combination	65
3-6 Diagram of a closed-loop SISO system with negative feedback	68
3-7 Generic Nyquist plot	70
3-8 Nyquist plots for single-arm arm-locking	71
3-9 Bode plot of $T_{sen}(f)$ with $\tau_{RT} = 33\text{s}$	72
3-10 Bode plot for a generic arm-locking controller	74
3-11 Closed-loop noise suppression for a generic arm-locking loop	75
3-12 Nyquist plot for common arm-locking	79
3-13 Magnitude of square-bracketed term in (3–38)	81
3-14 Combining pre-stabilization and arm-locking with a tuneable cavity	83
3-15 Combining pre-stabilization and arm-locking using a sideband cavity lock	84
3-16 Combining pre-stabilization and arm-locking using an offset PLL	85

4-1	The EPD technique applied to a single LISA arm	88
4-2	Optical layout of the UF LISA interferometry simulator	91
4-3	Frequency noise in the $L_1 - L_0$ beat note	92
4-4	Phase noise in the $L_1 - L_0$ beat note	93
4-5	Overview of the DSP system from Pentek Corporation	94
4-6	Schematic of a IQ phasemeter with feedback	99
4-7	Laplace domain model of the system in Figure 4-6	101
4-8	Bode plot of $G(s)$ for the software PM	104
4-9	Bode plot of $H(s)$ for the software PM	105
4-10	Expected behavior of the software PM	106
4-11	Observed behavior of the software PM	107
4-12	Schematic of the real-time hardware PM	108
4-13	Schematic of a Direct Digital Synthesizer	109
4-14	CIC decimation filter near the first aliasing band	110
4-15	Passband flatness of the CIC filter in the hardware PM	111
4-16	Feedback filter for hardware PM tracking loop	112
4-17	Packing format for PM data transferred over the VIM interface	113
4-18	VCO phase noise measured by four channels of the hardware PM	117
4-19	Linear spectral density of laser beat note phase	118
4-20	Qualitative amplitude spectrum of interfering beams with shot noise	119
4-21	Timeseries for entangled phase test using VCOs	122
4-22	LSD of an entangled phase test using VCOs	122
4-23	Analog electronics used to prepare beat signals for PM	124
4-24	Linearly-detrended phase for optical entangled-phase measurement	124
4-25	Quadratically-detrended phase for optical entangled-phase measurement	125
4-26	Linear spectral density for optical entangled-phase measurement	126
4-27	Noise suppression in optical entangled-phase measurement	127

4-28	Histogram of residual phase for laser beat note in hardware PM	128
4-29	Schematic of the NCO used in the 3 rd -generation EPD unit	132
4-30	Detrended phase of VCO signal in EPD test	134
4-31	Linear spectral densities in EPD test with VCO signals	134
5-1	Experimental setup for electronic arm-locking experiment	137
5-2	Transfer functions for electronic arm-locking experiment	138
5-3	Linear spectral density of arm-locked VCO signal	139
5-4	Closed-loop noise suppression for electronic arm-locking experiment	140
5-5	Experimental arrangement for the initial optical arm-locking experiments	141
5-6	Laplace-domain model of the system in Figure 5-5	142
5-7	Bode plot of controller for initial optical arm-locking system	143
5-8	Response of interferometer to phase modulation	144
5-9	Timeseries of $L_2 - L_0$ beat note for locked and unlocked cases	145
5-10	Detrended timeseries of $L_2 - L_0$ beat note for locked and unlocked cases	146
5-11	Close-up of locked case in Figure 5-9	146
5-12	Spectrum of locked and unlocked frequency noise	147
5-13	Timeseries of beat note phase for unlocked and locked cases	148
5-14	Phase noise spectra for the unlocked and locked cases	148
5-15	Closed-loop noise suppression for optical arm-locking	149
5-16	Error-point noise for locked and unlocked cases	150
5-17	Modification of electronics for improved optical arm-locking	151
5-18	Raw phase timeseries of arm-locking system characterization data	153
5-19	Detrended timeseries of arm-locking system characterization data	154
5-20	Linear spectral densities of arm-locking system characterization data	154
5-21	Magnitude of transfer function for arm-locking system	155
5-22	Timeseries of filtered and unfiltered frequency noise from VCO input	157
5-23	Measured and designed transfer functions of arm-locking control filter	157

A-1	An overview of the sampling process	163
A-2	The phenomenon of aliasing	165
A-3	An overview of the upconversion process	167
A-4	Assumed PDF for quantization error	172
A-5	Non-uniform, non-white quantization error	173
A-6	Comparison of Laplace and z domains	177
A-7	Windowed impulse response method for designing FIR filters	179
A-8	The direct-form I filter structure	181
A-9	The direct-form II transposed filter structure	182
A-10	Magnitude response of a generic two-stage CIC filter	183

KEY TO ABBREVIATIONS

ADC: Analog-to-Digital Converter
AU: Astronomical Unit
BH: Black Hole
BIFO: Bi-directional First-Input First-Output buffer
BS: Beam Splitter
CIC: Cascade Integrator Comb filter
DAC: Digital-to-Analog Converter
DC: Direct Current or zero Fourier frequency
DDS: Direct Digital Synthesizer
DF2T: Direct Form II Transposed
DMA: Direct Memory Access
DRS: Disturbance Reduction System
DSP: Digital Signal Processing
EMRI: Extreme Mass-Ratio Inspiral
EOM: Electro-Optic Modulator
EPD: Electronic Phase Delay
ESA: European Space Agency
FIR: Finite Impulse Response filter
FPGA: Field-Programmable Gate Array
FSR: Free Spectral Range
FTP: File Transfer Protocol
GR: General Relativity
GUI: Graphical User Interface
GW: Gravitational Wave
IEEE: Institute of Electrical and Electronic Engineers
IIR: Infinite Impulse Response filter
IMS: Interferometric Measurement System
IQ: In-phase/Quadrature
LIGO: Laser Interferometric Gravitational-Wave Observatory
LISA: Laser Interferometer Space Antenna
LO: Local Oscillator
LSB: Least-Significant Bit
LSD: Linear Spectral Density
LUT: Look-Up Table
MI: Michelson Interferometer
MSB: Most-Significant Bit
NaN: Not a Number
NASA: National Aeronautics and Space Administration
NCO: Numerically Controlled Oscillator
Nd:YAG: Neodymium-doped Yittrium Aluminum Garnet
NPRO: Non-Planar Ring Oscillator
NS: Neutron Star
OB: Optical Bench

PC: Personal Computer
PCI: Peripheral Component Interconnect interface
PD: Photodiode
PDF: Probability Density Function
PDH: Pound-Drever-Hall
PLL: Phase-Lock Loop
PM: Phase Meter
PZT: Piezoelectric actuator
RF: Radio Frequency
RMS: Root Mean Square
SC: Space Craft
SDRAM: Synchronous Dynamic Random Access Memory
SiC: Silicon Carbide
SMBH: Super-Massive Black Hole
SR: Special Relativity
TCP/IP: Transmission Control Protocol / Internet Protocol
TDI: Time-Delay Interferometry
TOA: Time Of Arrival
TT: Transverse-Traceless
UF: The University of Florida
ULE: Ultra-Low Expansion
VCO: Voltage-Controlled Oscillator
VIM: Velocity Interface Module
VME: Virtual Machine Environment
WD: White Dwarf
ZOH: Zero-Order Hold

KEY TO SYMBOLS

\square : D'Alembertian operator
 $A^{\mu\nu}$: tensor amplitude
 $A(t)$: amplitude
 b : bit width
 c : speed of light
 $c(s)$: control signal
 $C(t)$: Dirac delta-function comb
 ds^2 : differential spacetime interval
 $e(s)$: error signal
 $E(t)$: time-component of electric field
 E_{dig} : digitization noise energy
 E_{SN} : shot-noise energy
 $erfc(x)$: complimentary error function
 f : Fourier frequency
 f_c : cutoff frequency
 f_n : interferometer null frequency
 f_{Nyq} : Nyquist frequency
 f_{UG} : unity-gain frequency
 \mathcal{F} : Fourier transform
 fs : sampling frequency
 $g_{\mu\nu}$: metric tensor for general relativity
 G : Newton's gravitational constant
 $G(s)$: filter or system transfer function
 $G_{\mu\nu}$: Einstein curvature tensor
 h : gravitational wave strain, Planck's constant
 \hat{h}_+ : + polarization tensor
 \hat{h}_\times : \times polarization tensor
 h_+ : strain amplitude in + polarization
 h_\times : strain amplitude in \times polarization
 $h_{\mu\nu}$: metric perturbation
 $\bar{h}_{\mu\nu}$: trace-reversed metric perturbation
 $h_{\mu\nu}^{TT}$: metric perturbation in TT gauge
 h_{ij} : phase change due to GWs on light propagating from SC_i to SC_j
 $h(n)$: impulse response function
 $H(s)$: feedback transfer function
 I : quadrupole moment tensor
 $I(t)$: in-phase component, intensity
 k : wave number
 k^μ : 4-D wavevector
 \vec{k} : 3-D wavevector
 L : optical path length
 L^{GW} : gravitational wave luminosity
 \mathcal{L} : Laplace transform

M : total memory
 $M(t)$: mixer output
 M_{\odot} : Solar mass
 n : index
 N : photon number
 N_{chan} : number of channels
 $\tilde{p}(s)$: free-running noise
 P : signal power
 P_{slip} : probability of cycle slip
 Q : reduced quadrupole moment, quadrature component
 R : decimation rate, reference signal, rectangular impulse
 $R_{\alpha\beta\nu}^{\mu}$: Riemann curvature tensor
 s : Laplace complex frequency variable
 S_{ij} : PM signal from PD_{main} on OB_{ij}
 S_{+} : common-arm error signal
 S_{-} : differential-arm error signal
 S_{direct} : direct arm-locking error signal
 T : sampling interval, measurement time
 $T_{CL}(s)$: closed-loop transfer function
 $T_{OL}(s)$: open-loop transfer function
 $T_{sen}(s)$: arm-locking sensor transfer function
 $T_{\mu\nu}$: matter stress-energy tensor
 U_{dig} : digitization noise
 U_{dig1} : single-channel digitization noise
 U_{LSB} : LSB amplitude
 \tilde{U}_{SN} : shot-noise amplitude spectral density
 $w(n)$: window function
 $x(s)$: input signal
 $x(n)$: discrete input signal
 $y(s)$: output signal
 $y(n)$: discrete output signal
 z : z -domain variable
 \mathcal{Z} : z transform
 $\Gamma_{\nu\rho}^{\mu}$: Christoffel symbol
 $\delta(x)$: Dirac delta function
 δ_{ij} : Kronecker delta function
 δI_{dig} : digitization noise in in-phase component
 δQ_{dig} : digitization noise in quadrature component
 $\delta\nu_{dig}$: digitization noise in frequency correction
 $\Delta\tau$: arm length difference
 ϵ : electric field amplitude
 ϕ : signal phase
 ϕ_e : PM phase error
 ϕ_i : PM input phase
 ϕ_m : PM model phase

ϕ_o : PM output phase
 ϕ_r : PM residual phase
 ϕ_{ij} : phase difference between ϕ_i and ϕ_j
 $\eta_{\mu\nu}$: special relativity metric
 λ : wavelength
 ν : signal frequency
 ν_{corr} : PM correction frequency
 ν_{ij} : frequency difference between ν_i and ν_j
 ν_m : PM model frequency
 ν_{off} : PM offset frequency
 ρ : Laplace or z -domain pole
 $\bar{\tau}$: average arm length
 τ_{ij} : light travel time from SC_i to SC_j
 τ_{max} : maximum delay time in EPD unit
 τ_{RT} : round-trip light travel time
 ω : angular frequency
 Ω_{GW} : gravitational wave angular frequency
 ζ : Laplace or z -domain zero

Abstract of a Dissertation Presented to the Graduate School
of the University of Florida in Partial Fulfillment of the
Requirements for the Degree of Doctor of Philosophy

LABORATORY STUDIES OF ARM-LOCKING USING THE LASER
INTERFEROMETRY SPACE ANTENNA SIMULATOR AT THE UNIVERSITY OF
FLORIDA

By

James Ira Thorpe

December 2006

Chair: Guido Mueller

Major Department: Physics

The Laser Interferometer Space Antenna (LISA) is a collaboration between the National Aeronautics and Space Administration (NASA) and the European Space Agency (ESA) to design and build a space-based interferometric detector of gravitational waves. The LISA sensitivity band will range from 3×10^{-5} Hz to 100 mHz, a regime currently inaccessible to ground-based detectors.

The LISA detector will consist of a constellation of three identical spacecraft arranged in a triangular formation 5×10^6 km on a side. Each spacecraft will contain a pair of freely-falling proof-masses that will act as the geodesic-tracking test particles of general relativity. The separation between the proof-masses will be monitored using laser interferometry with a precision of ~ 10 pm, allowing for the detection of gravitational waves with strain amplitudes in the range of 10^{-21} .

The author is part of a group at the University of Florida that is developing a laboratory-based simulator of LISA interferometry. This dissertation describes the simulator in detail, emphasizing the electronic components designed and constructed by the author. These include a phase meter capable of measuring the phase of a cavity-stabilized laser beat-note with a noise floor of better than 10^{-5} cycles/ $\sqrt{\text{Hz}}$ from 1 Hz – 10 kHz and an electronic phase delay unit capable of delaying signals with frequencies up to 25 MHz for more than 300 s with or without a frequency offset.

Also described is a set of experiments made using the simulator that investigate arm-locking, a proposed method for reducing the phase-noise of the LISA lasers. A laser beat note was successfully stabilized to a 1.065 ms delay with a bandwidth of ~ 10 kHz. The residual frequency noise was less than $200 \text{ mHz}/\sqrt{\text{Hz}}$ from 10 mHz through 100 Hz.

CHAPTER 1

INTRODUCTION

1.1 Motivation for Gravitational Wave Astronomy

The history of science is rife with examples of new technologies leading to breakthroughs in our understanding of the natural world. This is particularly true in the fields of astronomy and astrophysics. The study of the heavens is undoubtedly one of the planet's oldest sciences and for countless millennia it proceeded with one instrument: the human eye. Fortunately, the eye is quite a good instrument and a great deal was learned about the universe using it.

The invention of the optical telescope brought an improvement in angular resolution over the eye, allowing Galileo to observe moons orbiting Jupiter. These observations helped to cement the Copernican model of a heliocentric universe. While telescopes improved the spatial resolution of astronomical observations, the advent of photography widened the spectral window of these observations to include wavelengths at which the human eye is insensitive.

The 20th century saw an explosion of new spectral windows opened to the heavens. The universe can now be observed in radio, microwave, infrared, optical, ultra-violet, X-ray, and γ -rays. Each of these new spectral windows produced surprising and significant results that altered our understanding of the universe. For example, microwave astronomy led to the detection of the cosmic microwave background and the validation of the Big Bang theory while X-ray observations provided the first evidence of the existence of black holes.

Nearly all of our information about the universe outside our own solar system comes from some form of electromagnetic radiation. Despite the wide range of observable frequencies (there are more than eighteen decades of frequency between a 1 MHz radio wave and a 10 GeV gamma ray), all electromagnetic observations have common characteristics. Electromagnetic radiation gives us direct information about the particles, atoms, or molecules that generate it and interact with it. It is fundamentally a probe

of microscopic physics. Only through the association of radiating material with a macroscopic object, such as the corona of a star or an accretion disk around a black hole, can we make observations of macroscopic objects.

Gravitational waves¹ (GWs) represent an entirely new potential source of information about the universe. A prediction of general relativity, GWs are disturbances in *spacetime*, the combined fabric of space and time that is the arena for physics in relativity. GWs are thought to be produced by a variety of astrophysical systems, ranging in mass from solar-mass neutron stars to black holes millions or billions of times larger at the centers of colliding galaxies. In addition, the Big Bang may have produced GWs which would exist today as a cosmological background. Unlike electromagnetic radiation, GWs couple directly to large-scale objects. This makes them ideal for probing gravity, the dominant force over macroscopic distances.

The ability to detect GWs will provide more than opening a new spectral window; it is more akin to providing an entirely new “sense” with which we can learn about the universe. If electromagnetic observations are our eyes, GWs are our ears. While it is risky to make grand predictions about what we may learn, it certainly seems that we ought to try and listen.

1.2 Gravitational Wave Detectors

The previous Section makes clear the motivations for trying to detect GWs. The reason why it has yet to be done is that it is extremely difficult. The effect of a GW passing through a detector is a tidal distortion characterized by a strain amplitude (change in length over length) on the order of 10^{-21} . A km-scale detector must detect length changes on the order of 10^{-18} m, 1000 times smaller than the classical radius of a proton.

¹ For two excellent extended introductions to gravitational wave astronomy, see Schutz [1] and Hughes [2].

Despite these difficulties, there has been significant effort over the past half-century to build GW detectors. The types and sizes of the detectors vary, with different detectors optimized to observe GWs in different frequency bands. While no confirmed direct detections have been made, measurements of orbital decay in binary pulsars [3] have provided extremely convincing circumstantial evidence that GWs exist and generally behave as expected.

One proposed GW detector is the Laser Interferometer Space Antenna (LISA), which will consist of three separate spacecraft forming a triangular detector with sides of $5 \text{ Gm} = 5 \times 10^9 \text{ m}$. To measure GWs, LISA must detect length changes in these arms with a precision of $\sim 10 \text{ pm}$. Achieving this level of precision over such vast distances requires a number of novel techniques and technologies.

1.3 LISA at the University of Florida

The author is part of a group in the Department of Physics at the University of Florida (UF) that is developing a laboratory-based simulator of LISA interferometry. The purpose of this simulator is to provide an arena in which the interferometric techniques of LISA can be studied and developed. It also provides a source of LISA-like signals with which to test prototype components. A long-term goal of the simulator is to have the ability to inject model GW signals into the apparatus and produce LISA-equivalent science signals with realistic instrumental noise. Such signals would be valuable for evaluating data analysis techniques.

The remainder of this dissertation is divided into four parts. Chapter 2 presents an overview of GWs including their theoretical origins, properties, likely sources, and potential detection methods. Chapter 3 describes LISA in detail, with an emphasis on the interferometry. Chapter 4 describes the development of the UF LISA interferometry simulator, focusing on the electronic components of the simulator that were designed and built by the author. Chapter 5 presents a series of experiments using the simulator that

investigate a laser phase-noise stabilization technique known as arm-locking that has been proposed for LISA.

CHAPTER 2 GRAVITATIONAL WAVES

2.1 Overview

2.1.1 Relativity

The theoretical framework of Special and General Relativity represents our best understanding of the macroscopic universe¹. In both cases, the three usual dimensions of space and the one dimension of time are combined into a single four-dimensional entity known as *spacetime*. In Special Relativity (SR), spacetime is a passive background in which physics occurs. Points in spacetime are known as *events*, and physics is concerned with the relation between events. For example, the set of events that mark the position of a particle in three-dimensional space as time evolves is known as that particle’s *worldline*. The worldline of a particular particle may be affected by non-gravitational phenomena such as electromagnetic or nuclear forces.

As with Newtonian physics, it is useful to define a coordinate system, or *frame*, which can be used to label and compare events. In Special Relativity (SR), there exists a special class of frames known as the Lorentz or inertial frames, in which free particles move in straight lines with uniform velocity. A set of coordinate transformations, known as the Lorentz transformations, relate the coordinates of an event in one inertial frame to the coordinates of the same event in another inertial frame. Physically observable quantities are independent of the particular frame used.

The four-dimensional Cartesian coordinate system $x^\mu = (t, x, y, z)$ can be used to describe an inertial frame in SR². The *interval*, or distance between events separated by

¹ Much of the theoretical development in this Chapter follows Schutz [4]. Other portions were adapted from Misner, et al. [5] and Shapiro & Teukolsky [6].

² Unless otherwise noted, I will adopt the “natural units” of $G = c = 1$ for this Chapter.

the four-vector V^μ , can be computed as

$$\Delta s^2 = \eta_{\mu\nu} V^\mu V^\nu, \quad (2-1)$$

where Δs is the interval, $\eta_{\mu\nu}$ is the metric tensor, and the Einstein summation convention ($x_\alpha y^\alpha = \sum_{\alpha=0}^3 x_\alpha y_\alpha$) applies. The metric tensor for an inertial frame in SR can be written using the Cartesian coordinates described above as

$$\eta_{\mu\nu} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \quad (2-2)$$

Spacetime intervals are split into three classes according to their sign. Four-vectors with negative intervals are known as time-like. The worldlines of all massive particles have time-like intervals. The magnitude of the interval is equal to the proper time, the time elapsed in an inertial frame comoving with the particle. Four-vectors with null intervals are known as light-like, because photons and other massless particles have worldlines with null intervals. Four-vectors with positive intervals are known as space-like.

In the language of differential geometry, spacetime is a four-dimensional Riemannian manifold for which the distance between points on the manifold is given by a rank-2 metric tensor. In SR the manifold is “flat,” meaning that every inertial frame is valid over all of spacetime and the metric $\eta_{\mu\nu}$ can be used to compute the spacetime interval between distant events. The straight lines that describe the worldlines of free particles are special curves known as *geodesics*, which have the property that the interval along them is extremal.

In General Relativity (GR), the manifold is curved in an additional dimension or dimensions. Because of the curvature, the global inertial frames of SR do not exist in GR. However, since spacetime is smooth, it appears to be flat over small distances. At each point in spacetime, a local inertial frame can be constructed in which the physics of SR

apply locally. The differential interval between two nearby points separated by dx^μ can be computed using a differential form of (2-1)

$$ds^2 = g_{\mu\nu}dx^\mu dx^\nu, \quad (2-3)$$

where the metric is labeled $g_{\mu\nu}$ to distinguish it from the SR metric. Since spacetime is curved, the local inertial frames will differ between adjacent patches and consequently $g_{\mu\nu}$ will be a function of position within spacetime. To compute the interval between distant events, (2-3) must be integrated along the path between the two events. As in SR, the worldlines of free particles in GR will follow geodesics. Due to the curvature of spacetime, geodesics in GR will not generally be straight lines in a local inertial frame. Variational principles can be used to produce equations describing geodesics in GR. If the worldline is described by a set of events parametrized by a scalar parameter λ , $x^\mu(\lambda)$, then the geodesic equation can be written as

$$\ddot{x}^\mu + \Gamma_{\nu\rho}^\mu \dot{x}^\nu \dot{x}^\rho = 0, \quad (2-4)$$

where the dot denotes derivation with respect to λ and $\Gamma_{\nu\rho}^\mu$ is a combination of derivatives of the metric known as a Christoffel symbol,

$$\Gamma_{\nu\rho}^\mu \equiv \frac{1}{2}g^{\mu\sigma}(g_{\sigma\nu,\rho} + g_{\sigma\rho,\nu} - g_{\rho\nu,\sigma}), \quad (2-5)$$

where

$$g_{\mu\nu,\rho} \equiv \frac{\partial g_{\mu\nu}}{\partial x^\rho}. \quad (2-6)$$

The geodesic equation describes the behavior of free particles in a curved spacetime. GR connects this with gravity by specifying that the source of the spacetime curvature is matter. More specifically it is the energy density of all forms of matter and non-gravitational forces. This relationship is expressed mathematically by the Einstein equations,

$$G_{\mu\nu} = 8\pi T_{\mu\nu}, \quad (2-7)$$

where $T_{\mu\nu}$ is the stress-energy tensor of matter and non-gravitational forces and $G_{\mu\nu}$ is the Einstein curvature tensor, a particular combination of metric derivatives. The form of $G_{\mu\nu}$ is imposed by the constraints placed by conservation laws for energy and momentum on $T_{\mu\nu}$. In order for (2-7) to be generally valid, $G_{\mu\nu}$ must satisfy the same constraints. The Einstein curvature tensor is a combination of metric derivatives that satisfies the constraints on $T_{\mu\nu}$.

GR is sometimes summed up by stating that spacetime tells matter how to move (geodesic equation) and matter tells spacetime how to curve (Einstein equations).

2.1.2 Weak-field GR and Gravitational Waves

The Einstein field equations are a set of ten coupled non-linear partial differential equations. Only a handful of analytic solutions are known. It is often useful to consider approximations to the full theory, which are more amenable to analytical study. One such approximation is the weak-field limit, in which the GR metric is equal to the SR metric plus some small perturbation, $h_{\mu\nu}$,

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \quad (2-8)$$

$$|h_{\mu\nu}| \ll 1 \quad (2-9)$$

The linearized theory is developed by truncating the full GR equations to first order in $h_{\mu\nu}$. In doing so, it is useful to exploit the gauge freedom of GR. Gauge freedom refers to the ability to make changes to tensors such as $h_{\mu\nu}$ without affecting the observable quantities, such as the spacetime interval between two events, that are computed from these tensors. The linearized Einstein equations are typically written using the Lorentz gauge, which requires

$$\bar{h}^{\mu\nu}{}_{,\nu} = 0. \quad (2-10)$$

With this condition, the linearized Einstein equations become

$$\square \bar{h}^{\mu\nu} = -16\pi T^{\mu\nu}, \quad (2-11)$$

where $\bar{h}^{\mu\nu} = h^{\mu\nu} - \frac{1}{2}\eta^{\mu\nu}h^\rho{}_\rho$ is the trace-reverse of $h_{\mu\nu}$ and \square is the D'Alembertian operator. The D'Alembertian operator is also known as the wave operator, since it gives the wave equation when applied to a function,

$$\square f = \left(-\frac{\partial^2}{\partial t^2} + \nabla^2\right)f, \quad (2-12)$$

where ∇^2 is the Laplacian in the three spatial dimensions. The form of (2-11) indicates that there will be a set of wave solutions to the linearized Einstein equations. These solutions are known as *Gravitational Waves* (GWs).

2.1.3 Properties of Gravitational Waves

The properties of GWs can be deduced from the homogeneous version of (2-11), which corresponds to linearized GR in a vacuum ($T^{\mu\nu} = 0$). The general solution of the homogeneous wave equation is a superposition of plane waves of the form

$$\bar{h}^{\mu\nu}(x^\mu) = A^{\mu\nu} \exp(ik_\mu x^\mu), \quad (2-13)$$

where $A^{\mu\nu}$ is an amplitude tensor and k^μ is the four-dimensional analog of the wave vector in classical radiation theory. In order for the expression in (2-13) to satisfy (2-11), k^μ must be a null or light-like four-vector,

$$k^\mu k_\mu = 0. \quad (2-14)$$

The dispersion relation for GWs can be found by expressing k^μ in a 3 + 1 (three spatial coordinates plus one time coordinate) coordinate system and identifying the time component, k^0 , as the angular frequency of the wave, ω . The condition in (2-14) can then be written as

$$\omega^2 = \left|\vec{k}\right|^2, \quad (2-15)$$

where \vec{k} is the spatial component of the wave vector. From (2-15) it can be seen that both the phase and group velocities of GWs are 1 in natural units, which corresponds to the speed of light.

In order to be a solution to (2-11), the expression in (2-13) must also satisfy the Lorentz gauge condition in (2-10). This places restrictions on $A^{\mu\nu}$; requiring it to be orthogonal to \vec{k} ,

$$A^{\mu\nu}k_\nu = 0. \quad (2-16)$$

The Lorentz gauge condition in (2-10) describes a class of gauges. Additional restrictions on $\bar{h}^{\mu\nu}$ can be obtained by choosing a particular gauge within this class. For GWs in vacuum, the *transverse-traceless* (TT) gauge is useful. The TT gauge conditions are

$$\bar{h}^\mu{}_\mu = 0 \quad (2-17)$$

and

$$\bar{h}_{\mu 0} = 0. \quad (2-18)$$

Within the TT gauge, there is an inertial frame of the background spacetime (the $\eta_{\mu\nu}$ in (2-8)) for which the wave is traveling in the z -direction. In this frame, there will only be two independent components of $\bar{h}^{\mu\nu}$,

$$h_{\mu\nu}^{TT} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_{xx} & h_{xy} & 0 \\ 0 & h_{xy} & -h_{xx} & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \quad (2-19)$$

The $h_{\mu\nu}^{TT}$ refers to transverse-traceless gauge and the overbar has been dropped since a traceless tensor is its own trace-reverse. The two independent components of $h_{\mu\nu}^{TT}$ are interpreted as two orthogonal polarization states for GWs. It is common to rewrite (2-19) using two scalar polarization states and two unit polarization tensors,

$$h_{\mu\nu}^{TT} = h_+ \hat{h}_+ + h_\times \hat{h}_\times \quad (2-20)$$

where the polarization tensors are

$$\hat{h}_+ = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad (2-21)$$

and

$$\hat{h}_\times = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \quad (2-22)$$

The two scalar polarization states, h_+ and h_\times each will obey a scalar wave equation analogous to (2-13),

$$h_{+/\times} = A_{+/\times} \exp(ik_\mu x^\mu). \quad (2-23)$$

2.1.4 Interaction with Matter

The preceding Sections described GWs as mathematical solutions of GR. How they manifest themselves physically can be examined by studying how these solutions effect the motion of free particles. Recall that the geodesic equation (2-4) describes the motion of free particles. A related equation, known as the geodesic deviation equation, describes the evolution of the 4-vector linking two nearby geodesics,

$$\ddot{\xi}^\mu = R^\mu{}_{\alpha\beta\nu} V^\alpha U^\beta \xi^\nu, \quad (2-24)$$

where ξ^μ is the 4-vector linking the two geodesics, the dots denote differentiation to a parameter of the geodesic (such as proper time), $R^\mu{}_{\alpha\beta\nu}$ is a combination of metric derivatives known as the Riemann tensor, and V^α and U^β are the four-velocities of the particles on the geodesics. Consider a frame in which there are two neighboring free

particles, both initially at rest, separated by a distance ε in the x direction. For this case,

$$\xi^\mu = (0, \varepsilon, 0, 0) \quad (2-25)$$

and

$$V^\alpha = U^\beta = (1, 0, 0, 0). \quad (2-26)$$

These expressions can be substituted into (2-24) to obtain

$$\ddot{\xi}^\mu = -\varepsilon R^\mu{}_{0x0}. \quad (2-27)$$

To first order in $h^{\mu\nu}$, the Riemann tensor is given by

$$R^\mu{}_{\nu\alpha\beta} = \frac{1}{2}\eta^{\mu\sigma}(h_{\sigma\beta,\nu\alpha} + h_{\nu\alpha,\sigma\beta} - h_{\sigma\alpha,\nu\beta} - h_{\nu\beta,\sigma\alpha}). \quad (2-28)$$

Using the expressions for $h^{\mu\nu}$ in the TT gauge, the four equations in (2-27) can be reduced to two

$$\ddot{\xi}^x = \frac{1}{2}\varepsilon\ddot{h}_+ \quad (2-29)$$

and

$$\ddot{\xi}^y = \frac{1}{2}\varepsilon\ddot{h}_\times. \quad (2-30)$$

All other components of $\ddot{\xi}^\mu$ are zero. The tidal effects of GWs on freely-falling particles will be restricted to the plane normal to the wave's propagation direction. The motion in this plane will be oscillatory, with an angular frequency equal to that of the GW. A similar analysis can be made for two particles initially separated by a distance ε in the y -direction. The results are

$$\ddot{\xi}^y = -\frac{1}{2}\varepsilon\ddot{h}_+ \quad (2-31)$$

and

$$\ddot{\xi}^x = \frac{1}{2}\varepsilon\ddot{h}_\times. \quad (2-32)$$

The expressions (2–29) through (2–32) can be used to determine the tidal distortions of an initially-circular ring of freely-falling test particles in the $x - y$ plane as a GW traveling in the z -direction passes by. The result is shown schematically in Figure 2-1.

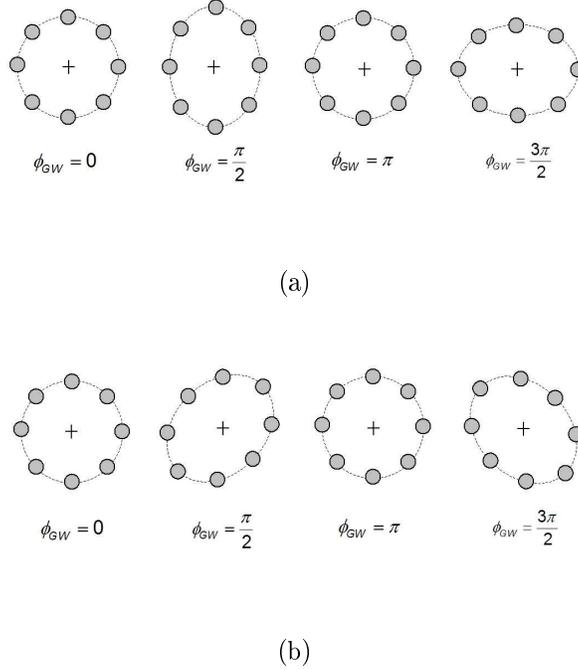


Figure 2-1. Distortion of an initially-circular ring of freely-falling test particles by a GW propagating into the plane for (a), the $+$ polarization and (b), the \times polarization. ϕ_{GW} refers to the phase of the GW.

2.1.5 Generation of Gravitational Waves

The general solution of the inhomogeneous wave equation (2–11) for $\bar{h}_{\mu\nu}$ at an event with coordinates (t, x^i) is given by the integral of the retarded Green's function over the past light cone of the event,

$$\bar{h}_{\mu\nu}(t, x^i) = 4 \int \frac{T_{\mu\nu}(t - |x^i - y^i|, y^i)}{|x^i - y^i|} d^3y. \quad (2-33)$$

The integrand can be simplified in the case that the source is compact and far from the field point x^i ,

$$\bar{h}_{\mu\nu}(t, x^i) \approx \frac{4}{r} \int T_{\mu\nu}(t - r, y^i) d^3y, \quad (2-34)$$

where r is the distance between the source and the field point and the integral is over the source where $|y^i| \ll r$. Further simplifications can be made by exploiting the restrictions placed on $T_{\mu\nu}$ by conservation laws. The laws of conservation of energy and conservation of momentum can be expressed together as

$$T^{\mu\nu}{}_{,\nu} = 0. \quad (2-35)$$

When (2-35) is applied to (2-34), the result to first order in $(1/r)$ is

$$\bar{h}^{\mu 0} \approx 0 \quad (2-36)$$

and

$$\bar{h}_{jk}(t, x^i) \approx \frac{2}{r} \ddot{I}_{jk}(t - r), \quad (2-37)$$

where I_{jk} is the quadrupole moment tensor of the source mass distribution and the dots indicate derivatives with respect to t . The quadrupole moment tensor is defined as

$$I^{jk} = \int \rho x^j x^k d^3x, \quad (2-38)$$

where $\rho \equiv T^{00}$ is the energy-density of the source distribution in its rest frame and the integral is over the entire source. Oftentimes, it is useful to express (2-36) and (2-37) in the TT gauge. In a coordinate frame at the observation point with the z -axis oriented along the propagation direction of the wave, the expressions become

$$h_+ = \frac{1}{r} \left[\ddot{Q}_{xx}(t - r) - \ddot{Q}_{yy}(t - r) \right] \quad (2-39)$$

and

$$h_\times = \frac{2}{r} \ddot{Q}_{xy}(t - r), \quad (2-40)$$

where Q_{ij} is the reduced quadrupole moment tensor, defined as

$$Q_{ij} = I_{ij} - \frac{1}{3} \delta_{ij} I^k{}_k \quad (2-41)$$

and δ_{ij} is the identity tensor. The expressions in (2-36) through (2-41) are known as the quadrupole approximation and can be used to describe the gravitational radiation in many physical systems. This is the topic of Section 2.2.

2.1.6 Energy Carried by Gravitational Waves

Like their electromagnetic counterparts, GWs carry energy. In GR, gravitational energy results from curvature, which is a global phenomenon. Consequently, the energy density cannot be assigned to a specific point and can only be computed as an average energy density over a region large enough to define the curvature. To compute the energy for a GW, the quadratic contributions from $h_{\mu\nu}$ to the left-hand side of the Einstein equation, (2-7) can be moved to the right-hand side and treated as a stress-energy source term. The result³ is

$$T_{\mu\nu}^{(GW)} = \frac{1}{32\pi} \left\langle \bar{h}_{\alpha\beta,\mu} \bar{h}^{\alpha\beta}{}_{,\nu} - \frac{1}{2} \bar{h}^{\alpha}{}_{\alpha,\mu} \bar{h}^{\alpha}{}_{\alpha,\nu} - \bar{h}^{\alpha\beta}{}_{\beta} \bar{h}_{\alpha\mu,\nu} - \bar{h}^{\alpha\beta}{}_{\beta} \bar{h}_{\alpha\nu,\mu} \right\rangle, \quad (2-42)$$

where the angle brackets denote an average over several wavelengths. The expression in (2-42) can be simplified if the TT gauge conditions ((2-17) and (2-18)) are applied,

$$T_{\mu\nu}^{(GW)} = \frac{1}{32\pi} \langle (h^{TT})_{ij,\mu} (h^{TT})^{ij}{}_{,\nu} \rangle. \quad (2-43)$$

Using (2-43) and the quadrupole approximation described in Section 2.1.5, the GW luminosity for a compact source can be estimated as

$$L^{GW} = \frac{1}{5} \langle \ddot{Q}_{ij} \ddot{Q}^{ij} \rangle. \quad (2-44)$$

To compute the luminosity in physical units, (2-44) is multiplied by the conversion factor

$$L_0 = c^5/G \approx 3.6 \times 10^{52} \text{ W}. \quad (2-45)$$

³ For details see Section 35.7 of Misner, et al. [5]

This luminosity is an upper limit that is never reached since $L^{GW} \ll 1$ in natural units. Nevertheless, in certain astrophysical systems, GWs carry a tremendous amount of energy and play an important role in the system dynamics.

2.2 Sources of Gravitational Waves

Section 2.1.5 described the quadrupole approximation for the generation of gravitational waves. It was found that a mass distribution with a time varying quadrupolar (or higher) moment would generate GWs. In theory, GWs could be generated by a laboratory apparatus such as the one in Figure 2-2. Consider a uniform beam of mass M and length L lying in the $x - y$ plane. The z -axis passes through the center of the beam and the beam makes an angle ϕ with the x -axis.

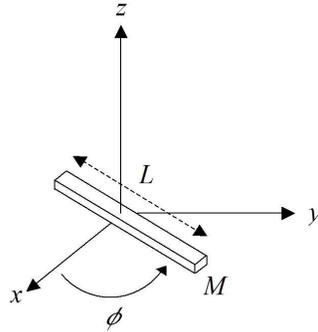


Figure 2-2. A hypothetical laboratory generator of GWs consisting of a bar of length L and mass M

Under the assumption that the cross-sectional dimensions of the beam are small compared to its length, the reduced quadrupole tensor can be written as

$$Q_{ij} = \frac{ML^2}{12} \begin{pmatrix} \cos^2 \phi - \frac{1}{3} & \cos \phi \sin \phi & 0 \\ \cos \phi \sin \phi & \sin^2 \phi - \frac{1}{3} & 0 \\ 0 & 0 & -\frac{1}{3} \end{pmatrix}. \quad (2-46)$$

If the beam is assumed to rotate about the z -axis with angular frequency ω , the second and third time derivatives of Q_{ij} are

$$\ddot{Q}_{ij} = \frac{ML^2\omega^2}{6} \begin{pmatrix} -\cos(2\phi) & -\sin(2\phi) & 0 \\ -\sin(2\phi) & \cos(2\phi) & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad (2-47)$$

and

$$\dddot{Q}_{ij} = \frac{ML^2\omega^3}{3} \begin{pmatrix} \sin(2\phi) & -\cos(2\phi) & 0 \\ -\cos(2\phi) & -\sin(2\phi) & 0 \\ 0 & 0 & 0 \end{pmatrix}. \quad (2-48)$$

The quadrupole formulas for h , (2-39) and (2-40) can be used to estimate the strain amplitude from (2-47). It is clear from (2-47) that the quadrupole moment oscillates with an angular frequency of 2ω . This is to be expected from the symmetry of the system under a rotation of π radians about the z -axis. The GWs will have this same frequency so that $\Omega_{GW} = 2\omega$.

Since the quadrupole approximation is valid only in the far-field, (2-39) and (2-40) must be evaluated at a source distance of at least one GW wavelength ($r = c/2\omega$). The strain amplitude at that distance can be estimated as

$$h^{(lab)} \approx \frac{2}{3}ML^2\omega^3. \quad (2-49)$$

For a beam with $M = 10^4$ kg, $L = 10$ m, and $\omega = 60$ rad/s, (2-49) gives $h^{(lab)} \sim 10^{-42}$. The GW luminosity from the beam can be estimated using (2-44) and (2-48) as

$$L_{GW}^{(lab)} \approx \frac{1}{15}M^2L^4\omega^6. \quad (2-50)$$

Using the same parameters, this gives $L_{GW}^{(lab)} \sim 10^{-33}$ W. It is clear from the small size of $h^{(lab)}$ and $L_{GW}^{(lab)}$ that GWs are not relevant for laboratory systems. What is needed to generate physically meaningful GWs is larger masses and higher velocities. Both can be found in astrophysical systems.

A binary star system is an example of an astrophysical system with a time-varying quadrupole moment. It is well-known that the 2-body problem in GR has no analytic solution. However, for most systems it is appropriate to use Newtonian mechanics to describe the orbital motion and then use the quadrupole approximation to compute the GW amplitudes and strains. Consider two point masses with masses m_1 and m_2 in a circular orbit of radius a in the $x - y$ plane, as shown in Figure 2-3.

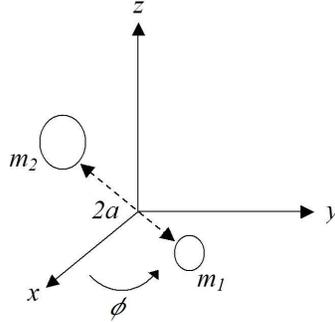


Figure 2-3. A binary star system as generator of GWs

If the system center of mass is placed at the origin, the position vectors of the two masses will be

$$\vec{x}_1 = a \frac{\mu}{m_1} (\cos \phi, \sin \phi) \quad (2-51)$$

and

$$\vec{x}_2 = a \frac{\mu}{m_2} (-\cos \phi, -\sin \phi), \quad (2-52)$$

where $\mu = m_1 m_2 / (m_1 + m_2)$ is the reduced mass and ϕ is the orbital phase angle, measured from the positive x -axis to m_1 . The reduced quadrupole moment for this system is

$$Q_{ij} = \mu a^2 \begin{pmatrix} \cos^2 \phi - \frac{1}{3} & \cos \phi \sin \phi & 0 \\ \cos \phi \sin \phi & \sin^2 \phi - \frac{1}{3} & 0 \\ 0 & 0 & -\frac{1}{3} \end{pmatrix} \quad (2-53)$$

This is the same form as that for the beam in (2-46) as should be expected from the similarity of the geometries. The time derivatives of Q_{ij} can be obtained from (2-47) and

(2-48) with the substitutions $M \rightarrow \mu$ and $L \rightarrow 2a$ as well as a division by 3 that results from the distribution of mass within the uniform beam. For the binary system, Kepler's law gives a relation between the orbital frequency and the binary separation,

$$\omega^2 = \frac{M}{a^3} \quad (2-54)$$

where $M = m_1 + m_2$ is the total mass. As with the rotating beam, the gravitational waves will be emitted at a frequency equal to twice the orbital frequency. Using the above relationships, the GW amplitude and luminosity for a binary system can be estimated as

$$h^{(binary)} \approx \frac{2 \mu M}{r a} \quad (2-55)$$

and

$$L_{GW}^{(binary)} = \frac{32 \mu^2 M^3}{5 a^5}. \quad (2-56)$$

The forms of (2-55) and (2-56) demonstrate that the largest and most energetic GWs will be generated in binaries with large mass and small separations. Ideal candidates for such binaries are binaries where one or both members is a compact object such as a white dwarf (WD), neutron star (NS), or black hole (BH).

For example, a NS-NS binary ($m_1 \approx m_2 \approx 1.4 M_{Sun} = 2.8 \times 10^{30}$ kg) with an orbital separation of $2a = 500$ km would produce a GW luminosity of $L_{GW}^{(binary)} \sim 10^{44}$ W at a frequency $\Omega_{GW} = 370$ Hz. At a distance of $1 \text{ Mpc} = 3 \times 10^{22}$ m, this would produce an energy flux of $\sim 9 \text{ mW/m}^2$ at Earth, about three times brighter than the visible light flux from the full Moon. The corresponding GW strain amplitude at Earth would be $h^{(binary)} = 2.3 \times 10^{-21}$, thought to be within the range of GW detectors.

The energy carried away from the binary in the form of GWs causes the overall energy of the binary to decrease with time. Consequently, the orbital radius must decrease while the orbital frequency increases. The decrease in orbital radius increases the GW energy output ($L_{GW}^{(binary)} \propto a^{-5}$), causing the system to radiate more strongly. The resulting GWs increase in both frequency and amplitude with time, a waveform known as

a *chirp*. So long as no other physical effects conspire to prevent it, the orbit will continue to decay until the two objects merge. The chirp waveform can be estimated using an adiabatic approximation in which the GWs are calculated from the Keplerian orbits and the orbital parameters are changed to match the energy loss. For a circular orbit, the total orbital energy is

$$E^{(binary)} = -\frac{1}{2} \frac{\mu M}{a}. \quad (2-57)$$

Taking a time derivative of (2-57) and equating it with (2-56) results in a differential equation for a ,

$$\dot{a} = -\frac{64}{5} \frac{\mu M^2}{a^3}. \quad (2-58)$$

This equation can be solved to yield

$$a(t) = a_0(1 - t/t_{merge})^{1/4}, \quad (2-59)$$

where a_0 is the orbital radius at time $t = 0$ and t_{merge} is the time of merger, given by

$$t_{merge} = \frac{5}{256} \frac{a_0^4}{M^2 \mu}. \quad (2-60)$$

As the orbital radius decreases, the accuracy of the Newtonian adiabatic approximation worsens. This is precisely the regime in which the GW luminosity is the largest, so it is important that more accurate methods be applied to predict GW waveforms. These include analytic treatments with relativistic corrections to the orbits as well as numerical simulations that incorporate the full Einstein equations [7].

In addition to binary systems, several other types of astrophysical sources of GWs are thought to exist. Rapidly rotating NSs with a slight asymmetry will produce GWs. The energy lost through GW emission will cause their rotation rate to decrease, much as the electromagnetic radiation from pulsars cause spin-down. The waveforms for such sources can be calculated in a manner similar to that for the binary systems [8]. Stellar core collapse associated with supernovae are also a likely source of GWs, although in order to generate GWs, there must be an asymmetric flow of mass. The difficulty in

modeling supernovae makes detailed predictions of their GW signatures hard to obtain [9]. Finally, there is also a possibility of a cosmological background of GWs analogous to the cosmic microwave background for electromagnetic radiation. This cosmic gravitational wave background would be a stochastic signal, the level of which can be estimated from cosmological arguments [2].

GW sources can be separated by frequency band, much as electromagnetic sources are separated into radio, visible, gamma-ray, etc. In general, larger masses translate to lower frequencies. The merger of a stellar-mass binary will occur in the ~ 1 kHz band, where as the merger of two super-massive black holes (SMBHs), with masses $10^6 M_{Sun} - 10^9 M_{Sun}$, will occur in the ~ 1 mHz band. A GW spectrum suggested by Hughes [2] is contained in Table 2-1.

Table 2-1. Suggested frequency bands for GWs

Band	Frequency Range	Persistent Srcs.	Transient Srcs.
Ultra-low Frequency	10^{-18} Hz \sim 10^{-13} Hz	?	?
Very-low Frequency	10^{-9} Hz \sim 10^{-7} Hz	SMBI CGBR	?
Low-Frequency	10^{-6} Hz \sim 1 Hz	BI EMRI CGBR	SMBM
High-Frequency	1 Hz \sim 10 kHz	RNS CGBR	BM SN

(SMBI = Super-Massive Binary Inspiral, BI = stellar-mass Binary Inspiral, EMRI = Extreme Mass-Ratio Inspiral, SMBM = Super-Massive Binary Merger, BM = stellar-mass Binary Merger, SN = Supernovae, RNS = Rotating/pulsating Neutron Stars, CGBR = Cosmic Gravitational Wave Background)

As with electromagnetic sources, the same physical object may radiate in different bands at different epochs within its evolution. A stellar-mass binary in the early stages

of inspiral will exist as a persistent source in the low-frequency band. As it evolves, the frequency will increase until it merges in the high-frequency band.

2.3 Detection of Gravitational Waves

With the knowledge that GWs represent a set of solutions to the Einstein equations (Section 2.1) and the existence of several plausible mechanisms for their generation (Section 2.2), it is a reasonable assumption that most of the universe is bathed in gravitational radiation. The obvious question is how can this radiation be detected. A number of techniques for detecting GWs have been proposed or implemented. These techniques generally fall into two categories: direct techniques which measure the amplitudes of the waves themselves and indirect techniques which infer the presence of the waves from their effects on a well-understood physical system.

2.3.1 Indirect Detection

Thus far, only indirect detections of GWs have been made. In 1975, Russel Hulse and Joseph Taylor discovered a pulsar known as PSR 1913+16 [3]. A pulsar is a rapidly-rotating NS with a highly-beamed radio emission. As the pulsar rotates, its radio beam sweeps across Earth, producing a pulse in a radio detector. Pulsars are some of the most stable oscillators in the universe and provide a unique opportunity for precision measurement of the motion of a distant compact object. After observing PSR 1913+16 for some time, it was determined that it was in a binary orbit with a radio-quiet companion, likely a second NS. The mass of the companion and the orbital parameters (radius, eccentricity, orbital phase, etc.) were extracted by fitting the pulse arrival times to an orbital model. Once the binary system was characterized, the expected GW luminosity could be computed from a modification of (2-56)⁴. With L_{GW} known, an energy balance

⁴ For elliptical orbits, (2-56) is modified by an enhancement factor $f(e)$ = $\frac{1+(73/24)e^2+(37/96)e^4}{(1-e^2)^{7/2}}$, that depends on the orbital eccentricity, e . The system also radiates preferentially at periastron, meaning that GW emission tends to circularize orbits [10].

could be used to determine the predicted rate of orbital decay as in (2–59). Figure 2-4 shows a plot of the observed shift in orbital phase (versus a non-decaying orbit) of PSR 1913+16 from 1975 to 1988 along with the predictions of GR.

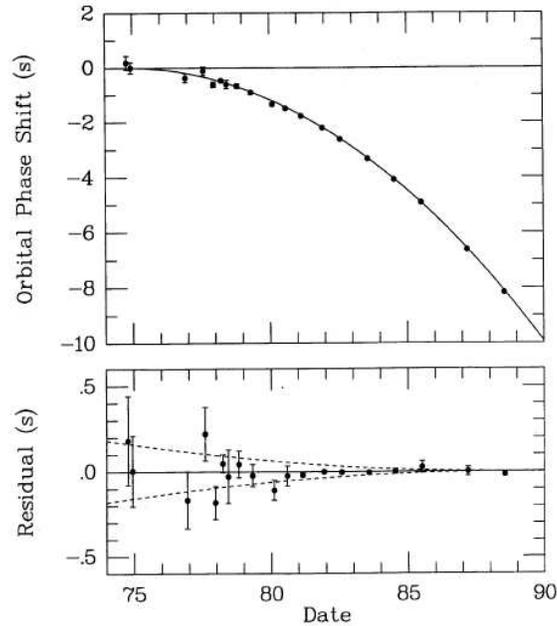


Figure 2-4. Observed shift of periastron for PSR 1913+16. The solid line is the predicted shift due to GW emission (Figure 5 from Taylor and Weisberg [11], used by permission of the American Astronomical Society)

The stunning agreement provides excellent circumstantial evidence for the existence of GWs and provided Hulse and Taylor with the 1993 Nobel Prize in physics. In addition to PSR 1913+16, several other binary pulsars have been observed. The observations of each have thus far been in agreement with the predictions of GR [12].

2.3.2 Direct Detection

While the measurements of binary pulsars provide extremely strong evidence for the existence of GWs, they do not allow the information carried by the waves themselves to be extracted. What is needed is a method to directly measure the GW strain $h(t)$. This will allow for comparison with predicted models of $h(t)$, providing tests of the models as well as providing a means to measure parameters of the systems generating the waves. In the

following subsections, the most common methods for measuring GWs are discussed. For a more exhaustive list of potential detectors see Misner, et al. [5], Chapter 37.

2.3.2.1 Doppler-tracking of spacecraft

Beginning in the 1960s, man-made probes began to leave Earth orbit and travel towards the outer planets. The radio communications systems on these spacecraft (SC) provide a means to make precise measurements of the spacetime interval between the SC and a receiver on Earth. Since both Earth and a SC in cruise phase approximate freely-falling particles, this is a measurement of geodesic separation and will be effected by GWs. The Doppler tracking technique [13, 14] begins with an ultra-stable oscillator of frequency ν_0 , which is used to drive an Earth-based transmitter. This signal travels to the distant SC, which receives a Doppler-shifted version of the signal a time τ_1 later. A phase-lock-loop (PLL) on-board the SC is used to fix the SC's local oscillator to the incoming signal. The SC then transmits this signal back to Earth, where it is received after an additional delay τ_2 and with a two-way Doppler shift $\Delta\nu$. For a coordinate system in which a GW propagates in the z -direction with the $x - y$ axes oriented parallel to the $+$ polarization (See Figure 2-5), the response of the Doppler shifts to GWs can be written as

$$\begin{aligned} \frac{\Delta\nu(t)}{\nu_0} = & -\frac{1-\mu}{2}h(t) - \mu h[t - (1+\mu)\tau_1] + \frac{1+\mu}{2}h(t - \tau_1 - \tau_2) \\ & + N_1(t) + N_2(t - \tau_1) + N_3(t - \tau_1 - \tau_2), \end{aligned} \quad (2-61)$$

where the polar angles to the SC are (θ, ϕ) , $\mu = \cos(\theta)$, N_1 , N_2 , and N_3 are noise terms and

$$h(t) \equiv h_+(t) \cos(2\phi) + h_\times(t) \sin(2\phi). \quad (2-62)$$

The response to GWs in (2-61) is sometimes called a three-pulse response, since an impulse in $h(t)$ will show up in the signal at three distinct times. For long wavelengths ($\lambda_{GW} \geq \tau_{1,2}$), the three pulses will interfere destructively. This sets the lower frequency limit for the Doppler-tracking technique.

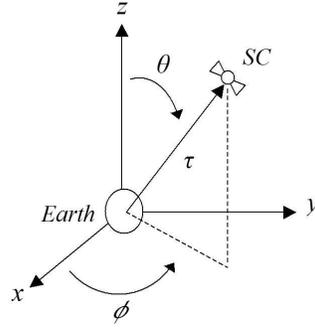


Figure 2-5. Concept for Doppler-tracking detection of GWs.

At high frequencies, the noise terms, which include scintillation in Earth’s atmosphere, scintillation in the interplanetary medium, mechanical motion of the antennae, and thermal noise in the receivers, begins to dominate the signal. For typical experiments, this sets the frequency range to be roughly $10^{-4} \text{ Hz} \leq f_{GW} \leq 10^{-1} \text{ Hz}$ [13]. The sensitivity of the signal is set by the remaining noise level in the system.

An improvement upon the simple Doppler tracking can be made by flying an additional ultra-stable oscillator on the SC and making a separate measurement of the one-way Doppler shift between Earth and the SC. Since the noise terms that enter into this measurement will be related to the noise terms in (2-61) under time shifts of τ , it is possible to create a linear combination of the Earth-SC and SC-Earth Doppler shifts that partially cancels the noise terms. It is expected that this two-way Doppler technique could provide amplitude sensitivity of 10^{-18} at frequencies around 1 mHz, corresponding to a strain spectral density amplitude of $3 \times 10^{-20} / \sqrt{\text{Hz}}$ [13]. Experiments using the one-way technique have been performed using the Pioneer SC, Galileo, Mars Surveyor, and most recently Cassini [15].

2.3.2.2 Pulsar timing

Another technique is to use a distributed array of pulsars as a timing network for GW detection [16, 17]. Pulsars are among the most precise clocks in the universe, a fact that made the indirect detection of GWs using binary-pulsars possible. In the direct

pulsar-timing technique, the pulse time-of-arrivals (TOAs) are measured for each pulsar and used to generate a model of the pulsar’s environment. This model includes all known effects on the TOAs, including relative motion between Earth and the pulsar, detector systematics, and the effect of GW emission on the pulsar orbit if it is in a binary system. A set of residual TOAs is then formed by subtracting the model TOAs from the observed TOAs. For a perfect model, the residuals will be zero. If a GW disturbs spacetime between the pulsar and Earth, the signal will show up in the residual TOAs. Roughly speaking, the relationship between TOA residuals and GW amplitude sensitivity is

$$h_{PT} \approx R \cdot \Omega_{GW} \tag{2-63}$$

where R is the rms of the TOA residuals in the absence of a GW, and Ω_{GW} is the frequency of the GW. The best current measurements have $R \sim 200$ ns, which allows for a GW sensitivity of $\sim 10^{-15}$ in the extremely low frequency band, $\Omega_{GW} \sim 1$ nHz. In theory, a single pulsar measurement would be capable of detecting GWs in this manner. In practice, multiple pulsars are needed to reduce the possibility of signals due to noise in the residuals. For N completely uncorrelated sets of residuals with equal R , the sensitivity in (2-63) will increase with \sqrt{N} . The Parkes Pulsar Timing Array [18] is a US-Australian collaboration with a goal of observing 20 pulsars with 100 ns residuals over a period of 10 yrs. This would give an amplitude sensitivity of $\sim 10^{-16}$ at $\Omega_{GW} = 1$ nHz. Sources in the extremely low frequency band include inspiraling SMBH binaries and stochastic sources such as the cosmological background.

2.3.2.3 Resonant mass detectors

The earliest GW detectors were resonant mass detectors, or “bars”, first conceived by Joseph Weber in the 1960s [19]. They consist of large masses (bars) suspended in such a way as to minimize damping. A passing GW will deposit some energy into mechanical vibrations of the bar. If this excess energy can be measured, the GW can be detected. The challenge for bar detectors is distinguishing the small amount of energy added to the bar

by a GW from the large amount of energy already present in the bar and the read-out system.

The amount of energy deposited in the bar can be increased by increasing the bar's mass. Today's bars have masses of several thousand kilograms. To prevent vibrations from the outside world from disturbing the bar, it must be mechanically well-isolated.

Cooling the bars to cryogenic temperatures reduces the thermal noise present in the bars. The remaining thermal noise can be mitigated by using materials with a high mechanical quality factor, or Q . This places most of the thermal energy in a narrow frequency band, leaving lower noise in the remainder of the band.

The second major challenge for bar detectors is measuring the energy within the bar without disturbing it. Most modern bar detectors use mechanically resonant read-out systems, which consist of smaller masses coupled to the main bar in such a way that they are resonant with the GW frequency of interest. The motion of these smaller masses is measured with electromechanical transducers built from SQUID (Sub-Quantum Interference Device) electronics.

Figure 2-6 shows a 1996 sensitivity curve for ALLEGRO, a bar detector in Baton Rouge, LA [20]. To produce the curve, the spectral density of the detector noise was scaled to equivalent GW strain amplitude. Detectable events would have a strain amplitude above the curve. The curve for ALLEGRO shows two narrow bands of maximum sensitivity, corresponding to resonances within the detector. In this sensitivity curve the maximum sensitivity reaches $10^{-21} / \sqrt{\text{Hz}}$ in a narrow band near 921 Hz. In general, size restrictions and limits on vibrational isolation limit the observational window of bar detectors to the high frequency regime ($f_{GW} \geq 100$ Hz, see Table 2-1). Persistent sources in this frequency band include rotating NSs and cosmological background. For GWs originating from optically-observed pulsars, the frequency of the GWs is known and the resonances of the bars can be tuned to search for it. Transient sources include the final merger of stellar mass binaries, supernovae, and other unmodeled sources.

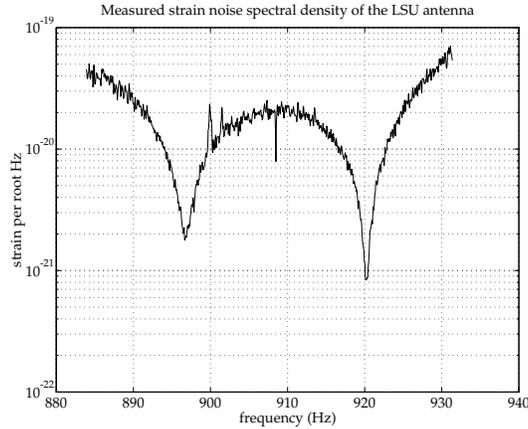


Figure 2-6. Sensitivity of ALLEGRO bar detector 1996 (courtesy of W.O. Hamilton)

To date, no confirmed detections of GWs have been made with bar detectors. A key technique for distinguishing GW signals from spurious noise bursts is coincidence measurements between multiple bar detectors. This can also help provide direction information for the observed GW, since bar detectors have nearly uniform antenna sensitivity patterns. A number of major bar experiments are underway around the world, some of which are listed in Table 2-2⁵.

Table 2-2. Operational GW bar detectors

Name	Location	Bar Temperature	Operational Date
ALLEGRO	Baton Rouge, USA	4.2 K	1991
ALTAIR	Frascati, Italy	2 K	1980
AURIGA	Lengaro, Italy	0.2 K	1997
EXPLORER	Geneva, Switzerland	2.6 K	1989
NAUTILUS	Rome, Italy	0.1 K	1994
NIOBE	Perth, Australia	5.0 K	1993

⁵ Table 2-2 is an adaptation of a similar table by Johnston [21].

In addition, there are plans to build larger detectors, including ones with spherical or nearly-spherical geometries. Spherical detectors both increase mass for a given volume as well as provide additional resonant modes that can be used to determine GW polarization and direction from a single detector [22].

2.3.2.4 Interferometric detectors

The tidal motion induced by GWs (Figure 2-1) is ideally suited to be detected with a Michelson Interferometer (MI). Consider a simple MI consisting of a light source, beam-splitter (BS), two mirrors (M_x, M_y), and a photodetector (D) oriented along the $x - y$ axes as shown in Figure 2-7. The optics are assumed to be free to move along the interferometer axes.

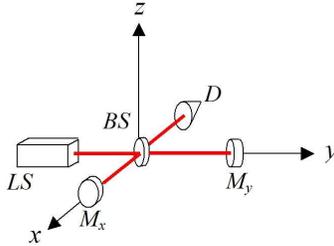


Figure 2-7. A Michelson interferometer as a detector of GWs. (LS = light source, BS = beam splitter, $M_{x,y}$ = mirrors, D = photodetector.)

If a GW propagating in the z -direction passes through the detector, the mirrors will respond as the masses in Figure 2-1, with the BS at the origin. The distances between the BS and the mirrors in the x, y arm will then be

$$L_x(t) = L_{x0}[1 + h_+(t) \cos(2\psi) - h_\times(t) \sin(2\psi)] \quad (2-64)$$

and

$$L_y(t) = L_{y0}[1 - h_+(t) \cos(2\psi) + h_\times(t) \sin(2\psi)] \quad (2-65)$$

where L_{x0} and L_{y0} are the nominal lengths of the x, y arms, $h_{+, \times}(t)$ are the polarization amplitudes of the GW, and ψ is the angle between the x -axis and the \hat{h}_+ polarization direction.

Interferometry is a technique for measuring changes in the arm-lengths given in (2-64) and (2-65). The light entering the BS can be described by an electric field oscillating at a given frequency with a phase $\phi(t)$. As each light beam makes its out-and-back trip along the arms, it will gain a phase of

$$\phi_i(t) = 2kL_i(t) \tag{2-66}$$

where $i = x, y$ and $k = 2\pi/\lambda$ is the wavenumber of the incoming light. When the light is recombined at the BS, the two beams will have a phase difference given by

$$\Delta\phi(t) = 2k\Delta L(t), \tag{2-67}$$

where $\Delta\phi(t) \equiv \phi_x(t) - \phi_y(t)$ and $\Delta L(t) \equiv L_x(t) - L_y(t)$. A number of techniques can be applied to measure the phase difference in (2-67), which can then be used to extract $h_+(t)$ and $h_\times(t)$.

For a more general relationship between the GW propagation direction, polarization, and the detector plane, the expressions in (2-64) and (2-65) will include a functions of sky position known as antenna patterns for each polarization. Except for certain orientations such as a GW propagating along the x or y axes, the antenna patterns are nearly uniform. Interferometric GW detectors are sensitive to GWs over the entire sky.

The observable frequency band for interferometric GW detectors is limited by their size and by noise sources. The expression for phase accumulation in (2-66) is valid only when the round-trip time is short compared to the GW period. For longer arms (or shorter GW periods), the sign of the GW strain will reverse as the light is propagating, causing the round-trip phase change to average towards zero. This effect acts as a low-pass filter with a $\text{sinc}(2Lf_{GW})$ transfer function. There is no fundamental limit to

the low-frequency response of a interferometric GW detector. However, practical limits are set by instrumental noise sources.

The many challenges that must be overcome in order to convert the concept in (2-7) into a working detector can be broken into two areas: the ability to build freely-falling test masses and the ability to make precision distance measurements between these test masses. The former is known in the ground-based interferometer community as displacement noise while the latter is called readout noise.

The chief source of displacement noise in interferometers is often seismic noise. An Earth-bound laboratory is not a freely-falling frame, but a free-fall condition in one dimension can be approximated by suspending the test masses on pendula. This constrains the test mass motion in the vertical and transverse directions. Longitudinally, for small displacements at frequencies above the natural frequencies of the pendula, the masses are free to move.

Vibrations can couple into the test masses through the pendula, spoiling the free-fall condition. Other sources of displacement noise are radiation pressure noise from the light on the mirrors, internal vibrations of the mirrors driven by thermal energy, and gas pressure noise. Significant effort is required to suppress these noise sources to a level sufficient for GW detection. The ultimate limit on displacement noise may be gravity gradient noise, which describes the time-dependent portions of the Newtonian gravitational field. For Earth-bound detectors with current technologies, gravity gradient noise limits the useful band to roughly 10 Hz and above.

The other class of noise sources for interferometric GW detectors is readout noise. Readout noise includes shot noise and phase noise in the light source. Shot noise can be reduced by increasing the power of the light source, however a penalty is paid in displacement noise through increased radiation pressure noise. The shot-noise/radiation-pressure noise limit represents the theoretical maximum sensitivity for a given detector operating over a broad frequency band.

Laser phase noise is a particularly important noise source for interferometric detectors. As shown in (2-67), the detector output is derived from the phase difference between the two arms, $\Delta\phi$. Changes in $\Delta\phi$ can result from two sources, a change in ΔL or a change in k

$$\delta(\Delta\phi) = 2k \cdot \delta(\Delta L) + 2\delta k \cdot \Delta L, \quad (2-68)$$

where δ denotes a time-dependant change from the nominal value. The first term corresponds to motion of the mirrors and a potential GW detection, while the second is caused by frequency (or phase) fluctuations in the light source. Note that the phase noise term is proportional to ΔL , so that in the case where the arm-lengths are exactly equal, it vanishes. To first order, an equal-arm interferometer is insensitive to phase noise of the light source.

In the past decade, several kilometer-scale interferometric detectors have been built around the globe (see Table 2-3, based on data from Johnston [21]). Multiple detectors are necessary for performing correlations in order to reduce the occurrences of false detections. In addition, the relative timing between events as measured by widely-separated detectors can provide information on the direction to the GW source.

Table 2-3. Major ground-based GW interferometers.

Project	Location	Arm Length	Status
LIGO	USA(2)	4 km	in operation
VIRGO	Italy	3 km	commisioning
GEO600	Germany	600 m	in operation
TAMA300	Japan	300 m	under construction

A representative example of a modern interferometric GW detectors are the LIGO (Laser Interferometric Gravitational Wave Observatory) detectors in the US [23]. Figure 2-8 shows an aerial photo of the 4 km LIGO detector at Hanford, WA. The effective length

of the arms is enhanced by placing Fabry-Perot cavities in each arm. A frequency and intensity stabilized Nd:YAG laser operating at 1064 nm provides ~ 6 W of power to the interferometer, which is enhanced by a factor of ~ 40 by a power-recycling mirror. The test masses are 10 kg fused silica optics suspended from a wire pendulum with a 0.75 Hz resonance frequency. Additional passive and active vibration isolation further reduce seismic coupling into the test masses.



Figure 2-8. Aerial photograph of the LIGO interferometric GW detector (courtesy LIGO Scientific Collaboration)

Figure 2-9 shows a sensitivity curve for LIGO during the science runs known as S4 (Spring 2005) and S5 (Nov. 2005 - present) along with the design goal for LIGO sensitivity. Aside from a few narrow peaks and a slight excess at low frequencies, the LIGO detectors are now performing at their design sensitivity. The peak sensitivity of $\sim 3 \times 10^{-23} / \sqrt{\text{Hz}}$ occurs just above 100 Hz. At higher frequencies the sensitivity degrades with a slope of roughly one (sensitivity is proportional to f^1). At lower frequencies the sensitivity decreases sharply, reaching $\sim 10^{-18} / \sqrt{\text{Hz}}$ at 10 Hz. This sharp decrease in sensitivity can be attributed to displacement noise, mainly from vibration coupling through the pendula as well as gravity gradient noise. This sets the LIGO observing band to the “high-frequency” band as defined by Table 2-1. LIGO is currently searching for GWs from many of the predicted sources in this band. While no confirmed detections have been made, the data has allowed important upper limits to be set [25–27].

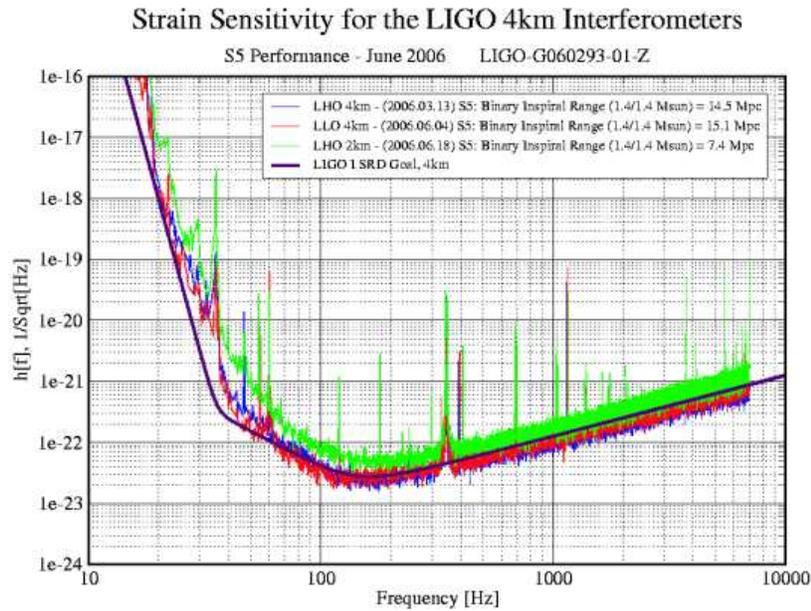


Figure 2-9. Sensitivity curves for the LIGO detectors during the S5 science run (courtesy LIGO Scientific Collaboration [24])

The sharp drop in sensitivity below ~ 30 Hz is sometimes called the “low-frequency wall” and represents a limit for current ground-based interferometric GW detectors. While future technology may allow for some improvement at low frequencies, it is unlikely that ground-based interferometers will be able to access GW sources in the low-frequency band. To do this, the detector must leave the noisy environment of Earth. One possibility is to place LIGO-type detectors on celestial bodies with less seismic activity, such as the Moon. Alternatively, the detector itself can be placed in space to avoid seismic noise completely. This is the goal of the Laser Interferometer Space Antenna (LISA), the subject of the remainder of this dissertation.

CHAPTER 3
THE LASER INTERFEROMETER SPACE ANTENNA

3.1 Introduction

As mentioned in Chapter 2, current understanding of GW detection suggests that gravity gradient noise will prevent Earth-based GW detectors from observing sources in the low-frequency band (10^{-6} Hz \sim 1 Hz, see Table 2-1). One way to access the many interesting sources in this band is to move the detector into space. The Laser Interferometer Space Antenna [28] (LISA) is a joint project of the National Aeronautics and Space Administration (NASA) and the European Space Agency (ESA) that plans to launch a space-based instrument capable of detecting gravitational radiation in the frequency band of 3×10^{-5} Hz to 10^{-1} Hz with a strain sensitivity in the range of $10^{-21}/\sqrt{\text{Hz}}$.

3.2 Sources

Three known types of sources populate the LISA observational window: galactic binaries, extreme mass-ratio inspirals (EMRIs), and SMBH mergers (see Figure 3-1). Galactic binaries refer to the early inspiral phase of stellar-mass compact objects. Since these objects will be far from merger while in the LISA band, they can be treated as persistent sources. So many of these sources are thought to exist that the LISA sensitivity will likely be limited by a confusion background of galactic binaries in some frequency regimes. The fact that GW sources are treated as “noise” to some in the LISA community is indicative of the difference in source abundance between the low-frequency and high-frequency GW bands. There are also a handful of binary systems that have been observed electromagnetically and should produce isolated GW signals in the LISA band. These “verification binaries” will provide an instant test of the instrument, as well as serving as valuable calibration sources throughout the mission lifetime [29].

An EMRI refers to a small ($1 M_{\odot} \sim 10 M_{\odot}$) compact object falling into a SMBH ($10^6 M_{\odot} \sim 10^9 M_{\odot}$). In an EMRI, the spacetime is dominated by the SMBH and the smaller object serves as a “test particle”, tracing out the geodesics of the spacetime

near the SMBH. This will allow for the first time precision tests of GR in highly-curved spacetimes, a major goal of LISA science [2].

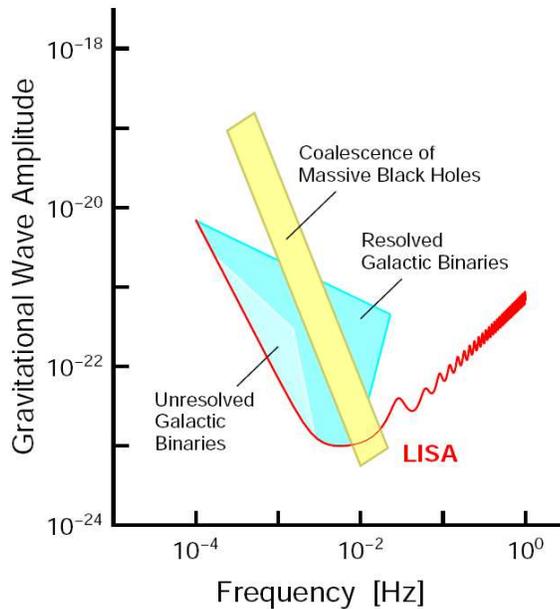


Figure 3-1. Sources in the LISA observational window (Courtesy NASA)

A merger of two SMBHs may occur during galactic collisions, when the SMBHs at the center of each parent galaxy inspiral into one another and merge. These events would be among the most energetic in the universe and would be visible to redshifts of $z \sim 5 - 10$ [2].

LISA will also have an opportunity to search for a cosmological background of GWs produced by the Big Bang. However, most models of such backgrounds that exist predict that they will lie below the LISA band.

3.3 Mission Design

The LISA mission concept calls for three individual spacecraft (SC) arranged in a triangular constellation approximately 5 Gm ($1 \text{ Gm} = 10^9 \text{ m}$) on a side. The center of the constellation will follow a circular heliocentric orbit with a radius of 1 AU, offset in orbital phase from Earth by approximately 20° , as shown in Figure 3-2. The plane of the

constellation is inclined with respect to the ecliptic plane by 60° and the constellation revolves in its plane with a period of one year.

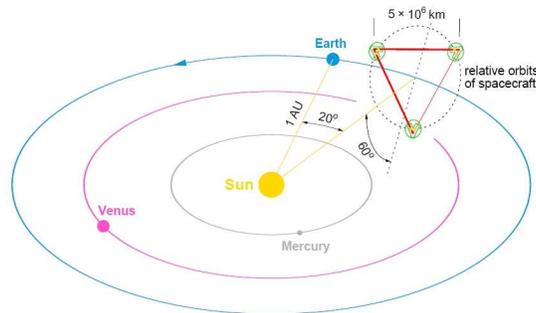


Figure 3-2. Orbital configuration of LISA constellation (Courtesy NASA)

Each SC contains two optical benches at the center of which is a 4-cm cube of gold-platinum alloy known as the *proof mass*. Like the mirrors in ground-based interferometers, the proof mass will represent the geodesic-tracking free particle in GR. Passing GWs will modulate the proper distance between the six proof-masses, an effect that will be measured using laser interferometry.

As with ground-based detectors, the challenges of LISA naturally divide into two areas: building a proof mass that approximates a freely-falling test particle and measuring the distance between the proof masses with a precision sufficient to detect the minute length changes caused by GWs. Accomplishing these tasks is the goal of the two major LISA instrumental systems, the Disturbance Reduction System (DRS), and the Interferometric Measurement System (IMS).

3.4 The Disturbance Reduction System (DRS)

A test mass in GR is completely isolated, its motion dictated solely by the geometry of the space-time in which it exists. Real objects in the universe can only approximate an ideal test mass; they are subject to electromagnetic interactions, particle interactions, and other spurious forces. Isolating the LISA proof mass from these other forces is the function of the disturbance reduction system (DRS). The LISA DRS is based on a

technology called drag-free control, in which the SC is used as a shield that flies around the proof mass. Upon reaching their designated orbits, the SC will carefully release the proof masses so that they are freely-floating inside a small enclosure. Capacitive and optical sensors will monitor the position and orientation of the proof-mass and feed this information to a controller. The controller will keep the proof masses centered in their enclosures by utilizing one of two actuators: electrostatic plates to push the proof mass and micro-Newton thrusters to move the SC. With an appropriately-designed controller, the proof masses will act much as the suspended mirrors in ground-based GW interferometers: free to move along the sensitive axis but constrained in other directions.

Design and construction of the DRS is an extremely challenging aspect of LISA technology. In order to reach the desired strain sensitivity, the residual acceleration of the proof-mass in the sensitive direction must be less than $\sim 10^{-15} \text{ (m/s}^2\text{)}/\sqrt{\text{Hz}}$. Dozens of potential noise sources such as electrostatic noises, thermal noises, and SC gravity gradient noise can spoil this goal and must be addressed. Much effort has been made to design ground-based experiments which can be used to investigate various aspects of the DRS [30]. Most of these use model proof masses suspended on torsion pendula, providing a similar system with a reduced number of degrees of freedom. In addition, an on-orbit test of the DRS technology will come with the LISA Pathfinder mission, a technology demonstrator mission planned for launch in 2009 [31].

3.5 The Interferometric Measurement System (IMS)

The other main function in LISA, measuring the distance between the proof-masses, is the role of the interferometric measurement system (IMS). The goal is to measure length changes on the order of $h \times L \approx 10^{-21}/\sqrt{\text{Hz}} \cdot 5 \text{ Gm} \approx 10 \text{ pm}/\sqrt{\text{Hz}}$ between pairs of proof-masses. While LISA is often colloquially referred to as “a Michelson interferometer in space”, in reality LISA operates quite differently from a Michelson interferometer or from any of the ground-based GW interferometers such as LIGO. Rather than utilizing many optical elements to generate a single electronic readout, LISA makes a series of

one-way interferometric measurements between pairs of optical components and then combines the results electronically to generate useful signals. The technologies required for this approach differ greatly from those required for traditional interferometry.

3.5.1 IMS Overview

Figure 3-3 contains a schematic of the LISA IMS, consisting of the three SC each with two identical optical benches. Each optical bench contains a proof mass, an infrared laser light-source, photoreceivers, and optics. Referring to the notation in the figure, OB_{ij} is the optical bench on SC_i oriented towards SC_j . The two optical benches on each SC are connected to one another via an optical fiber. Optical benches on opposite ends of a LISA “arm” are connected via a two 40 cm telescopes and a 5 Gm free-space link. Due to diffraction losses over the long arms, only about 100 pW of light are received from the ~ 1 W of light produced at the far SC.

Interferometry is used to make three types of measurements: distance between the proof mass and the optical bench, the distance between optical benches on different SC, and the phase difference between the lasers on adjacent optical benches.

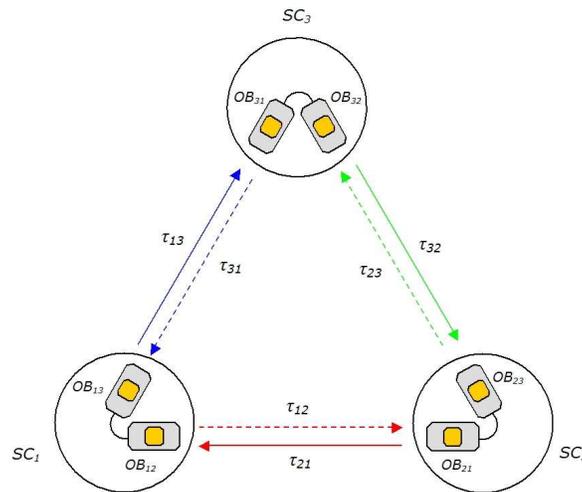


Figure 3-3. Diagram of the LISA IMS. OB_{ij} refers to the optical bench on SC_i oriented towards SC_j . τ_{ij} is the light travel time from SC_i to SC_j .

The basic premise of optical interferometry is to use the phase of a light field to make measurements of distance. In general, a light field at a particular frequency can be described by the real-part of a complex electric field given by

$$\vec{E}(t) = \vec{E}_0 \exp \{i[2\pi\nu t + \phi(t)]\} \cdot \exp[\vec{k} \cdot \hat{z}] \quad (3-1)$$

where $\Re[\vec{E}_0]$ describes the electric field amplitude and its transverse variations (spatial mode and polarization), ν is the optical frequency, and $\phi(t)$ is the optical phase, and \vec{k} is the wavevector. In (3-1), all of the variation of the light field is contained in $\phi(t)$, ν is assumed to be constant. An alternative view describes all variation in $E(t)$ as “frequency noise”, or a change in ν rather than ϕ . These two equivalent descriptions are easily related. If the frequency is described by $\nu(t) = \nu_0 + \delta\nu(t)$, the equivalent phase noise is

$$\phi(t) = 2\pi \int_0^t \delta\nu(\tau) d\tau + \phi(0) \quad (3-2)$$

and the equivalent noise spectra is

$$\tilde{\phi}(f) = \frac{\tilde{\delta\nu}(f)}{2\pi f} \quad (3-3)$$

where f is the Fourier frequency and the tilde indicates a frequency spectrum.

Since electromagnetic waves are linear, a superposition of two optical signals can be described by adding their electric fields. A photodiode (PD) can be used to measure the intensity of the combined beam, which is proportional to the squared magnitude of the total electric field. If the two signals have frequencies $\nu_{1,2}$ and phases $\phi_{1,2}(t)$, the PD output will be a signal of the form

$$S(t) \propto \sin[2\pi\nu_{12}t + \phi_{12}(t)] \quad (3-4)$$

where $\nu_{12} \equiv \nu_1 - \nu_2$ is the difference frequency between the two light beams and $\phi_{12}(t) \equiv \phi_1(t) - \phi_2(t)$ is the difference phase between the two beams. This signal is commonly referred to as a “beat note”.

In interferometry it is the phase of the beat note that contains the distance information that is of interest. Unfortunately, it also contains noise from the light source itself. In LISA the interfered light fields are produced by independent light sources, meaning that the size of the fluctuations in the beat note $\tilde{\phi}_{12}(f)$ will be roughly equivalent to the fluctuations in the individual lasers.

$$\tilde{\phi}_{12}(f) = \sqrt{\tilde{\phi}_1^2(f) + \tilde{\phi}_2^2(f)} \approx \sqrt{2}\tilde{\phi}_1(f) \quad (3-5)$$

The concept behind LISA interferometry is to measure these large phase fluctuations with high precision and then make combinations of different signals which will cancel the laser phase noise while leaving the phase fluctuations induced by GWs. The key to this approach is the ability to make phase measurements with a precision of $\sim 1 \mu\text{cycle}/\sqrt{\text{Hz}}$ of laser noise that may be greater than $10^6 \text{ cycles}/\sqrt{\text{Hz}}$ in the LISA band. This is the task of an instrument known as the phasemeter (PM). A detailed discussion of PMs, focusing on two prototypes designed for the UF LISA simulator, is presented in Section 4.5.

3.5.2 The Optical Bench

A conceptual LISA optical bench is shown in detail in Figure 3-4. Three different infrared laser beams enter the optical bench. The local beam (red), is produced by the laser associated with the optical bench. The adjacent beam (blue) is produced by the laser associated with the neighboring optical bench on the same SC, and reaches the optical bench through an optical fiber. The far beam (green), is produced by the laser associated with the optical bench on the far SC. These beams are interfered at three PDs, PD_{main} , PD_{back1} , and PD_{back2} .

In the baseline design of Figure 3-4, PD_{main} is used to interfere the incoming beam with the local beam, producing a signal containing the one-way motion between the local and far optical benches. In the “cross-over” design option, the incoming beam is interfered with the adjacent beam rather than the local beam. Since the local and adjacent beams will generally have different carrier frequencies, this will reduce the effect of stray light.

At PD_{back1} , the adjacent beam is interfered with a local beam that has reflected off the proof mass, producing a signal that contains the motion between the proof-mass and the local optical bench.

The signal at PD_{back2} is similar to that in PD_{back1} except that the local beam has not reflected off the proof mass. This signal serves as a reference signal to compare the two lasers on adjacent benches. It can also be used as the error signal in a phase-lock loop (PLL) in which the phase of the laser on one bench is forced to track the phase of the laser on the adjacent bench. This is the LISA equivalent of a beam-splitter in a traditional interferometer.

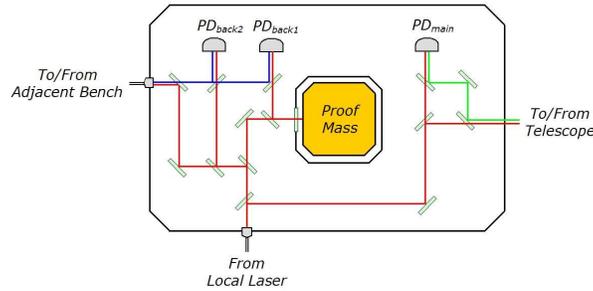


Figure 3-4. Diagram of a LISA optical bench. Light from the local laser (red) enters from the fiber coupler on the bottom, light from the adjacent optical bench (blue) enters from the left, and light from the far SC (green) enters from the right.

These three measurements are repeated on all six optical benches, resulting in eighteen independent measurements that must be properly combined in order to extract the distances between the proof masses.

Consider the two “backside” PDs (PD_{back1} and PD_{back2}) on the optical bench in Figure 3-4. A PM can be used to extract the phase of the beat signals. The PM outputs of the backside PDs are:

$$S_{back1}(t) = \phi_l(t) + kx_{pm}(t) - \phi_a(t) - n_{al}(t) \quad (3-6)$$

and

$$S_{back2}(t) = \phi_l(t) - \phi_a(t) - n_{al}(t), \quad (3-7)$$

where $\phi_l(t)$ and $\phi_a(t)$ are the phases of the light fields on the local and adjacent benches respectively, $k \equiv 2\pi\nu/c$ is the wavenumber of the light, $x_{pm}(t)$ is the motion of the proof mass in the sensitive direction relative to the optical bench, and $n_{al}(t)$ is the phase noise accumulated during travel from the adjacent bench to the local bench. This will include both noise in the fiber as well as the relative motion between the benches. By taking the linear combination $S_{back1}(t) - S_{back2}(t)$, the laser phase noise and fiber noise cancel, leaving only the term proportional to the proof mass motion.

The measurement between optical benches on opposite ends of an arm is accomplished using the signals from PD_{main} . The PM signals $S_{ij}(t)$ from PD_{main} on OB_{ij} are

$$S_{12}(t) = \phi_{12}(t) - \phi_{21}(t - \tau_{21}) + h_{21}(t) \quad (3-8)$$

and

$$S_{21}(t) = \phi_{21}(t) - \phi_{12}(t - \tau_{12}) + h_{12}(t), \quad (3-9)$$

where $\phi_{ij}(t)$ is the phase of the laser associated with OB_{ij} , $h_{ij}(t)$ is change in phase due to a GW for a beam traveling from SC_i to SC_j , and τ_{ij} is the light travel-time between SC_i and SC_j . The information of interest in $S_{ij}(t)$, the GW signal $h_{ij}(t)$, will be overwhelmed by the laser phase noise¹, $\phi_{ij}(t)$. Unlike the situation with the back-side interferometers, it is not possible to form a linear combination of $S_{ij}(t)$ and $S_{ji}(t)$ that eliminates $\phi_{ij}(t)$ while retaining $h_{ij}(t)$. The reason is that the phase noise terms enter with time delays due to the large separation between the SC.

3.5.3 Time Delay Interferometry

While it is not possible to create a signal free of laser phase noise using the PM signals on one arm, it is possible to do so, or nearly so, by using PM signals from multiple

¹ Other noise sources, such as relative motion of the SC due to non-gravitational effects, will also enter the PM signals at levels much higher than the GW signal. They are typically smaller than the laser phase noise and can be treated in a similar fashion.

arms with appropriate time delays. This process is known as time-delay interferometry [32] (TDI), and is key to the success of LISA. The three LISA arms provide a total of six one-way links, each of which can be potentially utilized to form a TDI signal. The resulting space of possible signals is large and is typically broken into several classes [33]. Some of the most basic TDI combinations are the Michelson-like combinations, typically referred to as X , Y , and Z . The three letters refer to the three-fold symmetry of LISA: the X combination is the Michelson-like combination with SC_1 as the “corner” SC, the Y combination is the Michelson-like combination with SC_2 as the corner SC, and the Z combination is the Michelson-like combination with SC_3 as the corner SC.

To form the X combination, the two lasers on-board SC_1 (Figure 3-3) are phase-locked using the signals on PD_{back2} so that $\phi_{12}(t) \approx \phi_{13}(t) = \phi_1(t)$. This is the LISA equivalent of the beam-splitter in a true Michelson interferometer.

To approximate the mirrors, the far SC (SC_2 and SC_3) are configured as optical transponders. The PM signals at PD_{main} on OB_{j1} will be

$$S_{j1}(t) = \phi_{j1}(t) - \phi_1(t - \tau_{1j}) + h_{1j}(t), \quad j = 2, 3. \quad (3-10)$$

A PLL is used to adjust the phase of the laser on the far SC so that $S_{j1}(t) \approx 0$.

Consequently,

$$\phi_{j1}(t) = \phi_1(t - \tau_{1j}) - h_{1j}(t). \quad (3-11)$$

This is the optical equivalent of the radio transponders used in the Doppler-tracking experiments discussed in Section 2.3.2.1. When the two beams return to SC_1 , the PM signal at PD_{main} on OB_{1j} will be

$$\begin{aligned} S_{1j}(t) &= \phi_1(t) - \phi_{j1}(t - \tau_{j1}) + h_{j1}(t) \\ &= \phi_1(t) - \phi_1(t - \tau_{1j} - \tau_{j1}) + h_{1j}(t - \tau_{j1}) + h_{j1}(t). \end{aligned} \quad (3-12)$$

The X combination is formed from $S_{12}(t)$ and $S_{13}(t)$ as follows:

$$X(t) = S_{12}(t) - S_{13}(t) - S_{12}(t - \tau_{13} - \tau_{31}) + S_{13}(t - \tau_{12} - \tau_{21}). \quad (3-13)$$

Substituting (3-12) into (3-13) and simplifying gives

$$\begin{aligned} X(t) = & h_{21}(t) - h_{31}(t) + h_{12}(t - \tau_{21}) - h_{13}(t - \tau_{31}) \\ & - h_{21}(t - \tau_{13} - \tau_{31}) + h_{31}(t - \tau_{12} - \tau_{21}) \\ & - h_{12}(t - \tau_{21} - \tau_{13} - \tau_{31}) + h_{31}(t - \tau_{12} - \tau_{31}). \end{aligned} \quad (3-14)$$

The X combination completely removes the contributions from laser phase noise, while leaving the contributions from the GW signals. The response of X to GWs is sometimes referred to as a “four-pulse” response since an impulse in $h(t)$ will result in an impulse in $X(t)$ at four distinct times: t , $t + \bar{\tau}$, $t + 2\bar{\tau}$, and $t + 3\bar{\tau}$, where $\bar{\tau}$ is an average of the four one-way light travel times. It is important to note that the quality of the PLLs in the transponders is not critical. In actual practice, the errors in the PLLs on each of the SC will be monitored and added into the TDI signals. A similar procedure can be used to measure and correct for any residual phase noise between $\phi_{12}(t)$ and $\phi_{13}(t)$ using the signal at PD_{back2} [34].

3.5.3.1 Visualizing TDI

The Michelson-like TDI variables can be visualized using the “rabbit-ear” diagram in Figure 3-5. Time delaying the PM signals by an amount τ can be interpreted as sending them along a virtual path with a light-travel time of τ . In the X combination, one beam (red) makes a physical trip (solid line) from SC_1 to SC_2 and back and then makes a virtual trip (dashed line) from SC_1 to SC_3 and back. The other beam (blue) does the reverse. When they return to SC_1 , both have traveled the same distance and consequently the phase noise is common and cancels out. This can be viewed as synthesizing an equal-arm Michelson interferometer or zero-area Sagnac interferometer from the individual PM signals.

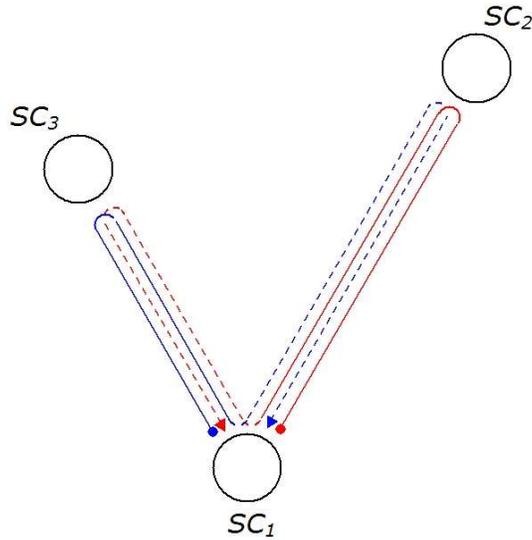


Figure 3-5. The “rabbit-ear” diagram for the first-generation TDI X combination. Each light beam originates at SC_1 and takes physical (solid) as well as virtual (dashed) trips to the far SC. The total round-trip path length for the two beams is identical and the laser phase noise in each beam at SC_1 is common.

3.5.3.2 Extensions to TDI

If the LISA constellation were static, the TDI variables such as X would perfectly cancel laser phase noise. However, as the orbits of the individual SC evolve, the constellation will change shape, causing the τ_{ij} to differ by up to one percent between arms. Additionally, the rotation of the constellation causes an asymmetry in the light travel time for a single arm, $\tau_{ij} \neq \tau_{ji}$. With time-dependant values of τ_i , laser phase noise is no longer completely canceled in the first-generation TDI variables such as X . Instead, it couples into the measurement at a level proportional to the relative velocity between the SC. One way to surmount this problem is to utilize the “second-generation” TDI variables [35], which include four additional terms and cancel out the relative velocities of the SC. These are sometimes referred to as the “eight-pulse” TDI variables, since an impulse in $h(t)$ will arrive at eight separate times in the signal. The cancellation of the second-generation TDI variables is also not perfect, since there is a relative acceleration

between the SC as well. However, the residual laser phase noise after second-generation TDI is applied is low enough so that satisfies the LISA error budget.

3.5.3.3 The zero-signal TDI variable

One important TDI variable for LISA is the symmetric Sagnac variable ζ [36]. A Sagnac interferometer consists of two beams from a common source propagating in opposite directions around a closed loop. The phase difference between the two signals is proportional to the area enclosed by the loop and the rotation speed of the loop. In LISA a Sagnac signal can be generated either physically through a combination of PLLs on appropriate benches or virtually through an appropriate combination of time-delayed PM signals. In either case, the rotation in the constellation results in a path length difference between the two beams of approximately 14 km, which is equivalent to a time-difference of $\approx 47 \mu\text{s}$. This will cause laser phase noise and other instrumental noise sources to couple into ζ . The GW signal, on the other hand, will *not* be present in ζ since the tidal distortions caused by GWs are area-preserving. What ζ provides is a measurement channel containing instrumental noise but no signal. This is essential for distinguishing between possible signals and instrumental noise since, unlike ground-based detectors, LISA will not have other detectors with which to perform correlations. In a sense, the three arms in LISA are equivalent to two co-located interferometers and the ζ variable measures the uncorrelated instrumental noise.

3.5.3.4 Limitations and noise sources

In addition to the restrictions placed on TDI by the motion of the constellation, there are sources of error that occur when the variables are formed in the first place. In order to form a variable such as X in (3–13), one must know the values of τ_{ij} . This requires an independent measurement of the range between the SC. Errors in this ranging measurement will degrade the noise cancellation in TDI. It is estimated that a ranging accuracy of 20 m to 200 m is needed to sufficiently suppress laser phase noise [34]. This

requirement is dependant on the initial laser phase noise, and can be relaxed if the input laser phase noise can be reduced through stabilization.

Another source of error in forming the TDI variables is obtaining the PM signals at the proper times. Once τ_{ij} is known from ranging measurements, signals such as $S_{ij}(t - \tau_{ij})$ must be formed. The original TDI concept called for a triggered PM fed by a real-time ranging system in order to obtain the time delayed signals at the appropriate times. This approach has since been abandoned in favor of a PM with a fixed sampling frequency of ~ 10 Hz. The PM data and ranging data are telemetered to the ground and the delayed combinations are formed by time-shifting the various PM signals. In order to achieve the required timing accuracy, the PM signals must be accurately interpolated within a small fraction of a sample period. This can be accomplished efficiently using a technique known as fractional delay filtering [37], which is discussed in Section A.4.6.

3.5.4 Arm-locking

In the previous Section, it was demonstrated how TDI will be utilized by LISA to suppress laser phase noise in the GW measurement. The requirements on TDI and its associated tasks (ranging and interpolation) are strongly tied to the input laser phase noise. LISA will use lasers that are intrinsically stable, but the laser phase noise will still be large compared with what is required. The stability of the lasers can be improved by several orders of magnitude by using an optical cavity [38, 39] or molecular transition [40, 41] as a frequency reference. The current LISA baseline calls for each laser to have an optical cavity system capable of providing a frequency stability of

$$\widetilde{\delta\nu}(f) \leq (30 \text{ Hz}/\sqrt{\text{Hz}}) \times \sqrt{1 + (1 \text{ mHz}/f)^4} \quad (3-15)$$

or better in the LISA band (3×10^{-5} Hz–0.1 Hz). With this frequency noise, second-generation TDI, and a ranging/interpolation error of ≤ 30 m, the IMS will meet its displacement noise requirements [34, 42]. Improvement in $\widetilde{\delta\nu}(f)$ over (3-15) will allow a relaxation of ranging/interpolation requirements and possibly the use of first-generation TDI variables.

Arm-locking [43] is a technique whereby some combination of the LISA arms is used as a frequency reference for laser stabilization. Although the lengths of the LISA arms change over a period of a year by as much as a few percent, in the LISA band they are extremely stable. If this stability can be transferred to the laser phase, the requirements on TDI can be relaxed considerably.

3.5.4.1 Closed-loop system dynamics

Before examining arm-locking in detail, it is useful to review the basics of closed-loop system dynamics. Figure 3-6 shows a Laplace-domain representation of a generic closed-loop system. The input $x(s)$ is combined with the control signal $c(s)$ to produce the error signal $e(s)$. The error signal propagates through the system, with transfer function $G(s)$, to produce the output signal $y(s)$. The controller, with transfer function $H(s)$, forms $c(s)$ from $y(s)$.

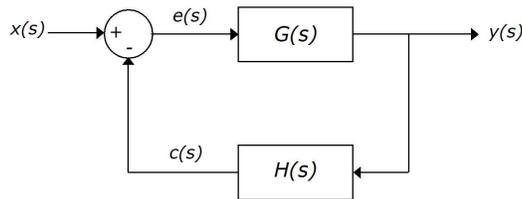


Figure 3-6. Diagram of a closed-loop SISO system with negative feedback. Signals: $x(s)$ = input, $e(s)$ = error, $y(s)$ = output, $c(s)$ = control. $G(s)$ is the system transfer function. $H(s)$ is the controller transfer function

This type of system can be referred to as a single-input single-output (SISO) closed-loop system with negative feedback [44]. The signals and the transfer functions $G(s)$ and $H(s)$ are complex-valued functions of the complex Laplace variable, $s = \sigma + 2\pi if$. The advantage of expressing the closed-loop system in the Laplace-domain is that the differential equations that relate the time-domain signals reduce to algebraic equations relating the Laplace-domain signals. The error, output, and control signals can be

expressed in terms of the input signal as

$$\frac{e(s)}{x(s)} = \frac{1}{1 + T_{OL}(s)}, \quad (3-16)$$

$$\frac{y(s)}{x(s)} = \frac{G(s)}{1 + T_{OL}(s)}, \quad (3-17)$$

and

$$\frac{c(s)}{x(s)} = \frac{T_{OL}(s)}{1 + T_{OL}(s)}, \quad (3-18)$$

where $T_{OL}(s) \equiv G(s)H(s)$ is the open-loop transfer function.

In the case of a control-loop used for stabilization, $x(t)$ is a noise input to the system and $y(t)$ is the noise in the system output. The goal of the control system is to reduce the magnitude of $y(t)$ for a given $x(t)$. Note that (3-17) contains the term $1 + T_{OL}(s)$ in the denominator. The magnitude of this term indicates the performance of the loop as a suppressor of noise. If $|1 + T_{OL}(s)| > 1$, the closed-loop value for $y(s)$ will be smaller than the open-loop value, given by (3-17) with $H(s) = 0$. If $|1 + T_{OL}(s)| \leq 1$, then the closed loop values for $y(s)$ will be greater than or equal to the open-loop value, a condition known as noise enhancement.

The performance of a particular stabilization system can be evaluated using a Nyquist plot, a plot of $T_{OL}(s)$ in the complex s -plane. As shown in Figure 3-7, the Nyquist plot has two regions separated by a circle of unit radius centered on the point $(-1, 0)$ marked with an \times . If $T_{OL}(s)$ lies inside the circle, the closed-loop system will enhance the noise. If it lies outside the circle, the closed-loop system will suppress the noise. The degree of noise enhancement or suppression is related to the distance from the point $(-1, 0)$. The closer $T_{OL}(s)$ lies to $(-1, 0)$ the larger the noise enhancement or smaller the noise suppression. If $T_{OL}(s)$ reaches the point $(-1, 0)$, known as a *pole*, the expressions in (3-16) - (3-18) become infinite.

In a Nyquist plot, $T_{OL}(s)$ is plotted as a curve parameterized by the Fourier frequency ($s \rightarrow 2\pi if$). For most systems, the curve will be a spiral with frequency increasing

clockwise. Since the gain of most systems eventually decreases with frequency, the spiral will be an inward spiral. Because of this tendency to spiral inwards, any curve that encircles the point $(-1, 0)$ will eventually become arbitrarily close to it, causing the system to become unstable. This behavior is summarized by the Nyquist stability criterion [44]: for closed-loop stability, the open-loop transfer function $T_{OL}(f)$ must not encircle the point $(-1, 0)$ in the complex plane.

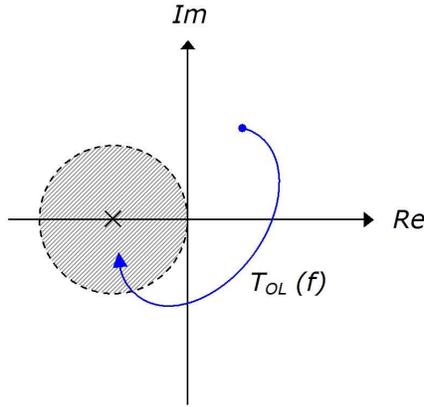


Figure 3-7. Generic Nyquist plot for open-loop transfer function $T_{OL}(f)$. The shaded region indicates noise enhancement. The pole at $(-1, 0)$ is marked by an \times . Frequency increases clockwise along the curve of $T_{OL}(f)$.

3.5.4.2 Steady-state arm-locking performance

In order to transfer the stability of the LISA arm to the laser phase, an error signal must be derived from the PM signals. Ideally, this signal would be directly proportional to the laser phase. The most basic arm-locking scheme is single-arm locking, where the round-trip length of a single LISA arm is used as a frequency reference. Consider the PM signal from PD_{main} on OB_{12} with SC_2 acting as a transponder, (3-12). In the absence of GW signals ($h_{12}(t) = h_{21}(t) = 0$), this signal can be simplified to

$$S_{12} = \phi_1(t) - \phi_1(t - \tau_{RT}), \quad (3-19)$$

where $\tau_{RT} \equiv \tau_{12} + \tau_{21}$ is the round-trip light travel time between SC_1 and SC_2 . The sensor's transfer function can be computed as

$$T_{sen}(s) \equiv \frac{\widetilde{S}_{12}(s)}{\widetilde{\phi}_1(s)} = 1 - e^{-s\tau_{RT}}. \quad (3-20)$$

The Nyquist plot of $T_{sen}(f)$, contained in Figure 3-8(a) is a circle, the result of a vector sum of the two terms in (3-20). The first term is a unit vector along the positive real axis. The second term is a rotating unit vector making an angle $2\pi f\tau$ with the negative real axis. As f increases, $T_{sen}(f)$ traces out a clockwise circle, reaching the origin at $f = f_n \equiv n/\tau_{RT}$, $n = 0, 1, 2, \dots$. These are the “null frequencies” for which the signals $\phi_1(t)$ and $\phi_1(t - \tau_{RT})$ are in phase and cancel. As the curve passes through the origin, the phase of $T_{sen}(f)$ shifts discontinuously from -90° to $+90^\circ$.

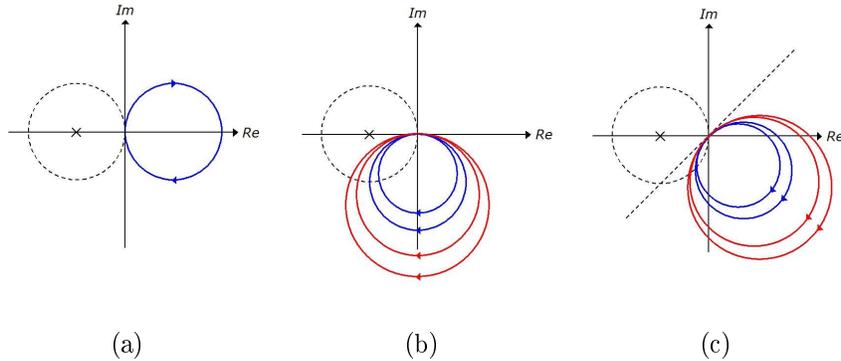


Figure 3-8. Representative Nyquist plots for single-arm arm-locking: (a) sensor $(1 - e^{-s\tau_{RT}})$, (b) system $\frac{1}{s}(1 - e^{-s\tau_{RT}})$, (c) open-loop $s^{1/2}\frac{1}{s}(1 - e^{-s\tau_{RT}})$. The red curves in (b) and (c) have higher gains than the blue curves.

In order to compute the system transfer function $G(s)$, $T_{sen}(s)$ must be combined with the actuator transfer function. Most laser actuators are frequency actuators, producing a change in laser frequency that is proportional to the controller input. A frequency actuator can be represented as a phase actuator in the Laplace domain with an additional $1/s$ in its transfer function. The system transfer function (sensor + actuator)

for single arm-locking is then

$$G(s) = \frac{G_0}{s}(1 - e^{-s\tau_{RT}}), \quad (3-21)$$

where G_0 is an overall constant gain factor. A Nyquist plot of $G(s)$ for two different values of G_0 is shown in Figure 3-8(b). Comparing the curves in Figure 3-8(b) with Figure 3-8(a), the effect of the actuator transfer function can be seen as a clockwise rotation of 90° coupled with a decrease in magnitude as frequency increases. This causes the system to enter the noise enhancement region (inside the dashed circle).

An alternative view of $G(s)$ is the Bode plot in Figure 3-9. Here the interferometer nulls and phase discontinuities are clearly seen at multiples of $1/\tau_{RT} \approx 30$ mHz.

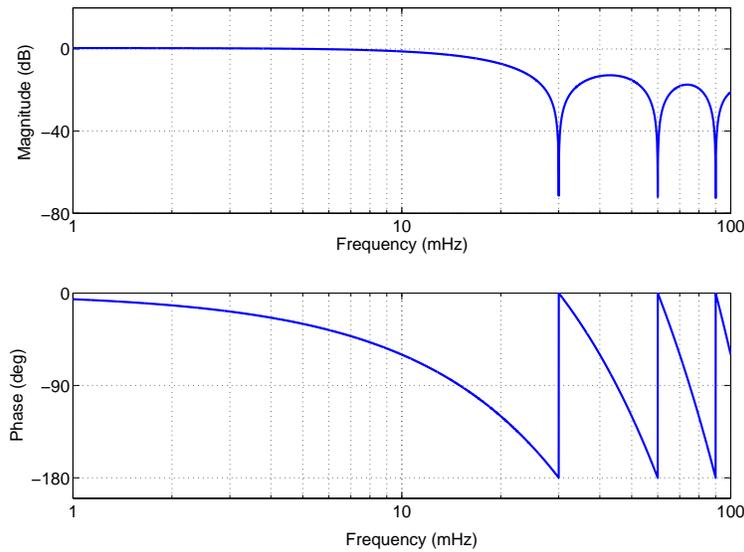


Figure 3-9. Bode plot of $T_{sen}(f)$ with $\tau_{RT} = 33$ s, the round-trip delay in a single LISA arm

As G_0 is increased (red curve vs. blue curve in Figure 3-8(b)), the system passes closer to the instability point at $(-1, 0)$. This is a marginally-stable condition where an increase in gain produces an increase in noise suppression at some frequencies but a corresponding increase in noise enhancement at other frequencies. Any additional phase loss will rotate the system further, allowing it to encompass the $(-1, 0)$ point and become unstable.

One solution to this problem is to design a controller with a transfer function whose magnitude drops below unity before $f = f_1 = 1/\tau_{RT}$. This allows large gains for $f \ll f_1$ while avoiding the instabilities at $f \geq f_1$. Unfortunately for single-arm locking, $f_1 \approx 30$ mHz lies right in the middle of the LISA measurement band.

An alternative approach is to utilize a controller that provides some phase advance in the vicinity of the interferometer nulls, rotating the Nyquist plot away from the point $(-1, 0)$. This can be achieved with a transfer function of the form

$$H(s) = H_0 s^p, \tag{3-22}$$

where H_0 is a gain constant and $0 \leq p \leq 1$. This form of $H(s)$ will produce a phase advance of $p \times 90^\circ$. A Nyquist plot of $T_{OL}(s) \equiv G(s)H(s)$ with $p = 1/2$ is shown in Figure 3-8(c). As compared to Figure 3-8(b), the curve is rotated 45° counter-clockwise. As the gain is increased (red curve vs. blue curve), the curve approaches the line $Re[T_{OL}(s)] = Im[T_{OL}(s)]$. With this type of controller, the gain can be arbitrarily increased (assuming no additional phase loss) without increasing the level of noise enhancement.

In the Bode representation (Figure 3-9), the phase response of the open-loop transfer function is equal to the sum of the system phase response and the controller phase response. The phase advance in the arm-locking controller lifts the phase minima at the null frequencies away from -180° . This provides some positive phase margin and hence stability.

The price paid for the phase advance of the controller in (3-22) is a reduced slope in the magnitude of $T_{OL}(f)$. For a given controller bandwidth, this limits the gain at low frequencies. Alternatively, a higher bandwidth is required to reach a given low-frequency gain.

A general arm-locking controller will have a transfer function similar to that in Figure 3-10. The frequency response can be divided into three distinct regions. For $f < f_1$, the controller can have a transfer function with a steep slope, allowing for large gains for

$f \ll f_1$. From slightly below f_1 to the unity-gain frequency², f_{UG} , the controller must be of the form in (3-22) so that it provides the necessary phase advance. For $f > f_{UG}$, the controller response can begin to roll off.

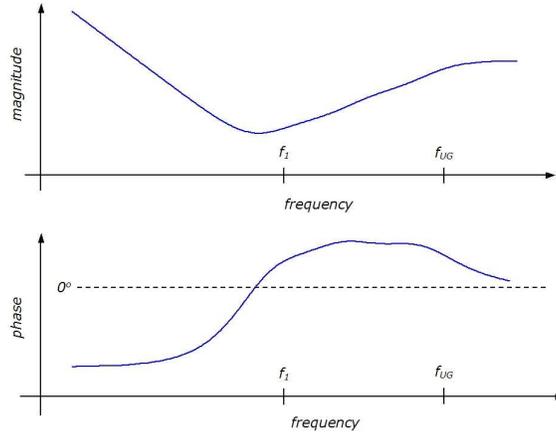


Figure 3-10. Bode plot for a generic arm-locking controller. The controller must provide a phase advance between f_1 and f_{UG}

For a general arm-locking controller with transfer function $H(s)$, the closed-loop suppression of phase noise can be computed from (3-16) and (3-21),

$$T_{CL}(s) \equiv \frac{\phi(s)}{p(s)} = \frac{1}{1 + H(s)G_0s^{-1}(1 - e^{-s\tau})}, \quad (3-23)$$

where $p(s)$ is the laser phase noise prior to arm-locking and $\phi(s)$ is the laser phase noise after arm locking. The magnitude of the closed loop suppression, $|T_{CL}(f)|$, is plotted for a generic arm-locking system in Figure 3-11.

For $f \ll f_1$, the suppression can be large, due to the steep roll-off of the controller. At a frequency just below f_1 , the closed-loop suppression crosses the 0 dB line, indicating noise enhancement. This corresponds to the curve of $T_{OL}(f)$ entering the dashed circle

² Strictly speaking, arm lock loops have many unity-gain frequencies, one for each time T_{OL} crosses the dashed circle in the Nyquist plot. There will be two of these points for each f_n where $|G_0H(f_n)/2\pi if_n| > 0$. Here we refer to the unity-gain frequency between the interferometer nulls $|G_0H(f_{UG})/2\pi if_{UG}| = 0$.

in Figure 3-8(c). The level of noise enhancement increases until it reaches a maximum, corresponding to the closest approach to the point $(-1, 0)$ in the Nyquist plot.

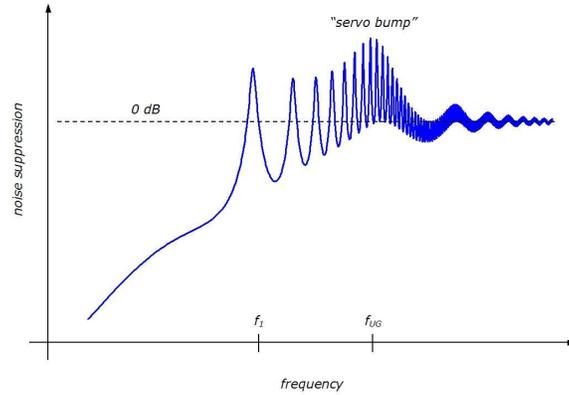


Figure 3-11. Closed-loop noise suppression for a generic arm-locking loop

At $f = f_1$, the closed-loop suppression is again 0 dB, corresponding to the curve passing through the origin in the Nyquist plot. This behavior is repeated at each f_n with the heights of the noise-enhancement peaks and the depths of the noise-suppression valleys decreasing as f increases. Near $f = f_{UG}$, a “servo bump” may occur, caused by the loss of phase in the controller as its magnitude flattens out.

3.5.4.3 Transient response

The noise suppression curve in Figure 3-11 represents a steady-state limit that is reached only after any transient responses decay. For a qualitative understanding of arm-locking transients, consider (3-19), the single-arm error signal in the time-domain. Under the assumption of high gain, the arm-locking controller will enforce the condition

$$\phi_1(t) \approx \phi_1(t - \tau_{RT}). \quad (3-24)$$

If the controller is suddenly switched on at $t = \tau_{RT}$, the phase noise $\phi(t)$ for $0 \leq t \leq \tau_{RT}$ will effectively become “frozen” in the system. Under the idealized conditions of infinite gain and no losses, the pattern of noise would repeat indefinitely. In the frequency domain,

this would correspond to all of the transient signal power being concentrated in the frequency bins near $f \approx f_n$, producing noise peaks similar to the ones in Figure 3-11.

For a real system, the relationship in (3-24) becomes more approximate as additional uncorrelated noise mixes into the system. As it does so, the transients decay and the system approaches its steady-state limit. The time constants for this decay are a critical measure of arm-locking performance. If the time-constants are too large, valuable observing time might be wasted while waiting for the noise to die down. It could also limit the flexibility of mission managers to unlock and re-lock the constellation as needed.

For a given controller, it is possible to make an analytical estimate of the transient frequencies and time-constants [45]. This is done using a Laplace-domain analysis that properly accounts for the integration constants in the Laplace-transform of the constituent signals. For the system described above,

$$f_n^{(trans)} \approx \frac{n}{\tau_{RT}} - \frac{Im\{\mathcal{G}(f_n)^{-1}\}}{2\pi\tau_{RT}}, \quad (3-25)$$

and

$$\tau_n^{(trans)} \approx \frac{\tau_{RT}}{Re\{\mathcal{G}(f_n)^{-1}\}}, \quad (3-26)$$

where $f_n^{(trans)}$ is the frequency and $\tau_n^{(trans)}$ is the time constant of the n^{th} transient, and $\mathcal{G}(s) \equiv G_0H(s)/s$ is the ratio of the open-loop transfer function to the interferometer transfer function. These expressions are valid to first order in $|\mathcal{G}(s)^{-1}|$. In order to compute the transient response for a given initial condition, the initial phase from $0 \leq t \leq \tau_{RT}$ is expressed as a Fourier sum of signals with frequencies $f_n^{(trans)}$ and amplitudes $A_n^{(trans)}$. After the loop is closed, the amplitudes will then decay with time constants $\tau_n^{(trans)}$.

Numerical simulations of arm-locking [43] suggest that the transient response may be suppressed by slowly ramping the loop gain to its steady-state value rather than suddenly turning it on at $t = \tau_{RT}$. The analytical treatment above does not apply for time-dependant gains, but the result can be intuitively understood as follows. For

the initially low gains (small G_0), the transient time constants will be small (3–26). The system will reach steady state quickly, but the steady state will only be a slight improvement over the free-running condition. The gain is then incrementally increased, increasing the time constants and decreasing the steady-state noise. However, the input noise is now slightly lower than the free-running case, corresponding to smaller initial $A_n^{(trans)}$. This process is then repeated until the final gain is reached. The overall time to reach steady state is decreased since

$$\left[1 - \exp \left(- \sum_i^N \tau_i \right) \right] > \sum_{i=1}^N [1 - \exp(-\tau_i)], \quad (3-27)$$

where τ_i represent the values of $\tau_n^{(trans)}$ for each quasi-steady-state value of the gain $G_0(i)$.

3.5.4.4 Alternative arm-locking schemes

The single-arm locking discussed in 3.5.4.2 is feasible, but its performance is less than ideal. Since $1/\tau_{RT}$ is in the LISA band, extremely large bandwidths would be needed to achieve sufficient suppression. A suppression in laser phase noise by a factor of 10^4 at 10 mHz would require a unity-gain frequency of 1 MHz for a controller with $p = 1/2$. Regardless of the controller shape or bandwidth, there will still be noise peaks in the “stabilized” spectrum that are larger than the un-stabilized noise in the same frequency bins.

The fundamental reason that these peaks are present is that the single-arm sensor signal, (3–19), contains no information about phase noise with Fourier frequencies $f = f_n$. Suppression of noise at those frequencies requires an error signal that is sensitive to noise at those frequencies. One way to obtain such an error signal is to utilize additional LISA arms. As mentioned in Section 3.5.3.4, the orbital dynamics of the constellation cause the LISA arms to differ by up to one percent. Consequently the values of f_n for one arm will be slightly different than those for another arm.

Consider the LISA constellation arranged as in the Michelson X TDI combination. SC_1 is designated the “master” SC and its two lasers are phase-locked to generate a single

light source with phase $\phi(t)$. The far SC are configured as transponders, so that the PM signals from PD_{main} on SC_1 are

$$S_{1j}(t) = \phi(t) - \phi(t - \tau_{1j}), \quad j = 2, 3. \quad (3-28)$$

If we define the average arm-length and arm-length difference as

$$\bar{\tau} \equiv \frac{\tau_{12} + \tau_{13}}{2} \quad (3-29)$$

and

$$\Delta\tau \equiv \tau_{12} - \tau_{13} \quad (3-30)$$

then (3-28) can be re-written as

$$S_{12}(t) = \phi(t) - \phi \left[t - \left(\bar{\tau} + \frac{\Delta\tau}{2} \right) \right] \quad (3-31)$$

and

$$S_{13}(t) = \phi(t) - \phi \left[t - \left(\bar{\tau} - \frac{\Delta\tau}{2} \right) \right]. \quad (3-32)$$

The transfer functions of S_{1j} in the Laplace domain can be found using (3-20) with $\tau \rightarrow \bar{\tau} \pm \Delta\tau/2$. The signals for each of the two arms can either be added or subtracted.

The former produces the common-arm error signal, S_+ , while the latter produces the difference-arm error signal, S_- . In the Laplace domain, these two signals are

$$S_+(s) = 2 \left[1 - e^{-s\bar{\tau}} \cosh \left(\frac{s\Delta\tau}{2} \right) \right] \quad (3-33)$$

and

$$S_-(s) = 2e^{-s\bar{\tau}} \sinh \left(\frac{s\Delta\tau}{2} \right). \quad (3-34)$$

Common-arm locking uses S_+ as an error signal. Figure 3-12 shows a schematic Nyquist plot of S_+ . The signal consists of two terms added vectorially, a vector of length 2 along the positive real axis and a vector of length $2|\cos(\pi f\Delta\tau)|$ making an angle of $2\pi f\bar{\tau}$ with the negative real axis. The curve will make its closest approach to the origin

for frequencies $f_n \equiv n/\bar{\tau}$. Unlike the case of single-arm locking, the error signal will not pass through the origin and consequently the error signal phase will not reach -90° . This may relax the constraints on the controller, allowing the gain to rolloff more steeply in the vicinity of the minima. For frequencies where $f = m/\Delta\tau$, $m = 1, 2, 3, \dots$, S_+ has a true null. At these frequencies, a controller must provide the same phase advance as the single-arm controller. For a one-percent arm-length difference in LISA, the first null would occur at 3 Hz, which is above the LISA measurement band.

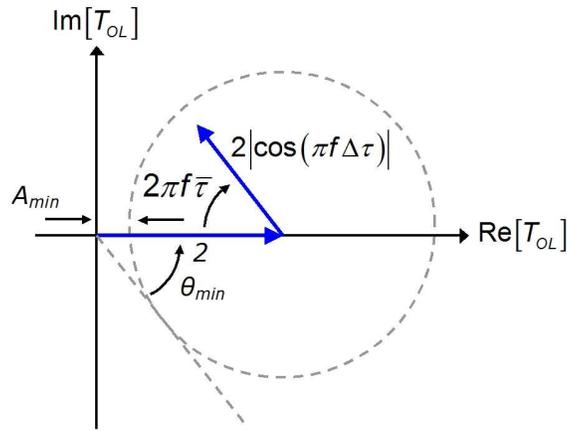


Figure 3-12. Nyquist plot for common arm-locking

The depth of the minima in S_+ and the corresponding minimum phase will depend on the value of cosine term as f approaches f_n . Since $\Delta\tau \ll \bar{\tau}$ in LISA, it is reasonable to approximate the Nyquist plot of S_+ as a series of circles centered on the point $(2, 0)$ with the radius given by $2|\cos(\pi f \Delta\tau)|$. The minimum amplitude and phase can then be estimated geometrically as

$$A_{min} \approx 2 [1 - \cos(\pi n \tau / \bar{\tau})] \quad (3-35)$$

and

$$\theta_{min} \approx \sin^{-1} [\cos(\pi n \Delta\tau / \bar{\tau})]. \quad (3-36)$$

The first concern is at $n = 1$, since the cosine function approaches unity at DC. For $\Delta\tau/\bar{\tau} = 0.01$, $A_{min} \approx 10^{-3}$ and $\theta_{min} \approx -88^\circ$. The additional 2° of phase margin provided

by common arm-locking would not cause a significant improvement in either stability or performance over single arm-locking. If $\Delta\tau$ were increased to $0.1\bar{\tau}$, the minimum amplitude would increase to ≈ 0.1 while the minimum phase would increase to -72° . This might provide some improvement, but it would also cause the location of the first true null to move from $\sim 3\text{ Hz}$ to $\sim 300\text{ mHz}$.

Given the preceding analysis of common-arm locking, it seems doubtful that it would be of much use in LISA. However, an improvement over common-arm locking known as direct arm-locking [42] looks to be more promising. The direct arm-locking error signal is formed from the common and difference error signals as

$$S_{direct}(t) = S_+(t) \pm \frac{2}{\Delta\tau} \int S_-(t) dt. \quad (3-37)$$

The sign of the integrated term depends on which arm is longer and is chosen so that the term is positive. In the Laplace domain, the direct arm-locking error signal can be written as

$$S_{direct}(s) = 2 \left\{ 1 - e^{-s\bar{\tau}} \left[\cosh\left(\frac{s\Delta\tau}{2}\right) - \frac{2}{s\Delta\tau} \sinh\left(\frac{s\Delta\tau}{2}\right) \right] \right\}. \quad (3-38)$$

An ideal arm-locking error signal would only retain the first term in the curly brackets. The second term will go to zero when the magnitude of the square-bracketed term is zero. Figure 3-13 contains a plot of the magnitude of the square-bracketed term versus Fourier frequency f . For $f \ll 1/\Delta\tau$, the term drops to zero. At these frequencies, the response of the direct arm-locking error signal is nearly flat, greatly relaxing the restrictions on controller shape. This allows for increased noise suppression at these frequencies and the removal of the frequency noise peaks.

As f increases, the square-bracketed term in (3-38) approaches unity, allowing the overall error signal to come closer to zero at the frequencies $f_n = n/\bar{\tau}$. This results in a loss of phase at these points and a corresponding increase in noise in the closed-loop system. Just before f reaches $1/\Delta\tau$, the term actually exceeds unity. In the Nyquist representation of Figure 3-12, this results in the system entering the noise enhancement

region on the left-hand side of the complex plane. The maximum amplitude of 1.06 is reached at $f \approx 0.87/\Delta\tau$. The direct arm-locking controller must provide additional phase advance at these frequencies in order to maintain the stability of the system.

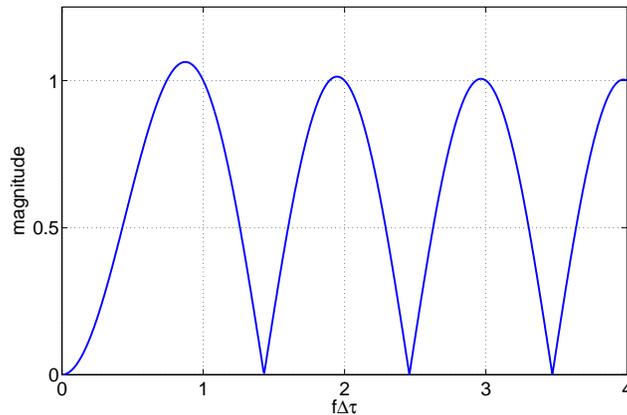


Figure 3-13. Magnitude of square-bracketed term in (3-38)

One minor problem with direct arm-locking arises in the fact that the arm-length difference is not constant and at times the arms are equal. This causes problems with scaling $S_-(t)$ by $1/\Delta\tau$ in the direct arm-locking error signal (3-37). The arms will only be equal occasionally (a few times a year) and it should be possible to address the problem by occasionally switching the location of the master SC, reverting to single arm-locking, or simply waiting for the arm-length difference to drift through zero while some scheduled maintenance activity is taking place.

A final possibility for an arm-locking error signal is to use the entire constellation in a Sagnac mode, much as for the ζ variable discussed in Section 3.5.3. Starting at the master SC, one beam is sent via phase-locks on the far SC on a counter-clockwise loop around the constellation. The other beam is sent on a clockwise loop. This produces two signals similar to those in (3-31) and (3-32), except that $\bar{\tau} \approx 50$ s and $\Delta\tau \approx 47$ μ s. These signals can also be combined to form the common and differential signals, which can be used to form a direct arm-locking error signal.

The small size of $\Delta\tau$ pushes the location of the first instability to 21 kHz, well above the LISA band. Furthermore, the rotation of the constellation is always in the same direction, meaning that $\Delta\tau$ will be constant. Sagnac arm-locking would allow for extremely large gains in the LISA band, reducing the residual phase noise considerably. The downside is that it would require all six links to be operational and quiet, whereas the other arm-locking configurations would still be viable if one or more links were inoperative or degraded.

3.5.4.5 GW signals

At first glance, there appears to be one major drawback to arm-locking: the arm-locking error signals are sensitive to phase changes caused by GWs as well as those caused by laser noise. If the signal is suppressed along with the noise, nothing is gained. However, it is important to remember that the GW signals in LISA are extracted from the TDI signals, not the individual PM outputs. The cancellation of laser frequency noise in TDI works for any distribution of laser noise. If some of the laser noise happens to be correlated with the GW signals, it will not effect the cancellation.

3.5.4.6 Interaction with pre-stabilization system

As mentioned earlier, the lasers on LISA will be externally stabilized to a local frequency reference, most likely a stable optical cavity. Ideally, arm-locking should cooperate with the local frequency stabilization, so that the maximum stability can be achieved. The problem in doing this is that the lock points for the local reference and the arm-locking reference will not generally be the same. For example, an optical cavity has a series of lock points separated by a free-spectral range (FSR) of $c/2L$, where L is the cavity length. For a 30 cm cavity, this corresponds to an FSR of 500 MHz. The linewidth of the cavity stabilized laser will be in the range of 10 Hz over 1000 s timescales. Over longer time scales, thermal and mechanical effects in the cavities can cause the lock points to drift by many MHz with slopes of 100 Hz/s or more [46].

A single-arm arm-lock controller will also have a series of lock points separated by an FSR of $c/(2 \times 5 \text{ Gm}) = 30 \text{ MHz}$, meaning that there will be an arm-locking point near the cavity lock point. Unfortunately the drifts in the cavity will cause the cavity lock point to move over thousands of arm-locking FSRs.

One possibility is to simply replace the pre-stabilization with the arm-locking loop. While this has the advantage of simplicity, it would require a much larger gain in the arm-locking loop. In addition, the PMs would have to be capable of handling the larger phase noise without losing accuracy.

Several possibilities for combining arm-locking with laser pre-stabilization have been proposed [42]. Each of these involve creating a “loop-within-a-loop” using an additional actuator. The tuneable-cavity approach, shown in Figure 3-14, replaces the fixed optical cavity with a PZT-actuated tuneable cavity. The laser is locked to the cavity using a standard locking scheme such as Pound-Drever-Hall (PDH) [47], providing the pre-stabilization. The arm-locking error signal is used to actuate the cavity, keeping it “locked” to the arm-length. The concern with this approach is that placing a PZT in the cavity will degrade the cavity’s length stability. It remains to be seen how severe this degradation will be.

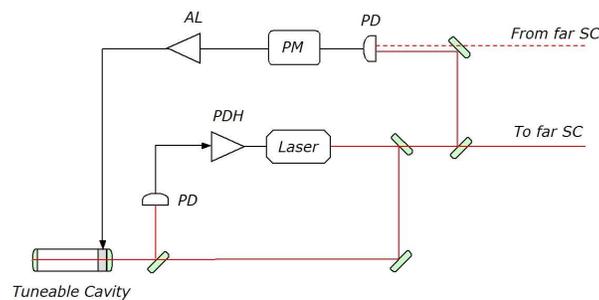


Figure 3-14. Combining pre-stabilization and arm-locking with a tuneable cavity

A second approach involves using a fixed cavity but modifying the PDH scheme to use a sideband lock. The PDH locking scheme utilizes an electro-optic modulator (EOM) to place RF sidebands on the laser beam entering the cavity. The error signal is typically

generated by making the carrier beam nearly resonant with the cavity and observing the phase of the reflected sidebands. It is also possible to generate an error signal by making one of the sidebands resonant in the cavity. In the sideband-lock approach (Figure 3-15), one RF sideband is locked to the cavity, providing stability to the carrier as well. The carrier is used to generate the arm-locking error signal, which is fed back to the local oscillator (LO) driving the EOM. This approach provides tunability while avoiding the need for a tuneable cavity. However, sideband locking is known to introduce additional noise sources and is not generally utilized in ultra-stable applications.

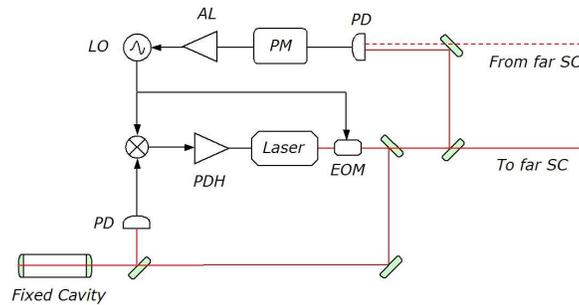


Figure 3-15. Combining pre-stabilization and arm-locking using a sideband cavity lock

A final option for combining arm-locking with laser pre-stabilization is the offset phase-lock approach, shown in Figure 3-16. Here an additional low-power laser is locked to a fixed cavity, providing a stable reference. The main high-power laser is phase-locked to this reference laser with an offset frequency provided by a tuneable LO. For a high-gain PLL, the phase noise of the main laser will be the same as that of the reference laser. The arm-locking error signal is derived from the main laser, and is used to generate a feedback to the LO. Since low-noise high-gain PLLs are an existing LISA technology requirement (transponder locks), no new technologies are needed. However the additional laser is a potential source of added mass and power consumption.

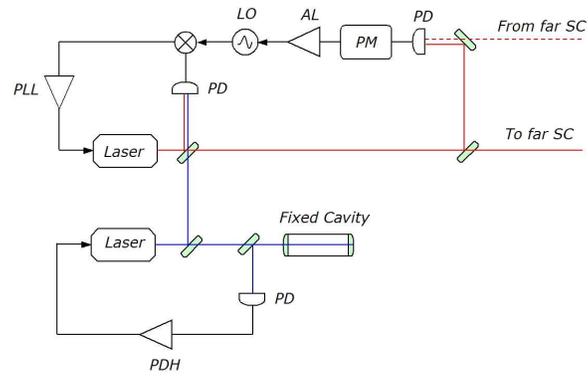


Figure 3-16. Combining pre-stabilization and arm-locking using a reference laser and an offset PLL

CHAPTER 4 THE UF LISA INTERFEROMETRY SIMULATOR

4.1 Background

A group of researchers at the University of Florida (UF), including the author, are developing a laboratory-based model of LISA interferometry. The purpose of this model is to provide an arena for studying measurement techniques and technologies under conditions similar to those present in LISA. Examples include laser pre-stabilization systems, phase-lock-loops, phasemeters, TDI, arm-locking, inter-SC ranging, and laser communication, all of which were discussed in Chapter 3. A long-term goal is to inject GW signals into the simulator and produce data streams with LISA-like noise sources that can be used by data analysis groups for mock data challenges.

There are two aspects of LISA interferometry that the simulator seeks to re-create: noise sources and transfer functions. For the most part, this can be accomplished by building a bench-top model of the IMS with a one-to-one correspondence between parts. For example, the optical cavity pre-stabilization system is built with a laser, an ultra-stable high-finesse cavity, and appropriate electronics. Standard components are used in favor of space-qualified versions to reduce cost, but the noise characteristics and transfer functions are essentially the same.

There is one aspect of LISA that cannot be simply copied on the ground: the size of the constellation. There is no viable way to optically delay a laser beam for 16 s. Delays in the microsecond or perhaps millisecond regime can be achieved using long optical fibers [48], but eventually loses and fiber noise will overwhelm the signal. To model this aspect of LISA interferometry, which is essential for studying TDI, arm-locking, and ranging, the UF simulator uses a novel technique known as Electronic Phase Delay (EPD) [49].

4.2 The EPD Concept

The motivation behind EPD comes from the fact that nearly all of the information of interest for interferometry is contained in the phase of the laser beams. Recall from (3-1) that the time-varying component of a light field can be described by a complex electric

field with a time-varying phase,

$$E(t) = E_0 \exp \{i[2\pi\nu t + \phi(t)]\}. \quad (4-1)$$

In LISA, the phase can include contributions from GW signals, laser phase noise, motion of the optical bench, etc. If $\phi(t)$ can be extracted from (4-1) and delayed, the result is interferometrically equivalent to delaying $E(t)$. Furthermore there is a limited bandwidth over which variations in $\phi(t)$ are important. For most components, this is the LISA band, but for some components, such as data communication and clock transfer, this may be as high as a few GHz.

The infrared lasers used in LISA will have wavelengths of $\sim 1 \mu\text{m}$, which corresponds to an optical frequency of $\nu \approx 280 \text{ THz}$. At such high frequencies, it is impossible to measure the oscillations in (4-1) directly. However, if two beams with different frequencies are interfered on a photodiode (PD) to produce a beat note, the resulting intensity signal is given by

$$I(t) \propto \sin[2\pi\Delta\nu t + \Delta\phi(t)], \quad (4-2)$$

where $\Delta\nu$ is the difference frequency between the beams and $\Delta\phi(t)$ is the difference in phase between the two beams. Unlike ν , $\Delta\nu$ can be made arbitrarily small by carefully tuning the lasers. If the two lasers have similar noise characteristics but are independent, then the noise spectrum of $\widetilde{\Delta\phi}(f)$ will be approximately equal to $\sqrt{2}\widetilde{\phi}(f)$. In terms of phase, $I(t)$ is an analogue of $E(t)$ at a lower frequency. If $\Delta\nu$ can be brought within the bandwidth of a photoreceiver and a digital signal processing (DSP) system, it will be possible to measure $I(t)$, store it in a digital delay buffer, and regenerate it at a later time. This is the concept behind the EPD technique.

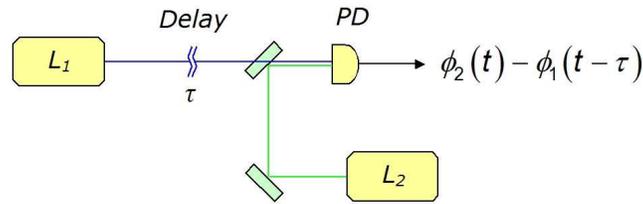
As an illustration of how EPD is applied to model LISA interferometry, consider the measurement of a single LISA arm between optical benches on different SC, shown in Figure 4-1(a). A laser on one SC (L_1), produces a light field with frequency ν_1 and phase $\phi_1(t)$. This light field traverses the 5 Gm to the other SC, incurring a delay of $\tau \approx 16 \text{ s}$ and

a frequency shift, $\nu_{Doppler}$, caused by the relative motion between the SC. At the other SC, the incoming light-field is interfered on a PD with the light from a local laser (L_2), which has a frequency ν_2 and phase $\phi_2(t)$. The frequency and phase of the resulting beat signal are

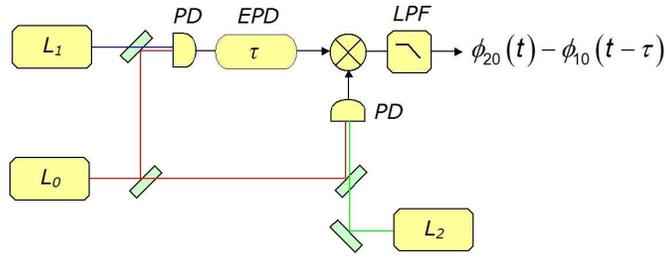
$$\nu_{LISA} = \nu_1 - \nu_2 + \nu_{Doppler} \quad (4-3)$$

and

$$\phi_{LISA} = \phi_2(t) - \phi_1(t - \tau) \quad (4-4)$$



(a) LISA



(b) EPD Analog

Figure 4-1. The EPD technique applied to a single LISA arm

The EPD equivalent of 4-1(a) is shown in Figure 4-1(b). The light field from L_1 is first interfered with a reference laser (L_0), which has phase $\phi_0(t)$. This beat signal has a frequency $\nu_{10} \equiv \nu_1 - \nu_0$ and phase $\phi_{10}(t) \equiv \phi_1(t) - \phi_0(t)$. So long as $\tilde{\phi}_0(f) \leq \tilde{\phi}_1(f)$ and the phases are uncorrelated, $\tilde{\phi}_{10}(f) \approx \tilde{\phi}_1(f)$. The $L_1 - L_0$ beat signal is the EPD analog of the optical signal from L_1 in Figure 4-1(a).

The $L_1 - L_0$ beat signal is used as an input to the EPD unit, which digitizes the signal, stores it in a memory buffer for a time τ , and regenerates the delayed signal. The output of the EPD unit is a signal with phase $\phi_{10}(t - \tau)$. Some versions of the EPD hardware (see Section 4.6) are also capable of putting a frequency shift on the beat signal, mimicking the Doppler shifts present in LISA.

The optical signal from L_2 in Figure 4-1(a) is modeled by a beat between L_2 and L_0 in Figure 4-1(b). This signal has a frequency $\nu_{20} \equiv \nu_2 - \nu_0$ and phase $\phi_{20}(t) \equiv \phi_2(t) - \phi_0(t)$. As with the $L_1 - L_0$ beat, the noise characteristics of this signal will be the same as that for the original optical signal so long as L_0 is independent from L_2 and has equal or lesser phase noise.

The PD in Figure 4-1(a) is replaced by an electronic mixer in Figure 4-1(b). The mixer performs a similar operation on the two electronic signals as the PD does on the two optical signals. The mixer output contains two terms, one with a frequency equal to the difference frequency of the two input signals and one with a frequency equal to the sum of the two input signals. A low-pass filter is used to remove the high-frequency term, leaving a signal with a frequency and phase given by

$$\nu_{EPD} = \nu_{10} - \nu_{20} + \nu_{Doppler} \quad (4-5)$$

and

$$\phi_{EPD} = \phi_{20}(t) - \phi_{10}(t - \tau). \quad (4-6)$$

A comparison of (4-5) and (4-6) with (4-3) and (4-4) indicates that the EPD model produces a signal that is of the same form as the LISA arm. With the restrictions on L_0 mentioned above, the noise characteristics will be similar as well. More complex models of LISA can be built up in a similar fashion. Table 4-1 lists the major components in LISA and their EPD equivalents.

As a final note, although Figure 4-1(b) is drawn with the same reference laser being used to generate $\phi_{10}(t)$ and $\phi_{20}(t)$, it is not a requirement of the EPD technique. Provided

all reference lasers had the appropriate noise characteristics, it would be possible to use a different reference laser for each beat note. Of course, using the same reference laser at all beat notes is more cost effective. Since the various beat notes are time-delayed before they are mixed, the noise from the common reference laser will not cancel out.

Table 4-1. Major LISA IMS components/signals and their EPD equivalents

LISA Component	EPD Equivalent
laser field	beat note with reference laser
optical delay	electronic delay
photodiode	electronic mixer
optical beat note	mixer output

4.3 Optical Components

The UF LISA interferometry simulator is designed to study primarily the SC to SC interferometry in LISA. While in theory it would be possible to include the backside interferometry by including a model proof mass suspended on a torsion pendulum [50], this would significantly complicate the experiment. Furthermore, such a combined experiment is not necessary at this stage since the SC to proof-mass and SC to SC interferometry are treated as separate measurements in LISA.

4.3.1 Layout

In the current optical layout of the simulator, shown in Figure 4-2, each of the three SC is modeled by an independent Nd:YAG non-planar ring oscillator (NPRO) laser, denoted as L_1 through L_3 in the figure. A fourth NPRO, L_0 , is used as a reference laser. The lasers L_1 and L_0 are each locked via the PDH method [47] to independent optical cavities housed in a thermally-isolated vacuum chamber. Beat notes between the far SC lasers (L_2 and L_3) are made with L_0 , allowing two complete LISA arms to be modeled [51].

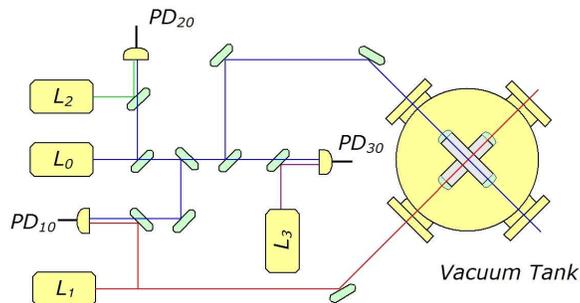


Figure 4-2. Optical layout of the UF LISA interferometry simulator. $L_1 - L_3$ represent $SC_1 - SC_3$ in LISA. L_0 is the reference laser.

4.3.2 Pre-stabilization

The function of the optical cavities in Figure 4-2 is to provide LISA-like laser phase noise for the simulator. The dominant noise source for cavity length in the LISA band is thought to be thermally-driven expansion. Consequently, the LISA cavities will likely use spacers of ultra-low expansion glass such as Dow-Corning’s ULE or Schott’s Zerodur. If the pre-stabilization requirements are relaxed due to improvements in TDI or arm-locking, it may be possible to utilize other materials such as Silicon Carbide (SiC). In parallel with the interferometry experiments, the UF group is studying the stability of various materials and bonding techniques [52]. Consequently, the optical cavities used for pre-stabilization in the simulator are occasionally changed.

Figure 4-3 shows a spectrum¹ of the beat note frequency between L_1 and L_0 for two pairs of optical cavities (Zerodur-Zerodur and Zerodur-SiC), along with the LISA pre-stabilization requirement given by (3-15). As can be seen, the Zerodur-SiC frequency noise rises steeply above the Zerodur-Zerodur frequency noise at frequencies below ~ 10 mHz.

¹ All spectra were computed using the MATLAB routine “myspd.m”, written by the author. It implements Welch’s method of overlapped-average periodograms as described by Heinzel et al. [53].

In addition, both beat notes exhibit long term drifts of tens or hundreds of MHz. For the Zerodur-Zerodur beats, the slope of the drift is typically ~ 1 Hz/s, comparable to the changing Doppler shift present in LISA. For the SiC-Zerodur beats, the slopes can be much larger, typically in the range of 100 Hz/s ~ 400 Hz/s. Over short time periods (hours to days) the drifts are monotonic but over longer time periods (days to weeks) the behavior becomes more complex. Similar drifts on the order of a few Hz/s were seen between a pair of ULE cavities in an other experiment [46], and the cause is not completely understood. For short-term interferometry experiments, the short-term drifts are of the most concern.

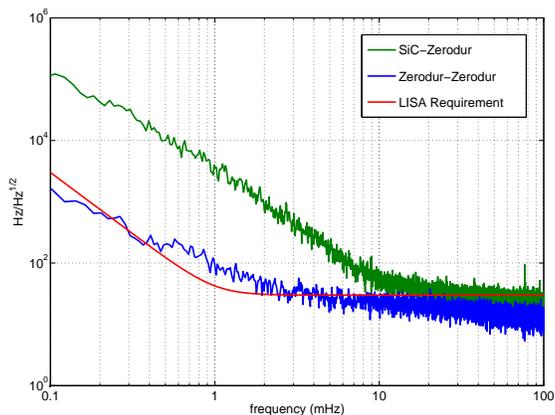


Figure 4-3. Frequency noise in the $L_1 - L_0$ beat note

As mentioned in Section 3.5, the noise in a sinusoidal signal can be placed either in the frequency or in the phase. Figure 4-4 shows the same data as Figure 4-3 converted to phase noise using (3-3).

4.4 Electronic Components

In addition to the optical components described in the previous Section, several electronic components are key to the operation of the simulator. These can be broken down into three primary categories: the control filters, phasemeters (PMs), and the EPD unit. Control filters are used to provide actuator signals for the laser-pre-stabilization, phase-lock loops, and arm-locking loops. PMs are used to measure the phase of the various

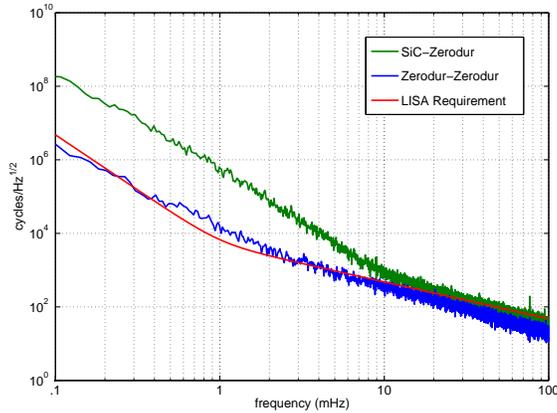


Figure 4-4. Data from 4-3 converted into phase noise by dividing by $2\pi f$, where f is the Fourier frequency

beat signals in the optical layout, providing inputs to the control loops as well as the science signals.

Both the control filters and the PMs are critical components for LISA as well. Since the simulator versions and their LISA counterparts are subjected to similar signals, the simulator provides an excellent arena for evaluating potential designs. The EPD unit is of course unique to the simulator, and must reproduce the LISA arm as faithfully as possible. The addition of noise not present in LISA or elimination of noise present in LISA would limit the accuracy of the simulator.

Several types of electronic architectures are employed in the UF simulator. These include both analog systems and digital systems, which are reviewed in Appendix A. For the most demanding electronic subsystems, such as the PMs and EPD unit, the UF group selected and purchased a high-speed digital signal processing system from the Pentek Corporation in Upper Saddle River, NJ. An overview schematic of the system is shown in Figure 4-5 below.

The system consists of three individual products, the model 4205 carrier board, the model 6256 digital downconverter, and the model 6228 digital upconverter. The model 4205 carrier board is housed in a VME crate and contains a 1 GHz, 32-bit PowerPC

microprocessor. The processor is connected via a PCI bus to several components, including SDRAM, serial ports, Ethernet ports, and the VME backplane. In addition the model 4205 provides four specialized high-speed interfaces known as velocity interface module (VIM) connectors.

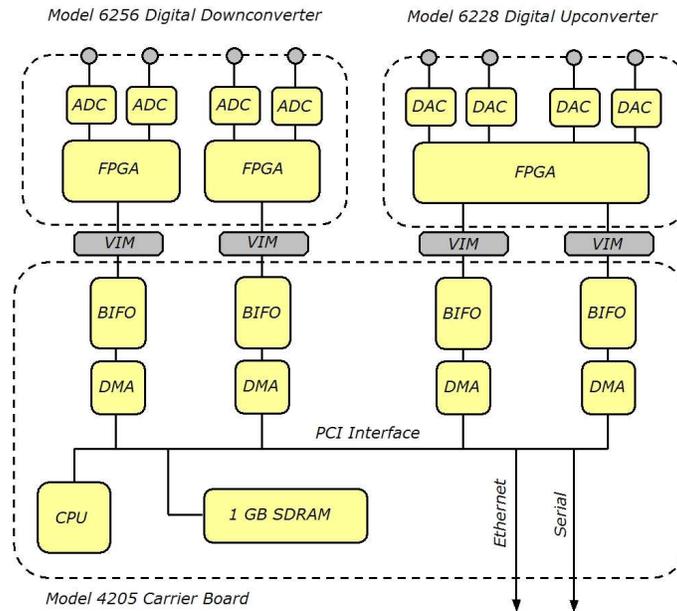


Figure 4-5. Overview of the DSP system from Pentek Corporation

The VIM interface consists of a 32-bit data interface and a 32-bit control/status interface. The data interface is connected to the main PCI bus via a bi-directional first-input first-output buffer (BIFO) and a direct-memory access (DMA) controller. The DMA controller allows data to be read from the BIFO and directly deposited into memory or another location on the PCI bus without the need for processor intervention. The speed of the VIM interface is limited by the speed of the PCI bus, which is clocked at 66 MHz. The control/status interface is connected to the processor via a separate 33 MHz PCI bus.

The 6256 digital downconverter contains four 14-bit ADCs that can be clocked at frequencies up to 105 MHz. The ADCs are connected to front-panel connectors via RF transformers with a high-pass frequency response. The -3 dB point of the transformers is at 400 kHz. Signals at lower frequencies cannot be measured with the 6256. The full-scale

input of the ADCs is reached with a +4 dBm signal at the front-panel input, which has an impedance of $50\ \Omega$.

Data from the ADCs is passed into one of two FPGAs, where it can be processed. The processed data is connected via two of the VIM interfaces to the BIFOs on the 4205. The VIM control/status interface also connects with the FPGA and can be used to configure the board.

The 6228 digital upconverter connects to the 4205 via the second pair of VIM modules. Data from the BIFOs passes into an FPGA where it is processed. The processed data is fed to two two-channel 16-bit DACs that can be clocked at frequencies up to 500 MHz. The DACs produce a full-scale output of $-2\ \text{dBm}$ and are coupled to the front-panel via RF transformers with a $50\ \Omega$ output impedance and a $-3\ \text{dB}$ point at 400 kHz. As with the 6256, the VIM control/status interface connects with the FPGA and can be used to configure the board.

This arrangement provides a powerful and flexible system for digital signal processing. The FPGAs on the 6256 and 6228 can be used to perform high-speed processing with fixed-point arithmetic, while the microprocessor can be used for floating-point processing at lower speeds. This system is used for the EPD units (Section 4.6), the hardware PM (Section 4.5.5), and one instance of the arm-locking control filter (Section 5.4).

4.5 Phasemeters

4.5.1 Overview

The phase of a sinusoidal signal can be specified by the time at which the signal has a specific value and specific first derivative (i.e. positive zero-crossing). Since this is fundamentally a timing measurement, all phasemeters (PM)s must be based on a reference clock. The phase measured by the PM is the difference phase between the signal and the reference clock. Any single phase measurement will be limited in accuracy by the phase stability of the reference clock. With multiple measurements using the same

reference clock, the phase noise of the reference clock can be measured and canceled. In the remainder of this Section, phase noise in the reference clock is not explicitly included.

A simple PM that is often used in the laboratory is the analog mixer. The signal to be measured, $S(t)$, can be described by a sinusoid with amplitude A , angular frequency ω , and phase $\phi(t)$,

$$S(t) = A \sin[\omega t + \phi(t)]. \quad (4-7)$$

This signal is mixed with a reference signal, $R(t)$, with the same frequency and a constant phase,

$$R(t) = \cos[\omega t]. \quad (4-8)$$

The mixer output, $M(t)$, is the product of (4-7) and (4-8),

$$\begin{aligned} M(t) &= S(t) \times R(t) \\ &= A \sin[\omega t + \phi(t)] \cos[\omega t] \\ &= \frac{A}{2} \{ \sin[\phi(t)] - \sin[2\omega t + \phi(t)] \}. \end{aligned} \quad (4-9)$$

The mixer output is then low-pass filtered to remove the 2ω term. If the phase noise is small and the signal amplitudes are constant, then the resulting signal is proportional to $\phi(t)$. Some analog mixers are specifically designed to compensate for the sinusoidal response and can produce a linear phase response for $|\phi(t)| \leq 70^\circ$.

The filtered mixer PM is an example of a general type of PM called in-phase/quadrature or IQ PMs. Mathematically, a sinusoidal wave at a given frequency contains two pieces of information. These can be expressed as the amplitude and phase, as was done in (4-7) or as the in-phase, $I(t)$, and quadrature, $Q(t)$, components:

$$A(t) \sin[\omega t + \phi(t)] = I(t) \sin[\omega t] + Q(t) \cos[\omega t]. \quad (4-10)$$

Trigonometric identities can be used to derive the following relationships between the IQ and amplitude-phase formalisms:

$$I(t) = A(t) \cos[\phi(t)], \quad (4-11)$$

$$Q(t) = A(t) \sin[\phi(t)], \quad (4-12)$$

$$A(t) = \sqrt{I(t)^2 + Q(t)^2}, \quad (4-13)$$

$$\phi(t) = \arctan[Q(t)/I(t)]. \quad (4-14)$$

Comparison of (4-12) with (4-9) reveals that the filtered mixer PM described above measures $Q(t)/2$ rather than $\phi(t)$. A complete IQ PM can be built by extending the filtered mixer concept to include two demodulations, one with $\cos[\omega t]$, which produces $Q(t)/2$, and one with $\sin[\omega t]$, which produces $I(t)/2$. The relations in (4-11) - (4-14) can then be used to compute $\phi(t)$ and $A(t)$.

In addition to direct mixing, $I(t)$ and $Q(t)$ can be measured in a number of other ways. One technique involves sampling the signal of interest with a sampling frequency equal to four times the carrier frequency of the signal [54]. Each set of four data points can be manipulated to measure the phase at a rate of one-half the carrier frequency. Another approach is to use integral transforms of the time series data to extract the phase [55].

One type of PM that is distinct from the IQ type is the counter/timer PM [56]. In this technique, the number of zero-crossings in a time interval T is counted, providing a crude estimate of the signal frequency. This estimate is then corrected by measuring the additional time between first and last zero-crossings and the time-interval boundaries. A combination of these two measurements gives an estimate of the phase accumulated during the interval T . This approach is used in frequency counters such as the ones used to measure the beat note stabilities discussed in Section 4.3.2.

4.5.2 Phasemeters for LISA-like signals

The accuracy requirement for the LISA PM is set by the error budget for the IMS to be $\sim 10^{-6}$ cycles/ $\sqrt{\text{Hz}}$ in the LISA measurement band (Shaddock [57] gives $3 \mu\text{cycles}/\sqrt{\text{Hz}}$ at 5 mHz). Reaching this level of performance requires careful suppression of other noise sources that present themselves as phase noise. Examples include phase and amplitude noise in the reference signal, residual signals from the second harmonic term in (4-9), and, for digital systems, digitization and quantization effects. It has been shown that commercial digital radio receivers are capable of meeting the LISA phase accuracy requirements for low-noise signals at fixed frequencies [59].

The input signals to the PM in LISA differ from those for a commercial radio receiver in two important ways: large shifts of the carrier frequency, and large intrinsic phase noise on the signal. The relative motion between the SC will cause the one-way Doppler shifts to vary by up to 30 MHz over the course of a year. If the frequencies of the SC lasers are held fixed, the beat frequencies will also vary over a range of 30 MHz. This range can be reduced by periodically adjusting the laser frequencies during the course of the year. One proposed frequency plan will keep the beat notes in the range 2 MHz – 20 MHz.

Of comparable size to the Doppler shifts are the drifts in the optical reference cavities discussed in Section 4.3.2. These drifts will be more difficult to model and must be taken into account when specifying the PM range requirements.

In addition to the frequency drifts, which can be considered as noise in the beat note below the LISA band, there is also a large amount of phase noise in the LISA measurement band, as evidenced by Figure 4-4. This large phase noise poses a problem to the IQ phasemeter. As shown in (4-11) and (4-12), $I(t)$ and $Q(t)$ are periodic in $\phi(t)$. Therefore the IQ phase measurement in (4-14) is a measurement of $\phi(t)$ modulo 2π and only gives $\phi(t)$ if $|\phi(t)| \leq \pi$ rad. In other words, the rms phase noise must be less than half a cycle in the measurement band or the phase noise measurement will “wrap”. The LISA laser phase noise requirement (Figure 4-4) corresponds to an rms phase noise of greater

than 10^6 cycles in the LISA band, meaning that the phase noise will constantly wrap between $\pm 1/2$ cycles. Addressing phase wrapping is a major challenge to the designers of the LISA PM.

4.5.3 An IQ phasemeter with a tracking LO

One way to address both the frequency drifts and the large laser phase noise is to demodulate the incoming signal with a local oscillator that tracks the phase of the incoming signal to within $\pm 1/2$ cycle. This is the approach taken by researchers at the Jet Propulsion Laboratory who are designing the LISA PM [57, 58]. A schematic of such a system is shown in Figure 4-6 below.

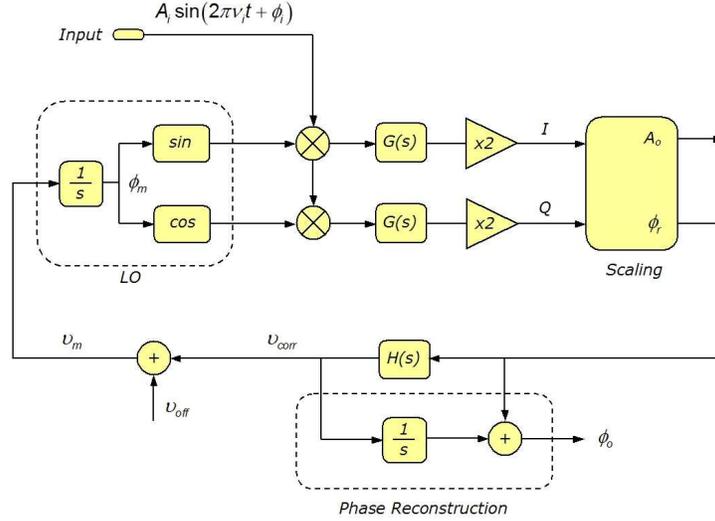


Figure 4-6. Schematic of a IQ phasemeter with feedback

The input signal is a sinusoid with frequency ν_i , amplitude $A_i(t)$, and phase $\phi_i(t)$. The frequency is assumed to be fixed and any frequency noise is converted to phase noise using (3-2). The input signal is demodulated with two signals from a local oscillator (LO), a cosine and a sine with model phase, $\phi_m(t)$. The sine term is filtered by a low-pass filter with a transfer function $G(s)$ and scaled by 2 to form the signal $I(t)$,

$$I(t) = 2G(t) \otimes \left\{ \frac{A_i(t)}{2} (\cos[\phi_i(t) - \phi_m(t)] - \cos[4\pi\nu_i t + \phi_i(t) + \phi_m(t)]) \right\}, \quad (4-15)$$

where $G(t)$ is the filter's impulse-response function, and \otimes denotes convolution. The cosine term is similarly filtered and scaled to form $Q(t)$,

$$Q(t) = 2G(t) \otimes \left\{ \frac{A_i(t)}{2} (\sin [\phi_i(t) - \phi_m(t)] + \sin [4\pi\nu_i t + \phi_i(t) + \phi_m(t)]) \right\}. \quad (4-16)$$

With a properly designed filter, the high-frequency terms in (4-15) and (4-16) can be eliminated while retaining the first terms. The additional scaling by a factor of 2 produces the standard definitions of $I(t)$ and $Q(t)$.

The $I(t)$ and $Q(t)$ signals can be used to compute the *residual phase*, defined as $\phi_r(t) \equiv \phi_i(t) - \phi_m(t)$, and the output amplitude, $A_o(t)$, using (4-13) and (4-14). If $\phi_r(t)$ is small enough to linearize the equations without introducing unacceptable errors, the relations become

$$\phi_r(t) \approx Q(t)/I(t) \quad (4-17)$$

and

$$A_o(t) \approx I(t). \quad (4-18)$$

The residual phase is used as an error signal for the LO tracking loop. It is filtered by a control filter with transfer function $H(s)$, forming the control signal for the LO. For most types of oscillators, the control signal is proportional to the oscillator's frequency. This adds an implicit factor of $1/s$ into the controller transfer function.

In most cases, the approximate frequency of the beat signal will be known, and an offset frequency, ν_{off} , can be added to the frequency correction, $\nu_{corr}(t)$, provided by the control filter. The model frequency, $\nu_m(t)$, is the sum of ν_{off} and $\nu_{corr}(t)$ and is integrated to form the model phase, $\phi_m(t)$,

$$\nu_m(t) = \nu_{off} + \nu_{corr}(t) \quad (4-19)$$

$$\phi_m(t) = \int \nu_m(t) dt. \quad (4-20)$$

To form the PM output phase, $\phi_o(t)$, $\phi_m(t)$ must be reconstructed and added to $\phi_r(t)$. The constant linear drift arising from the ν_{off} term in (4-20) does not contain any information of interest. Consequently, in the phase reconstruction only $\nu_{corr}(t)$ is included in the integration². The output phase is then

$$\phi_o(t) = \phi_r(t) + \int_0^t \nu_{corr}(t') dt' + \phi(0), \quad (4-21)$$

where $\phi(0)$ is an arbitrary reference phase, typically set to zero once the tracking loop has been closed. The output phase is a measurement of the phase of the input signal relative to a ideal (zero phase noise) sinusoid with frequency ν_{off} . A linear trend in $\phi_o(t)$ corresponds to a difference between ν_{off} and ν_i .

As with other closed-loop systems, it is useful to form a Laplace-domain model that can be used to analyze the PM's dynamics. Such a model for the system in Figure 4-6 is shown in Figure 4-7. To build the model, it is assumed that $|\phi_r(t)| \ll 1$, so that $Q(t) \approx A_i(t)\phi_r(t)/2$. The mixing and filtering process can then be replaced by a subtraction of phases followed by a filtering by $G(s)$.

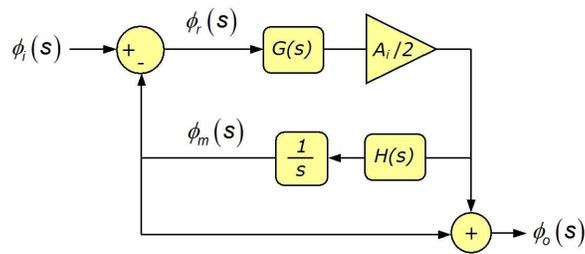


Figure 4-7. Laplace domain model of the system in Figure 4-6

² For discrete-time systems, a running sum of ν_{corr} scaled by the sample time of ϕ_m is equivalent to an integration.

Using the simplified model in Figure 4-7, the following relationships can be derived between the phase variables.

$$\phi_r(s)/\phi_i(s) = \frac{G'(s)}{1 + H'(s)G'(s)}, \quad (4-22)$$

$$\phi_m(s)/\phi_i(s) = \frac{H'(s)G'(s)}{1 + H'(s)G'(s)}, \quad (4-23)$$

$$\phi_o(s)/\phi_i(s) = \frac{G'(s) + H'(s)G'(s)}{1 + H'(s)G'(s)}, \quad (4-24)$$

where $G'(s) \equiv (A_i(t)/2)G(s)$ and $H'(s) \equiv (1/s)H(s)$. The PM error, defined as the difference between the output phase and the input phase,

$$\phi_e(t) \equiv \phi_o(t) - \phi_i(t), \quad (4-25)$$

will have a spectrum given by

$$\phi_e(s)/\phi_i(s) = \frac{G'(s) - 1}{1 + H'(s)G'(s)}. \quad (4-26)$$

Given an input phase noise spectrum, such as the laser noise in Figure 4-4, the relations in (4-22) and (4-26) can be used to design the filters $G(s)$ and $H(s)$ so that $|\phi_r(t)| < 1/2$ cycle and the phase accuracy requirements are met.

The chief function of $G(s)$ is to eliminate the effect of the second terms in (4-15) and (4-16). In frequency space, these terms will be a peaks centered at a frequency of $2\nu_i$ with linewidths related to the input phase noise spectrum, $\tilde{\phi}_i(f)$.

In general, the frequency $2\nu_i$ will be far away from the measurement band (> 10 MHz vs. 1 mHz for LISA), so the direct effect of the peak is not of much concern. The only requirement is that the peak be reduced to much less than $1/2$ cycle so that the residual phase will not wrap. The picture changes somewhat if the signals in (4-15) and (4-16) are sampled at a rate lower than 2ν (or down-sampled in the case of a digital PM). This

allows the frequency noise in the peak to alias into lower frequency bands³. If $G(s)$ also functions as the anti-aliasing filter, it must be carefully designed to minimize aliasing of the peak and other high-frequency phase noise into the measurement band.

The transfer function $G(s)$ should also be flat in the measurement band, to ensure that the PM output is constant for all phase-modulation frequencies of interest. This gain is typically scaled to 2 at DC so that a signal with amplitude $A_i(t) = 1$ will produce signals with the standard definitions of $I(t)$ and $Q(t)$.

The design of $H(s)$ is driven by the shape of $G(s)$ and the input phase noise. The goal of the design of $H(s)$ is to ensure that the residual phase does not wrap. This can be accomplished if

$$\int_{LF}^{\infty} \tilde{\phi}_r(f) df = \int_{LF}^{\infty} \left| \frac{G(f)}{1 + G(f)H(f)} \right| \tilde{\phi}_i(f) df \leq 0.5 \text{ cycle}, \quad (4-27)$$

where LF is the low-frequency limit of the measurement band. At low frequencies, the gain of $H(f)$ must increase at least as fast as the input phase noise. For the phase noises in Figure 4-4, this corresponds to a slope⁴ between f^{-2} and f^{-3} . The unity gain frequency must occur at a point past any peaks in $\phi_i(t)$ that are comparable to 0.5 cycles. The unity gain frequency is typically limited by phase loss associated with $G(s)$.

4.5.4 A Software Phasemeter

4.5.4.1 Design

As an initial investigation into IQ PMs with feedback, a software-based PM was built using MATLAB's SIMULINK environment. The PM was designed to analyze time series data from a laser beat note demodulated to approximately 10 kHz and sampled at a rate of 80 kHz. Building the PM in SIMULINK allowed for flexibility in the design of filters

³ For a discussion of the phenomenon of aliasing, see Section A.2.1.

⁴ Recall that one factor of f^{-1} is automatically present due to the conversion from frequency to phase in the LO.

and also limited the effects of digitization and quantization noise. Furthermore, it allowed for the PM to be tested by generating input signals with known phase noise from within SIMULINK.

The PM follows the form in Figure 4-6. The filter $G(s)$ is a 16-tap ($N = 16$) FIR filter with a passband of 10 kHz and a stopband of 18 kHz. It was designed in MATLAB using the equiripple design technique which specifies that the ripple in the passband match that in the stopband. A Bode plot of $G(s)$ is shown in Figure 4-8.

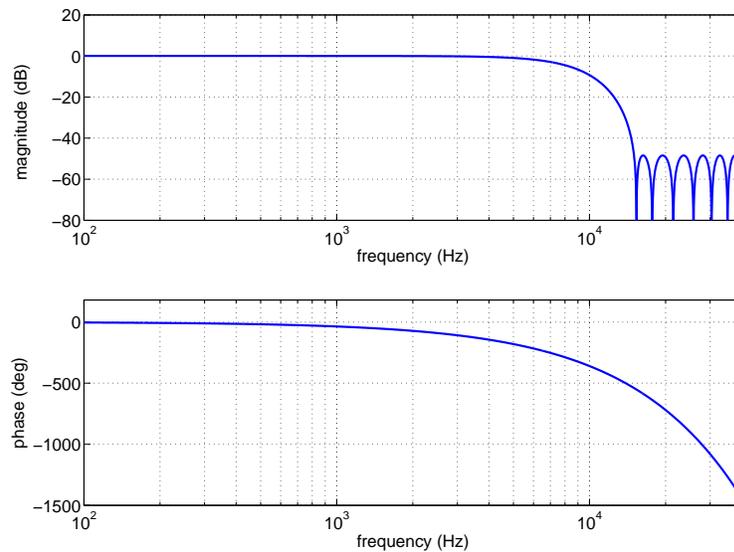


Figure 4-8. Bode plot of $G(s)$ for the software PM

The feedback filter $H(s)$ is designed to have an f^{-1} response for $200 \text{ Hz} \leq f \leq 20 \text{ kHz}$ and an f^{-2} response elsewhere, including the $1/s$ from the frequency to phase conversion in the LO. This can be accomplished with two poles at DC (one in the LO) and a pole-zero pair with the zero at 200 Hz and the pole at 20 kHz,

$$H(s) = 5 \cdot 10^8 \frac{(s - 2\pi \cdot 200 \text{ Hz})}{s^2(s - 2\pi \cdot 20 \text{ kHz})}. \quad (4-28)$$

A Bode plot of $H(s)$ is shown in Figure 4-9. The unity-gain frequency of the tracking loop in the software PM is approximately 2 kHz.

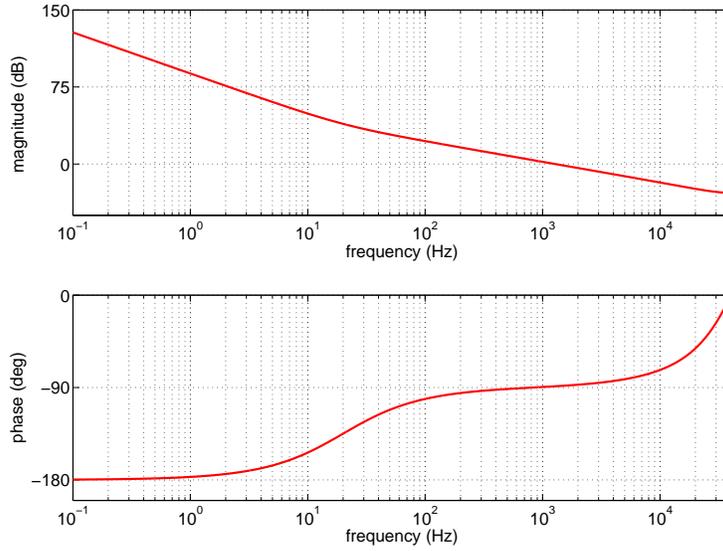


Figure 4-9. Bode plot of $H(s)$ for the software PM

4.5.4.2 Results

To test the software PM, an input phase noise signal with a phase noise spectral density of approximately $(1 \text{ cycle}/\sqrt{\text{Hz}}) \times (1 \text{ Hz}/f)$ was formed by filtering white noise with a digital filter having a single pole at 10 mHz. This phase noise was then placed on a carrier frequency of $\sim 10 \text{ kHz}$ using a second LO to produce an input signal. Using (4-22)-(4-26) with $G(s)$ and $H(s)$ as specified in Figure 4-8 and Figure 4-9, the spectral densities $\tilde{\phi}_r(f)$, $\tilde{\phi}_m(f)$, $\tilde{\phi}_o(f)$, and $\tilde{\phi}_e(f)$ can be predicted. These predictions are shown in Figure 4-10 .

To test these predictions, the simulation was run and $\phi_i(t)$, $\phi_r(t)$, $\phi_m(t)$, $\phi_o(t)$, and $\phi_e(t)$ were recorded for a 100s simulation time. To reduce the size of the generated data, the signals were down-sampled to a 10 kHz data rate using a cascaded-integrator-comb (CIC) decimation filter⁵ . The spectral densities of each signal are plotted in Figure 4-11.

⁵ For a discussion of CIC filters, see Section A.4.5.

A comparison of Figure 4-10 and Figure 4-11 reveals that, with the notable exception of the $\tilde{\phi}_e(f)$, the observed behavior matches that of the expected behavior. The fact that for frequencies below 20 Hz, $\tilde{\phi}_e(f)$ actually exceeds $\tilde{\phi}_r(f)$ indicates that the problem may lie in the reconstruction of the model phase.

The integration of the frequency correction $\nu_{corr}(t)$ used to form $\phi_m(t)$ is not exactly the same as the integration that takes place inside the LO. Inside the LO, the integrator uses a modulo- 1 cycle accumulator to compute the phase, since the phase is only used to compute sinusoids with values that repeat every 2π . Furthermore, the accumulator inside the LO includes the offset frequency ν_{off} . For ideal arithmetic operators, these differences do not matter, but it is possible that the numerical errors in the two cases differ. This seems unlikely in the case of the software PM since SIMULINK utilizes double-precision floating-point arithmetic.

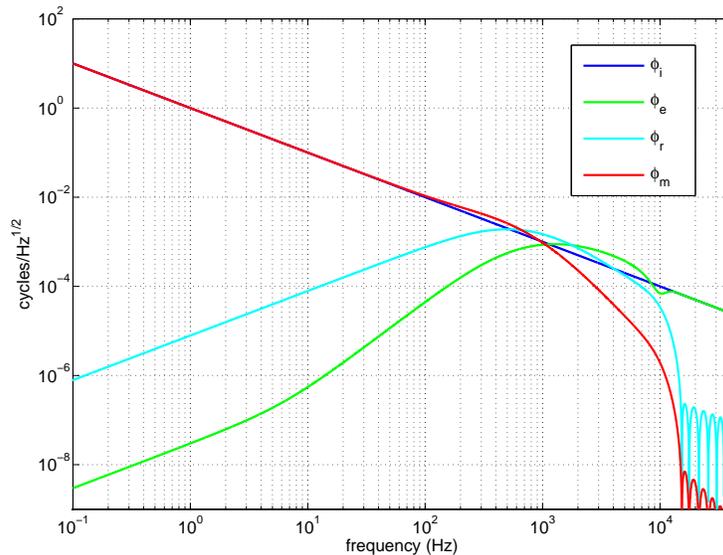


Figure 4-10. Expected behavior of the software PM for signal with input phase noise with linear spectral density $(1 \text{ cycle}/\sqrt{\text{Hz}}) \times (1 \text{ Hz}/f)$

Another possibility is an error in reconstruction due to a time shift between the reconstructed model phase and the residual phase. In order to accurately compute the output phase, the residual phase output must be delayed by whatever throughput

delay exists between the control filter and the output of the LO. In the SIMULINK environment, this delay should be well-known and has been corrected for.

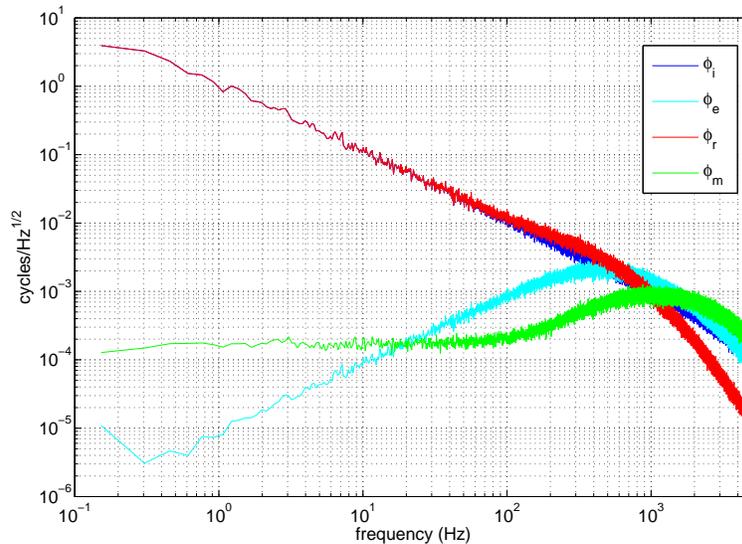


Figure 4-11. Observed behavior of the software PM

The software PM was successfully used for a number of simulator experiments, including investigations of TDI [51] and arm-locking [52]. These experiments were mainly proof-of-principle measurements, and the excess noise floor in Figure 4-11 was not much of a concern. One issue that did arise were the size of the data sets required to reach the LISA band. A signal digitized with 16-bit resolution at a frequency of 80 kHz produces data at a rate of 156 kB/s. For the TDI measurements, which required two signals recorded for several minutes, the data files reached 100 MB. This created issues with data storage as well as the length of time it took to process the data through the software PM.

It is likely that with further work, the problems with the software PM could have been addressed. However, once the initial success of the PM concept was verified, the focus was shifted to building a real-time hardware PM. This is the topic of the next Section.

4.5.5 A Real-time Hardware Phasemeter

The hardware PM was built using the DSP system described in Section 4.4. The design of the hardware PM, shown in Figure 4-12, is similar to the design described in 4.5.3. To take advantage of the different architectures present in the Pentek system, the PM function is broken into two parts. The *front-end*, implemented on the model 6256 daughterboard, tracks the incoming signal and generates the signals $I(t)$, $Q(t)$, and $\nu_{corr}(t)$. The *back-end*, implemented on the model 4205 carrier board, uses these signals to reconstruct the input phase and perform further processing.

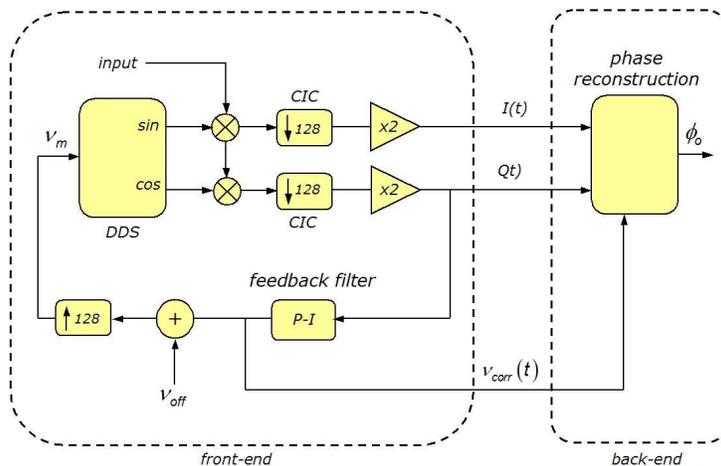


Figure 4-12. Schematic of the real-time hardware PM

4.5.5.1 Front-end design

The LO on the hardware PM is a type of numerically-controlled oscillator known as a direct digital synthesizer (DDS), shown in detail in Figure 4-13. The input to the DDS is the phase increment register, a fractional number equal to the fraction of cycles to advance per clock period. This corresponds to the model frequency, ν_m . For example, a phase increment of 0.1 with a 100 MHz clock frequency would correspond to $\nu_m = 10$ MHz. The frequency resolution of the DDS is

$$\delta\nu_m = \frac{2^{-N}}{T_{clk}}, \quad (4-29)$$

where T_{clk} is the clock period and N is the number of bits in the phase increment register. The phase increment is an unsigned binary fraction with 32 total bits and 32 fraction bits (U32.32)⁶. For a 100 MHz clock frequency, $\delta\nu_m \approx 23$ mHz.

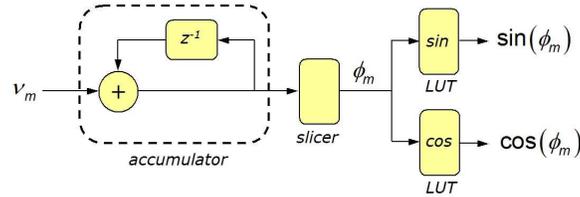


Figure 4-13. Schematic of a Direct Digital Synthesizer (DDS)

The phase increment is used as the input to an accumulator, which performs a running sum of the phase increment value at each rising clock edge. The accumulator variable is also U32.32 and wraps to zero on overflow. The value of the accumulator corresponds to the phase of the DDS in cycles. The accumulator output is “sliced” (re-quantized to a lower bit resolution) to a U10.10 and used to feed the address bits of two look-up tables (LUTs). The purpose of the slicer is to reduce the amount of memory needed for the LUTs while still preserving the frequency resolution of a larger phase increment width. The price paid is an increase in the amplitudes of spurious harmonics in the output of the LUT [60]. For the purposes of the PM, a 10-bit address depth in the LUT results in sufficiently low harmonics. Each LUT contains one cycle of sine or cosine waveforms in S14.13 format.

The input signals from the model 6256 ADCs, expressed as S14.14, are multiplied with the DDS outputs. These signals are then filtered by a 2-stage CIC filter with a decimation rate of 128. CIC filters, discussed in detail in Section A.4.5, are an efficient type of multi-rate digital filters often used to achieve large sample-rate changes with minimal aliasing. Section A.4.5 derives the magnitude response for a CIC decimator with

⁶ For an overview of binary fractions, see A.3.1.

a decimation factor of R . The magnitude response has nulls at the centers of the aliasing bands, $f = (k/R)f_s$, $k = 1, 2, 3 \dots$

Figure 4-14 contains a plot of $|G(f)|$ for the hardware PM around the first aliasing band, centered at $f_s/128 = 781.25$ kHz. The amplitude of the filter stays below 10^{-7} for a band of half-width ~ 250 Hz around the central null. No more than 10^{-7} of the phase noise at higher frequencies will be aliased into the band $0 \text{ Hz} \leq f \leq 250 \text{ Hz}$ by the CIC filter.

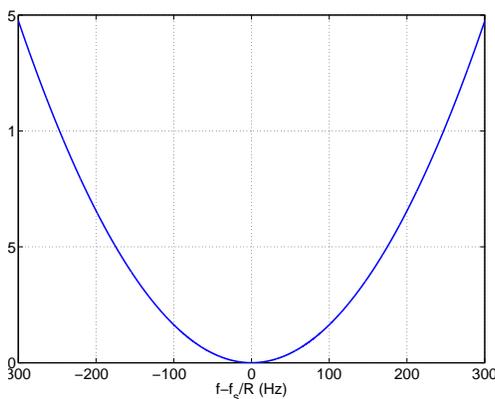


Figure 4-14. Magnitude response of the CIC decimation filter in the hardware PM near the first aliasing band at 781.25 kHz

The other concern for the magnitude response of the decimation filter is the passband flatness. Figure 4-15 contains a plot of $1 - |G(f)|$ near DC. As can be seen from the figure, the passband flatness is better than 10^{-7} up to ~ 135 Hz. This ensures that any coupling between phase noise frequency and measured phase noise amplitude will be minimal.

The phase of the CIC filter is linear, with an equivalent group delay of $\approx 1.3 \mu\text{s}$. After exiting the CIC filter, the data are scaled by a factor of two and formed into S16.16 words corresponding to the standard definitions of $I(t)$ and $Q(t)$ given in (4-11) and (4-12).

The hardware PM uses the $Q(t)$ signal rather than $\phi(t)$ as the error signal for the tracking loop driving the DDS. This is done to avoid the need to perform division or arctangent operations, both of which are problematic for fixed-point systems. The downside of this approach is that the tracking-loop gain scales with the signal amplitude.

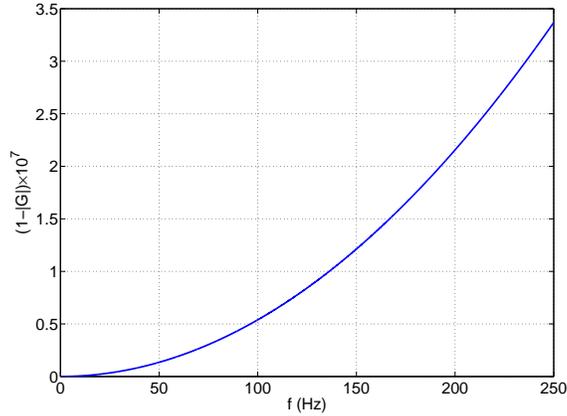


Figure 4-15. Passband flatness of the CIC filter in the hardware PM

This is not a serious problem so long as the gain in the tracking loop is easily adjustable and the signal amplitudes remain relatively constant during the measurement time.

The feedback filter consists of two parallel paths, a proportional path and an integral path. The integral path consists of an accumulator that performs a running sum of the error signal. The running sum operation is related to a true integral by T , the clock period of the accumulator

$$T \sum_{n=0}^N x(nT) \approx \int_0^{NT} x(t) dt. \quad (4-30)$$

A simple running sum is therefore equivalent to a continuous-domain integrator with an addition gain of $(1/T)$. The integral path of the feedback filter has an additional gain of 2^{-7} so that the overall transfer function can be written as

$$H_1(s) = 1 - 2^{-7} \frac{100 \text{ MHz}}{128} \frac{1}{s} \approx \frac{s - 970 \text{ Hz}}{s}. \quad (4-31)$$

In other words, the feedback filter is equivalent to a pole-zero pair with a pole at DC and a zero at 970 Hz. The output of the feedback filter is scaled by an overall gain, H_0 (U16.16), to produce an S32.32 signal containing the frequency correction, $\nu_{corr}(t)$. The frequency correction is in the same units as the phase increment for the DDS, cycles per clock period.

A frequency offset (U16.16) is added to $\nu_{corr}(t)$ and the result is up-sampled by 128 to generate the phase increment input to the DDS. Including the accumulator in the DDS, which has a clock frequency of 100 MHz, the overall feedback filter has the transfer function

$$H(s) = H_0(6 \times 10^{11} \text{ Hz}) \frac{(s - 970 \text{ Hz})}{s^2}. \quad (4-32)$$

For full-scale inputs ($A_i \approx 0.5$) a gain of $H_0 \approx 0.0005$ is used, corresponding to a unity-gain frequency of ~ 8 kHz. Figure 4-16 shows a plot of $|H(f)|$ computed using (4-32) for $H_0 = 0.0005$. Also shown is an observed value for $|H(f)|$ made from 5 s of laser beat note data. To compute $|H(f)|$, the measured model phase spectrum was divided by the measured residual phase spectrum. From (4-22) and (4-23) it is clear that $|\tilde{\phi}_m| / |\tilde{\phi}_r| = |H(f)|$.

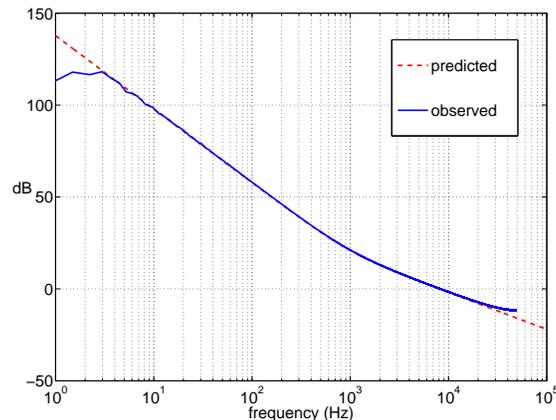


Figure 4-16. Feedback filter for hardware PM tracking loop. Predicted results were computed from (4-32), observed results were computed as a ratio of model phase to residual phase during a 5 s data run.

The outputs of the PM front end are $I(t)$, $Q(t)$, and $\nu_{corr}(t)$ signals ($IQ\nu$). These outputs are either transmitted as is at a rate of 781.25 kHz or are down-sampled in a second 2-stage CIC filter to 97.65625 kHz (an additional division by 8). This reduces the 10^{-7} anti-aliasing band to 30 Hz and the 10^{-7} flat passband to 17 Hz.

As shown in Figure 4-5, the model 6256 has two FPGAs, each of which is associated with two ADCs and one VIM interface. Consequently, each FPGA contains two PM

front-ends, one for each channel. The data for the two channels is then packed into a block of four 32-bit words using the packing format shown in Figure 4-17 and transferred across the VIM interface to the BIFO on the model 4205 carrier board. Due to the packing scheme, the average clock rate on the VIM is four times the sampling rate of the PM data.

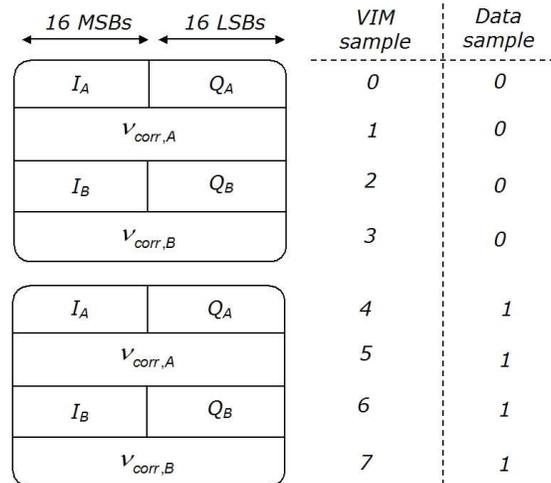


Figure 4-17. Packing format for PM data transferred over the VIM interface

The current front-end design occupies approximately 10% of the resources on the model 6256 FPGA (total for both channels). There is room to add several additional channels to allow for PM tracking of multiple tones for clock transfer, data transmission, etc. on the same input signal. Doing so would require a modification of the packing format in Figure 4-17. The primary bottleneck is the speed of the VIM and the processing in the PM back-end.

4.5.5.2 Back-end design

The first task of the PM back-end is to read the data placed on the BIFO by the PM front-end. This task must be accomplished with sufficient speed to prevent the BIFO from overflowing. The processor-side of the BIFO is connected to a 66 MHz PCI bus, which allows the processor to periodically read the BIFO at a high rate and then perform other tasks while waiting for the BIFO to re-fill. The raw data is read from the BIFO in large blocks of words, typically 1024. This corresponds to 256 samples of the $IQ\nu$ signals. This

block of data is then unpacked into six data streams, corresponding to the $IQ\nu$ signals for each of the two channels associated with the VIM.

The $IQ\nu$ signals are used as the inputs to a phase reconstruction algorithm. The first step in the algorithm is to convert the $IQ\nu$ signals into double-precision floating-point representation. This reduces quantization error in subsequent arithmetical operations and allows for operations such as division and arctangent. The reconstruction algorithm has three modes: arctangent, no arctangent, and model-only, summarized in Table 4-2. The arctangent mode is the most accurate, but is also the most computationally intensive. The no arctangent mode utilizes the approximation $\tan x \approx x$ for $x \ll 1$. The model-only mode, the simplest possible reconstruction mode, relies on the large gain of the tracking loop in the front-end. The error in the model-only mode is the residual phase, $\phi_r = \text{atan}(Q/I)$.

Table 4-2. Reconstruction algorithms for the hardware PM. R is the overall decimation rate between the DDS clock rate and the $IQ\nu$ data rate.

Mode	Algorithm
arctangent	$R \sum \nu_m + \text{atan}(Q/I)$
no arctangent	$R \sum \nu_m + Q/I$
model-only	$R \sum \nu_m$

The behavior of the PM back-end can be easily adjusted by running a different program on the model 4205. In the simplest versions, the user enters configuration data (offset frequencies, track-loop gains, etc.) via a serial terminal, the PM runs for a specified amount of time, and the data is transferred to a host computer via FTP. This data can be the raw data from the BIFO, the unpacked $IQ\nu$ data, the reconstructed phase data, or some processed version of the reconstructed phase. More sophisticated versions provide the user with a graphical user interface (GUI) for configuring the PM and continually stream the data via TCP/IP to a program running on a host PC.

4.5.5.3 Single-signal PM test with a VCO

Unlike the software PM, which could be provided with a known input phase, tests of the hardware PM necessarily involved real signals. Consequently effects such as clock jitter, ADC noise, amplitude noise, etc. were present in the tests, as they would be in any experimental environment. The first set of tests involved using a single oscillator, either a commercial function generator, analog voltage controlled oscillator (VCO), or laser beat note. The oscillator signal was split into multiple parts using RF splitters and run into separate PM channels.

The first test used a 25 MHz VCO demodulated to 10 MHz with a function generator as an input signal. The electronic beat signal was amplified using a $\times 10$ RF amplifier and split into four channels using three 50-50 RF splitters. The power of the function generator was adjusted so that the total signal power in each channel was ≈ 4 dBm, the full-scale input of the PM. Figure 4-18 contains a spectrum of the original signals and of three residuals from the pairwise subtractions. The lowest residual, labeled $\Delta\phi_{same}$ was obtained by subtracting pairs of signals that share a VIM interface (channel 1/2 and channel 3/4). The residuals between pairs on different VIM interfaces (channel 1/3, channel 1/4, channel 2/3, and channel 2/4) were larger by a factor of ~ 10 at high frequencies (curve labeled $\Delta\phi_{diff}$ in Figure 4-18).

One possibility for the higher noise would be a time-lag between the measured phase signal from each channel. The residual signal with a time delay is

$$\Delta\phi(t) = \phi(t) - \phi(t - \Delta t). \quad (4-33)$$

In the frequency domain, the magnitude of the residual noise is then

$$\left| \widetilde{\Delta\phi}(f) \right| = \left| \widetilde{\phi}(f) \right| 2 \sin(\pi f \Delta t). \quad (4-34)$$

This relationship can then be used to estimate Δt . For the curve labeled $\Delta\phi_{diff}$ in Figure 4-18, the estimate was $\Delta t \approx 6 \mu s$. The source of such a large Δt is not readily apparent. It

corresponds to many clock periods of the ADC and FPGA and is therefore unlikely to be due to any clock jitter between the two devices. One possibility is that when the back-end of the PM requests data from the front-end, there is a slight delay in requesting the data from the second VIM. If this is indeed the source of the delay, it might be reduced by synchronizing the data requests.

The data in Figure 4-18 were sampled at ~ 98 kHz, corresponding to a sample period of $\sim 10 \mu\text{s}$. The estimated delay Δt corresponds to 0.6 samples. Shifting the data by less than a sample period can be accomplished using fractional delay filtering, discussed in A.4.6. The curve labeled $\Delta\phi_{shift}$ in Figure 4-18 was obtained by shifting the data from channel 3 by 0.6 samples using a 51-point fractional-delay filter with a Lagrange window. As can be seen, the shift reduced the residual noise to nearly the level of $\Delta\phi_{same}$.

The curves labeled U_{dig} in Figure 4-18 represent an estimate of the digitization noise level present in the PM. The source of the digitization noise is the finite precision of the $IQ\nu$ data produced by the PM front-end. Section A.3.4 derives a formula (A-19) that estimates the linear spectral density of digitization noise for a specific sampling rate and bit resolution. Applying this formula for the $IQ\nu$ data gives

$$\delta\nu_{dig} = \frac{100 \text{ MHz} \cdot 2^{-32}}{\sqrt{6 \cdot 98 \text{ kHz}}} \approx 3 \mu\text{Hz}/\sqrt{\text{Hz}}, \quad (4-35)$$

$$\delta I_{dig} = \delta Q_{dig} = \frac{2^{-16}}{\sqrt{6 \cdot 98 \text{ kHz}}} \approx 2 \times 10^{-8}/\sqrt{\text{Hz}}, \quad (4-36)$$

where the units of I and Q are the full-scale amplitude of the ADC. Both of these noise sources contribute to phase noise. The frequency noise will scale to phase noise according to (3-3), producing a $1/f$ noise spectrum. The I and Q terms will remain flat. The total digitization noise for a single channel is given by

$$U_{dig1} \approx \sqrt{(\delta\nu_{dig}/2\pi f)^2 + (\sqrt{2}\delta Q_{dig}/2\pi)^2}. \quad (4-37)$$

For a two-channel measurement, the digitization noise floor is given by $\sqrt{2}U_{dig1}$. Two curves for digitization noise are plotted in Figure 4-18. The black curve corresponds to

the digitization noise for a system with a 32-bit DDS. Both $\Delta\phi_{same}$ and $\Delta\phi_{shift}$ follow the 32-bit version of U_{dig} , indicating that the PM performance is likely limited by digitization noise. Digitization noise prevents the current version of the hardware PM from meeting the LISA phase accuracy requirements at low frequencies. If the DDS accumulator width was increased to 42 bits (dashed red curve in Figure 4-18), the requirements could be met. Converting to a 42-bit accumulator will cause the front-end to occupy slightly more FPGA resources as well as require a re-design of the VIM data packing scheme in Figure 4-17.

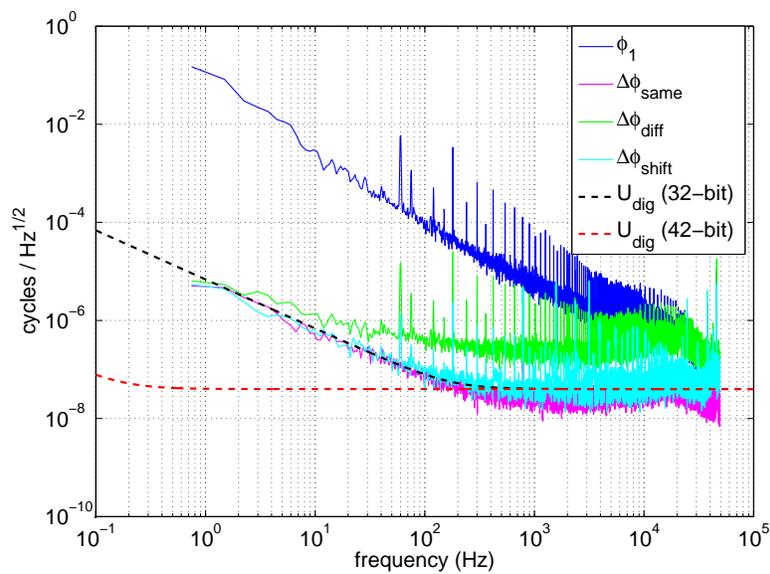


Figure 4-18. VCO phase noise measured by four channels of the hardware PM

4.5.5.4 Single-signal PM test with optical signals

A similar measurement was made using a beat between two lasers locked to independent cavities. A portion of the laser light from the laser locked to the optically-contacted Zerodur cavity on the main optical table (Figure 4-2) was focused into an optical fiber and transmitted to a second optical table. At the second optical table, the light from the fiber was superimposed on a photodiode with the light from a laser locked to a hydroxide-bonded Zerodur cavity. The light coming from the fiber had a power of $\approx 5 \mu\text{W}$ while the local light had a power of $56 \mu\text{W}$.

The resulting beat note had a frequency of 113.4 MHz and a power of -21.3 dBm. To bring the frequency in the range of the PM inputs, the beat note was demodulated using a local oscillator at 130 MHz. The local oscillator and the PM clock were both locked to a Rubidium-stabilized 10 MHz reference oscillator. The demodulated beat signal was split into two portions using a 50-50 RF splitter and amplified using an RF amplifier with an amplitude gain of 10. Figure 4-19 contains the linear spectral density of the input phase noise as well as the interchannel error and the digitization noise for both 32-bit and 42-bit DDS. At frequencies below ~ 100 Hz, the interchannel error lies on the 32-bit digitization noise, indicating that it is the limiting factor in the measurement.

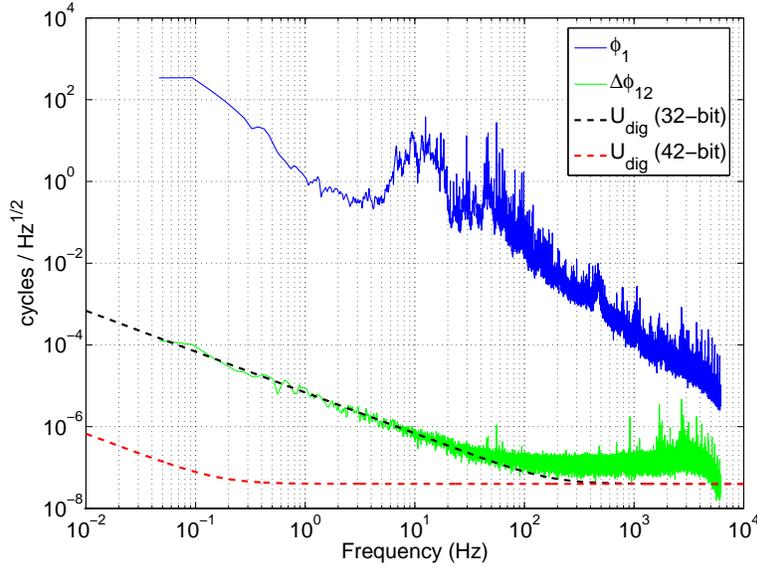


Figure 4-19. Linear spectral density of phase noise in laser beat note as measured by the hardware PM

Another potential source of noise in the PM is shot noise, the quantum-mechanical vacuum fluctuations of the light sources. Figure 4-20 shows a schematic representation of the amplitude spectrum of the electric fields of two light sources. Each source has a large peak at an frequency ω_i with an amplitude ϵ_i , where $i = 1, 2$. These peaks represent the coherent laser light, which has a phase ϕ_i .

Each light field also contains power in all frequency bins from quantum-mechanical vacuum fluctuations. The amplitude of these fluctuations is ϵ_{vac} and their phase is random and uncorrelated. In units where ϵ^2 equals photon number, $\epsilon_{vac}^2 = 1/2$. In these same units the carrier amplitudes can be computed from the signal power using

$$\epsilon_i = \sqrt{N_i} = \sqrt{\frac{P_i \lambda T}{hc}} \quad (4-38)$$

where N_i is the number of photons in the measurement, P_i is the signal power, λ is the carrier wavelength, T is the total measurement time, h is Planck's constant, and c is the speed of light.

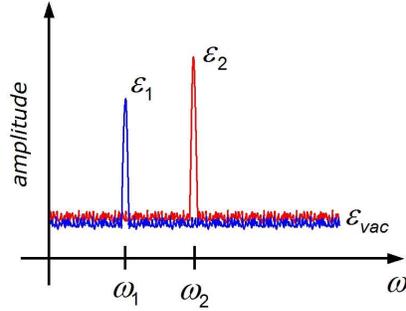


Figure 4-20. Qualitative amplitude spectrum of interfering beams with shot noise

If the two light sources in Figure 4-20 are superimposed, the total intensity will include a beat between the two carriers at $\omega_{12} \equiv \omega_1 - \omega_2$ with an amplitude $\sqrt{\epsilon_1 \epsilon_2}$. This beat note will have a phase $\phi_{12} \equiv \phi_1 - \phi_2$. In addition to this term, there will be several beat signals at ω_{12} that result from the superposition of one carrier and one shot noise bin. For example, the signal with amplitude ϵ_1 in source 1 will beat with shot noise at $\omega_1 \pm \omega_{12}$ in both source 1 and source 2. This results in a total of four uncorrelated terms, each with an amplitude of $\epsilon_1/\sqrt{2}$. The signal in source 2 will produce four similar terms from the vacuum noise in each signal at $\omega_2 \pm \omega_{12}$. Each of these terms will have a random phase. The total intensity at ω_{12} can be described by the real part of

$$I(t) = \left[\epsilon_1 \epsilon_2 e^{i\phi_{12}} + \epsilon_1 \sqrt{2} e^{i\phi_{r1}} + \epsilon_2 \sqrt{2} e^{i\phi_{r2}} \right] e^{i\omega_{12}t} \quad (4-39)$$

where ϕ_{ri} are random phase factors.

The phase of $I(t)$ is extracted in the PM by demodulating it with a signal at ω_{12} with a fixed phase. If the phase of the demodulating signal is chosen so that ϕ_{12} is small, the demodulated signal can be written as

$$S(t) = \epsilon_1 \epsilon_2 \left[\phi_{12} + \frac{\sqrt{2}}{\epsilon_2} e^{i\phi_{r1}} + \frac{\sqrt{2}}{\epsilon_1} e^{i\phi_{r2}} \right]. \quad (4-40)$$

The phase (in cycles) is extracted from $S(t)$ by dividing by $\epsilon_1 \epsilon_2$ and multiplying by 2π .

This yields the desired carrier phase, ϕ_{12} , plus two additional shot noise terms. The noise energy in these two terms is given by

$$E_{SN} = 4\pi^2 \left(\frac{2}{\epsilon_1^2} + \frac{2}{\epsilon_2^2} \right). \quad (4-41)$$

Since shot noise is white, this energy is distributed among all frequency bins from DC up to the Nyquist frequency⁷, given by $f_s/2$ where f_s is the sampling frequency. The shot-noise linear spectral density, \tilde{U}_{SN} , can consequently be related to the shot-noise energy using

$$\tilde{U}_{SN} = \sqrt{\frac{2E_{SN}}{f_s}}. \quad (4-42)$$

The shot noise limit for the two-signal measurement in Figure 4-19 can be estimated using (4-38), (4-41) and (4-42) with $P_1 = 5 \mu\text{W}$ and $P_2 = 56 \mu\text{W}$ and a measurement time of 85s. The resulting estimate is $\tilde{U}_{SN} \approx 2.5 \times 10^{-9}$ cycles/ $\sqrt{\text{Hz}}$, nearly two orders of magnitude below \tilde{U}_{dig} and the best results in Figure 4-19.

4.5.5.5 Entangled-phase PM test with VCO

The second type of test used for the hardware PM is a test of three entangled phases, designed to mimic the individual measurements and cancellations applied in TDI. The input signals for the test are derived from three oscillators, with frequencies ν_i and phases

⁷ For an introduction to the Nyquist frequency and sampling theory, see A.2.

$\phi_i(t)$. The oscillators are combined pair-wise to produce three beat signals:

$$S_{12}(t) = A_{12}(t) \sin [(\nu_1 - \nu_2)t + \phi_1(t) - \phi_2(t)], \quad (4-43)$$

$$S_{13}(t) = A_{13}(t) \sin [(\nu_1 - \nu_3)t + \phi_1(t) - \phi_3(t)], \quad (4-44)$$

$$S_{23}(t) = A_{23}(t) \sin [(\nu_2 - \nu_3)t + \phi_2(t) - \phi_3(t)], \quad (4-45)$$

where $A_{ij}(t)$ is the amplitude of $S_{ij}(t)$. These three signals are recorded by the PM in separate channels, producing the phases $\phi_{ij}(t)$. Each of the PM channels uses a separate offset frequency corresponding to the frequency of the individual beat note. The phases of the three signals should be related, such that a particular linear combination,

$$\phi_{123}(t) \equiv \phi_{12}(t) + \phi_{13}(t) - \phi_{23}(t), \quad (4-46)$$

should equal zero. The signs in this “null combination” will vary depending on the ordering of the oscillators in frequency space and any additional demodulations that are used to measure the beat signals.

Figure 4-21 shows the time-series of an entangled phase test using two VCOs and a function generator as oscillators. The beat frequencies were 4.86 MHz, 5.78 MHz, and 10.64 MHz and the amplitudes were scaled to match the 4 dBm full-scale inputs of the PM. The function generator has lower intrinsic phase noise than the VCOs, consequently the phase noise in the three signals is dominated by the VCO noise. The signals appear to be correlated, a fact which is confirmed by the plot of $\phi_{123}(t)$, which appears to be nearly zero on the scale of upper panel of the figure. The lower panel shows a blow-up of $\phi_{123}(t)$ with and without additional time shifting in $\phi_{23}(t)$ (see below).

Figure 4-22 contains the corresponding spectral densities of the three beat signals and the null combination. As with the single-signal tests in Section 4.5.5.4, there appeared to be a slight delay between the signal recorded on channel 3 and the signals recorded

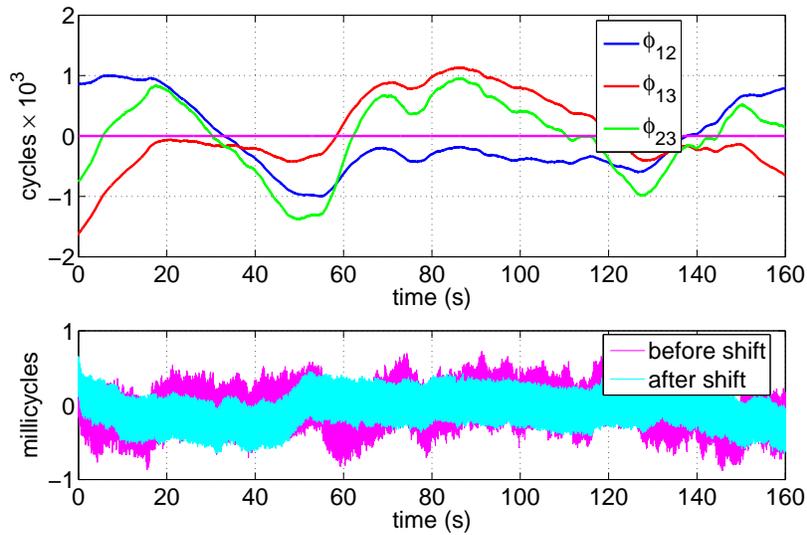


Figure 4-21. Linearly detrended time series for an entangled phase test using two VCOs and a function generator. Bottom panel shows close-up of residual noise in the null combination.

on channel 1 and channel 2. When a delay of $1.95 \mu\text{s}$ was removed using a 51-point Lagrange-windowed fractional delay filter, the residual noise in the null combination (cyan curve) was reduced to the level of the digitization noise given by $\sqrt{3}U_{dig1}$, with U_{dig1} being the single channel digitization noise given in (4-37).

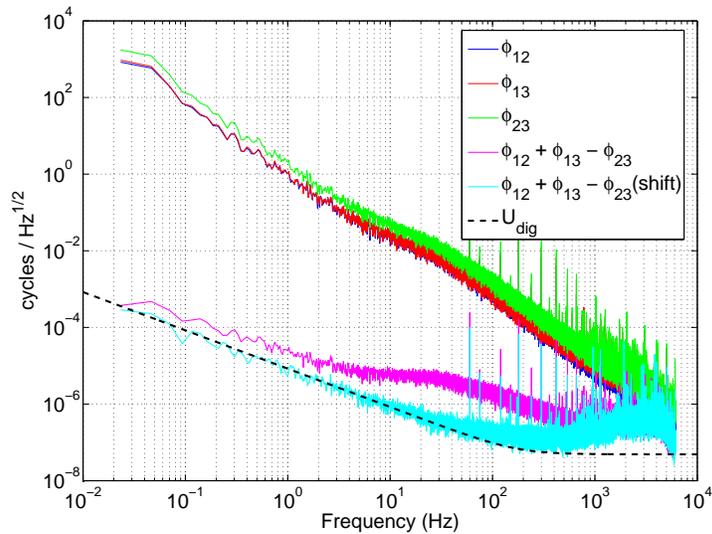


Figure 4-22. Linear spectral densities of an entangled phase test using two VCOs and a function generator

4.5.5.6 Entangled-phase PM test with optical signals

A second entangled-phase measurement was made using beat-notes between cavity-stabilized lasers as the source signals. A portion of the light from each of the two cavity-stabilized lasers on the main optical table (Figure 4-2) was focused into an optical fiber and transferred to a second optical table. The light exiting the fiber was combined with a third beam from an additional cavity-stabilized laser. The resulting photodiode signal contained three beat notes, one from each pair of lasers. The frequencies and powers of each signal are listed in Table 4-3.

Table 4-3. Beat note frequencies and amplitudes for optical entangled-phase measurement

Laser Pair	Frequency (MHz)	Power (dBm)
O-S	28.7	-23.3
O-H	113.4	-21.3
H-S	142.1	-17.3

O: Zerodur cavity with optically-contacted mirrors

H: Zerodur cavity with hydroxide-bonded mirrors

S: Silicon carbide cavity with optically-contacted mirrors

Before the beat signals in Table 4-3 could be read into the PM, they had to be conditioned so that their frequencies were ≤ 25 MHz and their amplitudes were approximately equal to 4 dBm, the full-scale input of the PM. This was accomplished using the signal conditioning arrangement in Figure 4-23. The PD signal was first amplified in a $\times 10$ RF amplifier and then split into two parts using a 50-50 RF splitter. Each of which was demodulated with a LO in order to bring the signals within the frequency range of the PM. One portion was demodulated using a LO at 22 MHz and filtered using a LPF with an 11 MHz corner frequency, producing a 6.7 MHz signal from the O-S beat. This signal was amplified in a second $\times 10$ RF amplifier and connected to a PM input channel.

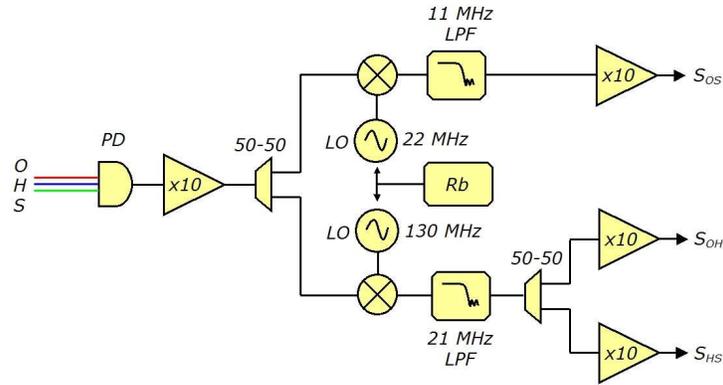


Figure 4-23. Analog electronics used to prepare beat signals in Table 4-3 for PM

The second portion was demodulated by another LO at 130 MHz and filtered with a 21 MHz LPF, placing the O-H beat at 16.6 MHz and the H-S beat at 12.1 MHz. Each of these signals was amplified in an additional RF amplifier and connected to PM input channels. The two LOs used for the demodulations as well as the PM clock were each locked to a Rubidium-stabilized 10 MHz reference signal in order to reduce the coupling of LO phase noise into the measurement.

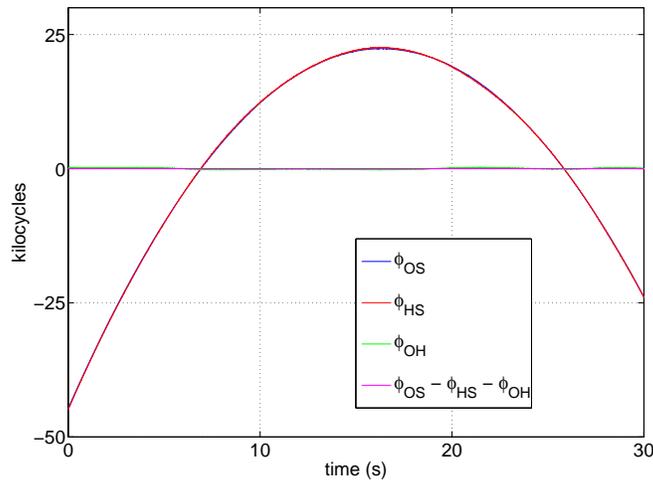


Figure 4-24. Linearly-detrended phase for optical entangled-phase measurement

Figure 4-24 shows the linearly-detrended phase output for the three channels and the null combination. The most obvious feature is the large parabolic tracks of ϕ_{OS} and ϕ_{HS} . This is a result of a frequency drift, discussed in 4.3.2, that is present between the SiC cavity and the Zerodur cavities. A fit to the data in Figure 4-24 gives a drift of ~ 248 Hz/s at the time of the measurements. Note that the drift between the two Zerodur cavities is much smaller, on the order of ~ 1 Hz/s, despite the fact that the two Zerodur cavities are located in different vacuum chambers on opposite ends of the laboratory.

The large frequency drifts do not pose a problem for data analysis, but they do limit the amount of time for which the PM can stay locked (see Section 4.5.5.7). For the three-signal measurements, the longest sets of data with three simultaneously-locked signals lasted around 30 s. Individual measurements using the O-H beat appeared to be able to last indefinitely.

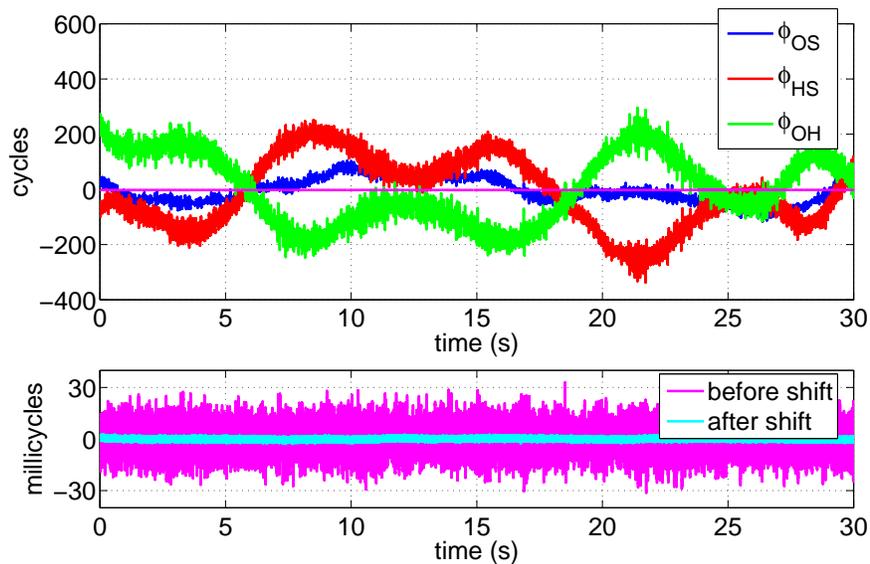


Figure 4-25. Quadratically-detrended phase for optical entangled-phase measurement

When the linear frequency drift of the beat notes is removed, the result is the time series in Figure 4-25. The noise of the three beat signals is approximately of the same amplitude, with that of ϕ_{OS} being slightly lower than the other two. This suggests that

the laser locked to the SiC cavity may have increased noise. More important is the lack of discernable noise in the null combination of the three beats, which is plotted with only a linear detrending (fixed frequency offset). The lower panel of Figure 4-25 contains a blow-up of the residual noise in the null combination both with and without an additional shift in $\phi_{OH}(t)$.

Figure 4-26 shows the linear spectral densities of the three individual beats as well as the null combination. As in the earlier measurements, a time delay was present in the third PM channel, which was used to measure ϕ_{OH} . The magnitude of this time delay was estimated by plotting the noise suppression in the null combination and making a fit using the expression in (4-34). The result was a time delay of $2.18 \mu\text{s}$, shown in Figure 4-27. To correct for this delay, the ϕ_{OH} data was shifted by $2.18 \mu\text{s}$ using a 51-point Lagrange-windowed fractional-delay filter. The results are the cyan curves in Figure 4-26 and Figure 4-27. The effect of the shift is significant, increasing the noise suppression by a factor of more than 100 near 10 Hz.

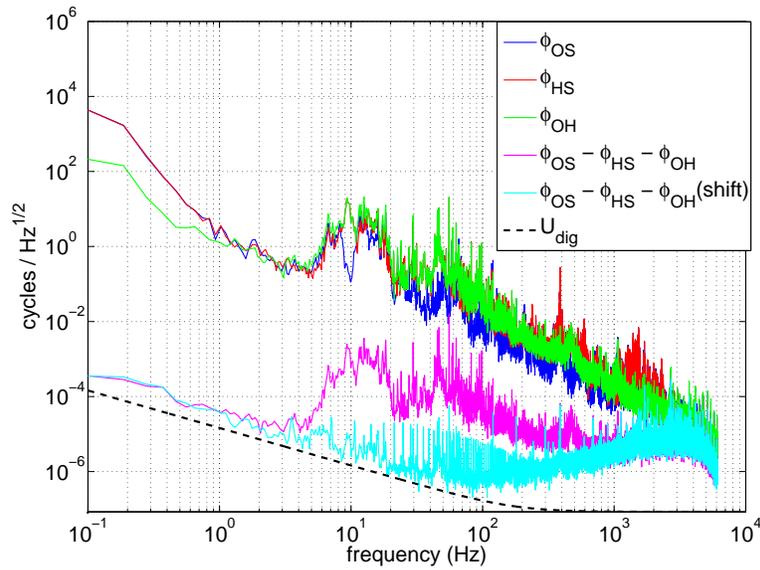


Figure 4-26. Linear spectral density for optical entangled-phase measurement

Also plotted in Figure 4-26 is the expected digitization noise, given by $\sqrt{3}U_{dig1}$ where U_{dig1} is the single-channel digitization noise given in (4-37). Unlike the VCO-measurement data, the residual phase-noise in the null combination does not reach the digitization noise. It instead follows a $1/f$ slope with an amplitude around 4.4 times larger than U_{dig} . One possible source for this noise would be relative phase noise between the two oscillators used to demodulate the beat signals.

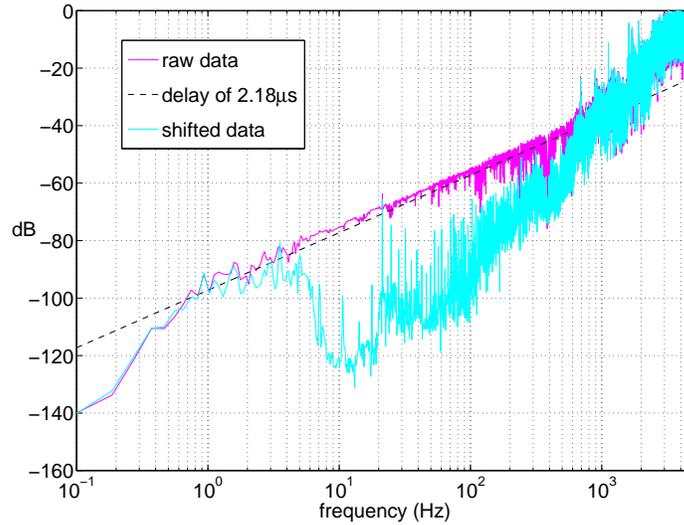


Figure 4-27. Noise suppression in null combination for optical entangled-phase measurement

4.5.5.7 Performance limitations

The previous two subsections describe the performance of the hardware PM in terms of error in measured phase noise. This performance can only be reached if the PM is operating normally and does not suffer from any failures.

One type of failure for a tracking IQ PM is a *cycle-slip*, which occurs when the amplitude of the residual phase exceeds 0.5 cycles. This can occur if the input noise is large or the gain of the tracking loop is low. One way to evaluate the likelihood of a cycle-slip is to plot a histogram of the residual phase and fit a Gaussian to it. Figure 4-28 shows a histogram for 5 s of residual phase in the hardware PM for a beat note between

the two Zerodur cavities. The data was fit with a Gaussian distribution of the form

$$P_{Gauss}(\phi_r) = A \exp \left[- \left(\frac{\phi_r - \mu}{\sigma} \right)^2 \right], \quad (4-47)$$

where A is the amplitude of the fit, ϕ_r is the residual phase, μ is the mean, and σ is the standard deviation. Using this distribution, the probability of a cycle slip can be computed using the complimentary error function

$$P_{slip} = A\sigma\sqrt{\pi} \operatorname{erfc} \left(\frac{0.5 - \mu}{\sigma} \right), \quad (4-48)$$

$$\operatorname{erfc}(x) \equiv \frac{2}{\sqrt{\pi}} \int_x^\infty e^{-y^2} dy. \quad (4-49)$$

For the residuals shown in Figure 4-28, the best fit Gaussian has an amplitude of 17%, a mean of zero, and a standard deviation of 7.4 millicycles. For such a small standard deviation, the probability of a cycle slip is effectively zero.

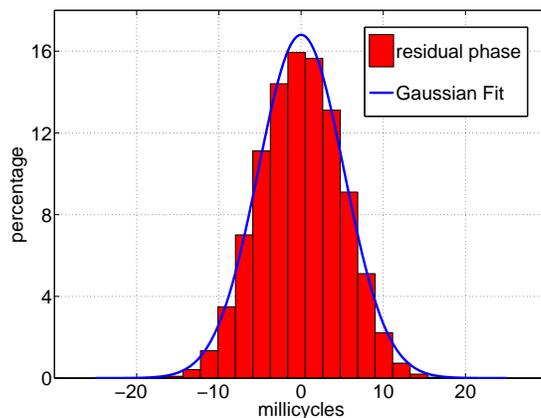


Figure 4-28. Histogram of residual phase for laser beat note in hardware PM

In addition to being caused by high-frequency phase noise, cycle slips can also be caused by large frequency drifts in the input signals, such as those on the SiC beat notes. At low frequencies, the gain of the tracking loop in the PM front-end increases as f^{-2} , with one power of f coming from the integrator in the feedback filter and the second power coming from the implicit integration in the DDS. A linear frequency drift corresponds to a quadratic increase of the phase with time, which in turn corresponds

to phase noise increasing with a slope of f^{-3} at low Fourier frequencies. As the Fourier frequency decreases, the phase noise will increase faster than the loop gain, meaning that the magnitude of the residual phase fluctuations will increase. Eventually this will cause a cycle slip to occur.

One way to reduce the effect of the frequency drifts is to increase the low-frequency gain of the PM by adding additional integrators at low frequency. Another approach would be to put a ramping signal into the frequency offset that would cancel most or all of the frequency ramp in the input signal. This feed-forward approach would not effect the stability of the PM, but would require that the frequency drift of the input signal be known in advance, as will be the case for the Doppler shifts in LISA. The best approach is probably a combination of the two with the majority of the drifts removed through feed-forward and the remaining noise suppressed through feedback. These features will likely be incorporated in future versions of the PM.

A more practical problem arises when trying to acquire lock with the PM using a signal with a large frequency drift. The locking range of the tracking loop on the PM is on the order of ~ 20 kHz, meaning that the difference between the signal frequency and the model offset frequency must be less than that amount in order for the PM to acquire the signal. This requires the user to anticipate the frequency of the signal if it is drifting, which can be difficult. The task becomes especially challenging when acquiring multiple signals, as in the entangled phase measurements. This problem could be solved by implementing an auto-lock function, where the PM front-end takes an initial sample of the time series, determines the frequency of the signal, and locks the tracking loop automatically. This feature will be implemented in subsequent versions of the PM.

4.6 EPD Unit

The purpose of the electronic phase delay (EPD) unit is simply to delay an electronic signal by a specified amount of time. This is accomplished using three steps: digitizing the signal, storing the digital information in a memory buffer, and regenerating the analog

signal. The important parameters for the EPD unit are the sampling rate, which sets the bandwidth of the EPD unit, and the amount of memory, which sets the maximum storage time. The maximum storage time can be computed as

$$\tau_{max} = \frac{M}{N_{chan}bf_s}, \quad (4-50)$$

where M is the amount of memory in bytes, N_{chan} is the number of delay channels, b is the width of the sample words in bytes, and f_s is the sampling rate.

The EPD units were developed in three generations, summarized in Table 4-4. The initial prototype was built using a DAP-5216a DSP card by Microstar, Inc. The DAP-5216a contains 16-bit ADCs and 16-bit DACs clocked at 200 kHz and 64 MB of SDRAM. This allowed signals in the ~ 10 kHz regime to be delayed by tens of seconds. This 1st-generation EPD unit was used for several early simulator experiments [61, 62].

Table 4-4. Progression of EPD units

EPD Unit	Hardware	Sampling Frequency	# of Chan.	Max. Delay
1 st -generation	Microstar	200 kHz	2	80 s
2 nd -generation	Pentek	≤ 25 MHz	4	2.6 s
3 rd -generation	Pentek	100 MHz (I/O) 97.66 kHz (phase)	4	340 s

4.6.1 Second-generation EPD unit

The 2nd-generation EPD unit was implemented in the Pentek hardware described in Section 4.4 and diagrammed in Figure 4-5. The input data is digitized by the ADCs on the model 6256 downconverter. The ADC data from each pair of channels is packed into a 32-bit word and transferred via the VIM interface to the BIFO on the model 4205 carrier board. The 4205 reads the data off the BIFO in blocks of 1024 words and places it in a memory buffer in the SDRAM. The output of the memory buffer is transferred via the

VIM interface to the model 6228 upconverter, where it is unpacked and used to feed the DACs.

The bandwidth of the 2nd-generation EPD unit has a lower limit given by the transformers that couple the ADCs and DACs to the signal inputs and outputs on the Pentek hardware. These transformers have a -3 dB corner frequency at 400 kHz. Inputs to the 2nd-generation EPD unit must have frequencies at least this large.

The bandwidth of the EPD unit is also limited at high frequencies. Although the ADCs and DACs are capable of being clocked at 100 MHz or more, the speed of the VIM interface is limited by the PCI bus. Since the incoming and outgoing data must share the 66 MHz PCI bus, the theoretical maximum data rate for continual two-way streaming across the VIM is 33 MHz. In practice this is reduced to ~ 25 MHz. For the four-channel case (two VIMs for input and two VIMs for output), the rate is further reduced to ~ 12.5 MHz. This limits the frequencies of delayed signals to $\leq 5 - 10$ MHz.

Signals passing through the 2nd-generation EPD unit experience an amplitude loss of -6 dBm due to the fact that the full-scale of the ADC input is $+4$ dBm while the full-scale of the DAC output is -2 dBm. This loss does not affect the signal-to-noise ratio of the input signal because both the signal and noise experience the same loss.

The timing resolution of the delay is given by the size of the blocks used to read from and write to the VIM interface. Typical values were blocks of 1024 words with a sampling rate of 25 MHz, corresponding to a timing resolution of $41 \mu\text{s}$.

The 2nd-generation EPD unit has been used in a number of experiments including measurements of TDI-like signals [51, 63] and investigations of arm-locking with optical signals [64].

The downside of the 2nd-generation system is that, despite having 1 GB of SDRAM, it is unable to delay signals for more than ~ 2.5 s. This is because the data rates for the high sampling frequencies are so large. This problem is especially annoying because there is very little information of interest in the LISA signals at high frequencies.

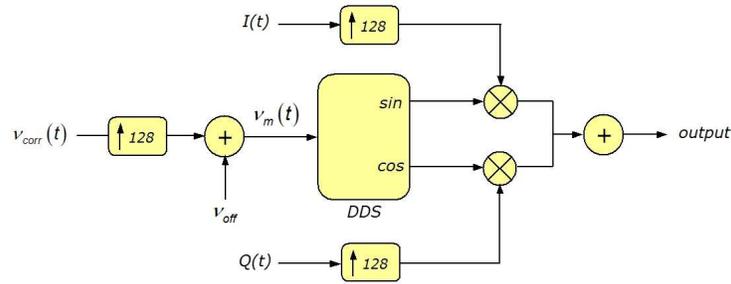


Figure 4-29. Schematic of the NCO used in the 3rd-generation EPD unit

The 3rd-generation EPD unit address this problem by incorporating a PM and an numerically-controlled oscillator (NCO) into the delay process.

4.6.2 Third-generation EPD unit

For the third-generation EPD unit, the model 6256 is configured as the front-end of a PM, as described in 4.5.5.1. The 100 MHz sampling frequencies allows for signals with frequencies of up to ~ 30 MHz. The $IQ\nu$ data for each channel is packed using the format in Figure 4-17 and transferred across the VIM at a rate of 97.65625 kHz. The 4205 reads this data off of the BIFO in blocks of 1024 and unpacks it into separate streams for each channel. These streams are stored in a memory buffer on SDRAM for the specified delay time. The delayed data is repacked and transferred across the VIM to the FPGA on the 6228. The 6228 unpacks the data and feeds the $IQ\nu$ data to a NCO, shown schematically in Figure 4-29.

The NCO converts the $IQ\nu$ data from the PM into sinusoidal signals with an offset frequency ν_{off} at a sampling rate of 100 MHz. The output of the NCOs is fed into the DACs, reproducing the signals.

By decimating the $IQ\nu$ data before storage, the 3rd-generation EPD unit greatly increases the maximum delay time while simultaneously increasing the maximum carrier frequency. The price paid is in the bandwidth of the EPD unit: any components of the signal above ~ 50 kHz will not be reproduced. This is above the bandwidth of most of the relevant signals with the exception of the modulation tones and possibly the active

bandwidths of the PLLs. The PLL bandwidth can easily be encompassed by reducing the decimation in the PM. Bypassing the additional divide-by-8 CIC filter in the PM will result in a EPD passband of ~ 400 kHz while still allowing for delays of ~ 40 s per channel.

The modulation tones are likely to be at tens of MHz and will have to be dealt with separately. One approach is to assign a separate PM channel for each modulation tone, delay each channel separately by the same amount, and re-generate the signal using a modulated NCO. So long as a limited number of modulation tones were included, this could be a viable approach.

An important feature of the 3rd-generation EPD unit is the ability to model the Doppler shifts produced by relative motion between the LISA SC. This is accomplished by using different values for ν_{off} in the PM and the NCO. For instance, an 11 MHz signal could be measured with a PM with $\nu_{off} = 10$ MHz and reconstructed by an NCO with $\nu_{off} = 5$ MHz. The delayed signal would have a carrier frequency of 6 MHz with the same phase information (after a delay) of the 11 MHz signal. This feature allows for more LISA-like arrangements to be used in the simulator.

As a test of the 3rd-generation EPD unit, a VCO was demodulated with a function generator to create a signal with a frequency of 8 MHz. This signal was split in two using a 50-50 RF splitter, with one portion passing through the 3rd-generation EPD unit with a time delay of 2 s and a Doppler shift of 4 MHz. The output of the EPD unit and the other portion of the input signal were connected to the inputs of a hardware PM running on a separate Pentek system. The linearly-detrended time series of the phase data for the two signals is shown in Figure 4-30.

To check the effectiveness of the delay unit, the delayed signal was shifted in time by various amounts using a combination of simple integer sampling point shifting and fractional delay filtering and then subtracted from the input signal. The delay time was optimized by minimizing the power spectrum of the shifted difference. For this data set, a delay time of $\tau = 2.039663$ s produced the minimal error.

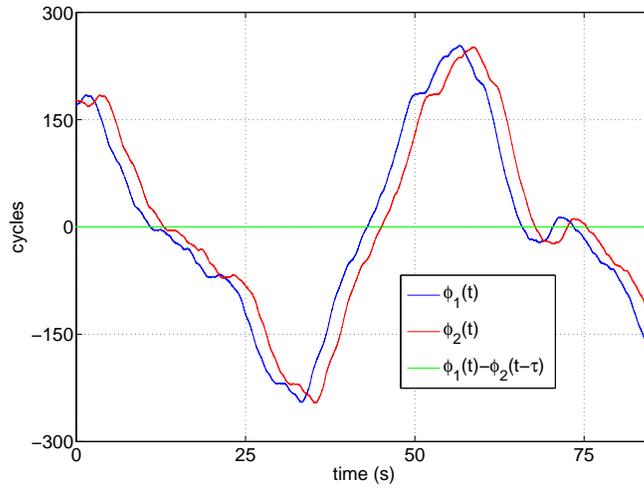


Figure 4-30. Detrended phase of VCO signal in EPD test. Expected time delay: $\tau = 2$ s.

The linear spectral density of the input signals and the shifted and subtracted combination is shown in Figure 4-31. For frequencies below ~ 100 Hz, the noise in the subtracted combination matches the digitization noise, computed as $\sqrt{2}U_{dig1}$, where U_{dig1} is the single-channel digitization noise given in (4-37). This indicates that the EPD unit does not add any additional noise to the signal beyond that added by the PM.

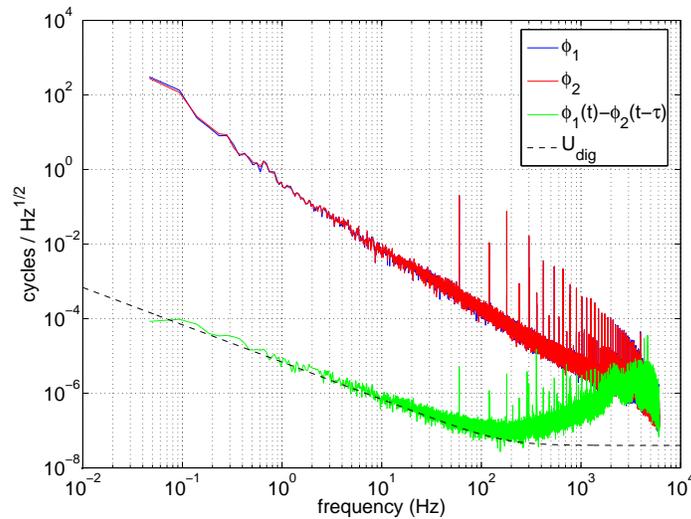


Figure 4-31. Linear spectral densities in EPD test with VCO signals. Optimal time delay: $\tau = 2.039663$ s

The disparity between the expected and observed time delays of ~ 40 ms is a bit puzzling. The majority of this extra time seems to be independent of the expected delay time, indicating that it is some sort of additive affect. The source is likely an error in the computation of the length of the storage buffer, which currently does not account for time the data spends in the various processing buffers. However it is estimated that the errors caused by the buffers should be of the order of $1 \text{ ms} \sim 5 \text{ ms}$. Regardless of its source, the timing error can be easily dealt with by adjusting the target delay times.

CHAPTER 5 ARM-LOCKING IN THE UF LISA INTERFEROMETRY SIMULATOR

5.1 Introduction

The laser phase-stabilization technique known as arm-locking plays an important role in the LISA IMS (Section 3.5.4). Arm locking is a function of the IMS as a whole, involving subcomponents such as phasemeters (PMs), laser pre-stabilization systems, and phase-lock loops (PLLs). To date, arm-locking has been studied analytically [43, 45], through time-domain simulations [42, 43, 65], and in several hardware-analog experiments [48, 61, 64, 66]. The particular advantage of the hardware experiments is that, in forcing one to actually build a working system, they can expose effects that have not been included in the analytic or numerical models.

The UF LISA interferometry simulator is ideally suited to studying arm-locking in a LISA-like environment. It is the only system in existence that can provide both realistic laser noise and realistic delay times. For studying arm-locking, the large delay times are essential, because they set the frequency scale for the controller.

In parallel with the development of the interferometry simulator (Chapter 4), the author has developed a series of hardware models of single-arm arm-locking based on the EPD technique. The initial model, described in Section 5.2, was a purely electronic model using a VCO in place of a laser beat note. This provided a proof-of-principle for the EPD technique. The next iteration, described in Section 5.3, incorporated improved electronics to allow locking of a pre-stabilized laser beat note to a ~ 1 ms delay.

The subsequent development of the hardware PM, described in Section 4.5.5, allowed for a change in the experimental topology that produced the improved optical system described in Section 5.4. This system is capable of generating LISA-like arm-locking error signals with delays of 1 s or more. Unfortunately, a technical issue with the implementation of the arm-locking filter for this system prevented the system from being locked.

5.2 Electronic Model

5.2.1 Method

The first EPD-based arm-locking experiment, described in detail by Thorpe & Mueller [61], was a purely electronic model. The experimental apparatus, shown in Figure 5-1, centers around a VCO with a nominal frequency of ~ 25 kHz. The VCO signal is split into two portions, one of which is delayed in an EPD unit. For these experiments, the EPD unit was the first-generation version described in Section 4.6, with a sampling frequency of 200 kHz. The delayed and prompt signals were mixed in an analog mixer and the mixer output was low-pass filtered by a single pole at 300 Hz, generating an error signal of the form

$$S(t) = 1.6 \frac{\text{V}}{\text{cycle}} [\phi(t) - \phi(t - \tau)] \quad (5-1)$$

where $\phi(t)$ is the phase of the VCO and τ is the delay time of the EPD unit. This error signal has the same form as the single-arm arm-locking error signal given in (3-19).

For these experiments, the delay time was set to $\tau \approx 500$ ms, so that the nulls in the interferometer response occurred at $f_n \approx n \cdot 2$ Hz.

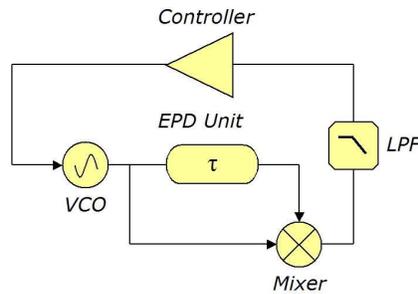


Figure 5-1. Experimental setup for electronic arm-locking experiment

The controller for the electronic model was implemented using a PC-based DSP system from National Instruments. The error signal was digitized at 1 kHz with 16-bit resolution and streamed into National Instruments LabVIEW software. The filter was an

IIR filter¹ designed using the bilinear transform method with Laplace-domain poles at 50 mHz, 3 Hz, and 30 Hz and Laplace-domain zeros at 500 mHz, 1 Hz and 10 Hz. A Bode plot of the controller is shown in Figure 5-2 along with a Bode plot of the arm-locking system (interferometer plus 1/s actuator). An overall gain of ≈ 200 has been removed from the controller transfer function to allow the two curves to overlap in the plot. Note from the lower panel of Figure 5-2 that the controller provides a phase advance in the vicinity of the interferometer nulls.

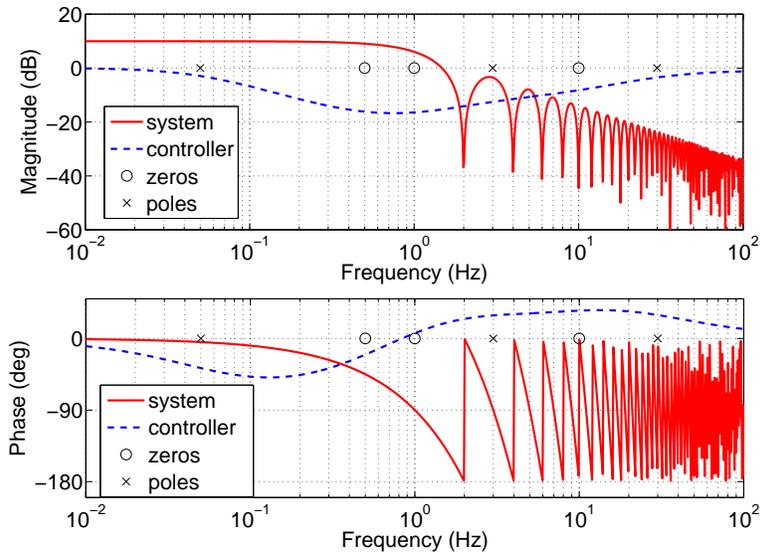


Figure 5-2. System and controller transfer functions for electronic arm-locking experiment

5.2.2 Results

To evaluate the electronic arm-locking system, a separate portion of the VCO signal was demodulated by a LO to DC. Since the rms phase noise of the arm-locked VCO was less than 1 cycle during the measurement band, the mixer output could be used as a PM, as described in 4.5. Figure 5-3 shows the linear spectral density (LSD) of the arm-locked

¹ For more information on IIR filters, see Section A.4. For details on the bilinear transform design method for IIR filters, see Section A.4.3.2.

VCO signal. The spectrum exhibits the characteristic noise spikes at Fourier frequencies just below the null frequencies.

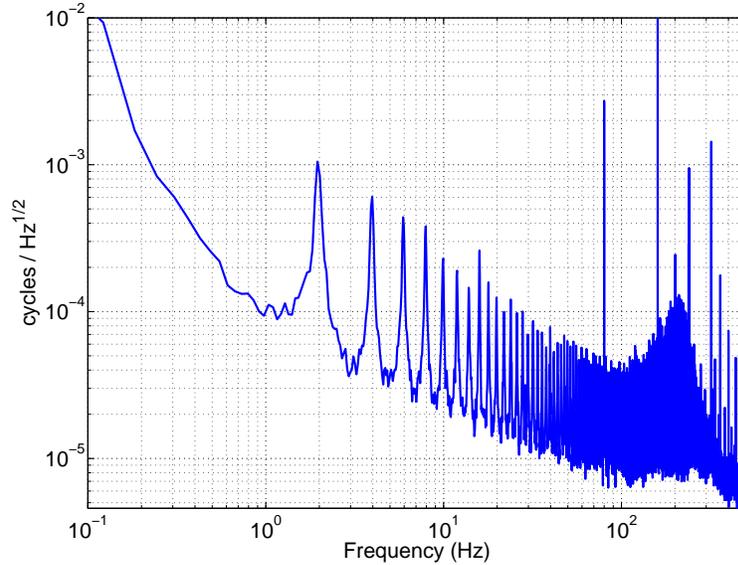


Figure 5-3. Linear spectral density of arm-locked VCO signal

Ideally, the unlocked spectrum would be included in Figure 5-3 as well. However, the rms phase noise in the unlocked VCO exceeds 1 cycle, preventing the crude mixer PM from working properly. As an alternative, the closed-loop noise suppression was measured in-loop by injecting a sinusoidal signal into the VCO control input and observing the corresponding signal in the error signal output. The results of this measurement are shown in Figure 5-4, along with a fit made using the known open-loop transfer functions of the components. The fit parameters were τ and an overall gain factor and the best fit values were $\tau = 500.9$ ms and $H_0G_0 = 200$.

5.2.3 Discussion

The good agreement between the measured and predicted responses in Figure 5-4 indicate that the EPD-based system is a reasonable hardware analogue for LISA. The primary limitations of this system is the short delay time, the low controller bandwidth, and the limited dynamic range of the mixer PMs. The latter two limitations precluded

this system from being used with optical signals, which have larger intrinsic phase noise than the VCO.

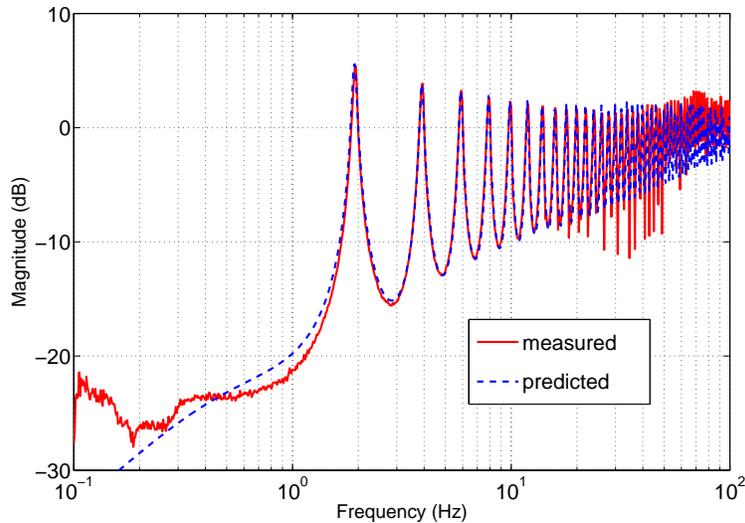


Figure 5-4. Closed-loop noise suppression for electronic arm-locking experiment

5.3 Initial Optical Model

5.3.1 Method

After the initial success of the electronic arm-locking model (Section 5.2), the next effort was building an arm-locking system incorporating optical signals [64]. The optical components of the system were composed from the optical bench shown in Figure 4-2. For the arm-locking experiments, these optics were arranged as shown in Figure 5-5 below. Laser 1 (L_1) is stabilized to an optical cavity consisting of a Zerodur spacer with optically-contacted mirrors. Laser 2 (L_2) is phase-locked to L_1 with a frequency offset provided by a VCO with a nominal frequency of 25 MHz. As discussed in Section 3.5.4.6, this is one approach that could be utilized on LISA to combine arm-locking with pre-stabilization.

The round-trip delay of the LISA arm is modeled using the EPD technique (Section 4.2). The reference laser is Laser 0 (L_0), which is locked to the SiC cavity. As mentioned

The arm-locking controller is a hybrid digital-analog system consisting of a second-order IIR digital filter with two additional integrators that can be switched in to provide additional low-frequency gain. Figure 5-7 contains a Bode plot for the controller with and without the integrators. The curve without the integrators is a measurement made using a network analyzer and includes the effects of latency in the filter. The curve with the integrators included combines the theoretical transfer function of the integrators and the measured latency in the filter system. With the integrators on, the filter has a slope of f^{-2} below 100 Hz. From 300 Hz to 30 kHz, the filter has a slope of $\sim f^{1/2}$, providing the required phase advance in the vicinity of the f_n frequencies.

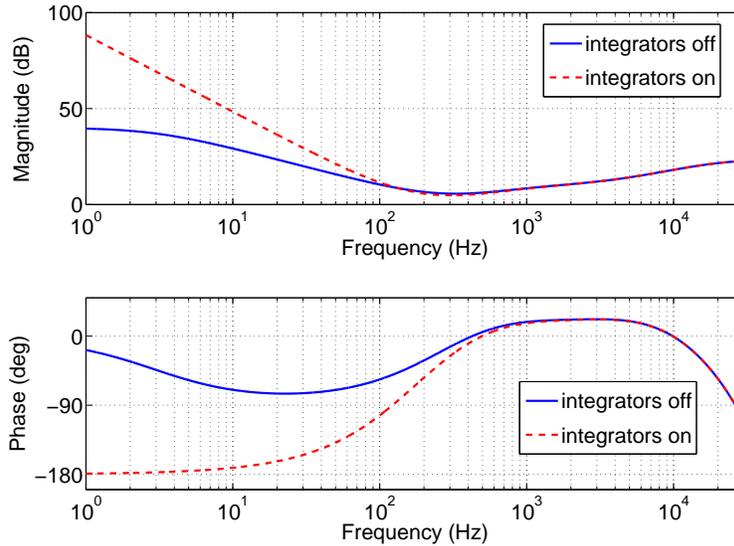


Figure 5-7. Bode plot of controller for initial optical arm-locking system

5.3.2 Results

An initial characterization of the system in Figure 5-5 was made by replacing the $L_2 - L_0$ beat note with a function generator. The response of the interferometer was measured by injecting sinusoidal phase modulation at a given frequency using the function generator and observing the mixer output signal. The results of these measurements are shown in Figure 5-8. The observed data was used to fit to the single-arm error signal

(3–19) to determine the actual time delay and an overall gain parameter. The best fit parameters were a delay of 1.065 ms and a gain of 8 mV/deg.

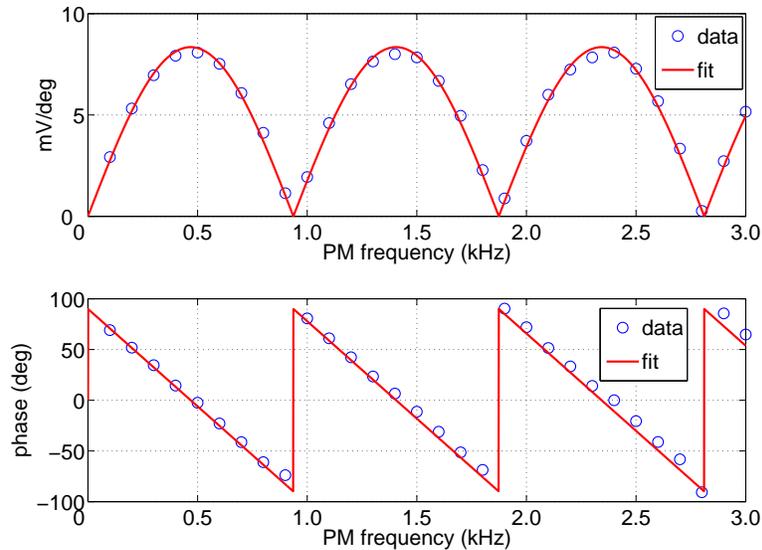


Figure 5-8. Response of interferometer to phase modulation

Once the system characterization was complete, the function generator signal was replaced by the L_2-L_0 beat note and the error signal was connected to the controller. The effectiveness of the arm-locking system was evaluated using an out-of-loop measurement system consisting of a beat note between L_2 and L_0 at a PD separate from the one used to generate the error signal. The stability of the beat note for both the locked and unlocked cases was measured using two instruments: a commercial frequency counter and the software PM described in Section 4.5.4. The frequency counter allowed for long-duration measurements to probe low-frequencies while the PM allowed the high-frequency regime to be studied.

5.3.2.1 Frequency counter measurements

Figure 5-9 contains a timeseries of the locked and unlocked frequency noise recorded using the frequency counter at a rate of 0.5 sample/s. The unlocked case clearly exhibits a

large frequency drift while the locked case remains nearly constant. Figure 5-10 shows the same data with a linear drift of 167 Hz/s removed from the unlocked data.

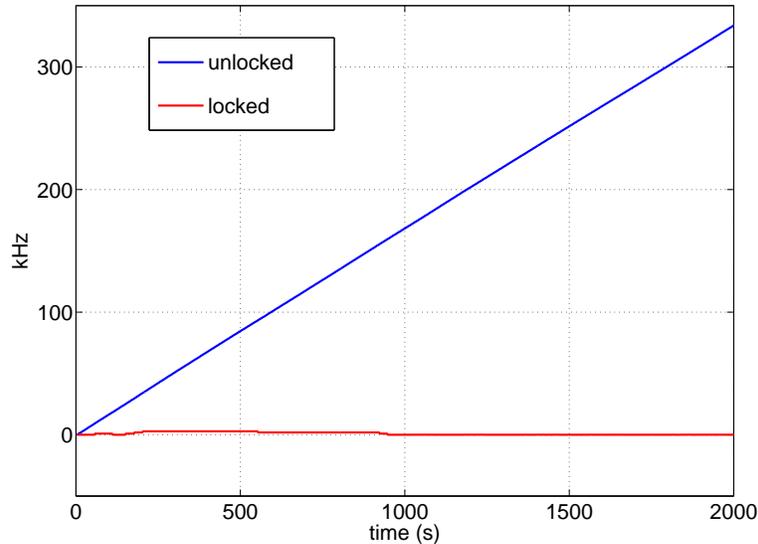


Figure 5-9. Timeseries of $L_2 - L_0$ beat note for locked and unlocked cases

The residual unlocked frequency noise drifts over approximately 1 kHz in 1000 s, with a smaller high-frequency component. On this scale, the locked timeseries consists of a series of flat plateaus separated by distinct vertical shifts. These plateaus are a result of the $\text{mod}(1 \text{ cycle})$ character of the error signal (5-2), which has lock points separated in frequency space by $1/\tau$. Excess noise can cause the arm-locking system to slip from one lock point to another, an event dubbed a “fringe-slip”. Close examination of the plateaus in Figure 5-9 shows that their levels differ by $1/1.065 \text{ ms} \approx 939 \text{ Hz}$.

Between the fringe-slips, the locked frequency appears nearly constant on the scale of Figure 5-9. Figure 5-11 shows a close-up of the locked frequency data between 1000 s and 2000 s. During this time period, the beat note frequency remained within roughly $\pm 250 \text{ mHz}$.

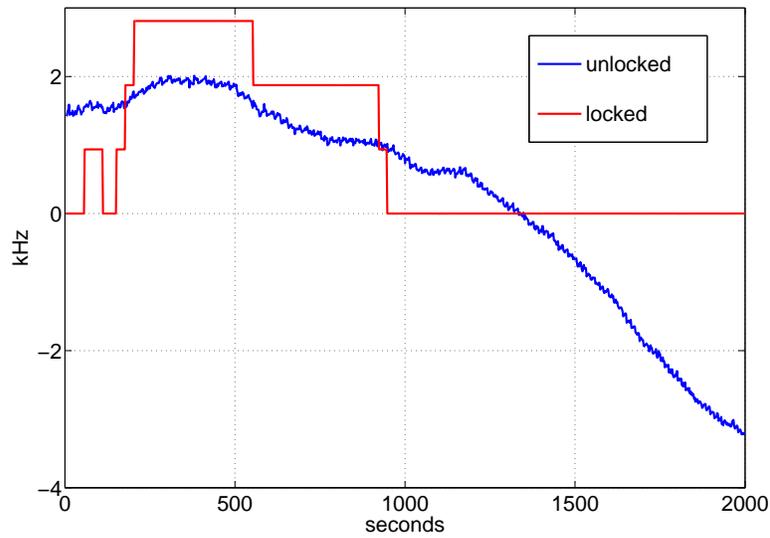


Figure 5-10. Timeseries of $L_2 - L_0$ beat note with linear trend removed from the unlocked case

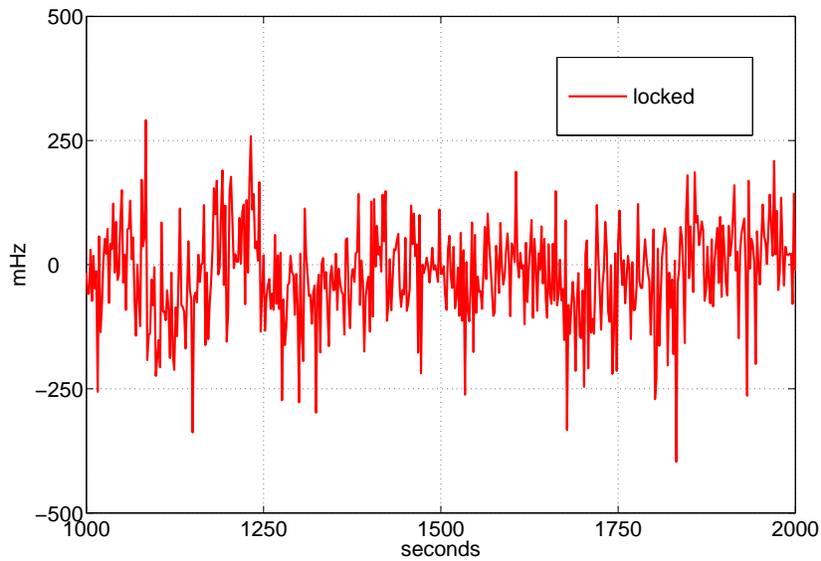


Figure 5-11. Close-up of locked case in Figure 5-9 from 1000 s to 2000 s

The spectra of the locked and unlocked frequency noise are shown in Figure 5-12. For the locked case, only data between fringe-slips was included. The locked frequency

noise lies below $200 \text{ mHz}/\sqrt{\text{Hz}}$ from 10 mHz to 250 mHz, a factor of ≈ 400 better than the unlocked case.

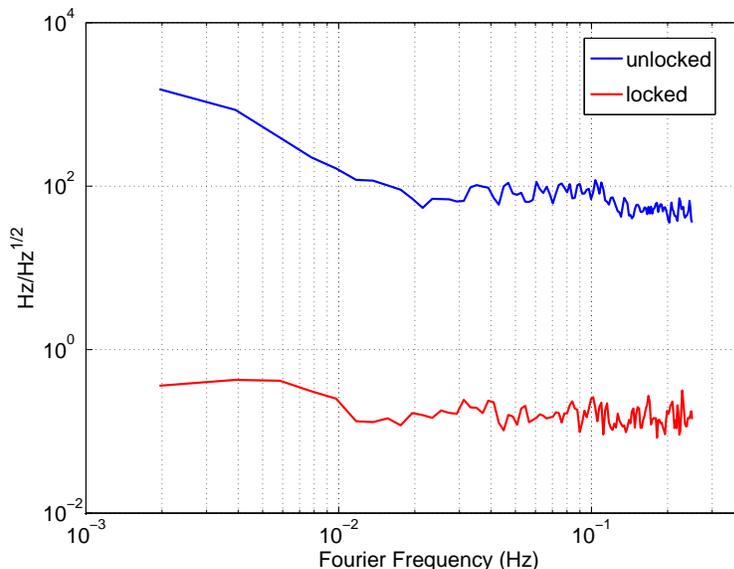


Figure 5-12. Spectrum of locked and unlocked frequency noise. Fringe-slips have been removed from the locked frequency noise.

5.3.2.2 Phasemeter data

The behavior of the arm-locking system at higher frequencies was studied using the software PM described in Section 4.5.4. The $L_2 - L_0$ beat signal from the out-of-loop photodiode was demodulated with a local oscillator to a frequency of $\sim 10 \text{ kHz}$. This signal was digitized at a rate of 80 kHz. The recorded data was then processed offline in the software PM, producing the phase timeseries shown in Figure 5-13. The reduction in phase noise in the locked case is clear.

The linear spectral densities of the PM signals are shown in the right-hand side of Figure 5-14. On the left-hand side of Figure 5-14 are the equivalent phase-noise spectra obtained from a scaling of the frequency noise spectra in Figure 5-12. Although the two data sets do not overlap, they are clearly consistent with one another.

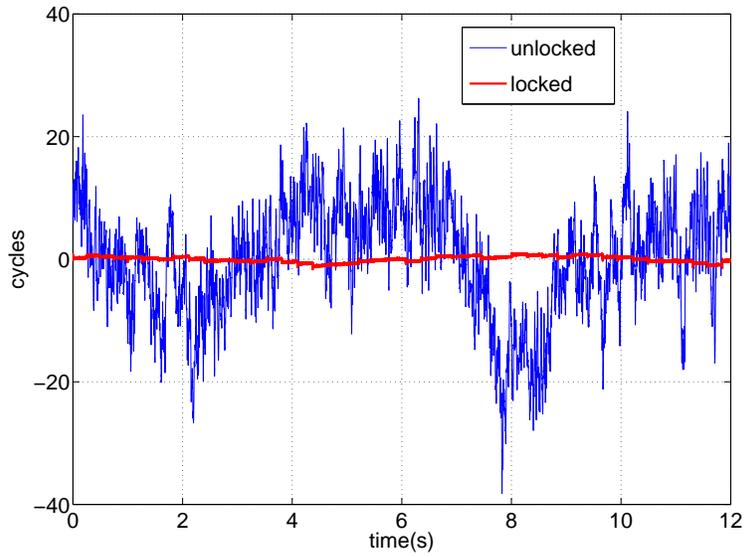


Figure 5-13. Timeseries of beat note phase for unlocked and locked cases

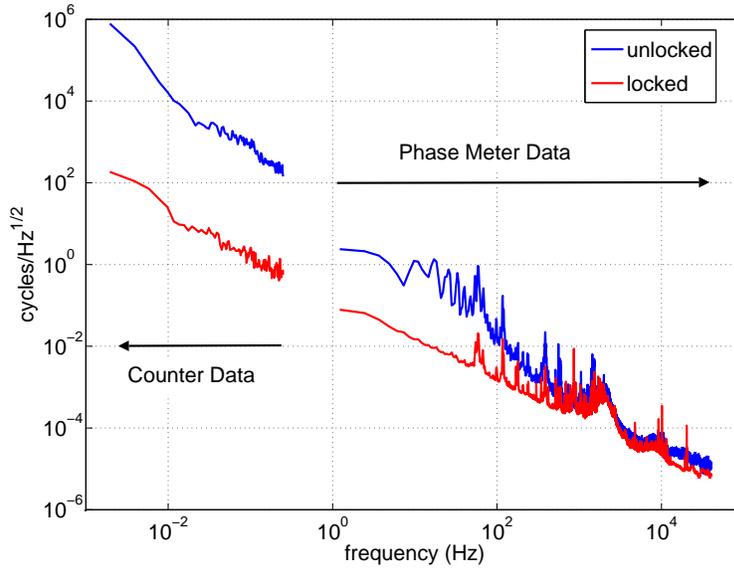


Figure 5-14. Phase noise spectra for the unlocked and locked cases. The spectra on the right-hand side are from the PM data. The spectra on the left-hand side are the frequency noise spectra from Figure 5-12 scaled to phase noise.

An estimate of the closed-loop noise suppression of the arm-locking loop can be made by dividing the locked spectral density by the unlocked spectral density, as shown in

Figure 5-15. Also shown in Figure 5-15 is the theoretical closed-loop suppression that was calculated using the measured open-loop transfer functions of the interferometer, control filter, and VCO. The measured closed-loop suppression generally follows the shape of the predicted transfer function, but deviates at both the upper and lower frequencies. The deviation at the upper frequencies is likely a result of limitations in the software PM, which has a limited resolution above ~ 1 kHz (see Section 4.5.4). At lower frequencies, the suppression appears to hit a noise floor around -40 dB.

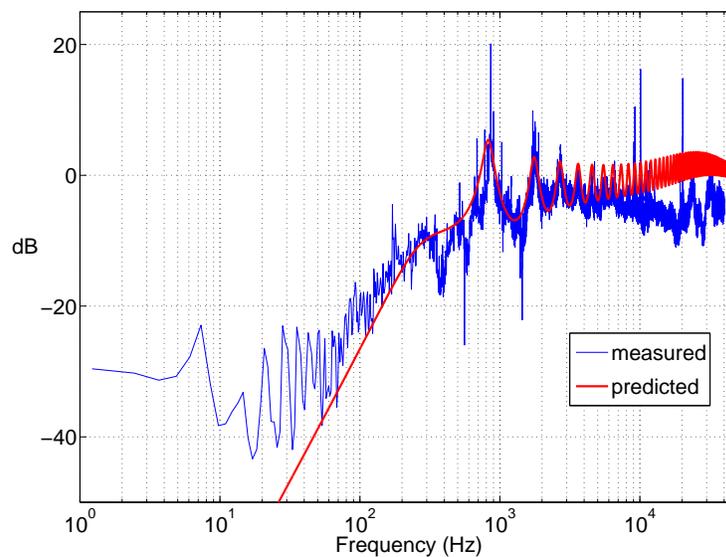


Figure 5-15. Closed-loop noise suppression for optical arm-locking

5.3.2.3 Error-point noise

As mentioned above, the response of the software PM degrades above ~ 1 kHz, making it difficult to estimate f_{UG} from Figure 5-15. A prediction based on the measured transfer functions of the system components estimates $f_{UG} \approx 12$ kHz. Figure 5-16 contains a measurement of the error point noise, recorded as the voltage out of the mixer in Figure 5-5, for the locked and unlocked cases. The error-point noise for the locked case is clearly lower than the unlocked case at low-frequencies. At higher frequencies, there is a broad

peak in the locked error-point noise. The peak is centered around 12 kHz, consistent with the predicted unity-gain frequency.

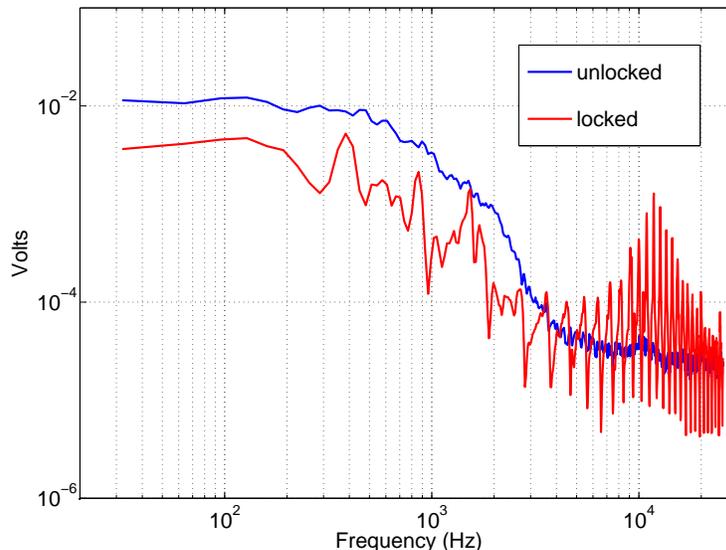


Figure 5-16. Error-point (mixer output) noise for locked and unlocked cases

5.3.3 Discussion

The results described above demonstrate that it is possible to incorporate optical signals with LISA-like noise into an EPD-based model of arm-locking. The primary limitation of this system is the small linear range of the interferometer coupled with the relatively large laser phase noise. Unlike the situation in LISA, there is no frequency offset between the two paths of the interferometer, resulting in a “beat note” at zero frequency. When coupled with the large laser phase noise, this distorts the error signal so that the measured signal is $[\phi(t) - \phi(t - \tau)]$ modulo 2π rather than $[\phi(t) - \phi(t - \tau)]$. This causes the fringe-slips, which prevent the system from being stable over long measurement times.

The closed-loop noise suppression could be increased by increasing the bandwidth of the control filter. For this particular filter implementation, the bandwidth is limited by latency in the digital portion of the loop.

5.4 Improved Optical Model

The development of the 3rd-generation EPD unit with the capability of providing frequency shifts (see Section 4.6) allows for an improvement on the optical model described above. The optical configuration remains the same as in Figure 5-5 while the electronics are modified as shown in Figure 5-17.

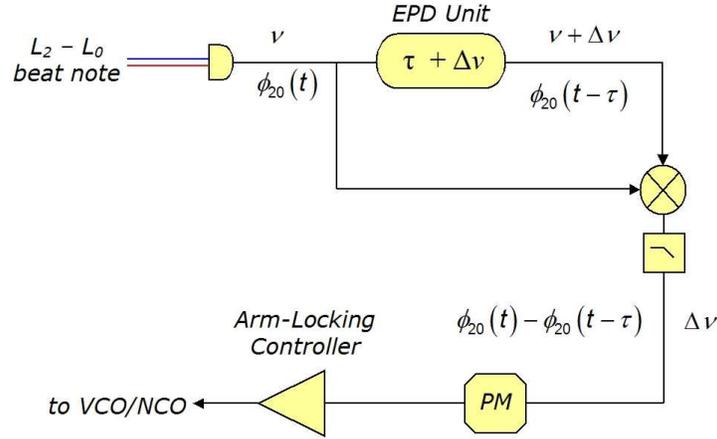


Figure 5-17. Modification of electronics for improved optical arm-locking

As in the original experiment, the $L_2 - L_0$ beat note is divided into two portions, one of which enters the EPD unit. The EPD unit delays the signal by an amount τ while also imparting a fixed frequency shift of $\Delta\nu$. The output of the EPD unit is mixed with the original beat note and the output is low-pass filtered, producing a signal with a frequency $\Delta\nu$ and a phase $\phi_{20}(t) - \phi_{20}(t - \tau)$. This is a direct analogue of the beat signal for single-arm locking in LISA, which will have a nearly constant frequency given by the Doppler shifts and any offset in the PLL at the far SC.

As in LISA, the phase of the mixer output signal can be read-out with a real-time PM set with an offset frequency equal to $\Delta\nu$. This provides an error signal proportional to $\phi_{20}(t) - \phi_{20}(t - \tau)$ so long as the phase difference remains in the linear range of the

PM. Consequently, the system is less susceptible to fringe-slips and the time delay can be increased to more LISA-like values.

5.4.1 System Characterization

To demonstrate the viability of the model in Figure 5-17, a measurement of the system transfer function was made. The $L_1 - L_2$ beat note was phase-locked to a 10 MHz LO signal using an analog PLL. This placed the $L_2 - L_0$ beat note at approximately 98 MHz, with an amplitude of -41 dBm. The $L_2 - L_0$ beat note was demodulated with a fixed 90 MHz LO signal and amplified using two RF amplifiers with amplitude gains of $\approx \times 10$ each. The amplified signal was split into three equal portions, each with an amplitude of roughly -4 dB.

One portion of the signal was fed into the EPD unit, which was set with a delay of 1 s and a Doppler shift of 4 MHz. The second portion of the signal was amplified by a third RF amplifier and fed into the LO port of an RF mixer. The other port of the mixer was connected to the output of the EPD unit, which had an amplitude of approximately -10 dBm. The mixer output was filtered with a 5 MHz corner frequency and amplified, producing a 4 MHz signal with an amplitude of roughly 0 dBm. This signal corresponds to the output of the mixer in Figure 5-17.

The mixer output signal and the third portion of the $L_2 - L_0$ beat note were each fed into the hardware PM described in Section 4.5.5. Figure 5-18 shows the raw timeseries of the phase data for both signals. The $L_2 - L_0$ beat note follows a quadratic trend

$$\phi_{20}(t) = at^2 + bt - c \quad (5-4)$$

with $a = 25.1$ cycles/s², $b = 1622$ cycles/s, and $c = -941$ cycles. The quadratic term is equivalent to a linear frequency drift of 25.1 Hz/s, caused by a drift between the L_1 and L_0 reference cavities, as discussed in 4.3.2. It happened that these measurements were taken during a time small drifts between the cavities, which allowed the PM to remained locked for long time periods.

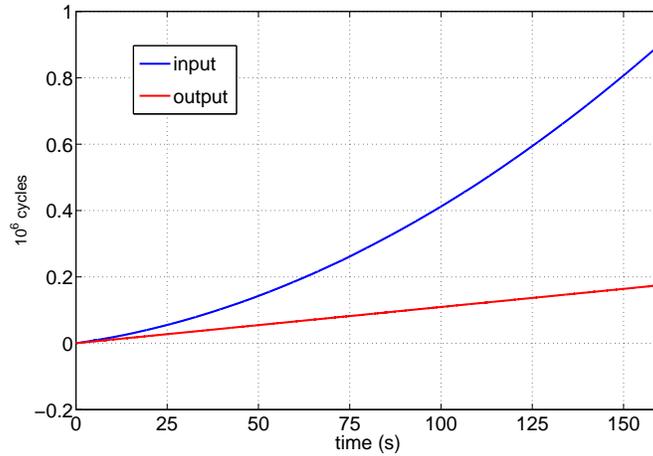


Figure 5-18. Raw phase timeseries of arm-locking system characterization data. Blue curve is $L_2 - L_0$ beat note, red is interferometer output, $S(t)$.

For an input with a quadratic trend in phase, the interferometer output, $S(t)$, will have a linear trend in phase. Using (5-4),

$$S(t) \equiv \phi_{20}(t) - \phi_{20}(t - \tau) = (2a\tau)t + (b\tau - a\tau^2). \quad (5-5)$$

This corresponds to a frequency offset of $2a\tau$ or ~ 50 Hz in the case of the data in Figure 5-18. $S(t)$ does indeed show a linear trend, although the slope is 1092 Hz rather than 50 Hz. Part of the discrepancy can be traced to a rounding error in the frequency offset register of the PM, which uses a U16.16 binary fraction to represent ν_{off} . The rounding error for a 4 MHz signal is approximately 671 Hz. The additional 370 Hz trend is unaccounted for.

Removing the quadratic trend from $\phi_{20}(t)$ and the linear trend from $S(t)$ yields the timeseries in Figure 5-19. Over long time periods, $S(t)$ is quieter than $\phi_{20}(t)$ since variations with periods longer than ~ 1 s are common to both the prompt and delayed signals.

Figure 5-20 shows the linear spectral densities $\tilde{\phi}_{20}(f)$ and $\tilde{S}(f)$. Clearly visible in $\tilde{S}(f)$ is a flattening at low frequencies as well as nulls at multiples of $1/\tau = 1$ Hz. An

estimate of the magnitude of the interferometer transfer function can be made by dividing the two curves in Figure 5-20, as shown in Figure 5-21.

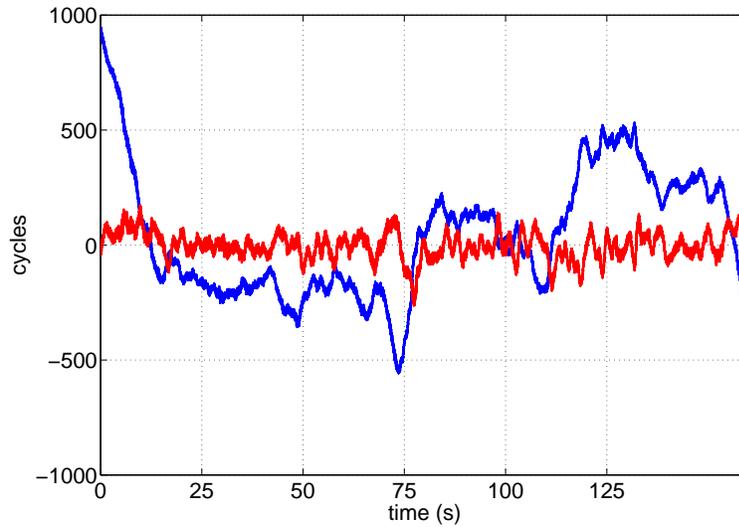


Figure 5-19. Detrended timeseries of arm-locking system characterization data. Blue curve is the quadratically-detrended $L_2 - L_0$ beat note, red is the linearly-detrended interferometer output, $S(t)$.

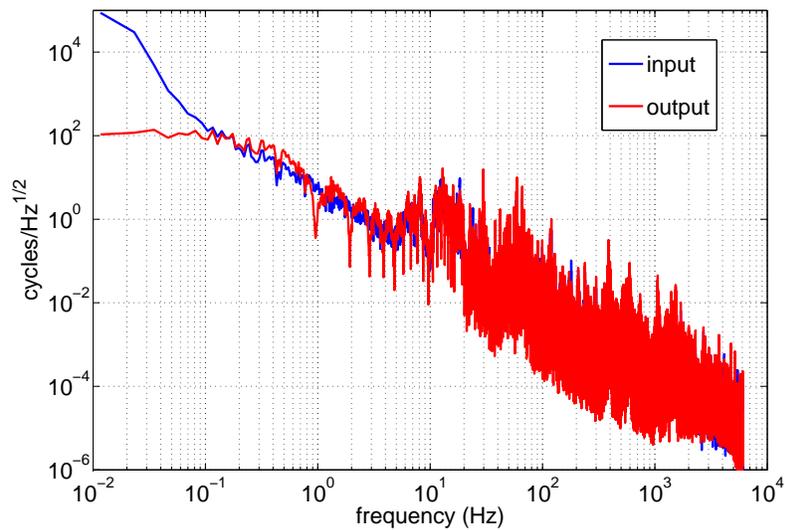


Figure 5-20. Linear spectral density of arm-locking system characterization data

Also shown in Figure 5-21 is the theoretical single-arm transfer function given in (3-20). The magnitude is given by

$$|1 - e^{-s\tau}| = 2 |\sin(\pi f\tau)| \quad (5-6)$$

The best fit using (5-6) for the data in Figure 5-21 obtained a delay of $\tau = 1.039$ s. The additional 39 ms of delay is consistent with the additional delay observed in the 3rd-generation EPD experiments described in 4.6.2. The additional roll-off in the measured response at low frequencies may only be due to a scarcity of points in the spectra in Figure 5-21. The measured depth of the interferometer nulls is affected by the frequency-resolution of the spectra. A longer measurement time would reduce both of these effects.

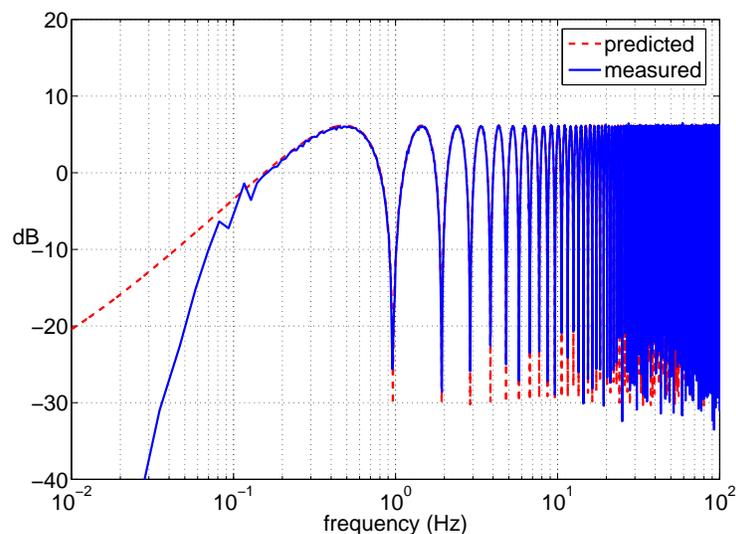


Figure 5-21. Magnitude of transfer function for arm-locking system. Fitted system has a delay of $\tau = 1.039$ s.

5.4.2 Filter Design

The presence of the PM in the loop as shown in Figure 5-17 greatly reduces the requirements on the arm-locking filter. Unlike the system described in Section 5.3, the error signal will be present regardless of the amount of noise in the system. Consequently,

the arm-locking loop does not have to suppress the noise in $\phi_{20}(t)$ to less than a cycle in order to avoid fringe slips.

Another advantage of the PM is that it can be used as a frequency-measurement device by using the ν_{corr} signal from the PM front-end as the frequency output. The measurement of frequency rather than phase introduces an additional factor of s in the arm-locking transfer function which can be balanced by a factor of $1/s$ in the control filter. This allows the control filter magnitude to decrease with increasing frequency, a situation that is generally more stable.

The phase error introduced in using ν_{corr} rather than the full $IQ\nu$ signals is equivalent to the residual phase, $\phi_r = \text{atan}(Q/I)$. So long as the PM remains locked, the rms value of ϕ_r will be less than 0.5 cycles. In the LISA band, ϕ_r should lie below the digitization noise in ν_{corr} .

For the system described in Section 5.4.1 above, a control filter with poles at 1 Hz, 10 Hz, 100 Hz, and 1 kHz and zeros at 3.163 Hz, 31.63 Hz, 316.3 Hz, and 3.163 kHz was designed using the bilinear transform method (Section A.4.3.2). This produces a transfer function that approximates $s^{-1/2}$ between 1 Hz and 10 kHz. The filter was implemented as a two-stage, second-order-section, direct-form II filter in the FPGA on-board the model 6256 downconverter².

Figure 5-22 shows a timeseries of filtered and unfiltered frequency noise for a VCO input. This data was used to generate a measurement of the magnitude of the filter transfer function by computing the linear spectral density for each signal and dividing the two spectra. This produces the measured result shown in Figure 5-23. It is clearly consistent with the predicted result, indicating that the filter is behaving as expected.

² For more information on filter structures see Section A.4.4.1. For an overview of the Pentek hardware, see Section 4.4.

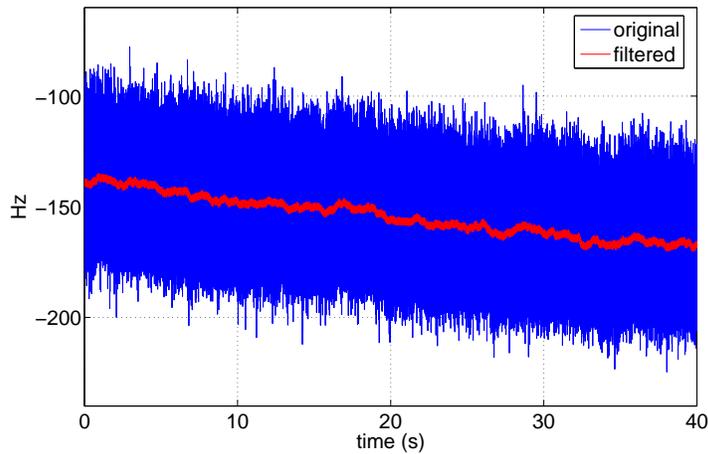


Figure 5-22. Timeseries of filtered and unfiltered frequency noise from VCO input

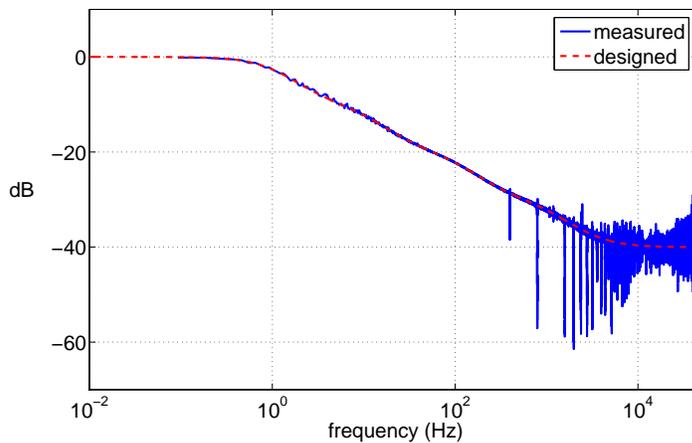


Figure 5-23. Measured and designed transfer functions of arm-locking control filter

As mentioned in Section 4.4, the DAC outputs on the 6228 upconverter are transformer-coupled with a -3 dB point at 400 kHz. Therefore they cannot be used to generate a DC control signal to pass to a VCO. To avoid this problem, the filter output can be used as a frequency input to an NCO running on the FPGA in the 6228. The NCO can then replace the VCO as the oscillator in the PLL between L_1 and L_2 in Figure 5-5. This should be similar to the situation in LISA, where the PM, control electronics, and NCO will all be part of a common avionics system.

Unfortunately, a problem with the Pentek system prevented the filter system from being completed. Although the PM/filter combination and the NCO can operate separately, there is a conflict in running them together which prevents the filtered output signal from reaching the NCO. The problem appears to result in the interrupt controller for the PCI bus on the model 4205 baseboard, but this has yet to be confirmed. A possible work-around involving bypassing the PCI bus has been conceptualized, and is currently in the implementation stage.

5.4.3 Results

Without a fully-operational filter, the improved optical model could not be locked. Nevertheless, the results from the transfer function measurements in [5.4.1](#) indicate that the system is capable of producing error signals with LISA-like noise and LISA-scale time delays. As potential PM/filter systems are developed at UF and elsewhere, the EPD-based optical arm-locking model will be available to evaluate them.

CHAPTER 6 CONCLUSION

Gravitational wave astronomy promises to be a major contributor to our understanding of the universe in the 21st century. It will provide a new method for gathering information that will compliment our existing ability for electromagnetic observations. GWs will be particularly important to our understanding of gravity, the dominant force over large distance scales.

Space-based detectors such as LISA will play an important role in GW astronomy. They are the only type of detectors currently under study that can access the low-frequency regime between $\sim 10^{-6}$ Hz and ~ 1 Hz. This band is rich with interesting sources such as galactic binaries, SMBH mergers, and EMRIs. In order to reach its design sensitivity, LISA relies on the precision and dynamic range of its phasemeters (PMs) as well as novel interferometric techniques such as time-delay interferometry (TDI) and arm-locking.

The UF LISA interferometry simulator provides an opportunity to study both technologies such as PMs and techniques such as arm-locking in a LISA-like environment. Realistic noise sources are generated using pre-stabilized lasers and other components similar to those in LISA. The electronic phase delay (EPD) technique allows the simulator to model the large optical path lengths in LISA, which is essential for studying system-level technologies such as TDI and arm-locking.

This dissertation has described the author's work in developing the simulator, especially the electronic subsystems that are critical to its success. It has also described a series of arm-locking experiments that gradually grew more sophisticated as the simulator improved. The results of these two efforts are summarized below.

6.1 Phasemeters and EPD Units

The success of the simulator hinges on the performance of the electronic subsystems, namely the EPD unit and the PM. Both of these technologies have been developed over several iterations. The current EPD unit is capable of delaying signals with carrier frequencies up to ~ 25 MHz for more than 300 s. It can also be used to shift the frequency

of the delayed signal, which can be used to model the Doppler shifts in the LISA arms. This is an essential feature for generating signals that are precise analogues of the LISA photodiode signals.

The current hardware PM is capable of measuring signals with LISA-like phase noise. The noise floor of the PM appears to be limited by digitization noise to a level of approximately

$$\left(10^{-7} \frac{\text{cycles}}{\sqrt{\text{Hz}}}\right) \sqrt{1 + \left(\frac{100 \text{ Hz}}{f}\right)^2} \quad (6-1)$$

where f is the Fourier frequency. While this noise does not meet the LISA requirements over the entire LISA band, the digitization noise could be brought below the LISA requirement by increasing the number of bits in the frequency-correction signal from 32 to 42. This should not be a major obstacle, but it will require some modifications to the PM software.

6.2 Arm-Locking

The experimental studies of arm-locking (Chapter 5) have demonstrated that the EPD technique can be used to create valid models of LISA interferometry. The results of the initial optical experiment with a ≈ 1 ms delay produced a beat note with a frequency noise of $\sim 200 \text{ mHz}/\sqrt{\text{Hz}}$ between 10 mHz and 250 mHz. This represents an improvement over the pre-stabilized lasers of more than three orders of magnitude. At higher frequencies, the noise suppression was consistent with the predictions of a Laplace-domain model of arm-locking. Based on measurements of the out-of-loop noise suppression and the in-loop error signal, the unity gain frequency was estimated to be $\approx 12 \text{ kHz}$. This demonstrates that the arm-locking loop was stable even though roughly ten interferometer nulls were included in the feedback bandwidth.

The improved optical model can produce a LISA-like error signal with delays up to the full LISA round-trip of 33 s. The magnitude of the model's transfer function was measured for a 1 s delay and found to be consistent with that of a single-arm error signal in LISA. A controller for locking the model was designed and built using the hardware PM

and a digital filter on an FPGA. While the controller is able to produce a signal with the correct transfer function, a problem with the digital signal processing hardware prevented the loop from being closed. Despite this drawback, the improved model stands as the first realistic model of a LISA arm. As potential controllers are developed at UF and elsewhere, it will be available to evaluate them.

APPENDIX A DIGITAL SIGNAL PROCESSING

A.1 Introduction

Digital signals differ from analog signals in two fundamental ways¹. They are discrete-time signals, meaning that they are only defined at particular times, typically multiples of a fixed clock period. The properties of discrete-time signals are discussed in Section A.2. Digital signals are also quantized in amplitude, they can only take on a limited number of values. The consequences of amplitude quantization are discussed in Section A.3.

Section A.4 provides an introduction to digital filtering, an extremely flexible and powerful technique for manipulating digital signals. Emphasis is placed on the types of digital filters used in developing the simulator electronics.

A.2 Sampling

Consider the continuous-time signal, $x(t)$, shown schematically in Figure A-1(a). By definition as a continuous signal, the value of $x(t)$ is specified for all values of t . The signal can be described in frequency space by the spectrum $\tilde{x}(f)$ shown in Figure A-1(b), which is related to the time series via the Fourier transform:

$$\tilde{x}(f) = \mathcal{F}[x(t)], \tag{A-1}$$

$$\mathcal{F}[x(t)] \equiv \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{2\pi i f t'} x(t') dt'. \tag{A-2}$$

Sampling the signal amounts to measuring $x(t)$ at specific values of t separated by the sample period T . The sampled signal can be represented as

$$S(t) = \begin{cases} x(t) & t = kT \\ 0 & t \neq kT \end{cases} \tag{A-3}$$

¹ For a more thorough treatment of digital signal processing, see Smith [67]

where k is an integer. This is equivalent to multiplying $x(t)$ by a comb of delta functions spaced by T ,

$$S(t) = x(t) \times C(t), \tag{A-4}$$

$$C(t) = \sum_{k=-\infty}^{\infty} \delta(t - kT). \tag{A-5}$$

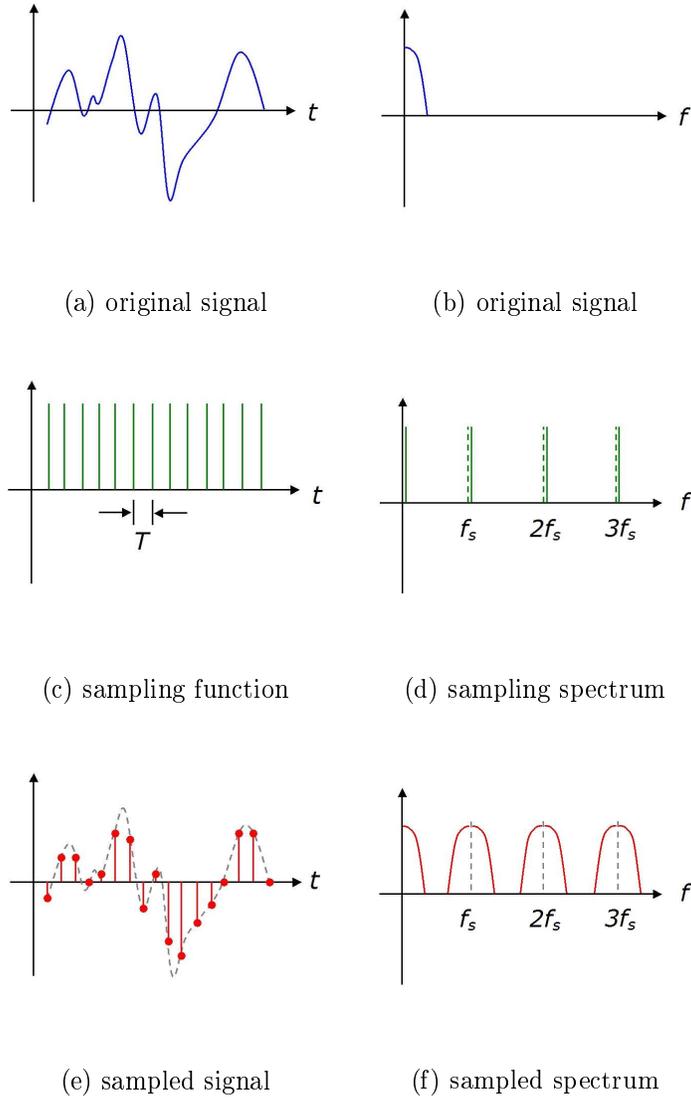


Figure A-1. An overview of the sampling process

This signal is shown in Figure A-1(c). The resulting signal, $S(t)$, shown in Figure A-1(e) is sometimes called an impulse train. The spectrum of the sampled signal can be determined using the fact that the Fourier transform of a product of two signals is equal to the convolution of the Fourier transforms of the two signals. In this case,

$$\tilde{S}(f) = \tilde{x}(f) \otimes \tilde{C}(f). \quad (\text{A-6})$$

The Fourier transform of comb of delta functions is another comb of delta functions.

Expressing all terms as positive frequencies, the spectrum of $C(t)$ is

$$\tilde{C}(f) = \frac{f_s}{\sqrt{2\pi}} \left\{ \delta(f) + \sum_{k=1}^{\infty} [\delta(f - kf_s) + \delta(-f - kf_s)] \right\}, \quad (\text{A-7})$$

where $f_s \equiv 1/T$, is the sampling frequency. The spectrum $\tilde{C}(f)$ is plotted in Figure A-1(d). The convolution in (A-6) will contain three terms. The first term, from the $\delta(f)$ in $\tilde{C}(f)$, places a copy of $\tilde{x}(f)$ at DC. The second term, from the $\delta(f - kf_s)$, places copies of $\tilde{x}(f)$ at each multiple of kf_s . The final term, from the $\delta(-f - kf_s)$, places frequency-reversed copies of $\tilde{x}(f)$ at each multiple of kf_s . The overall spectrum of the sampled signal, $\tilde{s}(f)$, is shown in Figure A-1(f).

A.2.1 Aliasing

With the exception of the spectral “images” in Figure A-1(f), the sampled signal is an exact replica of the time-domain signal. If the bandwidth of $\tilde{x}(f)$ is limited to $\leq f_s/2$, as show in Figure A-1, then the images do not effect the signal. If however, the bandwidth of $\tilde{x}(f)$ exceeds $f_s/2$, then the frequency-reversed image from $1/T$ will begin to overlap with the original spectrum. When this occurs, the contributions from the DC and the $1/T$ image overlap, as shown in Figure A-2.

This phenomenon is known as aliasing and is generally undesirable because high-frequency noise components in the original signal can map into lower frequencies in the sampled signal. To prevent aliasing, the input signal must be band-limited below

the Nyquist frequency,

$$f_{Nyq} \equiv f_s/2. \quad (\text{A-8})$$

If the input signal is not naturally band-limited, it is commonly filtered with an “anti-aliasing” filter before being sampled. Aliasing can also occur during the process of *downsampling*, the process of reducing the sample rate from a high rate to a lower rate in a multi-rate discrete-time system. The Nyquist sampling theorem applies in this case as well, although the anti-aliasing filters can be digital filters running at the high sample rate. In multi-rate filters, such as the CIC filters discussed in Section A.4.5, the anti-aliasing and downsampling functions can be combined into one operation, improving efficiency.

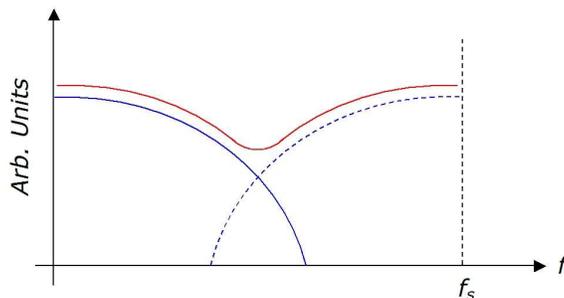


Figure A-2. The phenomenon of aliasing. The original spectrum (solid blue line) and the imaged spectrum (dashed blue line) overlap. The contributions from each cannot be separated in the resulting spectrum (red).

A.2.2 Upconversion

In some cases, it is necessary to convert a discrete-time signal to a continuous-time signal, a process known as upconversion. An example would be when using a digital filter to drive a PZT actuator on a laser. One way to do this would be to insert zeros between the known samples, re-creating the impulse train in Figure A-1(e). The resulting output spectra would contain the desired spectra plus the images, and could be filtered to eliminate the images.

This technique is often used in the related process of *upsampling*, increasing the sample rate from a low rate to a higher rate in a multi-rate discrete-time system. In

digital-to-analog converters (DACs), this is usually not practical, because it requires fast response in the analog components. Instead, most DACs utilize a technique known as zero-order hold (ZOH).

In a ZOH upconversion, the output is held constant between samples, producing a terraced effect. This can be described mathematically as a convolution in the time-domain of the impulse train and a rectangular impulse of width T

$$u(t) = s(t) \otimes R(t), \quad (\text{A-9})$$

$$R(t) \equiv \begin{cases} 1 & 0 \leq t < T \\ 0 & \text{otherwise} \end{cases}. \quad (\text{A-10})$$

This is shown in the left three panels of Figure A-3. In the frequency domain, the convolution implies a multiplication of the spectra

$$\tilde{u}(f) = \tilde{s}(f)\tilde{R}(f). \quad (\text{A-11})$$

The Fourier transform of a rectangular impulse is a sinc function,

$$|\tilde{R}(f)| = T \text{sinc}(fT), \quad (\text{A-12})$$

where $\text{sinc}(x) \equiv \sin(\pi x)/\pi x$. This is shown in Figure A-3(d). The spectrum of the upconverted signal, shown in Figure A-3(f) tracks $\tilde{s}(f)$ at low frequencies, but begins to fall off at higher frequencies due to the sinc response of $\tilde{R}(f)$. In theory, this passband droop can be corrected by a “reconstruction filter” with a sinc^{-1} frequency response. Such a filter is difficult to build in the analog domain, especially since for $f \geq f_{Nyq}$, the response of the filter should be zero to eliminate the spectral images at higher frequencies. In most systems the reconstruction filter is similar to the anti-aliasing filter, with a flat response in the pass-band. The effect of the sinc response can be reduced at a particular frequency by increasing the upsampling rate.

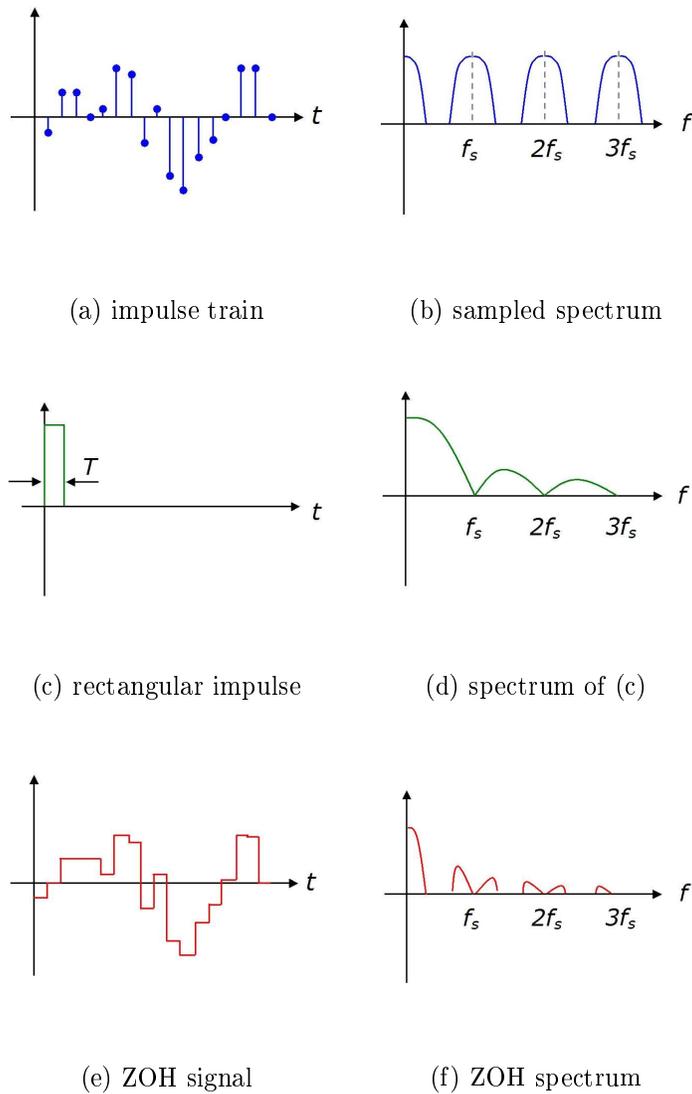


Figure A-3. An overview of the upconversion process

Associated with the sinc response in the magnitude of $\tilde{R}(f)$ is a linear phase response,

$$\angle R(f) = -\pi fT. \quad (\text{A-13})$$

This linear phase loss is equivalent to a delay of $T/2$ in the output. Graphically, it corresponds to the fact that the best-fit curve through the ZOH output is shifted from the original curve by $T/2$. This “delay” can be an important consideration when designing digital filters for real-time control systems.

A.3 Digital Signals

As mentioned in the introduction, digital signals are quantized in amplitude as well as in time. Each piece of information in a digital system is contained in a two-state system. Examples include two voltage levels, two current levels, or two magnetic states. The information can be represented by a binary digit, or *bit*, that can take on the values 0 or 1.

Digital signal processing refers to manipulating a bit or group of bits. For a single bit, the only operation is logical negation, abbreviated as NOT. For pairs of bits, the Boolean operations AND, OR, XOR (exclusive or), and their negated counterparts can be defined as well. These operations can be easily extended to groups of bits known as *words*.

An N -bit word can have 2^N possible values. These 2^N values can be chosen to represent any set of 2^N numeric values. For example, an 8-bit word can be used to describe the integers from 0 to 255 (a total of $2^8 = 256$ values). In this case, the mapping is simple: the k^{th} bit in the word represents the value of the 2^k digit. The word “10001101” with the most-significant bit (MSB) to the left, would represent $2^7 + 2^3 + 2^2 + 2^0 = 141$.

The Boolean logical operations can be combined to construct other digital operators. For example, an N -bit adder will generate the sum of two N -bit signals. In general, the sum of two N -bit signals can take on $2^{N+1} - 2$ different values, which requires $N + 1$ bits to represent. Some adders will include this extra bit, but many will express the output as an N -bit number. For sums larger than $2^N - 1$, the result will either be fixed at $2^N - 1$, known as saturation, or will “wrap” back into the 2^N values starting at zero. For example,

$$\begin{array}{r} 10001101 \\ +10001101 \\ \hline 00011011 \end{array} \Leftrightarrow \begin{array}{r} 141 \\ +141 \\ \hline 27 \end{array}. \quad (\text{A-14})$$

To compute the result using integers, the two integers are summed ($141 + 141 = 282$) and the result is expressed modulo 255, ($282 \bmod 255 = 27$).

The 2^N values in an N -bit word can also be used to represent signed integers. For example, an 8-bit word could be used to describe the integers from -128 to 127. In this

case there are several choices of mapping, the most common of which is known as *two's compliment*. In the two's compliment representation, the MSB corresponds to the $-(2^{N-1})$ digit. The word "10001101" would correspond to -115 in two's compliment. The name two's compliment refers to the property that the negative of a number can be generated by negating all of the bits and adding one. Flipping the bits gives "01110010" and adding one gives "01110011" = 115. Efficient adders can be built using two's compliment representation to add or subtract signed integers.

A.3.1 Binary Fractions

Digital words can be used to describe non-integer values as well. A common way of doing so is using binary fractions, where a word is divided into an integer portion and a fractional portion by a "binary point". For example, an 8-bit word could be divided into 4 bits of unsigned integer and 4 bits of fraction by placing the binary point after the fourth bit. The bit to the left of the binary point represents the 2^0 place while the bit immediately to the right represents the 2^{-1} place. The word "1000.1101" would represent $2^3 + 2^{-1} + 2^{-2} + 2^{-4} = 8.8125$.

Binary fractions can also be used with two's compliment to formed signed values. In two's compliment, "1000.1101" would represent $-2^3 + 2^{-1} + 2^{-2} + 2^{-4} = -7.1875$. The negation procedure for two's compliment still works with binary fractions. Hence "0111.0011" represents $2^2 + 2^1 + 2^0 + 2^{-3} + 2^{-4} = 7.1875$.

The various representations are often abbreviated using the notation *Type N.B. Type* is either "U" for unsigned representations or "S" for signed representations. N is the total number of bits in the word and B is number of bits to the right side of the binary point. For example S16.15 refers to a 16-bit signed integer with 15 bits of fraction, which can represent the values from -1 to $1 - 2^{-15}$ in increments of 2^{-15} .

A.3.2 Multiplication and Other Operations

The multiplication of two digital words is similar to addition. The product of two N -bit words can span $2^{2N} - 2^{N+1} + 1$ possible values. As with adders, it is either possible

to use enough bits for the product to cover all possible values or to use a smaller number of bits and throw away some information. This information can either be taken from the MSBs by saturation or wrapping or it can be taken from the least-significant bits by rounding or truncation.

As an example, consider two U4.2 words, each representing the numbers from 0 to 3.75 in increments of 0.25. The product can range from 0 to 14.0625 in increments of 0.0625. The entire range of possible products can be described by a U8.4 word. However, consider the case when only 4 bits are available to store the output. One choice is to use a U4.0 word, which ranges from 0 to 16 in increments of 1. This will ensure that the largest possible values of the product can be represented, but will result in a loss of precision.

Another choice is to use a U4.4 word, which covers the range from 0 to 0.9375 in increments of 0.0625. This provides all of the precision, but has a limited range and will overflow (through saturation or wrapping) for some values. Using a U4.2 word for the output represents a compromise between these two extremes.

In systems with many additions and multiplications, it is important to carefully manage the overflow and rounding conditions of each operator. Undesired overflows or rounding errors can cause otherwise well-designed systems to perform poorly or not at all.

Division of two digital words is not easily accomplished. Division by a constant value can be accomplished by multiplication by the reciprocal value, but division of two arbitrary words is typically avoided in fixed-point systems.

Nonlinear operations such as *log*, *sin*, *atan*, etc. can be accomplished using look-up tables (LUTs). A LUT is a list of words stored in memory or hard-coded in digital logic. A particular word in the list is selected by providing an address word to the input of the LUT. The *depth* of the LUT refers to the number of bits in the address word while the *width* of the LUT refers to the width of the words at each address. This is sometimes abbreviated as *depth* \times *width*. For example, an 8×10 LUT could be used to store the function $\sin(\phi)$. The address word would correspond to ϕ from 0 to $(2\pi)(1 - 2^{-8})$ rad

in increments of $(2\pi)2^{-8}$ rad. The word at the corresponding address would contain the function represented as a 10-bit integer, say as S10.8. A LUT will generate errors due to both the finite length of the address word and the finite width of the output.

A.3.3 Floating-point Representations

The binary representations of integers and fractions discussed above are generally called *fixed-point* representations, because the binary point is fixed at a particular location in the word. There are also a set of *floating-point* representations, where the binary point is non-stationary. The most common floating-point representations use the sign, mantissa, exponent representation:

$$x = (\pm)(mantissa) \exp(exponent). \tag{A-15}$$

One bit, typically the MSB, is used to indicate the sign of the number. A second sub-word is used to describe a mantissa using an unsigned fixed-point representation. The remainder of the bits are used to describe the exponent using a signed fixed-point representation. The IEEE maintains standards for floating-point representations of numbers, some of which are summarized in Table A-1. In addition to the numbers explicitly covered, they also have representations for $\pm\infty$ and not-a-number (NaN), which signifies a mathematical error such as divide by zero.

Table A-1. IEEE standard floating point representations

type	total bits	mantissa	exponent	range
float	32	23	8	$\pm 10^{38.53}, \pm\infty, \text{NaN}$
double	64	52	11	$\pm 10^{308.25}, \pm\infty, \text{NaN}$

For numbers of roughly the same magnitude, the precision of the representation is constant. As the size of the number increases or decreases, the precision is correspondingly decreased or increased. This happens in discrete steps, which can sometimes produce strange results.

It is much more difficult to build floating-point adders and multipliers using basic logic gates. As a result, floating-point representation is typically employed only on microprocessors or specialized DSP platforms. With these platforms, it becomes possible to perform a wide variety of operations including non-linear operations, with minimal error.

A.3.4 Digitization Noise

The combined effects of amplitude quantization and sampling lead to a noise source known as digitization noise. The quantization error can be defined as the difference between the true value of a sample and the quantized value. In most situations it is a reasonable assumption that the quantization error for a particular point is equally likely to lie between $U_{LSB}/2$ and $-U_{LSB}/2$, where U_{LSB} is the level of the most significant bit. This probability distribution function (PDF) is pictured in Figure A-4.

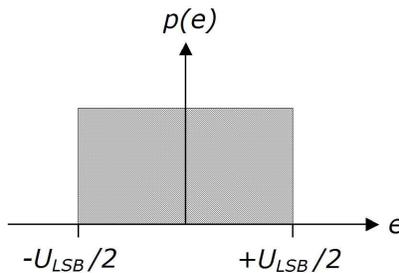


Figure A-4. Assumed PDF for quantization error

In general, the energy contained in the quantization error is equal to the second moment of the PDF,

$$E_{dig} = \frac{\int_{-\infty}^{\infty} e^2 (p(e) - \bar{p}) de}{\int_{-\infty}^{\infty} p(e) de} \quad (\text{A-16})$$

where $p(e)$ is the probability for a particular error e and \bar{p} is the mean of $p(e)$. For the uniform PDF in Figure A-4, (A-16) can be evaluated as

$$E_{dig} = \frac{\int_{-U_{LSB}/2}^{U_{LSB}/2} e^2 de}{\int_{-U_{LSB}/2}^{U_{LSB}/2} de} = \frac{U_{LSB}^2}{12} \quad (\text{A-17})$$

The spectral distribution of the quantization error is assumed to be white. For white noise sampled with a sampling frequency f_s , the total energy can be written as

$$E_{dig} = \tilde{U}_{dig}^2 \frac{f_s}{2} \quad (\text{A-18})$$

where \tilde{U}_{dig} is the spectral amplitude of the white noise. For white noise with a constant energy, \tilde{U}_{dig} will decrease as f_s is increased because the energy can spread over more frequency bands. Substituting $U_{LSB}^2/12$ for E_{dig} in (A-18) and solving for \tilde{U}_{dig} gives

$$\tilde{U}_{dig} = \frac{U_{LSB}}{\sqrt{6f_s}}. \quad (\text{A-19})$$

This formula for digitization noise makes two assumptions: the PDF for the quantization error is the one in Figure A-4 and the spectrum of the noise is white. For signals with a random noise component on the order of U_{LSB} , this is a reasonable assumption.

For signals with lower noise (or fewer bits), one or both of these assumptions can break down. Figure A-5(a) shows the PDF of the quantization error for a pure sinusoidal signal with unit amplitude and frequency $f = f_s/10$ quantized as a S6.5 binary fraction.

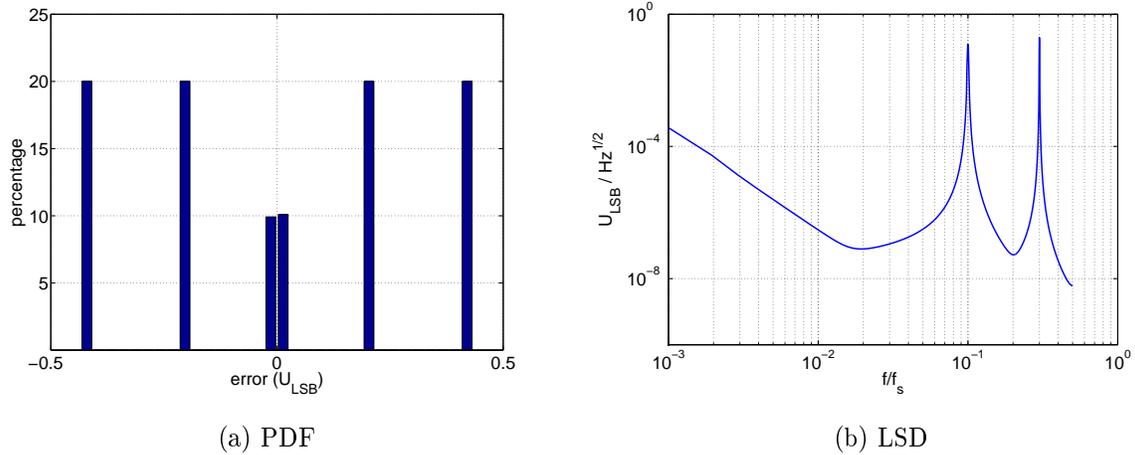


Figure A-5. Non-uniform, non-white quantization error from a sinusoid with unit amplitude and $f = f_s/10$ quantized as a S6.5 binary fraction

The quantization error in Figure A-5(a) falls into a limited range of bins and is no longer uniform. The linear spectral density (LSD) of the quantization error is shown in Figure A-5(b). The peaks at $f_s/10$ and $2f_s/10$ result from the fact that the timeseries of the quantization error is a repetitive sequence.

A.4 Digital Filtering

In the language of signal processing, a filter is a device with a specific transfer function. Typically filters refer to devices which are designed to alter the overall transfer function of a physical system. In the case of closed-loop systems, the transfer function of the feedback element is often divided into a control filter, which can be specified, and an actuator, which is a fixed attribute of the system.

Digital filters² are operators which convert a set of sampled input values into a set of sampled output values. A generic digital filter is specified by a recursion relation relating inputs and outputs at different sample times. For an input $x[n]$ and output $y[n]$ where n denotes the n^{th} sample, a generic recursion relation for a digital filter is

$$y[n] = \sum_{i=0}^N b_i x[n-i] + \sum_{j=1}^M a_j y[n-j]. \quad (\text{A-20})$$

The recursion coefficients a_j and b_i specify the behavior of the filter. The form in (A-20) is known as a causal filter, since the output only depends on previous outputs and the current and previous input. Filters that operate in real-time, such as control filters must be causal. Filters used for offline data processing can be acausal, including negative indicies for i that amount to knowledge of future samples.

² See Jackson [68] for a more complete treatment of digital filtering

A.4.1 Time-domain Response

In the continuous-time domain, filters can be described by their impulse response function, $h(s)$. The impulse response function describes the time-series output of the filter for an impulsive input, $x(t) = \delta(t)$.

Discrete-time filters can also be described by an impulse response function, $h[n]$. The discrete-time impulse response is the output of the filter for a unit impulse in the input, $x[n] = \delta_{n0}$. Digital filters are divided into two types according to their impulse response functions. For finite-impulse response (FIR) filters, $h[n]$ goes to zero after a certain number of samples. Referring to (A-20), an FIR is a filter with $a_j = 0$, or no feedback. For an FIR,

$$h[n] = \begin{cases} b_n & 0 \leq n \leq N \\ 0 & \textit{otherwise} \end{cases} . \quad (\text{A-21})$$

By contrast, for an infinite impulse response (IIR) filter, $h[n]$ never reaches zero. This occurs when the filter uses feedback, $a_j \neq 0$ for some j .

A.4.2 Frequency Response

The frequency response of a continuous-time filter is often expressed in the Laplace domain, which is related to the continuous time-domain by the Laplace transform,

$$x(s) = \mathcal{L} \{x(t)\} , \quad (\text{A-22})$$

$$\mathcal{L} \{x(t)\} \equiv \int_0^{\infty} e^{st'} x(t') dt' , \quad (\text{A-23})$$

The transfer function of a filter with input $x(t)$ and output $y(t)$ can be written in the Laplace domain as

$$H(s) \equiv \frac{y(s)}{x(s)} . \quad (\text{A-24})$$

The transfer function is also the Laplace transform of the impulse response function, $\mathcal{L}\{h(t)\} = H(s)$. In general, $H(s)$ can be described using *zeros* (locations in the s -plane where $|H(s)| = 0$) and *poles* (locations in the s -plane where $|H(s)| = \infty$). In equation

form,

$$H(s) = K \frac{\prod_i^N (s - \zeta_i)}{\prod_j^M (s - \rho_j)}, \quad (\text{A-25})$$

where K is a constant, ζ_i is the location of the i^{th} zero, and ρ_j is the location of the j^{th} pole in the complex s -plane.

The Laplace variable, s is related to the Fourier frequency by the relation

$$s = \sigma + 2\pi if, \quad (\text{A-26})$$

where σ is a real constant. The frequency response of a filter with Laplace-domain transfer function $H(s)$ can be determined by evaluating $H(s)$ along the imaginary s -axis, as shown in Figure A-6(a).

The frequency response of a digital filter can be computed using the z -transform, a discrete analogue of the Laplace transform.

$$x(z) = \mathcal{Z} \{x[n]\}, \quad (\text{A-27})$$

$$\mathcal{Z} \{x[n]\} \equiv \sum_{k=0}^{\infty} z^k x[k]. \quad (\text{A-28})$$

The z -transform has the property that $\mathcal{Z}\{x[n-k]\} = z^{-k} \mathcal{Z}\{x[n]\}$. Using this relationship, the recursion relation in (A-20) can be expressed in the z -domain as

$$y(z) = x(z) \sum_{i=0}^N b_i z^{-i} + y(z) \sum_{j=1}^M a_j z^{-j}. \quad (\text{A-29})$$

This can be rearranged to define the z -domain transfer function of the digital filter as

$$H(z) \equiv \frac{y(z)}{x(z)} = \frac{\sum_{i=0}^N b_i z^{-i}}{1 - \sum_{j=1}^M a_j z^{-j}}. \quad (\text{A-30})$$

It can be shown that the z -domain transfer function is the z -transform of the impulse response function, $H(z) = \mathcal{Z}\{h[z]\}$. An alternative expression for $H(z)$ can be obtained by factoring the numerator and denominator of (A-30). The pole-zero representation of $H(z)$ is

$$H(z) = K \frac{\prod_{i=0}^N (z - \zeta_{z,i})}{\prod_{j=0}^M (z - \rho_{z,j})}, \quad (\text{A-31})$$

where K is a constant, $\zeta_{z,i}$ the i^{th} zero and $\rho_{z,j}$ is the j^{th} pole in the complex z -plane.

It is important to note that the z -domain poles and zeros are not generally the same as the Laplace domain poles and zeros given in (A-25). The two can be related using the relationship between the z -transform variable is related to the Laplace variable,

$$z = e^{sT}, \quad (\text{A-32})$$

where T is the sampling period of discrete-time system. The z -transform variable can also be related to the Fourier frequency by combining (A-26) with (A-32),

$$z = e^{2\pi i f T}. \quad (\text{A-33})$$

Geometrically, the frequency axis in the Fourier domain is mapped onto a unit circle centered at the origin in the z -plane, as shown in Figure A-6(b). The frequencies k/T where $k = 0, 1, 2, 3, \dots$ all map to the point $(1, 0)$ on the z -plane. For a filter with a known $H(z)$, the frequency response can be computed using (A-33). In addition, the left-hand side of the complex s -plane, where poles must be located in order for a system to be stable, maps to the region inside the unit circle in the z -plane.

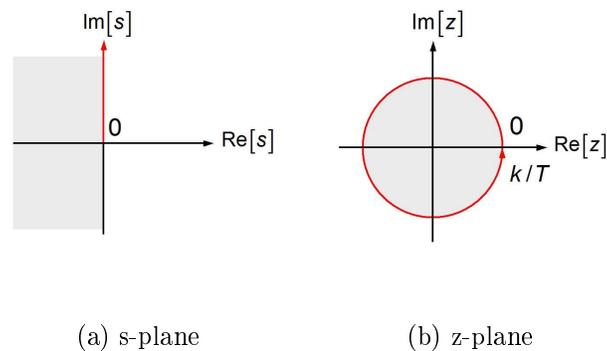


Figure A-6. Comparison of Laplace and z domains. The red line represents the frequency axis.

A.4.3 Design Methods

The previous Section described how to compute the frequency response of a given digital filter. The task of the filter designer is to invert this problem: determine the digital filter for a desired frequency response. A number of standard methods for filter design exist. The following Sections describe a few used by the author. A more thorough treatment can be found in references on the subject such as Jackson [68].

A.4.3.1 FIR Filters - windowed impulse response method

A useful method for designing FIR filters is the windowed transfer function method. This method takes advantage of the relationship between the filter transfer function in the frequency domain and the impulse response function in the time domain. The starting point for the design is the magnitude response of the filter at a series of discrete points in the frequency domain. For example, an ideal low-pass filter will have the transfer function in Figure A-7(a),

$$|H(f)| = \begin{cases} 1 & 0 \leq f < f_c \\ 0 & f_c \leq f \leq f_{Nyq} \end{cases} \quad (\text{A-34})$$

where f_c is the cut-off frequency. The impulse response of the filter can be computed by performing a discrete inverse Fourier transform on the magnitude response. In general, the impulse response function will be an infinite series. For the ideal low-pass in Figure A-7(a), the impulse response is given by,

$$h[n] = \text{sinc}(2f_c n / f_s) \quad -\infty \leq n \leq \infty, \quad (\text{A-35})$$

where f_s is the sampling frequency and $\text{sinc}(x) \equiv \sin(\pi x) / \pi x$. The impulse response is plotted for $-30 \leq n \leq 30$ in Figure A-7(b). This infinite impulse response can be approximated by truncating the series after a total of N points. The edge effects caused by the truncation can be mitigated by multiplying the infinite impulse response with a window function that goes to zero at the end points. For example, the infinite impulse response in Figure A-7(b) can be truncated to a 31-point response using the Hamming

window

$$h[n] = w[n]\text{sinc}(2f_c n/f_s) \quad -15 \leq n \leq 15, \quad (\text{A-36})$$

$$w[n] = 0.54 + 0.46 \cos\left(\frac{2\pi n}{31}\right) \quad (\text{A-37})$$

This is shown in Figure A-7(d). The frequency response of this windowed filter can be determined by performing a discrete Fourier transform of the windowed impulse response function. The frequency-response of the 31-point Hanning-windowed sinc is shown in Figure A-7 (c).

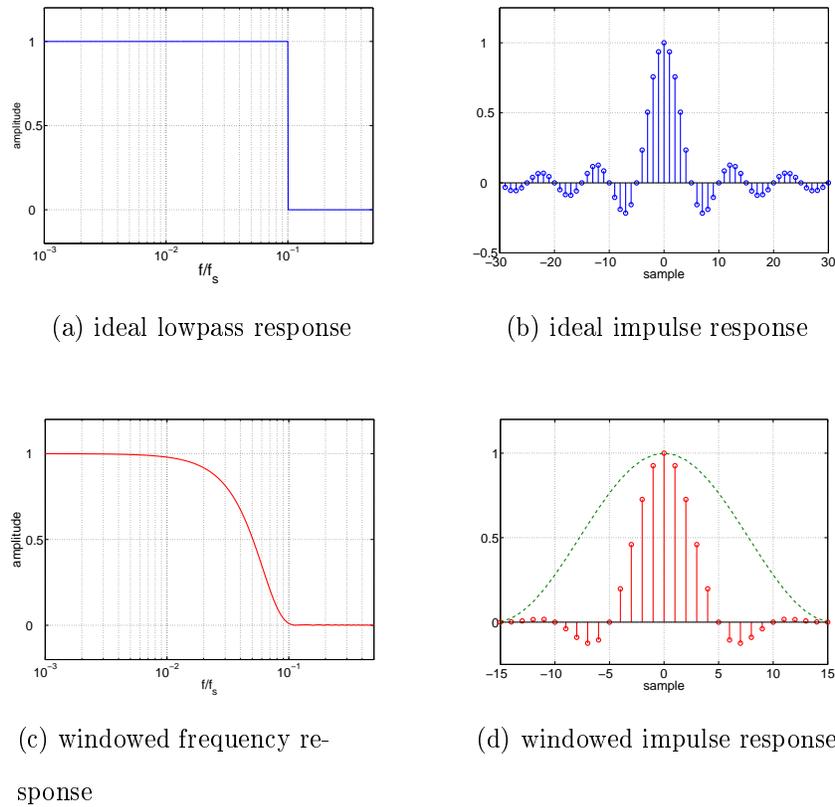


Figure A-7. Windowed impulse response method for designing FIR filters

The actual filter response in Figure A-7(c) exhibits similar low-pass characteristics as the ideal response in Figure A-7(a). It differs from the ideal response in several important ways: the slope of the magnitude response for $f \approx f_c$ is finite, the passband ($0 \leq f < f_c$)

does not have a perfectly flat magnitude response, and the stopband ($f_c \leq f \leq f_{Nyq}$) response is not zero. As the number of points used in the filter is increased, each of these differences decrease. Tradeoffs between the different effects can also be made by adjusting the window function. In general, window functions which produce low ripple (flat response in the passband and stopband) will have shallower slopes while window functions which lead to steep slopes will produce larger ripple.

For an FIR filter, the impulse response function gives the filter coefficients, via (A-21). The windowed impulse response will generally include negative values of n , corresponding to an acausal filter. An equivalent causal filter can be constructed by shifting $h[n]$ to positive values of n . This will produce a filter with the same magnitude response but a linear phase response corresponding to a delay equal to the number of points shifted. For example, the 31-point filter in (A-36) will be shifted to the right by 15 points, corresponding to a delay of $15/f_s$. This will produce a phase response of $-2\pi f(15/f_s)$.

A.4.3.2 IIR Filters - bilinear transform method

Oftentimes a desired filter response is known in the Laplace domain. One technique for converting from a Laplace domain representation, such as the pole-zero form in (A-25), to a set of recursion coefficients for an IIR digital filter is the bilinear transform method. The bilinear transform method starts with the relationship between s and z , given in (A-32). This relationship can be inverted to give

$$sT = \ln z. \tag{A-38}$$

This non-linear relationship makes it difficult to make an analytic conversion between the Laplace-domain poles and zeros and the IIR recursion coefficients. This can be done more easily if (A-38) is simplified using the bilinear approximation,

$$s \approx \frac{2}{T} \frac{1 - z^{-1}}{1 + z^{-1}}. \tag{A-39}$$

which is valid for $|sT| \ll 1$. The steps for generating the recursion coefficients for an IIR filter that will approximate a given Laplace-domain filter are as follows. Use (A-39) to replace s in (A-25). Expand all of the terms in the numerator and denominator, and algebraically manipulate it into the form of (A-30). The recursion coefficients can then be extracted using (A-30). For all but the most simple filters, this manipulation becomes extremely cumbersome to perform analytically. However, it is easy to perform numerically.

A.4.4 Realization and Practicalities

A.4.4.1 Filter structures

Once the filter coefficients are known, a data set can be filtered using (A-20). One way to realize this equation in digital logic would be the *direct-form I* structure in Figure A-8. This implementation requires $N + M$ multipliers, $N + M + 1$ adders, and $N + M$ registers for storing the prior input and output data.

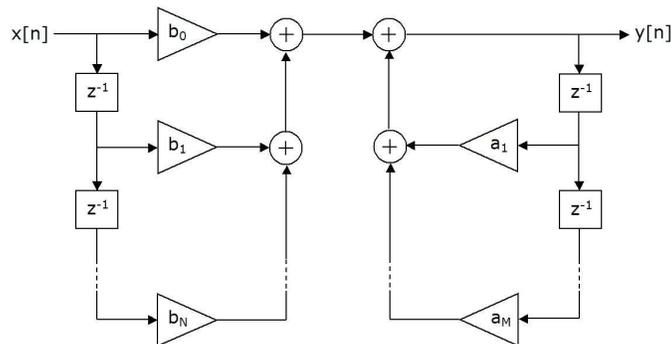


Figure A-8. The direct-form I filter structure

The same result can be achieved with a number of other filter structures, some of which are more computationally efficient. For example, the *direct-form II transposed* (DF2T) filter structure is shown in Figure A-9. The DF2T structure reduces the number of registers to $\max(N, M)$.

Algebraically, all filter structures are identical, but because the operations occur in different order, effects such as quantization noise can vary from structure to structure. This is most pronounced in IIR filters, which are made more sensitive to quantization

error by their feedback. An IIR filter implemented with one filter structure may perform well while the same filter implemented with a different structure may not perform at all.

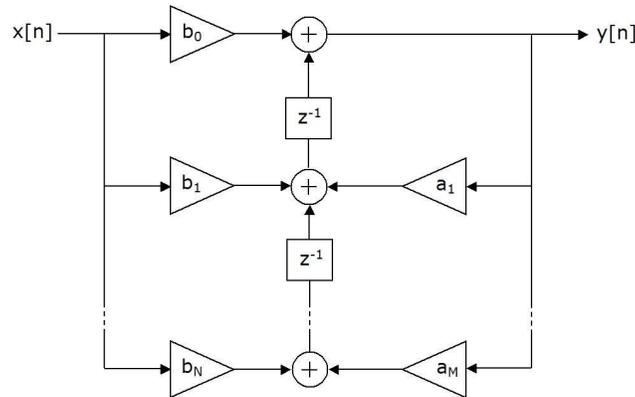


Figure A-9. The direct-form II transposed (DF2T) filter structure

A good choice for IIR filters is a cascade of second-order-sections (SOS). A second-order section is a filter with $N = M = 2$ and can provide up to two pole-zero pairs. Higher-order filters can be formed by cascading a series of SOSs. By adjusting which poles and zeros are placed in which section and the overall order of the sections, tradeoffs can be made between the stability of the filter and its dynamic range. This is similar to the process of designing an analog filter using multiple operational amplifier stages.

A.4.4.2 Latency

A given filter implementation will take a certain amount of time, known as the latency, to produce an output from a given input. This means that the phase response of the filter will have an additional linear phase lag. For filters used in out-of-loop analysis this delay is typically not important and it is sufficient to require only that it remain constant. For filters used in closed-loop systems, such as control filters, the phase lag caused by latency can affect the stability of the loop. One source that is often overlooked is the phase lag due to a ZOH in a system with an analog output. The ZOH phase lag, (A-13), is equivalent to a delay of $T/2$.

A.4.5 CIC Filters

Cascade Integrator-Comb (CIC) filters, also called Hoghenouer filters [69], are a special type of filter used for making large sample rate changes in DSP systems. They use no multipliers, which makes them computationally efficient. An N -stage CIC decimator consists of N integrators (single pole at DC) followed by a decimation of R and N differentiators (single zero at DC). The transfer function of a CIC filter in the z -domain is given by

$$G(z) = \frac{(1 - z^{-R})^N}{(1 - z^{-1})^N}, \quad (\text{A-40})$$

where $z = \exp(s/f_s)$ is the z -variable relative to the fast clock frequency, f_s . Mathematically, a n -stage CIC decimator is equivalent to a cascade of N running averages (boxcar filters) of length R . The frequency response of a CIC filter can be computed from (A-40) with $z \rightarrow \exp(2\pi i f/f_s)$. The magnitude response is

$$|G(f)| = \left| \frac{\sin(\pi R f/f_s)}{\sin(\pi f/f_s)} \right|^N. \quad (\text{A-41})$$

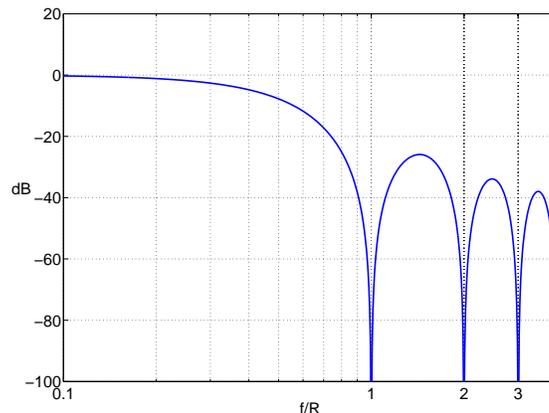


Figure A-10. Magnitude response of a generic two-stage CIC filter

The CIC magnitude response has nulls at $f = (k/R)f_s$, $k = 1, 2, 3 \dots$, which are at the centers of the aliasing bands for resampling at f_s/R . A Bode plot for a CIC filter is

shown in Figure A-10. The phase of the CIC filter is linear, with an equivalent group delay of $(R - 1)/f_s$ for $N = 2$, the case used in the hardware phasemeter.

A.4.6 Fractional-Delay Filters

The ability to interpolate between the data points of a sampled signal is a key capability for the LISA data analysis systems. In TDI, the phase-noise canceling data combinations require knowledge of the phasemeter signals at precise times. A ranging accuracy of 20 m in TDI corresponds to a timing accuracy of ~ 70 ns. Achieving this accuracy directly by sampling the signal at 15 MHz would be impractical, given the amount of data that would be generated. The alternative is to sample the data at a lower rate, around 10 Hz, and then interpolate between data points to produce the delayed versions of the phasemeter signals.

Fractional delay filters [37] are one technique used to interpolate between sampled data points. The basis for a fractional delay filter is the ideal low-pass filter in Figure A-7(a). As discussed in Section A.4.3.1, the corresponding impulse response for an ideal low-pass is the sinc function given in (A-35). For the case where the cutoff frequency is the Nyquist frequency, (A-35) simplifies to

$$h[n] = \text{sinc}(n) = \delta_{n0}, \quad (\text{A-42})$$

where δ_{ij} is the Kronecker delta function. A convolution with this impulse response function will result in an exact replica of the input signal. Fractional delay filters use a modification of (A-42) to produce a filter which does not effect the magnitude of a signal but produces a linear phase response which corresponds to a delay. For a delay of D samples, where $-0.5 \leq D \leq 0.5$, the fractional-delay filter kernel is given by

$$h[n] = \text{sinc}(n - D). \quad (\text{A-43})$$

Since fractional delay filters are FIR filters, the infinite series in (A-43) must be truncated to a finite length. This is generally accomplished with a window function. The truncation

and windowing will cause the filter to affect the magnitude of the input signal as well as the phase. These effects can be limited by using a sufficiently-long filter and employing a window function with a flat passband, such as the Blackman or Lagrange windows [37].

REFERENCES

- [1] B. Schutz, “Gravitational wave astronomy,” *Class. Quantum Grav.*, vol. 16, pp. A131–A156, 1999.
- [2] S. Hughes, “Listening to the universe with gravitational-wave astronomy,” *Annals of Physics*, vol. 303, 2003.
- [3] J. Taylor and J. Weisberg, “A new test of general relativity: gravitational radiation and the binary pulsar PSR 1913+16,” *The Astrophysical Journal*, vol. 253, pp. 908–920, February 1982.
- [4] B. Schutz, *A first course in general relativity*, Cambridge University Press, Cambridge, UK, 1985.
- [5] C. Misner, K. Thorne, and J. Wheeler, *Gravitation*, W.H. Freeman and Company, San Francisco, CA, 1973.
- [6] S. Shapiro and S. Teukolsky, *Black Holes, White Dwarfs, and Neutron Stars: The Physics of Compact Objects*, John Wiley and Sons, Toronto, Canada, 1983.
- [7] L. Lehner, “Numerical relativity: a review,” *Class. Quantum Grav.*, vol. 18, pp. R25–R86, 2001.
- [8] D. Jones, “Gravitational waves from rotating strained neutron stars,” *Classical Quantum Gravity*, vol. 19, pp. 1255–1265, April 2002.
- [9] C. Fryer, D. Holz, and S. Hughes, “Gravitational wave emission from core collapse of massive stars,” *The Astrophysical Journal*, vol. 565, pp. 430–446, January 2002.
- [10] P. Peters and J. Mathews, “Gravitational radiation from point masses in a Keplerian orbit,” *Physical Review*, vol. 131, no. 1, 1963.
- [11] J. Taylor and J. Weisberg, “Further experimental tests of relativistic gravity using the binary pulsar PSR 1913+16,” *The Astrophysical Journal*, vol. 345, pp. 434–450, October 1989.
- [12] I. Stairs, Z. Arzoumanian, F. Camilo, A. Lyne, D. Nice, and J. Taylor, “Measurement of relativistic orbital decay in the PSR B1354+12 binary system,” *The Astrophysical Journal*, vol. 505, pp. 352–357, September 1998.
- [13] M. Tinto, “Spacecraft Doppler tracking as a xylophone detector of gravitational radiation,” *Physical Review D*, vol. 53, no. 10, pp. 5354–5364, May 1996.
- [14] M. Tinto and J. Armstrong, “Spacecraft Doppler tracking as a narrow-band detector of gravitational radiation,” *Physical Review D*, vol. 58, 2002.
- [15] B. Bertotti, L. Iess, and P. Tortora, “A test of general relativity using radio links with the Cassini spacecraft,” *Nature*, vol. 425, pp. 374–376, 2003.

- [16] F. Jenet, A. Lommen, S. Larson, and L. Wen, “Constraining the properties of supermassive black hole systems using pulsar timing: an application to 3C 66B,” *The Astrophysical Journal*, vol. 606, pp. 799–803, May 2004.
- [17] F. Jenet, G. Hobbs, K. Lee, and R. Manchester, “Detecting the stochastic gravitational wave background using pulsar timing,” *The Astrophysical Journal*, vol. 625.
- [18] R. Manchester, “The Parkes pulsar timing array,” *Chinese Journal of Astronomy and Astrophysics*, astro-ph/0604288 in press.
- [19] J. Weber, “Detection and generation of gravitational waves,” *Physical Review*, vol. 117, no. 1, pp. 306, January 1960.
- [20] Mauceli et al., “The ALLEGRO GW detector,” *Physical Review D*, vol. 54, pp. 1264–1275, July 1996.
- [21] W. Johnston, “List of gravitational wave detectors,” <http://www.johnstonsarchive.net/relativity/gwdtable.html>, September 2003, Accessed October 2006.
- [22] S. Merkowitz and W. Johnson, “Spherical gravitational wave antennas and the truncated icosohedral arrangement,” *Physical Review D*, vol. 51, pp. 2546, 1995.
- [23] A. Abromovici et al., “The laser interferometer gravitational-wave observatory,” *Science*, vol. 256, pp. 325–333, April 1992.
- [24] A. Lazzarini and D. Shoemaker, “LIGO laboratory home page for interferometer sensitivities,” http://www.ligo.caltech.edu/~jzweizig/distribution/LSC_Data/, n.d., Accessed November 2006.
- [25] B. Abbot et al.(LSC), “Search for gravitational waves from binary black hole inspirals in LIGO data,” *Phys. Rev. D*, vol. 73, 2005.
- [26] B. Abbot et al.(LSC), “Search for gravitational waves from galactic and extra-galactic binary neutron stars,” *Phys. Rev. D*, vol. 73, 2005.
- [27] B. Abbot et al.(LSC), “Upper limits on gravitational wave bursts in LIGO’s second science run,” *Phys. Rev. D*, vol. 73, 2005.
- [28] P. Bender, K. Danzmann, and the LISA Study Team, “Laser interferometer space antenna for the detection of gravitational waves, pre-Phase A report,” Tech. Rep. MPQ233, Max-Planck-Institut für Quantenoptik, Garching, 1998, 2nd ed.
- [29] A. Stroerer and A. Vecchio, “The LISA verification binaries,” *Class. Quantum Grav.*, vol. 23, pp. S809–S818, 2006.
- [30] L. Carbone et al., “Characterization of disturbance sources for LISA: torsion pendulum results,” *Class. Quantum Grav.*, vol. 22, pp. S509–S519, 2005.

- [31] S. Anza et al., “The LTP experiment on the LISA pathfinder mission,” *Class. Quantum Grav.*, vol. 22, pp. S125–S138, 2005.
- [32] M. Tinto and J. Armstrong, “Cancellation of laser noise in an unequal arm interferometer detector of gravitational radiation,” *Physical Review D*, vol. 59, no. 102003, 1999.
- [33] S. Durandhar and K. Nayak, “Algebraic approach to time-delay data analysis in LISA,” *Physical Review D*, vol. 65, no. 102002, 2002.
- [34] M. Tinto, D. Shaddock, J. Sylvestre, and J. Armstrong, “Implementation of time-delay interferometry for LISA,” *Physical Review D*, vol. 67, no. 122003, pp. 1–17, 2003.
- [35] D. Shaddock, M. Tinto, F. Estabrook, and J. Armstrong, “Data combinations accounting for LISA spacecraft motion,” *Physical Review D*, vol. 68, no. 061303, 2003.
- [36] D. Shaddock, “Operating LISA as a Sagnac interferometer,” *Physical Review D*, vol. 69, no. 022001, 2004.
- [37] D. Shaddock, B. Ware, R. Spero, and M. Vallisneri, “Post-process time delay interferometry for LISA,” *Physical Review D*, vol. 70, no. 081101(R), 2004.
- [38] C. Salomon, D. Hils, and J. Hall, “Laser stabilization at the millihertz level,” *Journal of the Optical Society of America*, vol. 5, no. 8, pp. 1576–1587, February 1988.
- [39] B. Young, F. Cruz, W. Itano, and J. Bergquist, “Visible lasers with subhertz linewidths,” *Physical Review Letters*, vol. 82, no. 19, pp. 3799–3802, May 1999.
- [40] M. Eickhoff and J. Hall, “Optical standard at 532 nm,” *IEEE Transactions on instrumentation and measurement*, vol. 44, no. 2, pp. 155–158, 1995.
- [41] J. Hall, L. Ma, M. Taubman, B. Tiemann, F. Hong, O. Pfister, and J. Ye, “Stabilization and frequency measurement of the I_2 -stabilized Nd:YAG laser,” *IEEE Transactions on instrumentation and measurement*, vol. 48, no. 2, pp. 583–586, April 1999.
- [42] D. Shaddock, “Direct arm-locking,” LIMAS 2006-001 v1.1, May 2004.
- [43] B. Sheard, M. Gray, D. McClelland, and D. Shaddock, “Laser frequency stabilization by locking to a LISA arm,” *Phys. Lett. A*, vol. 320, no. 1, pp. 9–21, 2003.
- [44] J. Shearer, B. Kulakowski, and J. Gardner, *Dynamic Modeling and Control of Engineering Systems*, Prentice Hall, Inc., Upper Saddle River, NJ, 2nd edition, 1997.
- [45] M. Tinto and M. Rakhmanov, “On the laser frequency stabilization by locking to a LISA arm,” gr-qc/0408076, 2004.
- [46] G. Mueller, P. McNamara, J. Thorpe, and J. Camp, “Laser frequency stabilization for LISA,” Tech. Rep. NASA/TM-2005-212794, NASA, December 2005.

- [47] E. Black, “An introduction to Pound-Drever-Hall laser frequency stabilization,” *Am. Jour. Phys.*, vol. 69, no. 1, pp. 79–87, January 2001.
- [48] B. Sheard, M. Gray, D. Shaddock, and D. McClelland, “Laser frequency noise suppression by arm locking in LISA: progress towards a bench-top demonstration,” *Class. Quantum Grav.*, vol. 22, pp. S221–S226, 2005.
- [49] J. Thorpe, R. Cruz, S. Sankar, and G. Mueller, “Electronic phase delay — a first step towards a bench-top model of LISA,” *Class. Quantum Grav.*, vol. 22, pp. S227–S234, 2005.
- [50] M. Hueller, A. Cavalleri, R. Dolesi, S. Vitale, and W. Weber, “Torsion pendulum facility for ground testing of gravitational sensors for LISA,” *Class. Quantum Grav.*, vol. 19, pp. 1757–1765, 2002.
- [51] R. J. Cruz, *Development of the UF LISA Benchtop Simulator for Time Delay Interferometry*, Ph.D. dissertation, The University of Florida, May 2006.
- [52] A. Preston and et al., “Dimensional stability of hexology SA Silicon Carbide and Zerodur glass using hydroxide-catalysis bonding for optical systems in space,” in *Proceedings of SPIE*, E. Atad-Ettedgui, J. Antebi, and D. Lemke, Eds., 2006, vol. 6273 of *Optomechanical Technologies for Astronomy*, pp. 649–656.
- [53] G. Heinzel, A. Rudiger, and R. Schilling, “Spectrum and spectral density estimation by the discrete fourier transform (DFT), including a comprehensive list of window functions and some new flat-top windows,” AEI Internal Report, February 2002.
- [54] D. A. Shaddock, “Digital phasemeter using in-phase and quadrature sampling,” unpublished, August 2003.
- [55] A. Cruise, D. Hoyland, and S. Aston, “Implementation of the phasemeter for LISA LTP,” *Class. Quantum Grav.*, vol. 22, pp. S165–S169, 2005.
- [56] S. Pollack, O. Jennrich, R. Stebbins, and P. Bender, “Status of LISA phase measurement work in the US,” *Class. Quantum Grav.*, vol. 20, pp. S192–S199, 2003.
- [57] D. Shaddock, B. Ware, P. Halverson, R. Spero, and B. Klipstein, “An overview of the LISA phasemeter,” in *Proceedings of the Sixth International LISA Symposium*, in press.
- [58] B. Ware, W. M. Folkner, D. Shaddock, R. E. Spero, P. G. Halverson, I. Harris, and T. Rogstad, “Phase Measurement System for Inter-Spacecraft Laser Metrology,” in *Proceedings of the sixth annual NASA Earth Science Technology Conference*, June 2006.
- [59] M. R. Marcin, “Digital receiver phase meter for LISA,” *IEEE Transactions on Instrumentation and Measurement*, vol. 54, pp. 2466–2453, 2005.

- [60] “Direct Digital Synthesizer (DDS) v4.1,” *product specification* Xilinx, Inc., San Jose, CA, 2001.
- [61] J. Thorpe and G. Mueller, “Experimental verification of arm-locking for LISA using electronic phase delay,” *Physics Lett. A*, vol. 342, pp. 199–204, 2005.
- [62] J. Thorpe, R. Cruz, S. Sankar, G. Mueller, and P. McNamara, “First step toward a benchtop model of the laser interferometer space antenna,” *Optics Letters*, vol. 29, no. 24, pp. 2843–2845, December 2004.
- [63] R. Cruz and et al, “The LISA benchtop simulator at the University of Florida,” *Class. Quantum Grav.*, vol. 23, pp. S761–S767, 2006.
- [64] J. Thorpe, R. Cruz, M. Hartmann, and G. Mueller, “Arm-locking in a LISA-like hardware model: A status report,” in *Proceedings of the Sixth International LISA Symposium*.
- [65] J. Sylvestre, “Simulations of laser locking to a LISA arm,” gr-qc/0408076, 2004.
- [66] A. F. G. Marin, G. Heinzel, R. Schilling, V. Wand, F. G. Cervantes, F. Steier, O. Jennrich, A. Weidner, and K. Danzman, “Laser phase locking to a LISA arm: Experimental approach,” *Class. Quantum Grav.*, vol. 22, pp. S235–S242, 2005.
- [67] S. W. Smith, *The Scientist and Engineer’s Guide to Digital Signal Processing*, California Technical Publishing, San Diego, CA, 1997.
- [68] L. Jackson, *Digital Filters and Signal Processing*, Springer, 2 edition, 1988.
- [69] E. Hogenauer, “An economical class of digital filters for decimation and interpolation,” *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 29, no. 2, pp. 155–162, 1981.

BIOGRAPHICAL SKETCH

I was born in Santa Fe, New Mexico, on June 2nd, 1979. I spent my childhood and adolescence in Santa Fe, graduating from the Santa Fe Preparatory School in 1997. In the fall of 1997, I enrolled at Bucknell University in Lewisburg, PA as a mechanical engineering student. While taking my freshman and sophomore physics classes, I discovered I enjoyed physics as much as mechanical engineering and decided to pursue a double major in the two fields. In 2001, I graduated summa cum laude with a B.S. in mechanical engineering and a B.A. in physics. In the fall of 2001, I enrolled in the Ph.D. program at the physics department at the University of Maryland, College Park as a NASA Laboratory for High-Energy Astrophysics (LHEA) fellow. As part of the LHEA fellowship, I was required to work at the NASA Goddard Spaceflight Center (GSFC) during the summer of 2002. I began working with Guido Mueller at GSFC on laser stabilization for the Laser Interferometer Space Antenna (LISA). After the summer, I continued to work on LISA as a LHEA fellow, obtaining a M.S. in physics in December of 2002. In order to continue my work on LISA, I elected to transfer to the University of Florida in January of 2004, again working with Guido Mueller. Along with Guido and Rachel Cruz, I worked to build up the LISA lab in the physics department at UF. In August of 2004, I was awarded the University Alumni Fellowship, which has continued to support me for my tenure at UF. In December of 2005, I was awarded the Tom Scott Memorial Prize for best experimentalist graduate student by the physics department.