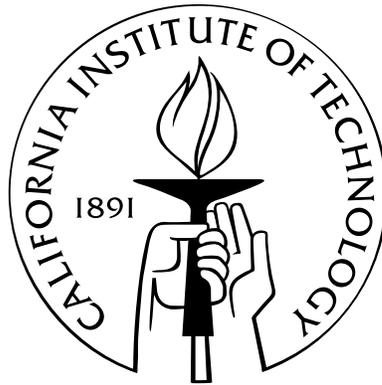


Topics in gravitational-wave astronomy

Thesis by
R. O'Shaughnessy

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



California Institute of Technology
Pasadena, California

2003
(Defended 29 July 2003)

© 2003

R. O'Shaughnessy
All Rights Reserved

Acknowledgements

Foremost, I would like to thank my advisor, Kip Thorne, for suggesting research problems, for assisting me with many of those problems, and most of all for his guidance on how to be a better scientist (e.g., using clearer prose and more effective communication). I am glad he spent the many long hours required to demonstrate to me how helpful a clear, cogent, and well-organized presentation can be in conveying the essential physics underlying an argument.

Aside from Kip Thorne, who deserves thanks for all my work, I must thank the many people who collaborated with and assisted me:

- *Compact-object inspiral*: Dan Kennefick deserves my heartfelt thanks for the kindness with which he provided invaluable numerical data.
- *Thermoelastic noise in LIGO*: My close collaborator, Erika d’Ambrosio, provided a wonderful sounding board and independent check on many of my calculations, as well as providing my main channel to the practicalities underlying my “theorist’s perspective” of LIGO. I also thank Sergey Vyatchanin (and his collaborator Sergey Strigin) for pleasant discussions and for his independent work on optical simulations and thermoelastic noise.

I am also happy to acknowledge the assistance of Peter Fritschel, Phil Willems, and GariLynn Billingsley, who all tried (not always successfully) to educate me about practical concerns and limitations in advanced LIGO designs. Bill Kells, Ray Beausoleil, and Bob Spero all provided very helpful comments on optics and interferometry. Finally, Lee Lindblom provided a key insight regarding my thermoelastic postprocessing code, allowing me to remove many kludges and obtain reliable answers.

- *Geometric-optics analysis of stability of Einstein’s equations*: Both Mark Scheel and Lee Lindblom provided critical assistance throughout the prolonged development of my geometric-optics analysis: Mark provided many analytic expressions (in Mathematica) and numerical tools with which I developed and tested my expressions; while Lee provided invaluable critical feedback on my argument, encouraging me to develop a more readable and comprehensive presentation.

I also thank the many participants in Caltech’s numerical relativity visitor’s program (too numerous to list here) for listening to and commenting on my early efforts on this subject.

During my time as a grad student, I had more interesting conversations than I could possibly remember; rather than miss someone, I thank everyone I met during my time here for making it an interesting and educational experience.

Finally, I would like to thank my family for all the kindness they've shown me through my time here, kindness that I've always appreciated but never acknowledged enough.

In addition, this work was supported in part by NSF grant PHY-0099568.

Abstract

Over the next decade, both ground-based (e.g., the Laser Interferometer Gravitational-Wave Observatory, or LIGO) and space-based (the Laser Interferometric Space Antenna, or LISA) gravitational-wave detectors should detect gravitational waves emitted by the motion of compact objects (e.g., black hole and neutron star binaries). These waves potentially contain useful information about the structure and dynamics of the compact-object systems that emit them. Since gravity-wave signals are inherently weak, any efforts to detect them naturally involve high-quality detectors and good models for expected signals. This thesis presents methods to improve (i) LIGO detector quality, (ii) our knowledge of waveforms for certain LIGO and LISA sources, and (iii) models for the rate of detectability of a particular LISA source. More specifically, this thesis studies:

1. *Plunge of a compact object into a supermassive black hole:* LISA is likely to detect many inspirals of compact objects (i.e., neutron stars or \sim few-stellar-mass black holes) into supermassive black holes ($\sim 10^5 - 10^7 M_\odot$, on the small end of what one expects to find in the center of a galaxy). Because these compact objects very slowly lose energy through their emitted gravitational waves, their motion is well-approximated at any instant by a stable geodesic orbit, and over long periods by a succession of stable geodesic orbits. Eventually, each inspiralling compact object will reach its last stable orbit, and will subsequently plunge rapidly into the hole. The location of this last stable orbit provides a sensitive probe of strong-field geometry near the supermassive black hole. Since this entire process (i.e., the transition from inspiral to plunge) takes place within the frequency band of greatest LISA sensitivity, LISA could conceivably observe this transition and thereby constrain the location of the last stable orbit and hence strong-field general relativity.

Previous computations by Ori and Thorne have suggested that, while the entire inspiral – by virtue of its long duration – can be easily seen by LISA, the transition from inspiral to plunge cannot (albeit just barely so). In Chapter 2, I perform a more generic computation – I estimate the chances that LISA could detect the transition from an eccentric equatorial inspiral to plunge – that comes to the same conclusion: LISA can expect a maximum signal-to-noise ratio of order ~ 1 from transition-to-plunge events. Therefore, the present LISA design cannot be expected to reliably measure the last stages of inspiral. However, a LISA design with a

slightly lower noise floor could potentially encounter a few events with strong enough signal to observe the transition.

2. *Scheme to reduce thermoelastic noise in advanced LIGO*: After its first upgrade, LIGO will use sapphire mirrors. Because sapphire expands considerably when heated, the upgraded LIGO detector (advanced-LIGO) will have high *thermoelastic noise*. [Thermoelastic noise occurs because millimeter-scale thermal fluctuations in the mirror bulk induce expansion and contraction, causing the mirror surface to shimmer.] The advanced-LIGO interferometer’s sensitivity could be significantly enhanced by merely reducing thermoelastic noise. In collaboration with Kip Thorne, Erika d’Ambrosio, Sergey Vyatchanin, and Sergey Strigin, I developed a proposal to reduce thermoelastic noise in advanced LIGO by switching LIGO cavity optics from simple spherical mirrors to a new, *Mexican-hat* shape. If advanced LIGO were redesigned to use these mirrors, it would have significantly greater effective range (i.e., for binary neutron star inspirals, an increase by a factor ~ 1.4) and would thus detect more inspiral events (for binary NS inspirals, a rate increase by a factor ~ 2.5).
3. *Geometric-optics-based analysis of stability of symmetric-hyperbolic formulations of Einstein’s equations*: The late stages of binary black hole inspiral and merger should produce some of the strongest and potentially most detectable signals for LIGO. But the waveforms emitted by the late stages of inspiral and merger remain poorly known; the simulations which would provide accurate predictions for these waveforms – full numerical evolutions of Einstein’s equations – cannot yet be successfully completed. More generally, no matter how they are performed – Einstein’s equations admit many possible representations – numerical simulations of sufficiently generic spacetimes always fail after a relatively short interval.

Different simulations fail for different reasons: each representation of Einstein’s equations presents its own difficulties. For example, when solving first-order symmetric hyperbolic formulations of Einstein’s evolution equations, errors naturally present in simulations accidentally excite ill-behaved *exact but unphysical solutions* to the equations we evolve (i.e., solutions which both grow exponentially and violate the Einstein energy and momentum constraints). Eventually, these solutions grow so large that the simulation fails.

Certain special types of these ill-behaved solutions (i.e., short-wavelength wave-packet solutions) are particularly easy to analyze; and by understanding their properties, I can (and did) make definite predictions about which simulations (based on first-order symmetric hyperbolic formulations of Einstein’s equations) will be particularly ill-behaved. [These predictions have not yet been systematically tested in full nonlinear simulations.]

Contents

Acknowledgements	iii
Abstract	v
1 Introduction	1
1.1 Late stages of the inspiral of compact objects into supermassive black holes	3
1.1.1 Last stable orbits as probe of strong-field gravity	3
1.1.2 LISA detection of the last stable orbit	4
1.1.3 Prior and present work: Circular and eccentric equatorial inspirals	4
1.2 Methods to reduce thermoelastic noise in advanced-LIGO designs	6
1.2.1 Review of thermoelastic noise	6
1.2.2 Using a larger beam to lower thermoelastic noise	7
1.2.3 Methods behind supporting computations	9
1.3 Stability analysis of various formulations of Einstein’s evolution equations, based on geometric optics techniques	10
1.3.1 Geometric optics approach to stability	11
1.4 Appendix: Writing Einstein’s equations in a manner suitable for numerical solution .	12
1.4.1 Basic Einstein’s equations	12
1.4.2 3+1 perspective	12
1.4.2.1 Gauge freedom	13
1.4.2.2 Constraint and evolution equations	13
Bibliography	15
2 Inspiral of compact objects into black holes	17
2.1 Introduction	17
2.1.1 Outline of this paper and summary of conclusions	19
2.2 Physical framework underlying the transition length estimate	20

2.2.1	Why we may still approximate the potential as static when computing the radial orbit	21
2.2.2	Adiabatic approach to estimating the duration of the transition from inspiral to plunge	22
2.2.3	Explicit conditions under which we may continue to approximate the potential as static	25
2.2.4	Inputs necessary for estimating the transition length	25
2.2.4.1	Estimating ΔI	26
2.2.4.2	Estimating $I_{\text{ad,min}}$	27
2.2.4.3	Choosing δr_{ref}	27
2.3	Estimating the probability for LISA to observe a transition	28
2.3.1	Estimating LISA's signal-to-noise for a given transition	29
2.3.2	Method for estimating the probability of detecting some transition during LISA's operation	31
2.3.3	Probability of detecting a transition during LISA's operation	32
2.4	Schwarzschild supermassive black hole (SMBH)	33
2.4.1	Choosing parameters	33
2.4.2	Dependence of transition parameters on eccentricity	34
2.4.3	Transition duration	35
2.4.4	Prospects for LISA detecting a given transition	36
2.5	Kerr SMBH	37
2.5.1	Parameterizing orbits, potential	37
2.5.2	Dependence of transition parameters on e,p	37
2.5.3	Transition duration	38
2.5.4	Prospects for LISA detecting a given transition	40
2.5.5	On probability of detection	40
2.6	Summary	41
	Bibliography	42
3	Reducing thermoelastic noise in advanced LIGO by flattening the light beams	44
3.1	Introduction and summary	45
3.1.1	The context: noise in advanced LIGO interferometers	45
3.1.2	The Physical Nature of Thermoelastic Noise; Motivation for Reshaping Beams	46
3.1.3	Summary of analysis and results	47
3.1.4	Notation	51
3.2	Mexican-mat mirrors and the mesa modes they support	53

3.2.1	Mesa fields	53
3.2.2	Gaussian fields	55
3.2.3	Diffraction losses	55
3.2.4	Mirrors and normalized flux for mesa and gaussian modes	56
3.3	Thermoelastic noise and neutron-star binary range for mesa-beam interferometers	58
3.3.1	Thermoelastic noise	58
3.3.1.1	Quantifying the thermoelastic noise: the noise integral and the NS/NS range	58
3.3.1.2	Baseline test masses with spherical mirrors and Gaussian beams	60
3.3.1.3	Cylindrical test masses with MH mirrors and mesa beams	60
3.3.1.4	Conical test masses	61
3.4	Sensitivity to mirror tilts, displacements and figure errors	61
3.4.1	Foundations for investigation	62
3.4.1.1	Our tools of analysis	62
3.4.1.2	Parameter values for comparing mesa and Gaussian configurations	63
3.4.1.3	Reflection and transmission coefficients	63
3.4.2	Driving a mesa-beam interferometer with a gaussian beam	63
3.4.3	Parasitic modes in arm cavities	64
3.4.4	Mirror tilt in arm cavities	64
3.4.4.1	Parasitic mode mixing in arm cavities	64
3.4.4.2	Diffraction losses	67
3.4.4.3	Arm cavity gain	67
3.4.4.4	Dark port power	67
3.4.5	Transverse displacement of arm cavities' mirrors	68
3.4.6	Errors in the arm cavities' mirror figures	69
3.4.6.1	Billingsley's worst-case figure error	69
3.4.6.2	Mode mixing by figure errors without compensating tilt	70
3.4.6.3	Mode mixing by figure errors with compensating tilt	72
3.4.7	Influence of mirror tilt and figure errors on thermoelastic noise	74
3.4.8	Mesa vs. Gaussian beams in recycling cavities	75
3.4.9	Increase in shot noise due to mirror tilts	77
3.4.9.1	Foundations	77
3.4.9.2	Carrier light in arm cavity	78
3.4.9.3	Signal light in arm cavity	79
3.4.9.4	Transmission of signal light through SR cavity	81
3.4.9.5	Reference light for baseline readout	82

3.4.9.6	Transmission of RF reference light through a power recycling cavity without an output mode cleaner	82
3.4.10	Increase in shot noise due to mirror figure errors	83
3.4.10.1	Carrier light in arm cavity	84
3.4.10.2	Signal light in arm cavity	84
3.4.10.3	Transmission of signal light through SR cavity	85
3.4.10.4	Reference light for baseline readout	86
3.4.10.5	Transmission of RF reference light through power recycling cavity without an output mode cleaner	86
3.5	Conclusions and Future Research	86
Bibliography		88
4	Calculations of thermoelastic noise and interferometer sensitivity to perturbation for the mesa-beam proposal for advanced LIGO	90
4.1	Introduction	91
4.1.1	Outline of paper	91
4.1.2	Connection with other published work	93
4.1.3	Notation	93
4.1.3.1	Symbols	93
4.1.3.2	Tensor notation	95
4.1.3.3	State and operator notation	95
4.2	Theory of an idealized Mexican-hat arm cavity	96
4.2.1	Mexican-hat mirrors and mesa beams	96
4.2.2	Thermoelastic noise	97
4.2.2.1	Effective static elasticity model for I_A	99
4.2.3	Propagation, the eigencondition, and diffraction losses	100
4.2.3.1	Principles of paraxial optics	101
4.2.3.2	Describing paraxial propagation through an arm cavity	102
4.2.3.3	Diffraction losses: Exactly	103
4.2.3.4	Diffraction losses: The clipping approximation	104
4.2.3.5	Diffraction losses of the ideal mesa beam, computed via the clipping approximation and corrected to estimate the true diffraction losses	104
4.2.4	Advanced LIGO design constraints	105
4.2.5	Summary: Exploring mesa-beam arm cavity configurations	106
4.3	Driving a mesa-beam arm cavity with a Gaussian beam	107
4.3.1	Beam and cavity used in this section	107

4.3.1.1	Fiducial mesa-beam cavity	107
4.3.1.2	Gaussian beam used to drive the fiducial arm cavity	108
4.3.2	Operation of fiducial cavity when driven by a Gaussian beam	108
4.3.2.1	Decomposing light along the cavity eigenstates	108
4.3.2.2	Amplification of light entering cavity	109
4.3.2.3	Output: Transmission of light out of cavity, and recombination with reflected light	109
4.4	Effect of mirror perturbations on arm cavities and the interferometer	110
4.4.1	Basis states for perturbation theory	110
4.4.2	Effect of perturbations on light in the arm cavities	111
4.4.2.1	Results of perturbation expansion	111
4.4.2.2	Estimating convergence of perturbation expansion	112
4.4.3	Implications of change in optical state for other quantities	113
4.4.3.1	Change in diffraction losses under perturbation of one mirror	113
4.4.3.2	Change in cavity gain due to perturbations of one mirror	113
4.4.3.3	Change in dark port power due to perturbations of one mirror	114
4.4.3.4	Influence of perturbations of one mirror on thermoelastic noise	115
4.4.4	Observations which simplify our computations	115
4.4.4.1	Symmetry and the influence of first-order changes in state	115
4.4.4.2	Assuming (roughly) equal contributions from the two mirrors to changes in I_A and diffraction losses	116
4.4.5	Summary: Using perturbation theory to explore the sensitivity of the arm cavity and interferometer to perturbations	116
4.5	Numerical implementations of the thermoelastic noise integral, the optical eigenequation, and optical perturbation theory	117
4.5.1	Thermoelastic noise for perfect (undeformed) mirrors	117
4.5.1.1	General approach	118
4.5.1.2	Special case: cylindrical mirror	118
4.5.1.3	Approximate technique: half-infinite mirrors	118
4.5.2	Numerically solving for the resonant optical eigenmodes of a cavity bounded by arbitrary axisymmetric mirrors	119
4.5.2.1	Setting up the problem to be solved; preliminary analytic simplifications	119
4.5.2.2	Method of numerical solution	120
4.5.2.3	Interpreting and applying the numerical solution	120
4.5.3	Numerical implementation of perturbation theory	121

4.5.4	Numerical exploration of changing diffraction losses	121
4.5.5	Numerical investigation of changes in thermoelastic noise when one mirror shape is perturbed	121
4.6	Dependence of thermoelastic noise on mirror and beam shape	122
4.6.1	Baseline advanced LIGO configuration	122
4.6.2	Improved baseline advanced LIGO configuration	123
4.6.3	Mesa beams reflecting off identical 40 kg cylindrical mirrors	123
4.6.4	Mesa beams reflecting off identical 40 kg frustum mirrors	124
4.7	Interferometer sensitivity to mirror perturbations	125
4.7.1	Frequency distribution of parasitic modes as a measure of arm cavity sensitiv- ity to perturbations	126
4.7.2	Effect of displacement on cavities bounded by spherical and Mexican-hat mirrors	128
4.7.2.1	Displacement of spherical mirrors	128
4.7.2.2	Displacement of Mexican-hat mirrors	129
4.7.2.3	Power in parasitic modes	129
4.7.2.4	Power out the dark port	129
4.7.3	Effect of tilt on the resonant eigenstate of an arm cavity bounded by spherical and Mexican-hat mirrors	130
4.7.3.1	Tilt-induced changes in the arm cavity ground state	130
4.7.3.2	Parasitic mode power excited by tilt	131
4.7.3.3	Tilt-induced changes in the diffraction losses of the ground state of the resonant arm cavity (Mexican-hat only)	131
4.7.4	Effect of tilt on arm cavities and interferometers using Mexican-hat mirrors but driven by Gaussian beams	132
4.7.4.1	Tilt and thermoelastic noise	132
4.7.4.2	Tilt and cavity gain, when driven by a Gaussian beam	133
4.7.4.3	Tilt and dark port power, when driven by a Gaussian beam	133
4.7.5	Effect of mirror figure error on mesa-beam interferometers	134
4.7.5.1	GLB's worst-case figure error	134
4.7.5.2	Effect of mirror figure errors, without compensating tilt	136
4.7.5.3	Mirror figure errors, with compensating tilt	137
4.7.5.4	Influence of mirror figure errors on thermoelastic noise	138
4.8	Conclusions	139

Bibliography

141

5	Numerical relativity	143
5.1	Introduction	143
5.1.1	Outline of remainder of paper	144
5.1.2	Connection with prior work	146
5.1.2.1	Study a short-time, rather than long-time, instability mechanism . .	146
5.1.2.2	Study an instability mechanism, not necessarily the dominant one .	146
5.1.2.3	Short-wavelength approximations	147
5.1.2.4	Energy norms	147
5.2	Ray optics limit of first-order symmetric hyperbolic systems	148
5.2.1	Definitions	148
5.2.2	Form of ray-optics solution	150
5.2.3	Interpreting the geometric-optics equations	150
5.2.3.1	Path of the ray	151
5.2.3.2	Propagating polarization along ray	151
5.2.4	When do geometric-optics solutions exist?	152
5.3	Defining wave packets	153
5.3.1	Constructing wave packets	153
5.3.2	Specialized wave packets I: Coherent wave packets	154
5.3.3	Specialized wave packets II: Prototypical coherent wave packets (PCWPs) . .	154
5.4	Describing and bounding the growth rate of wave packets	155
5.4.1	Energy norms and the magnitude of geometric-optics solutions	155
5.4.2	Energy norms and the growth rate of wave packets	156
5.4.3	Energy norms and the growth rate of coherent wave packets	157
5.4.4	Energy norms and the growth rate of PCWPs	157
5.5	Geometric optics limit of KST: Rindler	157
5.5.1	Generating the FOSHLS using the background Rindler space	158
5.5.2	Describing local plane waves by diagonalizing $A^a \hat{x}_a$	158
5.5.3	Deriving the polarization and energy equations, for propagation in the x direction on the light cone	159
5.5.3.1	Essential tool: Diagonalizing $P_s F P_s$	159
5.5.3.2	Polarization equation for general geometric-optics solutions	160
5.5.3.3	Energy equation for general geometric-optics solutions	161
5.5.3.4	Energy equation in a special case: PCWPs	162
5.5.4	Comparing growth rate expressions to simulations of prototypical coherent wave pulses	162
5.5.4.1	Specific simulations we ran	163

5.5.4.2	Results	163
5.6	Geometric optics limit of KST: PG	164
5.6.1	Generating the FOSHLS using a background Painleve-Gullstrand space	164
5.6.2	Local plane waves and diagonalizing $A^a \hat{r}_a$	165
5.6.3	Deriving the polarization and energy equations, for radial propagation on the light cone	165
5.6.3.1	Essential tool: Diagonalizing $P_s F P_s$ with $s = \pm 1$	166
5.6.3.2	Polarization equation for $s = \pm 1$	167
5.6.3.3	Energy equation for $s = \pm 1$	167
5.6.4	Deriving the polarization and energy equations, for radial propagation against the shift vector	168
5.6.4.1	Growth rate of PCWPs	169
5.6.4.2	Essential tool: Diagonalizing $P_o F P_o$	169
5.6.4.3	Aside: Why can't we follow the previous pattern?	170
5.7	Transients and limitations on numerical simulations: Rindler	171
5.7.1	Transients studied	171
5.7.2	Amplification expected	172
5.7.3	Relevance of our computation to numerical simulations	173
5.8	Transients and limitations on numerical simulations: PG	174
5.8.1	Transients studied	175
5.8.2	Amplification conditions	176
5.8.3	Results: Some KST parameters which have transients which amplify by 10^{32}	177
5.8.4	Generalizations of our method which could generate stronger constraints on KST parameters	177
5.9	Conclusions	179
	Bibliography	180
	A Appendix for compact object inspiral	181
A.1	Evolution of the maximum	181
A.2	Kerr Parameters and Constraints	182
A.2.1	Separatrix	183
A.2.2	Minimum of potential	183
A.3	Results from Ori and Thorne	183
	Bibliography	185

B Appendix for thermoelastic noise	186
B.1 Approximate formulae for mesa modes	186
B.2 Sapphire material parameters: Notation and values	188
B.3 General description of mesa beams	189
B.3.1 Constructing mesa beams	189
B.3.1.1 Canonical mesa beams: mesa beams for symmetric cavities	189
B.3.2 Canonical mesa beams at the mirror surfaces	190
B.4 Converting the fluctuation-dissipation model problem to a static model problem at low frequencies	190
B.4.1 General quasistatic approach for the elastic response to an oscillating surface stress	191
B.4.1.1 Step 1: Express the problem in the accelerated frame	191
B.4.1.2 Step 2: Factor out all sinusoidal dependence	192
B.4.1.3 Step 3: Approximate the problem as static in the accelerated frame	192
B.4.1.4 Summary	192
B.4.2 Using the quasistatic elastic solution to simplify our thermoelastic noise integral	193
B.5 On the completeness of basis states for an arm cavity with two identical, infinite mirrors	193
B.6 Second-order perturbation theory for eigenstates of an individual arm cavity	194
B.6.1 Preliminaries: Setting up notation for the expansion	194
B.6.2 Perturbation theory expansion, expressed using operators	195
B.6.3 Perturbation theory expansion, expressed using basis states	196
B.7 Perturbation theory for cavities bounded by two identical spherical mirrors	196
B.7.1 Background: Notation and definitions	197
B.7.1.1 Coordinate representation of eigenfunctions	197
B.7.1.2 Basis as eigensolutions of cavity	197
B.7.1.3 Establishing a quantum-mechanical correspondence	198
B.7.1.4 Creation and annihilation operators for optical states	198
B.7.2 Tabulating useful matrix elements	199
B.7.3 Effects of tilt on beam state	199
B.7.4 Effect of displacement on beam state	200
B.8 Postprocessing performed on finite-element computations used in thermoelastic noise integrals	200
B.8.1 Caltech postprocessing code	200
B.8.1.1 Postprocessing technique	201
B.8.1.2 Testing the result	202
B.9 Thermoelastic noise of half-infinite mirrors	202

B.9.1	Elastic solutions for the expansion (θ)	203
B.9.2	Thermoelastic integral I_A	203
B.10	Configurations with nonidentical frustum mirrors and mesa-like beams	203
B.10.1	Thermoelastic noise integral for ETM designs	204
B.10.2	Aside: Constructing mesa-like beams appropriate to nonidentical mirrors	206
Bibliography		207
C Appendix for numerical relativity		209
C.1	Useful identities used in the text	209
C.1.1	An alternative approach to the group velocity	209
C.1.2	The no-rotation condition	209
C.1.3	Reorganizing inner products for the polarization equation	210
C.2	Demonstrating that the ray-optics approach provides high-quality approximate solutions to the FOSHLS	211
C.2.1	Review: Writing equations in terms of ϕ and $d_{l,\alpha}$	212
C.2.2	Natural scales used in order-of-magnitude estimates	212
C.2.3	Degree to which ray-optics solution satisfies the FOSHLS	213
C.2.4	Length of time ray-optics solution can be trusted	214
C.3	When do PCWPs exist?	215
C.3.1	Rewriting polarization equation in the basis of eigenvectors of O_j	215
C.3.2	Sufficient conditions for PCWP to be exact solution	216
C.3.3	PCWP as limit of arbitrary rapidly growing coherent wave packet	216
C.4	Bounding the energy growth rate	216
C.5	Rays optics and KST 2-parameter formulation	217
C.5.1	Generally	217
C.5.1.1	Principal part and symmetrizer	217
C.5.1.2	Eigenvalues and group velocity	218
C.5.1.3	Eigenfields and projection operators	218
C.5.2	Special case: Flat spatial metric	219
C.5.2.1	Simplifying general polarization equation	219
C.5.2.2	Simplifying the energy equations	219

List of Figures

2.1	The effective potential $V(r, E, L)$ for radial geodesic motion, near the transition from inspiral to plunge.	21
2.2	The change ΔI of the radial potential's maximum during the last radial orbit.	26
2.3	Range of the number of angular cycles expected for a transition from inspiral to plunge into a Schwarzschild hole, with mass ratio 10^{-5}	36
2.4	A contour plot of the shortest possible length (in number of cycles) for a transition from inspiral to plunge into a Kerr black hole, for mass ratio 10^{-5}	38
2.5	A contour plot of the longest possible length (in number of cycles) for a transition from inspiral to plunge into a Kerr black hole, for mass ratio 10^{-5}	39
2.6	A contour plot of the fiducial "longest likely" length (in number of cycles) for a transition from inspiral to plunge into a Kerr black hole, for mass ratio 10^{-5}	39
2.7	A contour plot of S/N for LISA for a fiducial inspiral at 1Gpc.	40
3.1	Noise curves for advanced LIGO.	46
3.2	The power distributions of gaussian and mesa beams.	57
3.3	The shape of spherical and mexican-hat mirrors.	57
3.4	The range of the LIGO network for NS/NS inspiral, as a function of the thermoelastic noise power spectrum.	59
3.5	Contour diagram maps of the power in the first- and second-order parasitic modes excited by a tilt of the ETM of an MH cavity.	66
3.6	Contour diagram of the worst-case figure error.	70
3.7	The power distribution in parasitic modes when the ETM is deformed by a fraction 0.2 of the worst-case figure error.	71
3.8	A contour map of the worst-case figure error, reduced by a factor 0.2 and tilted by the optimal tilt angle.	73
3.9	The power distribution in parasitic modes when the ETM is deformed by a fraction 0.2 of the worst-case figure error and is tilted by the optimal tilt angle.	74
4.1	The power distributions of Gaussian and mesa beams.	97

4.2	The shape of spherical and mexican-hat mirrors.	98
4.3	A figure demonstrating notation for symmetric optical cavities.	98
4.4	A figure demonstrating a model problem which we used in conjunction with the fluctuation-dissipation theorem to calculate thermoelastic noise.	101
4.5	This figure illustrates our notation for propagation operators in paraxial optics (i.e. demonstrates the action of certain operators).	103
4.6	A contour plot of the clipping-approximation diffraction losses of the mesa beam, as a function of mirror radius and beam radius parameter D	105
4.7	The worst-case figure error.	135
5.1	The growth rates for certain wave packets in Rindler space, computed (i) from full numerical simulations of Einstein's equations and (ii) from the expressions in this paper.	164
5.2	The radially-propagating characteristics in the KST formulation of Einstein's equations, linearized about Painleve-Gullstrand coordinates.	166
5.3	The future domain of dependence of the region $0.01 \leq x \leq 1$ for the KST 2-parameter formulation of Einstein's equations, linearized about Rindler space.	172
5.4	A plot of the most rapid growth rate of prototypical coherent wave packets in Rindler space.	173
5.5	The future domain of dependence of the region $2 \leq r \leq 10$ for the KST 2-parameter formulation of Einstein's equations, linearized about Painleve-Gullstrand coordinates.	175
5.6	A plot showing which values of the two KST parameters are excluded by the mechanism presented in this paper.	178
B.1	Comparison of the moduli of the exact flat-topped mode $U(4, r)$, the zero-order approximation $U_0(4, r)$ [Eq. B.4], the third-order approximation $U_3(4, r)$ [Eq. (B.5)], and the asymptotic approximation to the zero-order approximation, $U_{0a}(4, r)$ [Eq. (B.7)]. Top: linear plot; bottom: logarithmic plot.	187

List of Tables

3.1	Optimized light-beam configurations, their thermoelastic noise compared to the baseline, their neutron-star binary range, and their event rate for NS/NS inspiral divided by the baseline rate.	61
4.1	The thermoelastic noise integral I for a cylindrical test mass and a Gaussian beam, relative to the baseline advanced LIGO configuration.	123
4.2	The thermoelastic integral I for a cylindrical test mass and a mesa beam, relative to the value for the baseline advanced LIGO configuration.	124
4.3	The thermoelastic integral I for a frustum input test mass (ITM) and a Mexican-tat Beam, relative to the value for the baseline advanced LIGO configuration.	125
4.4	Optimal thermoelastic noise integrals for arm cavity configurations involving (i) identical cylinders and (ii) identical frustums.	126
4.5	The distribution of low-lying parasitic modes in an arm cavity bounded by Mexican hat mirrors.	128
B.1	The thermoelastic integral I for a frustum end test mass (ETM) and a Mexican-Hat Beam, in units of $I_{\text{BL}} = 2.57 \times 10^{-28} \text{s}^4 \text{g}^{-2} \text{cm}^{-1}$. The values of I/I_{BL} are estimated to be accurate to within one per cent.	205
B.2	Optimized test-mass and light beam configurations, their thermoelastic noise compared to the baseline.	206

Chapter 1

Introduction

Classical astronomy is a mature field that uses the electromagnetic spectrum to explore the universe. But by using gravitational rather than electromagnetic waves—by using waves produced only by the most energetic events in the universe; waves that, unlike their electromagnetic counterparts, cannot easily be obscured by dust or other intervening low-density matter—the budding field of gravitational wave astronomy will provide an entirely new view of the universe [1, 2, 3].

Because gravitational-wave signals will be very weak, both sensitive detectors (to maximize the potential for detection) and knowledge of the most likely signals (to permit detection of the signal in the presence of detector noise) are needed. Thus, gravitational-wave astronomy has consisted of efforts in four closely related areas:

1. *Source statistics and estimates of detectability:* Each class of source has been (often roughly) assessed, to determine whether its members are worth looking for (and potentially to simplify the process of finding them). Estimates of the frequency and distribution of signal strengths of each likely source (if they can be obtained) have therefore been developed.¹
2. *Detector design, construction, and commissioning:* Detectors sensitive to likely sources—such as the ground-based Laser Interferometer Gravitational-Wave Observatory (LIGO)—have been designed and built, and are presently being commissioned.² Further, future detectors—such as planned upgrades to the LIGO interferometer, along with the Laser Interferometer Space Antenna (LISA)—are being planned, to further increase the overall sensitivity and bandwidth of the present gravitational-wave-detection network.
3. *Source simulation, waveform modelling, and template extraction:* Each source produces some associated waveform, which must be modelled (more or less accurately) to provide a reference

¹For example, these estimates can help guide the design of both future detectors and signal-detection algorithms.

²The LIGO interferometer is an extremely complicated combination of interlocking systems; problems with any subsystem can easily cause problems with the whole interferometer. During the “commissioning” phase, the interferometer’s components are successively installed and integrated, and these components and the entire interferometer are systematically tested, until the interferometer finally reaches design specifications (or some acceptable approximation thereof).

against which data can be compared. Ideally, either accurate waveforms are unnecessary (e.g., for burst sources, whose waveforms are strong short pulses whose internal structure need not be modelled) or are fairly well-determined (i.e., for equatorial inspiral of a compact object into a supermassive hole). When more accurate waveforms are needed than are available, the waves are modelled by (for example) some plausible parameterized functions (e.g., [4, 5]).

4. *Detection and parameter estimation*: Finally, statistical analyses can be developed to determine whether signals are present in the detector’s data and what sources these signals most likely come from, based on knowledge of what signals are likely, of the the detector sensitivity, and of the waveforms associated with various sources. In practice, the optimal statistical analyses may be difficult or impossible to implement (i.e., because of the large number of templates needed, as compared to the computing power available, or because the waveforms needed for optimal searching simply aren’t well-enough known); suboptimal alternatives are being developed to handle these cases.

This thesis surveys three topics, drawn from each of the first three categories of research:

- *Compact inspirals into supermassive black holes* (source statistics and detectability; some source modelling): If data analysis challenges are addressed, each year LISA should potentially detect several inspirals (possibly more) of compact objects (i.e., a neutron star or \sim few stellar-mass black hole) into supermassive ($M \sim 10^5 - 10^7 M_\odot$) black holes. The last few waves from each inspiral tell (e.g., through their frequency) where the compact object made its last few orbits, so these waves provide a simple probe of strong field gravity. Section 1.1 and Chapter 2 describe whether LISA could be reasonably expected to detect these last few orbits, when the compact object’s inspiral is confined to the equatorial plane of the capturing hole.
- *Thermoelastic noise in advanced LIGO* (detector design): If the current plans for the next-generation LIGO interferometer (advanced LIGO) are used (and various materials can indeed be produced at material specifications), thermoelastic effects (i.e., the effects of stochastic temperature fluctuations, and their associated expansions/contractions) will dominate the noise in the detector. Section 1.2 and Chapters 3 and 4 present a scheme, relying primarily on modified optics, to reduce the effect of this noise and increase the sensitivity and range (both by a factor ~ 1.5 in amplitude) of the interferometer.
- *Growth of undesirable analytic solutions in numerical simulations of Einstein’s equations* (source simulation): Numerical simulations of Einstein’s equations are necessary to obtain waveforms for binary black hole inspiral which are accurate enough to be used for signal detection. For the purposes of numerical simulation, one can devise many formulations of Einstein’s evolution

equations. As these formulations differ in their treatment of functions which are *not* solutions of Einstein’s equations (e.g., which violate the energy and momentum constraints), they possess different stability properties. Section 1.3 and Chapter 5 present a novel analytic technique to analyze the stability properties of various formulations of Einstein’s evolution equations.

1.1 Late stages of the inspiral of compact objects into supermassive black holes

One of the many goals of gravitational-wave astronomy is to test strong-field general relativity in the neighborhood of a black hole. Specifically, while we expect the weak- and distant-field geometry of a black hole spacetime to closely correspond to the predictions of general relativity (as weak-field relativity has been extensively tested [6]), closer to the black hole the stronger fields involved could possibly unmask any potential differences between general relativity and the true theory of gravity.

1.1.1 Last stable orbits as probe of strong-field gravity

Perhaps the simplest (but not the only³) measure of strong-field gravity in a black-hole spacetime are the *last stable orbits* in that spacetime [9]. The last stable orbits are those orbits (i.e., geodesics) which lie at the boundary between i) bound, stable motion about the black hole and ii) capture by the hole.⁴

This surface plays a critical role when a compact object (i.e., a neutron star or stellar-mass black hole) orbits a supermassive black hole ($M \sim 10^5 - 10^7 M_\odot$). Because gravitational radiation is weak, over any short period the compact object’s orbit may be well-described by a geodesic. Over longer periods, gravitational radiation causes slow changes in the conserved constants parameterizing the geodesic; effectively, the orbit slowly move through a sequence of instantaneously geodesic orbits.⁵ Eventually, this radiation-reaction-induced flow through the space of stable orbits encounters the surface of last stable orbits; or, in other words, eventually under the action of gravitational radiation the orbit draws so close to the black hole that the object’s motion can no longer protect it from the pull of gravity. In short, the particle plunges into the black hole. Therefore, the waves emitted during the inspiral abruptly cease.⁶ The abrupt termination of the wavetrain can be used to determine the

³Fintan Ryan [7, 8] has performed a suggestive computation, arguing that measurements (by LISA) of gravitational waves emitted from a compact object (i.e., a neutron star or stellar-mass black hole) orbiting a supermassive black hole ($M \sim 10^5 - 10^7 M_\odot$) can be used to partially reconstruct some of the multipole moments of the supermassive black hole’s spacetime geometry. [Because the LISA detector has noise, the multipole moments are not known perfectly, and higher-order moments are harder to constrain.]

⁴The space of geodesics inherits a natural topology from the six-dimensional space that parameterizes it (four initial values for the coordinate location and four values of the coordinate momentum, modulo the conservation-of-rest-mass constraint, and modulo equality of two geodesics that overlap entirely).

⁵Sam Finn, Scott Hughes, Dan Kennefick, and others have implemented codes to model these compact-object inspirals using precisely this technique.

⁶To take a concrete example, consider a particle in the late stages of inspiral into a Kerr black hole. At some well-

radius of the last stable orbit, a simple probe of strong-field physics.⁷

1.1.2 LISA detection of the last stable orbit

Inspirals of compact objects into supermassive black holes are common enough [10, 11], and emit strong enough waves [12], that they can potentially be seen (both high signal to noise and frequency greater than \sim few/year) by the planned detector sensitive to their waves (the Laser Interferometer Space Antenna, or LISA) [13].⁸ Unfortunately (for our purposes), while these inspirals produce a high S/N , much of the signal is accumulated while the compact object is relatively far from the last stable orbit [12]. Estimates of the strength of these last few waves and of the overall rate of inspirals must be combined to determine whether LISA can measure the waves emitted during the last stable orbit from *any* inspiral.

1.1.3 Prior and present work: Circular and eccentric equatorial inspirals

Ori and Thorne [14] performed the first computations to estimate, for waves emitted from a compact object orbiting a supermassive black hole, how effectively LISA could detect the transition from inspiral to plunge. For simplicity, they limited attention to circular equatorial inspirals. For those limited inspirals, they found a disappointing result: for plausible sources ($10M_{\odot}$ orbiting 10^6M_{\odot}) LISA could expect only $S/N \sim 1$ —not even enough for reliable detection, let alone parameter estimation.

But since Ori and Thorne limited attention only to circular equatorial inspiral, possibilities still remained that some more generic (yet still physically plausible) combination of orbital parameters could produce a longer or stronger signal (i.e., larger S/N). For example, *eccentric* equatorial inspirals offered the potential for significantly longer transition times (because the transition involved an unstable equilibrium of the effective one-dimensional radial potential).⁹

My own work, presented in Chapter 2, generalized the Ori-Thorne approach to include the case of eccentric equatorial inspiral into a Kerr black hole. As with Ori and Thorne, the analysis was based upon how an effective one-dimensional equation for the radial motion¹⁰ [i.e., an equation for

defined Boyer-Lindquist radius (which depends on the inclination and eccentricity of the orbit)—the radial location of the transition from inspiral to plunge is exceedingly well-determined, with errors that go as an inverse power of the mass ratio.—the particle ceases to orbit the hole in a stable fashion and plunges rapidly into the hole.

⁷For example, for circular equatorial orbits, the wave are predominantly quadrupole, emitted at twice the orbital frequency of the particle’s orbital motion. Therefore, when the waves terminate, the frequency of the last few waves seen provides a simple and direct measure of the orbital frequency of the last stable orbit.

⁸Unfortunately, the signals emitted by compact-object inspirals are very complicated, making data analysis difficult. Thus, while *potentially* LISA could detect many inspirals, practical data analysis issues may significantly reduce the effectiveness of LISA in detecting those inspirals. Indeed, the LISA design may be modified (i.e., lower noise floor) specifically to compensate for limitations in our ability to optimally detect these signals with matched filtering.

⁹As will be discussed in greater detail in Chapter 2, for an eccentric orbit near the transition from inspiral to plunge, the transition occurs about a turning point of the one-dimensional radial effective potential $V(r)$. The particle naturally stays a very long time near that turning point during the transition.

¹⁰As noted in standard references [24], the geodesic equation for test-particle motion in a Kerr black hole separates into equations for the radial and angular motions. The radial equation depends on the conserved constants E and L .

$r(t)$] changed with time as the conserved constants (energy and orbital angular momentum of the compact object) gradually changed. Using this potential, I . . .

1. *Developed models for the last stable orbits:* Because the effective potential varies slowly under the action of gravitational radiation, I realized that, for all except nearly circular transitions from inspiral to plunge, the geodesic orbit is a “zoom-whirl” orbit. In a zoom-whirl orbit, the particle spends much of its orbit circling (“whirling”) many times around the central black hole; the object then “zooms” out to its outer turning point and back. In the case of the transition, the particle comes in, whirls several times, and then plunges rapidly into the hole.
2. *Predicted the possible range of durations of the transition:* The number of times the particle “whirls” around the hole is a simple measure of how long the transition lasts. Unfortunately, the number of whirls a particular particle will perform depends sensitively on initial conditions. I found a simple formula to express the range (and even probability distribution) of possible transition durations. For example, for inspiral of a $10M_{\odot}$ hole into a 10^6M_{\odot} Schwarzschild hole, I found that the transition could last from ~ 5 to (in very unlikely cases) ~ 10 orbital cycles.
3. *Estimated the signal to noise from transition waves seen by LISA:* Using a simple model to estimate how effectively LISA might detect the simple sinusoidal waves emitted during the transition, I determined how strong a signal LISA could see from a characteristic source. For a $10M_{\odot}$ source at 1Gpc inspiralling with eccentricity $e = 1/3$, I expect most inspirals should have signal-to-noise ratio between 0.9 and 1.
4. *Estimated the rate at which strong signals could be seen:* Finally, using simple statistics, I tried to estimate the strongest signal LISA could plausibly see. Because I had to use relatively poorly-known statistics for the capture rate of compact objects into supermassive holes, I could not make definitive statements. However, even with relatively optimistic assumptions about the number of compact-object captures [11], I found LISA had only a 50% chance of detecting one strong signal (i.e., with $S/N > 4$); for more realistic parameters, LISA would have a 50% chance of detecting no signal with strength $S/N > 2.3$.

In short, my conclusions remain the same as Ori and Thorne’s: unless significantly more black hole inspirals occur than I assumed, LISA stands little chance of even detecting waves emitted during the transition from inspiral to plunge, let alone of using those transition waves to significantly constrain the innermost stable orbit and (more generally) strong-field gravity.¹¹

It is this equation we solve.

¹¹The waves emitted during the transition are merely the *simplest* measure of strong-field gravity; they are not the only measure. For example, Kip Thorne (private communication) suspects that the location of the last stable orbit can be accurately (to \sim one percent) determined using waves emitted before the transition.

1.2 Methods to reduce thermoelastic noise in advanced-LIGO designs

According to current plans for the first LIGO upgrade (i.e., advanced LIGO, or LIGO-II) [15, 16], the interferometers will substantially increase in sensitivity. If all goes as currently planned (e.g., if sapphire test-mass mirrors are used and if the mirror coatings can be created with sufficiently low mechanical losses), the sensitivity of the interferometer will be limited by *thermoelastic noise*. The next largest expected sources of noise, radiation pressure noise and shot noise, will produce a significantly lower amount of noise. Therefore, the sensitivity of the advanced-LIGO design could be easily increased by merely reducing thermoelastic noise.

1.2.1 Review of thermoelastic noise

Thermoelastic noise is one of many forms of noise that enters the LIGO output through motions of the test-mass mirrors. Physically, gravitational waves cause the centers-of-mass of the four LIGO test-mass mirrors to move; therefore, ideally, the LIGO output should be related to measurements of the mirror center-of-masses. Practically, however, the LIGO interferometer output is directly related not to the distances between mirror center-of-masses, but rather between mirror *surfaces*, where the effective mirror surface location x_{eff} is determined in terms of the true mirror surface location $x_{\text{mirr}}(r)$ by an average, weighted by the power in the beam:

$$x_{\text{eff}} = \left[\int d(\text{area}) x_{\text{mirr}}(r) I(r) \right] / P ; \quad (1.1)$$

here $I(r)$ is the beam intensity at radius r and where $P = \int I d(\text{area})$ is the beam power. When the mirrors move and deform, the effective location x_{eff} changes, in a manner not necessarily correlated to the center-of-mass mirror location or to any passing gravitational wave. In particular, stochastic motions of the mirror surfaces produce noise in the LIGO output.

Many different processes cause the mirror surfaces to move and deform randomly, with different names given to the noise produced by different processes.¹² For thermoelastic noise, the relevant mirror deformations are produced by elastic deformations in the bulk of the mirror, deformations which in turn arise due to temperature fluctuations in various small regions in the bulk.

To be very explicit, at any given instant thermal fluctuations, via their bulk expansion and the elastic properties of the mirror, produce many small deformations on the surface of the mirror. Loosely, these deformations are very small bumps (or dips) in the surface of the mirror, with bump

¹²For example, *seismic noise* arises because of assorted motions of the ground near the LIGO site, motions which cause the LIGO site (and hence the mirror suspensions and eventually the mirrors themselves) to move. Also, *brownian (thermal) noise* is the name given to noise associated with oscillations in the low-frequency vibrational eigenmodes of the mirror. These modes are thermally excited and evolve stochastically.

width of order the diffusion length of heat inside the mirror over a gravity wave timescale, or $\sim 0.3\text{mm}$.¹³ At each instant, the LIGO laser approximates the true surface location via Eq. (1.1); because bumps and dips are equally likely, most bumps and dips cancel (i.e., they average out). But because the beam has a finite extent and because therefore Eq. (1.1) involves an average over a large but finite number of possible bumps and dips, the bumps and dips do not perfectly cancel: the mirror location estimate x_{eff} retains some small error, related to whether, at the given instant, more bumps or more dips happened to lie within the beam cross-section.¹⁴ Therefore, since the bumps and dips change stochastically, so too does LIGO's estimate of the mirror surface location. This random process is the source of thermoelastic noise.

1.2.2 Using a larger beam to lower thermoelastic noise

To reduce thermoelastic noise without significant complications (i.e., cooling the mirrors or changing the mirror substrate), one must improve the average used in Eq. (1.1) so more bumps and dips tend to cancel. But the framework of conventional optics (i.e., spherical mirrors) and conventional mirror shapes (i.e., cylinders) does not allow for substantial reductions in thermoelastic noise. To improve the average one necessarily must broaden the beam (i.e., increase the width of the gaussian beam). In order to significantly broaden the beam without losing significant amounts of power off the mirror's edges, the cylindrical mirror's radius must be increased. [For thermoelastic noise, a bigger mirror is always better than a smaller mirror of the same proportions.] But unfortunately technical problems associated with the manufacture of large sapphire mirrors in effect limit the radius of a sapphire cylinder one can buy; also, practical limits of the advanced LIGO suspension limit the mass of the mirror advanced LIGO can use. Therefore, we cannot substantially improve upon the baseline advanced LIGO design by simply requiring larger beams and mirrors.

Indeed, as one would expect, the current advanced-LIGO design is already a near-optimal choice, given the limitations of conventional mirrors and optics. Therefore, as groups at Caltech and Moscow independently realized, to obtain lower thermoelastic noise without cooling the mirrors, LIGO must use unconventional mirrors and optics:

1. *Mirror reshaping:* At Moscow, Sergey Vyatchanin and Sergey Strigin proposed studying mirror shapes more generic than simple cylinders (e.g., frustum).¹⁵ With a broader class of mirrors to explore, they could potentially find certain combinations which could both be manufactured and possess a larger front face size than the baseline cylinder. Because the mirror front face size

¹³This number should be taken only as a rough guide to the scale of the relevant thermoelastic deformations. In fact, just as thermal fluctuations occur over all time and length scales, so the associated thermoelastic deformations exist over a *range* of scales.

¹⁴As described in greater detail in Chapter 3 and below, the error in the mirror location estimate also depends on the beam *shape*. A flatter beam performs a more equitable (and therefore better) average.

¹⁵A frustum is the shape resulting when an axisymmetric cone is cut along two planes perpendicular to the cone axis.

is larger, the thermoelastically produced bumps and dips should cancel out more completely, thus lowering thermoelastic noise.

2. *Nonstandard optics*: At Caltech, Kip Thorne and I realized that a LIGO laser with a flatter intensity profile should produce a more equitable average in Eq. (1.1) and thus lower thermoelastic noise. Moreover, Kip realized how to explicitly construct a simple form for a flat-topped beam (*mesa beam*) and for mirrors (*Mexican-hat mirrors*) that reflected that beam back into itself.

Since these two approaches seemed to offer independent methods to reduce thermoelastic noise, these two collaborations—augmented at Caltech by Erika d’Ambrosio—joined together, to discover how effective these two ideas would be when combined. But after consulting with GariLynn Billingsley—who told us that, because of mirror fabrication issues, any large sapphire block we could cut to make a frustum could just as well be cut to create a larger cylinder—the collaboration focused its attention primarily on the Caltech *mesa beam* proposal.

Chapter 3, adapted from a paper composed by Kip Thorne, summarizes the principal results of our collaboration:

- *Configurations with lower thermoelastic noise exist*: If mesa beams are used with otherwise unchanged cylindrical sapphire advanced-LIGO test-mass mirrors, the thermoelastic noise power will be lower, by a factor 0.34, than the corresponding noise produced with conventional optics.
- *Recycling cavities are insensitive to the arm cavity mirror shape*: Because the power and signal recycling cavities are short and have low finesse, just about anything can resonate inside them: they are highly insensitive to the light used. (Thus, an interferometer using mesa beams can operate with conventional advanced LIGO recycling mirrors; no redesign is required.)
- *Each arm cavity, and the overall interferometer, is not overly sensitive to small tilts and displacements*: Even though the Mexican-hat mirrors needed in a mesa-beam arm cavity are much flatter in the center than their conventional counterparts, each arm cavity is only marginally more sensitive to tilt and displacement than their gaussian counterparts (as measured by the power going out the dark port of an interferometer with one mirror perturbed, and by the change in the power spectrum of noise as a function of tilt and displacement).¹⁶
- *Each arm cavity, and the overall interferometer, is not overly sensitive to mirror figure error*:

We determined how accurately the Mexican-hat mirrors need to be fabricated.¹⁷ For example,

¹⁶Since mesa-beam arm cavities are roughly as sensitive to displacement and tilt as the planned advanced-LIGO arm cavities, roughly the same feedback control servo system planned for advanced-LIGO (albeit with the feedback control loops adapted for the appropriate mesa-beam-induced weights) can be used to control mirror displacement and tilt.

¹⁷Since the Mexican-hat mirrors are very flat in the center (and of an unfamiliar shape), they pose a potential machining challenge: Mexican-hat mirrors might have greater surface figure errors than a corresponding conventional mirror.

to keep changes due to mirror figure error in the power spectrum of noise below 1%, the peak-to-valley variations Δz in the mirror shapes must be less than ~ 2.0 nm (for a wideband interferometer).

1.2.3 Methods behind supporting computations

Where Chapter 3 provides the overall results of the entire collaboration, in particular providing a broad survey of those results of practical significance to LIGO design, in Chapter 4 I describe the three techniques I used to obtain many of the results quoted in Chapter 3:¹⁸

- *Numerical solution for optical modes in the cavity:* I formulated a (standard) integral eigenequation for the modes of an optical cavity bounded by arbitrarily shaped (but finite and axisymmetric) mirrors. I then wrote numerical code to find the optical eigenmodes of these cavities.
- *Optical perturbation theory:* I wrote out (and designed code to rapidly evaluate) second-order optical perturbation theory for the cavity.
- *Numerical approach to thermoelastic noise:* I constructed both semianalytic (series) and fully-numerical (finite-element) elasticity models to determine the thermoelastic noise associated with a given beam shape.

As described in detail in Sections 4.6 and 4.7, I combined these tools to obtain almost all the results presented in Chapter 3 (with the exception of Sections 3.4.8 and 3.4.9). For example, to describe how the contribution to thermoelastic noise from an individual mirror increased due to mirror figure errors in another, I (i) used a basis of numerically tabulated eigensolutions to (ii) construct an explicit form for optical perturbation theory; applying that expansion, I (iii) deduced how the ground state of the cavity changed due to mirror defects; and then I (iv) evaluated the thermoelastic noise associated with the deformed beam configuration.

Erika d’Ambrosio also independently corroborated many of our results, using a standard code (the “FFT code”)¹⁹ designed to simulate the full optical properties of the whole LIGO interferometer [25].

¹⁸In our collaboration, we tried to provide several independent checks of all key results. While I wrote Chapter 3 and performed all the computations mentioned in it, Strigin and Sergey Vyatchanin are coauthors because they corroborated many of my results, using similar (but independent) methods.

¹⁹The code simulates the optical fields (on a grid) at key locations of the interferometer. The name follows from the Fourier transforms used to accelerate computations.

1.3 Stability analysis of various formulations of Einstein’s evolution equations, based on geometric optics techniques

In order to provide models for waveforms produced by the late stages of comparable-mass binary black-hole inspiral (e.g., two stellar mass black holes) sufficiently accurate to be used as templates for LIGO detection waveforms, we must solve Einstein’s equations numerically.

Usually, Einstein’s equations are solved by rewriting them in a “3+1” form (i.e., a space and time decomposition), with 6 evolution equations and 4 constraint equations (the energy and momentum constraints) for the spatial metric g_{ab} and extrinsic curvature K_{ab} (both 3×3 tensors). [Section 1.4 of this introduction, an appendix, provides a brief review of the ADM 3+1 decomposition designed for the nonspecialist.] To find solutions to Einstein’s equations, one selects initial data for K_{ab} and g_{ab} which satisfy the constraints; one then uses only six of the equations and constraints to deduce the behavior of the metric and extrinsic curvature at later times. If the equations are solved without error, the remaining four equations must hold automatically. [For example, in *unconstrained evolution* one uses only the six evolution equations to evolve the system forward in time; the four constraints hold automatically.]

Unfortunately, numerical simulations of realistic black hole binaries—and even of simple static black hole spacetimes²⁰—fail dramatically after only a comparatively short time. Moreover, the problems presently limiting simulations do not arise merely from numerical problems (e.g., poor code or lack of resolution) or insufficient computer time. Rather, considerable evidence suggests that fundamental features of the *continuum equations and boundary conditions* lead to most observed problems with simulations. For example, various authors have discovered that certain ways of writing Einstein’s equations (and various ways of adding boundary conditions to Einstein’s equations) are ill-posed (e.g., [17, 18]).²¹ In addition, Kidder, Scheel, and Teukolsky found evidence suggesting that, when they evolved various well-posed forms of the 3+1 equations, their simulations were always limited by the presence of constraint-violating solutions which were excited by small (i.e., roundoff-level) errors in their initial data [19, 20].

Since the equations themselves largely determine the stability properties of numerical simulations that employ them, many different, continuously parameterized formulations of Einstein’s evolution equations have been introduced [19, 21], in part motivated by the hope that—somewhere among the many formulations—better-behaved formulations might be found. And somewhere among the vast array of presently proposed formulations a well-behaved formulation might well exist; the problem is

²⁰If symmetries are exploited, such as $O(3)$ symmetry or even quadrant symmetry for a static Schwarzschild black hole, then most serious, debugged numerical relativity codes run forever. However, full 3-d simulations run into problems.

²¹Technically, ill-posedness means that the growth rate of errors cannot be bounded by a constant, independent of the solution being simulated. In practice, ill-posedness implies errors associated with shorter scales grow faster. Ill-posedness causes terrible problems for numerical simulations; typically, simulations run at successively higher resolutions do not converge (i.e., they have errors which increase with increasing resolution).

finding it. Even putting aside the difficulties inherent in a blind search through a high-dimensional parameter space of formulations, we still have trouble *rapidly* determining how “well-behaved” a given formulation is.

Of course, we can always simply run full simulations to deduce the stability properties of a given formulation. Unfortunately, numerical evolutions are slow, often expensive, and, furthermore, at present require some human oversight to produce believable answers (e.g., through convergence testing and other sanity checks). In other words, full numerical simulations are far from suitable for a blind search in a high-dimensional parameter space. Therefore, we would much prefer analytic insight into the stability properties of various formulations of Einstein’s equations.

1.3.1 Geometric optics approach to stability

Chapter 5 presents a simple quasianalytic technique, based on the growth of analytically tractable solutions, which can rapidly discover the most ill-behaved well-posed formulations of Einstein’s equations. Generally speaking, this technique consists of . . .

1. *Constructing (approximate) solutions*, based on the geometric-optics approximation, to a broad class of well-posed partial differential equations which include many formulations of Einstein’s equations;
2. *Measuring the amplification of certain linearized transients*—or, more precisely, measuring how certain geometric-optics solutions of the PDE, when linearized about a known solution, grow in the future domain of dependence of the initial data slice;²² and finally
3. *Determining if the amplification is practically unacceptable*, based on a reasonable conjecture about how a numerical simulation of a nonlinear PDE will behave in the presence of exceedingly large, generic initial data (i.e., the simulation will rapidly fail).

To provide a clear demonstration of these methods, Chapter 5 also applies this technique to the KST 2-parameter formulation of Einstein’s equations [19], when those equations are linearized about two elementary spacetimes:²³ (i) flat space in Rindler coordinates and (ii) a Schwarzschild black hole in Painleve-Gullstrand coordinates. In each case, this method reveals a large region of KST parameter space permit enormous (i.e., a factor of 10^{32}) magnification of some transients within a light-crossing time.

²²Loosely speaking, the future domain of dependence of a slice is that part of spacetime which depends *only* on information supplied on the data slice; that part of spacetime depends in no way on the boundary conditions.

²³These spacetimes are well-understood, and are used as model problems for generic numerical simulations.

1.4 Appendix: Writing Einstein's equations in a manner suitable for numerical solution

In this appendix, we briefly review the ADM 3+1 decomposition, a technique for decomposing Einstein's equations in a manner suitable for numerical time evolution—that is, for slicing up a 4-manifold into timeslices, and for expressing the metric on each timeslice in terms of the metric on preceding timeslices. This review is intended as a brief summary for the nonspecialist, so that nonspecialist can better appreciate the significance of Chapter 5. For those seeking a more systematic treatment, the literature contains many excellent surveys of 3+1 decomposition of Einstein's equations.²⁴

Notation

Following traditional notation in this field (cf., e.g., [24]), we use a coordinate-based tensor notation $(x^0, x^1, x^2, x^3) = (t, x, y, z)$, with Greek indices (α, β, \dots) that run from 0 to 3 and latin indices that run from 1 to 3. The time coordinate will always be indexed by zero (i.e., $x^0 = t$).

1.4.1 Basic Einstein's equations

Einstein's equations in vacuum may be expressed as

$$G_{\alpha\beta} = R_{\alpha\beta} - \frac{1}{2}Rg_{\alpha\beta} = 0,$$

where the Ricci tensor R_{ab} is a contraction of the Riemann tensor (i.e., $R_{\alpha\beta} = g^{\mu\nu} R_{\alpha\mu\beta\nu}$). In terms of an explicit family of coordinates, the Riemann tensor can be expressed in terms of first and second partial derivatives of the metric tensor [cf., e.g., MTW (8.24) and (8.44) [24]]. Thus, because the Ricci tensor is symmetric and because spacetime has four dimensions, Einstein's equations are 10 equations involving the metric $g_{\alpha\beta}$.

The Riemann tensor (and hence $G_{\alpha\beta}$) is linear in the second derivatives.

1.4.2 3+1 perspective

As a practical matter, we usually solve physical equations by evolving them forward in time—that is, we subdivide spacetime into slices, and we express fields on future slices in terms of their values on earlier slices. We can do the same with Einstein's equations, expressing the value of the metric $g_{\alpha\beta}$ on later slices in terms of its values on earlier slices.

²⁴KST Section II A [19] provides a brief summary of the essential equations. York's review article [22] is the classic treatment of the subject. Chapter 10 of Wald's book [23] provides an excellent survey to 3+1 methods for relativity; note, however, his sign for the extrinsic curvature differs from the usual numerical relativity convention.

1.4.2.1 Gauge freedom

Strictly speaking, the 10 fields of the metric $g_{\alpha\beta}$ cannot be fully specified with Einstein's 10 equations ($G_{\alpha\beta} = 0$) alone: because we have freedom to choose coordinates (i.e., four spatial functions) completely freely, in reality we have only 6 true degrees of freedom, which we usually express as the spatial components of the metric g_{ab} . The remaining metric components must be specified by some additional procedure (i.e., equivalent to choosing coordinates on the 4-manifold); typically, these remaining components are expressed in terms of the lapse (α) and shift vector (β^a):

$$ds^2 = -\alpha^2 dt^2 + g_{ab}(dx^a + \beta^a dt)(dx^b + \beta^b dt) .$$

Conversely, Einstein's 10 equations contain only 6 equations which involve second time derivatives of the metric. Only the components G_{ab} for both a and b spatial involve second time derivatives; the components $G_{\alpha\beta}n^\beta$ (for n^β the normal to the timeslice) involve only first time derivatives. In other words, if we express these equations in first-order form, defining a quantity K_{ab} to be closely related to first time derivatives of the metric²⁵, only the 6 equations $G_{ab} = 0$ for a, b spatial involve first time derivatives of K ; the expressions $G_{\alpha\beta}n^\beta$ involve no time derivatives at all.

1.4.2.2 Constraint and evolution equations

To summarize, we may express Einstein's equations in first-order form (that is, using as variables the spatial metric g_{ab} and the extrinsic curvature K_{ab}) as two separate sets of equations:

- *Constraint equations:* The equations $G_{\alpha\beta}n^\beta = 0$ do not involve any time derivatives K_{ab} and g_{ab} ; they depend only on values on the present timeslice. These four expressions are denoted the *constraint equations*.
- *Evolution equations:* The equations $G_{ab} = 0$ for a and b spatial are linear in first time derivatives of K_{ab} . Combined with the definition of K in terms of the first time derivative of the spatial metric, which may be thought of as an evolution equation for K_{ab} , these equations provide a coupled first-order system of evolution equations for K_{ab} and g_{ab} equivalent to the original second-order system for g_{ab} alone. These six expressions are therefore denoted the *evolution equations*.

One can show that the evolution equations preserve the constraints: if initial data for K_{ab} and g_{ab} are chosen to satisfy the constraints, then at each later time, the metric g_{ab} and extrinsic curvature K_{ab} also always satisfy the constraints. Therefore, to solve Einstein's equations, one need only choose some gauge convention for the lapse α and shift β^a ; choose initial data for g_{ab} and K_{ab} which

²⁵Strictly, $K_{ab} = -L_n g_{ab}/2$ is the Lie derivative of the spatial metric along the normal direction n^a to the timeslice.

satisfies the constraint equations; and then evolve that initial data according to the 6 evolution equations to find the metric at all later times.

Bibliography

- [1] K. S. Thorne, gr-qc/9704042
- [2] K. S. Thorne, gr-qc/0506086
- [3] K. S. Thorne, in *300 years of gravitation*, eds. S. Hawking and W. Israel, Cambridge University Press, 1987.
- [4] A. Buonanno, Y. Chen, and M. Vallisneri, Phys. Rev. D **67**, 104025 (2003)
- [5] A. Buonanno, Y. Chen, and M. Vallisneri, Phys. Rev. D **67**, 024016 (2003)
- [6] C. M. Will, *Theory and experiment in physics*, Cambridge University Press, 1993. gravitational
- [7] F. D. Ryan, Phys. Rev. D **56**, 1845 (1997)
- [8] F. D. Ryan, Phys. Rev. D **52**, 5707 (1995)
- [9] E. Poisson, Phys. Rev. D **54** 5939 (1996)
- [10] S. Sigurdsson and M. J. Rees, MNRAS **284** 318
- [11] M. Freitag, astro-ph/0107193
- [12] S. Finn and K. S. Thorne, Phys. Rev. D **62** 124021 (2000)
- [13] Several websites provide public-relations-oriented (<http://lisa.jpl.nasa.gov>, at NASA) and science-oriented (<http://www.srl.caltech.edu/lisa>, at Caltech) descriptions of the LISA project.
- [14] A. Ori and K.S. Thorne, Phys. Rev. D. **62** 124022
- [15] LIGO II Conceptual Project Book, available at <http://www.ligo.caltech.edu/docs/M/M990288-A1.pdf>
- [16] The LIGO Scientific Collaboration (LSC) maintains a website describing plans for the first-generation upgrade: <http://www.ligo.caltech.edu/~ligo2/>
- [17] G. Calabrese et al, Phys. Rev. D **66** 041501 (2002)

- [18] G. Calabrese and O. Sarbach, gr-qc/0303040
- [19] L. E. Kidder, M. A. Scheel, and S. A. Teukolsky, Phys. Rev. D **64** (2001) 064017
- [20] M. A. Scheel and L. Lindblom, Phys. Rev. D **66** (2002) 084014
- [21] L. Lindblom and M. A. Scheel, Phys. Rev. D **67** (2003) 124005
- [22] J. W. York, “Kinematics and Dynamics of General Relativity,” in *Sources of Gravitational Radiation*, edited by L. L. Smarr, Cambridge University Press, 1978
- [23] R. Wald, *General Relativity*, University of Chicago Press, 1984.
- [24] C. Misner, K. S. Thorne, and J. A. Wheeler, *Gravitation*, W. H. Freeman and Co., 1973
- [25] E. D’Ambrosio, Phys. Rev. D **67**, 102004

Chapter 2

Inspiral of compact objects into black holes

[This chapter, along with the associated appendices, is precisely the text of R. O’Shaughnessy, Phys. Rev. D **67** 044004 (2002).]

Abstract

Ori and Thorne have discussed the duration and observability (with LISA) of the transition from circular, equatorial inspiral to plunge for stellar-mass objects into supermassive ($10^5 - 10^8 M_\odot$) Kerr black holes. We extend their computation to eccentric Kerr equatorial orbits. Even with orbital parameters near-exactly determined, we find that there is no universal length for the transition; rather, the length of the transition depends sensitively—essentially randomly—on initial conditions. Still, Ori and Thorne’s zero-eccentricity results are essentially an upper bound on the length of eccentric transitions involving similar bodies (e.g., a fixed). Hence the implications for observations are no better: if the massive body is $M = 10^6 M_\odot$, the captured body has mass m , and the process occurs at distance d from LISA, then $S/N \lesssim (m/10M_\odot)(1\text{Gpc}/d) \times O(1)$, with the precise constant depending on the black hole spin. For low-mass bodies ($m \lesssim 7M_\odot$) for which the event rate is at least vaguely understood, we expect little chance (probably [much] less than 10%, depending strongly on the astrophysical assumptions) of LISA detecting a transition event with $S/N > 5$ during its run; however, even a small infusion of higher-mass bodies or a slight improvement in LISA’s noise curve could potentially produce $S/N > 5$ transition events during LISA’s lifetime.

2.1 Introduction

The gravitational waves emitted during inspiral and infall of a body (mass m) into a black hole (mass M) should reveal detailed information about the orbital geometry and the hole’s spacetime

geometry, thereby providing high-precision tests of general relativity [1]. While the scattering of waves off background curvature implies that waves emitted at any time during the inspiral provide some small measure of even the smallest scale variations in the background spacetime geometry, the waves emitted as a particle passes through a region provide the most sensitive tests of that region: they reveal what path the particle has followed, and therefore constrain the spacetime to permit such a path. Therefore, to provide a sensitive probe of the innermost regions of black-hole spacetimes, we want to study orbits that pass as near as possible to the hole itself. Unfortunately, this means that the signals that are potentially the most informative about the hole's innermost structure are typically the briefest: they arise from the end of the bound portion of the orbit and from the transition from inspiral to plunge. Since the relevant fraction of the orbit persists for only a small fraction of the overall detectable inspiral, we have significantly less probability to resolve waves during this interval than to resolve earlier, longer portions of the inspiral. One therefore wants to roughly characterize the waves emitted during these intervals (in the case of LISA sources, the goal of this paper). If this characterization suggests that planned observatories such as LIGO or LISA could detect them, one should then carry out much more detailed studies of these last few orbits and the waves they emit.

For inspirals appropriate to the LIGO band ($\sim 10\text{-}10^3$ Hz) and which LIGO can plausibly detect ($\eta = m/M \sim 0.01$ to 1), order-of-magnitude computations (say, by post-Newtonian methods) suggest the last few waves are detectable [2]. But because in this regime simple approximation techniques (such as post-Newtonian [2, 3, 4] or test-particle approximations) break down, and because numerical relativity [5] codes remain incapable of evolving orbits accurately enough to find the waves, the community does not yet possess a waveform trusted for any purpose beyond detection.

LISA's band ($\sim 10^{-3}\text{-}10^{-1}$ Hz) will prove more sensitive to extreme mass-ratio infalls—that is, to stellar-mass black holes, white dwarfs, and neutron stars falling into supermassive [$M = O(10^5\text{-}10^8)$, so $\eta \sim 10^{-4}$ to 10^{-8}] black holes [6]. With such extreme mass ratios, the computation of detailed waveforms for purposes beyond mere detection should prove much simpler: to understand evolution, we need do nothing more than solve the classical radiation-reaction problem, albeit on a curved spacetime and with a gravitational, rather than electromagnetic, field [7]. While this problem hasn't been solved to the accuracy required to construct long-integration-time coherent detection templates, one can employ adiabatic approximations to address most preliminary investigations. For example, as Ori and Thorne [8] have discussed in the context of circular inspiral, to understand the $\eta \ll 1$ transition's duration—measured in experimentally observable gravitational wave cycles—we do not need a precise knowledge of the reaction force. An averaged reaction force—one we can easily deduce from the radiation of conserved constants—suffices for the short interval we will coherently employ it. Applying this reaction force, we can follow the particle through transition and thereby predict roughly how long this transition will last.

The goal of this paper is to extend the Ori-Thorne analysis to eccentric Kerr orbits, in an effort to estimate the prospects of LISA detecting a transition from inspiral to plunge.

This analysis relies on using the radiation of two conserved constants E, L to compute the effect of radiation on the orbit. But for Kerr inclined orbits there is an additional constant—the Carter constant—whose evolution has not yet been related to fluxes at infinity. Since we lack the necessary tools, we leave the Kerr inclined case to a future paper.

2.1.1 Outline of this paper and summary of conclusions

In Sec. 2.2, we will outline the basic physical framework behind our approach. In particular, we will introduce an explicit procedure to estimate the time duration of a transition. This procedure takes as input the net (time-averaged) fluxes of energy and angular momentum from the particle’s instantaneously geodesic orbits, input one obtains from a solution of the Teukolsky equation given a geodesic orbit as source. This procedure also takes as input some observationally-defined interpretation of what “the transition region” is. As the latter is ambiguous, and depends on exactly what sorts of templates one uses to find it, the exact length of the transition will depend on the convention one uses.

Ideally, one should define some unambiguous set of templates and match those against the simulated emitted waves to both define the transition duration and deduce the resulting signal-to-noise ratio for a given source. But for brevity and simplicity, as discussed in Sec. 2.3, we will use a much cruder scheme—based on a purely sinusoidal, quadrupolar model for the waves—to characterize the expected LISA signal-to-noise ratio from a specific transition crossing. Given S/N for an event and loosely-understood rates for transition events, we then develop, in Sec. 2.3, a scheme for estimating the probability that LISA will see an event with S/N greater than some detection threshold.

With this complete scheme for estimating the signal-to-noise of a characteristic source and determining the probability that LISA, in its currently-planned configuration, will see something, in Sec. 2.4 and Sec. 2.5 we will apply it to inspirals into Schwarzschild and Kerr holes, respectively. We find in Sec. 2.5.4 that Ori and Thorne’s zero-eccentricity results are essentially an upper bound on the length of eccentric transitions involving similar bodies (e.g., a fixed). It follows, in Sec. 2.5.5, that if we accept current (rough) astrophysical estimates of the masses and numbers of inspiralling stellar-mass black holes and if we employ only the current LISA design, we expect LISA will not see any transitions from inspiral to plunge during its lifetime, though it may come close.

Slight changes in LISA could make some transitions detectable. Dramatic improvements would be required to render LISA sensitive to prograde inspirals of stellar-mass black holes into rapidly-spinning ($a > 0.9$) supermassive holes. But assuming such inspirals are a small proportion of all inspirals, if the LISA noise curve is lowered by a factor of 3 (as is under currently discussion for

other reasons), or if nature provides black holes more massive than $10M_\odot$ (say $30M_\odot$) in numbers approaching current estimates for $10M_\odot$, LISA would have a good chance of seeing one or two transitions sometime during its lifetime.

2.2 Physical framework underlying the transition length estimate

In the (formal) absence of radiation reaction, a particle in equatorial orbit about a Kerr hole moves along a geodesic. Its radial motion can be determined from a first integral of the geodesic equation (equivalent to conservation of rest mass; see comments in Appendix A.1) [9]:

$$\left(\frac{dr}{d\tau}\right)^2 + V[r(\tau), E, L] = 0, \quad (2.1)$$

$$V \equiv -(E^2 - 1) - \frac{2}{r} + \frac{(L^2 - a^2(E^2 - 1))}{r^2} - \frac{2(L - aE)^2}{r^3}. \quad (2.2)$$

Here and throughout this paper all quantities are, for simplicity, made dimensionless using the particle's mass m and the hole's mass M : $E = (\text{orbital energy})/m$, $L = (\text{orbital angular momentum})/mM$, $r = (\text{orbital boyer-lindquist radius})/M$, $\tau = (\text{particle's proper time})/M$, and $a = (\text{hole spin angular momentum})/M^2$. Physical solutions may be specified by (E, L) or by any other pair of equivalent orbital parameters. It is conventional in the inspiral literature to employ as alternatives the parameters p (a relativistic generalization of semi-latus rectum) and e (a relativistic generalization of orbital eccentricity) [10, 11]; these parameters are discussed in more detail in Appendix A.2.

We concern ourselves with a region of parameter space for which the maximum V_{\max} of the potential is nearly 0 (Fig. 2.1) and which therefore nearly admits a circular orbit at the radius r_{\max} of the maximum. The geodesic equation Eq. (2.1) implies that particles can spend an extremely (logarithmically) long time near the maximum; i.e., the particle can “whirl” several times about the hole in angle without moving significantly in r . It is conventional to call this portion of the orbit the “whirl.”

In the presence of radiation reaction (henceforth assumed weak), we must add to the geodesic equation (gauge-dependent) time-varying terms which reflect the (gauge-dependent) influence of gravitational radiation on the test particle's path. These gauge-dependent terms oscillate on the same characteristic timescale as the radiation field. Since the radiation field is predominantly produced during the whirl part of the orbit, the radiation field predominantly oscillates at harmonics of the angular frequency Ω of circular orbits at the maximum. By averaging these reaction forces over a few cycles (e.g., over times $\approx 2\pi/\Omega$) to obtain their secular effect, we in principle find expressions

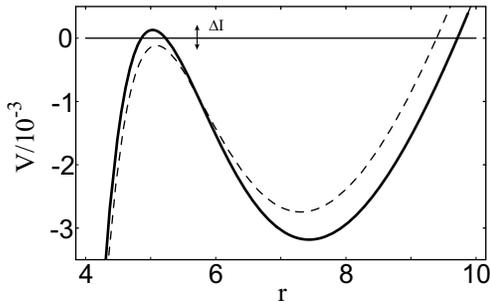


Figure 2.1: The effective potential $V(r, E, L)$ for radial geodesic motion gradually evolves during inspiral. If the initial eccentricity is nonzero, the maximum of the potential will gradually decrease until it passes below zero, thus permitting the particle to fall into the hole. By way of illustration, we plot $V(r, E, L)$ for $a = 0$ and (solid) $(E, L) = (0.948157, 3.53038)$ and (dashed) $(E, L) = (0.947454, 3.52092)$. Under the action of radiation reaction, if $\eta = 2 \times 10^{-3}$ (an exaggerated mass ratio) the first system will evolve into the second after one radial period (of the first potential).

for $E(\tau)$ and $L(\tau)$. Since the averaging time can still (particularly when the particle whirls several times about the hole) be shorter than the time the particle spends whirling around the central hole, we can to a good approximation employ Eq. (2.1) with *time-varying* $E(\tau)$, $L(\tau)$ to follow the orbit when the particle is near the maximum of the potential—and in particular the interval in which the particle goes from nearly-geodesic bound orbit into rapid plunge into the hole.

In this paper, we do not compute $E(\tau)$ and $L(\tau)$ in the thorough, general manner described here. See Sec. 2.2.4 for a discussion of what information about $E(t)$ and $L(t)$ is required for our estimate and how that information is obtained.

2.2.1 Why we may still approximate the potential as static when computing the radial orbit

According to Eq. (2.1), a test particle on an approach to a black hole will fall into the hole if the local maximum of V is negative, or equivalently if

$$I \equiv -V(r_{\max}) = -V_{\max} , \quad (2.3)$$

is positive. Radiation reaction reduces the local maximum faster than this potential V flattens out. As the maximum decreases, the particle spends ever-more of its radial cycle near the local maximum. Eventually, we reach configurations such as those shown in Fig. 2.1, where the particle can slip over the maximum and fall into the black hole.

While configurations with $I \approx 0$ appear delicately balanced and therefore highly sensitive to small changes in (E, L) , in fact under weak conditions (conditions made more explicit in Sec. 2.2.3) one may ignore radiation reaction when computing a radial orbit and treat the potential V as static.

Suppose a particle starts its whirl with some values for (E, L) (and therefore I). During the whirl, even though (E, L) change, the peak of the potential (I) will not change significantly (see Appendix A.1). Moreover, the location of the peak usually moves slowly relative to the particle. A condition for when the latter holds is presented in Sec. 2.2.3. Therefore, during the whirl one can ignore radiation reaction.

After the particle finishes whirling about the hole, it moves outward to its outer radial turning point and back. During this period, the maximum *does* change, from I to $I' = I + \Delta I$. Unless the potential is nearly flat, however, the potential away from the neighborhood of the hole will not change much as (E, L) change. Again, a condition for when the latter holds is presented in Sec. 2.2.3. Therefore, in this interval one can again ignore radiation reaction.

When we attempt to evolve the particle through the next “whirl,” we need the correct value of the height of the maximum (now $I' = I + \Delta I$) to determine how long the particle whirls around the hole. Therefore, when we start the cycle anew, we must use a potential with parameters $(E + \Delta E, L + \Delta L)$, with ΔE and ΔL the change in these constants over the preceding full radial period. If $I' < 0$, the particle will “bounce” off the maximum and we repeat the cycle above once more. But eventually we will have $I' > 0$, at which point the particle will move across the maximum during its whirl and will subsequently “plunge” into the hole.

To summarize: so long as the potential is approximately static (cf. Sec. 2.2.3), we expect we can understand transitions from inspiral to plunge by way of examining the geodesic equation [Eq. (2.1)] in the neighborhood of the local maximum, using $I \in [0, \Delta I]$.

2.2.2 Adiabatic approach to estimating the duration of the transition from inspiral to plunge

So long as we can treat the potential as static, we can approximate Eq. (2.1) in the neighborhood of the potential’s maximum at $r = r_{\max}$ by the form

$$\gamma^2 \left(\frac{d\delta r}{dt} \right)^2 + \frac{\delta r^2}{\tau_o^2} = I. \quad (2.4)$$

Here $\delta r \equiv r - r_{\max}$; r_{\max} is the instantaneously static location of the local maximum of V , and also the point about which we have expanded the potential;

$$\tau_0 \equiv (V''/2)^{-1/2} \quad (2.5)$$

is a constant related to the curvature of the potential at the transition location; t is the (dimensionless) time at infinity; and γ is the redshift factor relating proper time τ to Boyer-lindquist coordinate

time at $r = r_{\max}$:

$$\gamma = dt/d\tau_{r=r_{\max}} . \quad (2.6)$$

Estimating the duration of a given transition: Solutions to Eq. (2.4) give hyperbolic motion; for example, the solution appropriate to $I > 0$ (and therefore to a particle crossing over the maximum and falling into the hole) is

$$r(t) - r_{\max} \approx \sqrt{I}\tau_0 \sinh t/(\tau_0\gamma) . \quad (2.7)$$

Using this solution, we conclude that the transition time going from $r - r_{\max} = -\delta r_{\text{ref}}$ to $r - r_{\max} = \delta r_{\text{ref}}$ is

$$T_c(\delta r_{\text{ref}}, I) \approx 2\gamma\tau_0 \sinh^{-1} \left[\delta r_{\text{ref}}/(\tau_0\sqrt{I}) \right] \approx 2\gamma\tau_0 \ln \left[2\delta r_{\text{ref}}/(\tau_0\sqrt{I}) \right] . \quad (2.8)$$

Hence given δr_{ref} , a quantity which defines what we mean by “the transition extent,” we can estimate the length of any transition (characterized by I) at any transition location (characterized by the explicit values that go into γ, τ_0).

Estimating the distribution of transition durations: There is no unique transition duration. Rather, we have a distribution of durations, depending on the distribution of I at the start of the particle’s final whirl. But that distribution is simple: since an initial configuration of particles will have some distribution of I , since this distribution evolves smoothly with no “knowledge” of the preferred scale ΔI , and since ΔI will be smaller than any scale in the distribution function, a test particle on its final, plunge-triggering whirl has an approximately equal probability to have any $I \in [0, \Delta I]$. Therefore, the probability density for a test particle to have a given duration between T_c and $T_c + dT_c$ is $dP \propto dT_c(dI/dT_c) \propto dT_c \exp[-T_c/\gamma\tau_0]$; see Eq. (2.8). Denoting by

$$T_{c-} = T_c(\delta r_{\text{ref}}, \Delta I) \quad (2.9)$$

the minimum possible transition duration, and ignoring the tiny regime of transitions which are nonadiabatic (see Sec. 2.2.3 below), we conclude that

$$dP \approx \Theta(T_c - T_{c-}) e^{-(T_c - T_{c-})/\gamma\tau_0} dT_c/\gamma\tau_0 . \quad (2.10)$$

[where $\Theta(x) = 1$ when $x > 0$, 0 otherwise].

We can also characterize distribution of crossing times by a function $T_c(p)$ such that only a fraction p of particles could (assuming the conditions of Sec. 2.2.3 hold) have longer crossing times. For example, only a fraction 10^{-n} of particles will have duration longer than

$$T_{c,n} \equiv T_c(\delta r_{\text{ref}}, \Delta I 10^{-n}) \approx 2\gamma\tau_0 \ln \left[\frac{2\delta r_{\text{ref}}}{\tau_0 10^{-n/2} \sqrt{\Delta I}} \right] . \quad (2.11)$$

Additional comments:

- **Converting to number of cycles:** As the particle passes through the transition region, the particle “whirls” about the black hole a few times. Since its radial location is largely fixed while it whirls around the hole, so is its angular frequency $d\phi/dt \equiv \Omega$; therefore, we can re-express any duration T_c in terms of a “number of orbital cycles” the particle “whirls” around the hole N_c , defined by

$$N_c = \frac{T_c \Omega}{2\pi}. \quad (2.12)$$

Since we concern ourselves with only Kerr equatorial orbits, we have

$$\Omega(r) = \frac{\text{sign}(a)}{r^{3/2} + a}. \quad (2.13)$$

- **Characteristic duration and variation of T_c with e :** By examining the quadratic approximation to the potential [Eq. (2.4)], or equally well from Eq. (2.8), we see that the transition duration is always $T_c \sim (\text{few}) \times \gamma \tau_o$ —that is, the crossing time is around the natural timescale of the effective potential. Admittedly, since $O(\Delta I) = O(\eta)$, the quantity labeled (few) could be—and will be—significant; therefore, the logarithmic correction in Eq. (2.8) is necessary. But for purposes of understanding the variation of crossing time with orbital parameters, largely we can regard $T_c \sim \gamma \tau_o$. For example, we expect T_c to increase monotonically with decreasing orbital eccentricity e —that is, as the maximum possible energy barrier decreases and the potential flattens out—simply because τ_o does. [By way of example, see Eq. (2.41), an expression for $\gamma \tau_o$ appropriate to Schwarzschild.]

- **On variation of T_c with η :** Similarly, we can loosely characterize the dependence of the duration distribution—or, for clarity, T_{c-} —on η by noting i) $\sinh^{-1}(x) \approx \ln 2x$ when x is large and ii) $\Delta I \propto \eta$, so we can characterize variation with η by $H(\eta_o)$, defined by

$$\frac{T_{c-}(\eta)}{T_{c-}(\eta_o)} - 1 \approx \frac{\ln \sqrt{\eta_o/\eta}}{\ln \left(\delta r_{\text{ref}}/\tau_o \sqrt{\eta_o \Delta I/\eta} \right)} \equiv \ln(\eta_o/\eta) H(\eta_o)$$

[where we have used the fact that $\Delta I/\eta$ is independent of η to justify writing the denominator as $2/H(\eta_o)$]. In other words, while the minimum transition duration will grow slightly shorter with larger mass ratios, the dependence (like the dependence on $\Delta I/\eta$) is weak; typically (e.g., for Schwarzschild) we find $H \in \sim [0.1, 0.4]$.

2.2.3 Explicit conditions under which we may continue to approximate the potential as static

Throughout our analysis, we have approximated the potential as static. As outlined in Sec. 2.2.1, there are two ways in which this approximation could fail.

First, the potential away from the maximum could change significantly during one whole radial orbit. Generally the change of V at any specific location is small. Such changes therefore matter only if the potential is delicately balanced near zero at every point in which the particle orbits. More explicitly, we expect problems if the change ΔI of the potential's maximum during one whole radial orbit is comparable to the difference between the maximum and minimum of V . Therefore, we conservatively require

$$I_{\max} \equiv V(r_{\max}) - V(r_{\min}) \gg \Delta I . \quad (2.14)$$

An explicit form for I_{\max} is presented in Eq. (A.11). Since the potential gets very flat as $e \rightarrow 0$, our approximations will break down at eccentricities below e_{\min} , defined by solutions to

$$\Delta I = I_{\max}(e_{\min}) . \quad (2.15)$$

Second, the radial location r_{\max} of the maximum could move significantly while the particle is in its last whirl about the hole. Based on Eq. (2.1), to prevent against this we require $I = (dr/d\tau)^2 \gg (dr_{\max}/d\tau)^2$, i.e., that

$$I \gg I_{\text{ad},\min} \equiv \left(\frac{dr_{\max}}{d\tau} \right)^2 = \left(\gamma \frac{dr_{\max}}{dt} \right)^2 . \quad (2.16)$$

The precise procedure that we will use to estimate $dr_{\max}/d\tau$ will be discussed in Sec. 2.2.4. In summary, so long as $I \gg I_{\text{ad},\min}$, or equivalently so long as the crossing duration significantly shorter than

$$T_{c,\text{ad},\min} = T_c(\delta r_{\text{ref}}, I_{\text{ad},\min}) , \quad (2.17)$$

gradual motion of the potential will not significantly alter the transition length estimates presented earlier.

2.2.4 Inputs necessary for estimating the transition length

In the above we have outlined a computational procedure which takes as input δr_{ref} and knowledge about radiation reaction (namely, about ΔI and about dr_{\max}/dt) and which gives us in return an estimate of the length of any specific transition from inspiral to plunge. We now describe the explicit approximations we shall use to estimate ΔI and $I_{\text{ad},\min}$ from known information about $E(\tau)$ and $L(\tau)$. We also make an explicit choice for δr_{ref} .

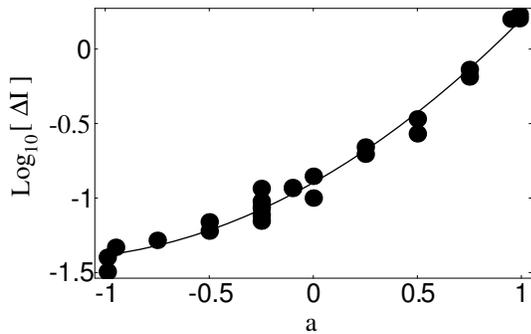


Figure 2.2: The (log of the) change of the potential’s maximum during the last radial orbit ($\log_{10}[\Delta I/\eta]$) versus a . The points show our ΔI for various cases with parameters (p, e) close to those consistent with circular orbits (the boundary between stable and unstable orbits); these points were obtained from numerical solutions of the Teukolsky equation by Glampedakis and Kennefic, using the procedure described in Sec. 2.2.4.1. For each a , solid circles show values of $\log_{10}[\Delta I/\eta]$ for several different values of e ; that these points are all consistent with a single fit demonstrates that ΔI is approximately independent of e . The solid curve is the quadratic fit Eq. (2.18). [In cases where insufficient data was present for extrapolation to the line $p = p_s(e, a)$, solid circles also indicate rough upper and lower limits expected of $\log_{10}[\Delta I/\eta]$.]

2.2.4.1 Estimating ΔI

As described in Sec. 2.2.1, we obtain ΔI by comparing the potential V when the conserved constants are (E, L) to the potential V when they are $(E + \Delta E, L + \Delta L)$, where ΔE and ΔL are the change in the appropriate conserved constants over one radial orbit. We obtain ΔE and ΔL from numerical solutions to the Teukolsky equation. From their code, Glampedakis and Kennefic have kindly provided time-averaged fluxes $\langle dE/dt \rangle$ and $\langle dL/dt \rangle$ [11], which, when combined with an expression for the radial period $T(E, L)$ as given in any classical relativity text [e.g., Eq. (33.37) of MTW [9]], yields ΔE and ΔL , and thus ΔI [Eqs. (2.2) and (2.3)]. In this fashion, for each black hole (parametrized by spin parameter a), we can find $\Delta I(p, e, a)$ for any equatorial geodesic with parameters (p, e) .

We need ΔI only for the particle’s last whirl. Since the maximum is extremely close to zero, the orbital parameters (p, e, a) nearly satisfy a condition for the existence of (unstable) circular orbits $p = p_s(e, a)$ (see Appendix A.2). This curve also necessarily serves as the boundary between stable orbits and plunge. In the vicinity of this boundary line, ΔI is well-approximated by its nonzero values on the boundary. So for our computation we seek an expression $\Delta I(e, a)$.

In practice, from the values of ΔI at points near this boundary line, we extrapolate to estimate ΔI on the boundary surface itself. Figure 2.2 shows the results of our extrapolation.

One can argue that $\Delta I(p, e, a)$ on the last-stable-orbit boundary $p = p_s(e, a)$ should largely be independent of e at moderate eccentricity ¹. For this reason, Figure 2.2 shows results only as a

¹Since the orbit is nearly circular, radiation of conserved constants should be nearly uniform in time, so assume $E \approx E_o + t \times dE/dt$ and similarly for L . Take a third-order approximation to the potential. Find an explicit expression for dI/dt in terms of the solution $r(t)$ and the motion of the maximum r_{\max} . Use an approximate (sinusoidal+constant)

function of one parameter (a). Numerical data over the range $e \in [0, 0.5]$ support this conjecture. Therefore, so long as we avoid $e \approx 1$, where this conjecture has not yet been tested and likely fails, we can approximate $\Delta I/\eta$ by a function independent of e . Fitting a relatively simple function (exponential form in a , independent of e) to the data in Fig. 2.2 we find

$$\log_{10} \frac{\Delta I}{\eta} = -0.8972 + 0.7911a + 0.3047a^2. \quad (2.18)$$

2.2.4.2 Estimating $I_{\text{ad},\text{min}}$

To evaluate $I_{\text{ad},\text{min}}$, we need no more than i) knowledge of the potential (which tells us r_{max} as a function of E, L) and ii) knowledge of $dE/dt, dL/dt$ when the particles are in nearly-circular orbit near the hole.

In principle, we could approximate the latter by the appropriate values for an exactly circular (unstable) orbit. As a practical matter, comprehensive tabulation of the physically appropriate instantaneous dE/dt and dL/dt for all transitions of interest—namely, the values appropriate to a circular unstable orbit—proves time-consuming and technically challenging. Furthermore, because the crossing time depends only weakly (logarithmically) on I , and because exceedingly few particles will have $I \lesssim I_{\text{ad},\text{min}}$, we only need $I_{\text{ad},\text{min}}$ to order of magnitude.

Therefore, for practical purposes, when estimating $I_{\text{ad},\text{min}}$ by way of Eq. (2.16) we will i) perform the computation for $dr_{\text{max}}/d\tau$ analytically in terms of dE/dt and dL/dt , ii) simplify under the assumption $dE = \Omega dL$, which would be valid if we used the true forms for dE/dt and dL/dt , and then iii) insert for dL/dt the Peters-Mathews expression (an estimate obtained using linearized, quadrupolar emission from newtonian orbits) [12]

$$\frac{dL}{dt} \approx \eta \frac{32}{5} \frac{1}{p^{7/2}} (1 - e^2)^{3/2} \left[1 + \frac{7}{8} e^2 \right]. \quad (2.19)$$

2.2.4.3 Choosing δr_{ref}

To complete our procedure, we must define “the” transition duration. Unfortunately, because “the” transition from inspiral to plunge occurs at no definite location, has no well-defined start or finish, the transition duration remains a matter of convention². We shall adopt a convention motivated by a simple model of gravitational-wave data analysis.

The key feature of waves emitted during the transition is their considerable simplicity: they are emitted from a nearly-circular-equatorial orbit at r_{max} , and hence are characterized by the angular

solution for $r(t)$ in the previous expression to show that ΔI over one radial period is approximately independent of eccentricity.

²The closest “natural” definition would be some fraction, defined some way or another, of the length of the binding region. But since the binding region goes to zero length, when the potential gets flat, the length of the transition would go to zero. We therefore would have the unusual result that the transition from circular inspiral to plunge took no time. This result is inconsistent with the Ori & Thorne value.

frequency Ω associated with circular orbits there. If we were to try to detect these gravitational waves—for simplicity, focusing on the dominant frequency component, $\omega = 2\Omega$ —we would want to insure that our model $\phi = 2\Omega t$ for the gravitational wave phase agrees, within π , with the true wave phase.

The true rate of change of orbital phase is

$$\frac{d\phi}{dt}(r, E, L) = \frac{g^{\phi\phi}L - g^{\phi t}E}{-g^{tt}E + g^{t\phi}L}, \quad (2.20)$$

where $g^{\alpha\beta}$ are known Kerr metric functions in Boyer-Lindquist coordinates, and E, L are consistent with the circular orbit at $r = r_{\max}$ (use standard expressions for E, L appropriate to circular orbits, such as Eqs. (2.12) and (2.13) of Bardeen, Press, and Teukolsky [13]). Demanding that the difference between the true angular phase and our fiducial reference $2\Omega t$ be no more than π over the length of the transition, we find a constraint on the crossing duration T_c :

$$\pm 2\pi/4 = \int_{-T_c/2}^{T_c/2} \left[\frac{d\phi}{dt}(r(t)) - \frac{d\phi}{dt}(r_{\max}) \right] dt. \quad (2.21)$$

When we insert $r(t) - r_{\max} = A \sinh[t/\tau_o\gamma]$ with $A = \sqrt{I}\tau_o$ [Eq. (2.7)] into the above, we find an expression we can invert for $\delta r_{\text{ref}}(\Delta I)$:

$$\pi/4 \approx \left| \frac{d(d\phi/dt)}{dr} \right| \gamma\tau_o \sqrt{\delta r_{\text{ref}}^2 + A^2} \approx \left| \frac{d(d\phi/dt)}{dr} \right| \gamma\tau_o \delta r_{\text{ref}} \quad (2.22)$$

[where the constant A has been neglected in this expression, as it is always much smaller than δr_{ref}].

Solving for δr_{ref} , we obtain

$$\delta r_{\text{ref}} = \frac{\pi}{4\gamma\tau_o |d(d\phi/dt)/dr|}. \quad (2.23)$$

We will use this form even when it predicts $\delta r_{\text{ref}} = O(1)$ [in other words, when $\delta r_{\text{ref}} = O(M)$ when we convert to physical units]. Notice this δr_{ref} is independent of mass ratio.

2.3 Estimating the probability for LISA to observe a transition

We wish to estimate, for each choice of the supermassive hole's angular momentum and distance from earth, and for each choice of test particle orbital parameters, the signal-to-noise (S/N) LISA would obtain from waves emitted during the transition. By combining this S/N with the (poorly-known) statistics of black-hole inspirals, we can estimate the probability LISA will see a transition event (e.g., have $S/N > 5$).

2.3.1 Estimating LISA's signal-to-noise for a given transition

Since the transition waves are emitted by a circular orbit of frequency

$$f_{\text{orb}} = \Omega(r_{\text{max}})/2\pi M, \quad (2.24)$$

the gravitational waves will be at that frequency and its harmonics. For simplicity, assume that LISA detects only the strongest waves, the waves emitted from the second harmonic $f_{\text{tr}} = 2f_{\text{orb}}$. These waves will last for an interval

$$\Delta t = N_c/f_{\text{orb}} (= MT_c). \quad (2.25)$$

We can approximate their characteristic rms (source-orientation-averaged) amplitude [following OT equation (4.7)] as a Peters-Mathews-style quadrupole term (averaged over all orientations) times a relativistic correction:

$$h^{\text{rms}} = \frac{8}{\sqrt{5}} \frac{M}{d} \eta \Omega(r_{\text{max}})^{2/3} \sqrt{\dot{\mathcal{E}}_{\infty,2}}. \quad (2.26)$$

Here d is the distance to the source and $\dot{\mathcal{E}}_{\infty,2}$ is a relativistic correction factor defined explicitly in OT equation (2.3).

LISA has a spectral density of noise S_h for waves incident on it with optimal propagation direction and polarization; it has spectral density $5S_h$ for typical directions and polarizations. Therefore, on average, LISA should accumulate a signal-to-noise from the transition event given by

$$(S/N)_{\text{tr}} = \frac{h^{\text{rms}}}{\sqrt{5S_h(f_{\text{tr}})/\Delta t}}. \quad (2.27)$$

Particularly special sources could have significantly higher S/N . For example, we can pick up an increase of $\sqrt{5}$ if the source is ideally positioned on the sky, and a similar increase if the source itself is optimally oriented. But overall, the above scheme suffices to estimate the signal-to-noise LISA would see from the transition between inspiral and plunge for any capture m into M with any specific source parameters (e.g., e , a) at any distance d .

Explicit expressions needed to compute LISA's signal-to-noise for a given transition

To evaluate Eq. (2.27), we need in addition to N_c and Ω [which enter into S/N via δt and f_{tr}] the LISA noise curve S_h and the relativistic correction factor $\dot{\mathcal{E}}_{\infty,2}$. The LISA noise curve may be modeled by [OT equation (4.9)]

$$S_h(f) = \left[(4.6 \times 10^{-21})^2 + (3.5 \times 10^{-26})^2 \left(\frac{1\text{Hz}}{f} \right)^4 + (3.5 \times 10^{-19})^2 \left(\frac{f}{1\text{Hz}} \right)^2 \right] \text{Hz}^{-1}. \quad (2.28)$$

The appropriate relativistic correction factor $\dot{\mathcal{E}}_{\infty,2}$ can in principle be extracted from simulations of waves emitted by particles in unstable circular orbits. As in practice the latter proves time-consuming to evaluate and tabulate for all possible eccentric orbits and for all a , for simplicity we will assume that the appropriate relativistic correction factor is i) fixed for all orbits close to a black hole of angular momentum a and ii) given explicitly by the value appropriate to the innermost stable circular orbit (ISCO). This latter expression has been tabulated by Ori and Thorne (see the $\dot{\mathcal{E}}_{\infty,2}$ column in their Table II); we approximate their results by

$$\log_{10} \left(\dot{\mathcal{E}}_{\infty,2} \right)_{\text{OT}} \approx -0.0473 + 0.211x - 0.053x^2 + 0.034x^3 + 0.010x^4 \quad (2.29)$$

where $x = \log_{10}(1 - a)$.

Dominant terms in the signal-to-noise estimate

As written, the signal-to-noise estimate Eq. (2.27) disguises what kinds of effects predominantly influence it—for example, whether changes in the strength of radiation emitted prove more important or less than changes in the duration T_c of the transition. To clarify the dominant contributions to our estimate, fix some a and compare the signal-to-noise between two transitions (1,2) involving otherwise arbitrary parameters (e.g., m , M , d , e). Substituting expressions for Δt [Eq. (2.25)], h^{rms} [Eq. (2.26)], and $f_{\text{tr}} = 2f_{\text{orb}}$ [Eq. (2.24)] into Eq. (2.27); assuming $\dot{\mathcal{E}}_{\infty,2}$ is a fixed function of a ; and comparing the resulting S/N at two sets of orbital parameters, we find

$$\frac{(S/N)_2}{(S/N)_1} = \frac{m_2 d_1}{m_1 d_2} \sqrt{\frac{N_{c,2}}{N_{c,1}}} \left(\frac{\Omega_1}{\Omega_2} \right)^{1/6} \times \sqrt{\frac{M_1 S_h \left(\frac{\Omega_2}{\pi M_2} \right)}{M_2 S_h \left(\frac{\Omega_1}{\pi M_1} \right)}}. \quad (2.30)$$

The first two terms reflect the natural m/d scaling of emitted waves. The third term reflects the fact that more orbits around the hole during the transition mean more gravitational wave cycles seen by LISA. The fourth term, which combines the fact that gravitational waves emitted closer to the hole are stronger and yet last for less time, is to a good approximation constant. Finally, the last term reflects LISA's sensitivity. The only term which depends explicitly on M (ignoring the weak variation in N_c), this last term selects black hole masses which have their transition close to optimally positioned in the LISA band, or $M \approx (\text{few}) \times 10^6$. So long as the mass is so, this term varies comparatively little.

2.3.2 Method for estimating the probability of detecting some transition during LISA's operation

Above we gave a procedure for computing the S/N for any given source. But the sources which produce the strongest signals (inspirals very close by) are rare. Therefore, for any given $(S/N)_o$ we have a certain probability that, during the entire operation time T_L of LISA, we detect no inspirals with $S/N > (S/N)_o$.

Since the relevant statistics for supermassive black holes and compact objects are poorly known, we will not attempt a detailed calculation that allows for all possible factors (e.g., source-orientation effects). Instead, for a first-pass estimate of the likelihood that LISA will see a transition, we will i) fix $M = 10^6$, ii) approximate LISA's noise curve as flat (in other words, ignore variations in S/N due to the emitted radiation being slightly off LISA's peak sensitivity), iii) ignore any orientation-related increase in the emissivity of the source or the sensitivity of LISA, iv) approximate N_c as independent of m , v) further replace N_c at each a by some characteristic number of cycles (the precise value to be chosen later, when we understand how N_c varies), and vi) assume all black holes have the same value of a (again, to be chosen later). To be particularly explicit, we assume the S/N varies with m , d , and a in the following manner:

$$\left(\frac{S}{N}\right)(m, d, a) \approx \frac{m}{10M_\odot} \frac{1\text{Gpc}}{d} \left(\frac{S}{N}\right)_A = K \frac{m}{d}. \quad (2.31)$$

Here $(S/N)_A \equiv (S/N)_A(10M_\odot, 1\text{Gpc}, a)$ is a fiducial approximation to the signal-to-noise ratio for an inspiral with $m = 10M_\odot$, $d = 1\text{Gpc}$, and a .

Suppose we have a discrete family of possible compact objects of masses m_k with rates (per galaxy containing a $10^6 M_\odot$ hole) r_k ; suppose the number density of galaxies containing a $10^6 M_\odot$ hole is ρ_g . Subdividing the universe into cubes of cell size Δr , we find the probability a given cell has an inspiral of mass m_k into a $10^6 M_\odot$ hole at some time during the lifetime T_L of LISA is $p_k = \rho_g r_k T_L \Delta r^3$. Suppose we're concerned with a threshold S/N level $S/N = s_o$. At such a level we could see a source of mass m_k out to a distance $d_k = K m_k / s_o$. If no inspirals have $S/N > s_o$, then for every cell in range, we have no inspirals of any mass type. Therefore, the probability that no inspirals occur with $S/N < s_o$ is

$$P(\text{no } S/N > s_o) = \prod_k (1 - p_k)^{4\pi d_k^3 / 3\Delta r^3} \approx \exp \left[-\frac{4\pi R_{\text{net}} T_L}{3} \frac{K^3 \langle m^3 \rangle}{s_o^3} \right] \quad (2.32)$$

where in the last line we use $p_k \ll 1$, $R_{\text{net}} \equiv \sum \rho_g r_k$ (the net event rate per unit volume for all inspirals), and $\langle m^3 \rangle \equiv \sum \rho_g r_k m^3 / R_{\text{net}}$ (the mean cubed mass of inspiralling bodies, where weights are by event rate). Note that $4\pi K^3 \langle m^3 \rangle / (3s_o^3)$ is the volume of space in which an inspiral involving a mass $\langle m^3 \rangle^{1/3}$ can be seen with a signal-to-noise ratio $> s_o$. Necessarily, the probability that *some*

source has $S/N > s_o$ is $P(\text{some } S/N > s_o) = 1 - P(\text{no } S/N > s_o)$.

We can reorganize this expression to tell us, for a given probability P_{no} , at what S/N we will have a probability P_{no} of having no signals of stronger strength:

$$\left(\frac{S}{N}\right)_{\text{no}}(P_{\text{no}}) \equiv \frac{\langle m^3 \rangle^{1/3}}{10M_{\odot}} \frac{1\text{Gpc}}{\left(\frac{4\pi}{3}R_{\text{net}}T_L\right)^{-1/3}} \frac{1}{[\ln(1/P_{\text{no}})]^{1/3}} \times \left(\frac{S}{N}\right)_A. \quad (2.33)$$

2.3.3 Probability of detecting a transition during LISA's operation

The S/N threshold [Eq. (2.33)] depends very sensitively (through $\langle m^3 \rangle^{1/3}$) on low-probability high-mass inspirals. By way of illustration, a family of $0.6M_{\odot}$ white dwarfs inspiralling with rate R and a black hole family of mass $30M_{\odot}$ and rate $10^{-4}R$ contribute in similar proportions to $\langle m^3 \rangle^{1/3}$. At present, the astrophysical community lacks sufficiently understanding of the high-mass tail to be able to reliably compute $\langle m^3 \rangle^{1/3}$. Therefore, we will neglect such objects and focus on the slightly better understood problem of capture of conventional compact objects. Doing so, we will underestimate the true $(S/N)_{\text{no}}$.

Even disregarding the high-mass holes, event rates for capture [14, 15] remain very loosely determined, ranging from rates of $\sim 2 \times 10^{-6}/\text{yr}/\text{galaxy}$ to $\sim 10^{-4}/\text{yr}/\text{galaxy}$. We take two cases as characteristic:

- *Freitag* (F) Based on astrophysical discussion by Miralda-Escude and Gould [16], Freitag allows for three species: white dwarfs ($m_{\text{WD}} = 0.6M_{\odot}$, $r_{\text{WD}} \sim 10^{-5}/\text{yr}$); neutron stars ($m_{\text{NS}} = 1.4M_{\odot}$, $r_{\text{NS}} \sim 2 \times 10^{-6}/\text{yr}$); and black holes ($m_{\text{BH}} = 7M_{\odot}$, $r_{\text{BH}} \sim 10^{-6}/\text{yr}$). In this case, the net event rate R_{net} is dominated by low-mass WD inspirals, but black holes dominate the events seen by LISA. Using a LISA lifetime $T_L = 3$ yr and (based on Sigurdsson and Rees's estimate that the density of 10^6M_{\odot} holes at their cores is around the density of spirals, since spirals have low mass and ellipticals high mass supermassive holes [15]) $\rho_g \sim 0.003/\text{Mpc}^3$, we find

$$\left(\frac{S}{N}\right)_{\text{no,F}} = 2.67 \left(\frac{\ln(2)}{\ln(1/P_{\text{no}})}\right)^{1/3} \times \left(\frac{S}{N}\right)_A(10M_{\odot}, 1\text{Gpc}).$$

- *Sigurdsson and Rees* (SR) They consider two types of galaxies—spirals and dwarf ellipticals—but only the latter leads to significant event rates. In that case, they use the following masses and rates for the three species: WD ($m_{\text{WD}} = 0.6M_{\odot}$, $r_{\text{WD}} \sim 3 \times 10^{-8}/\text{yr}$); NS ($m_{\text{NS}} = 1.4M_{\odot}$, $r_{\text{NS}} \sim 10^{-7}/\text{yr}$); and BH ($m_{\text{BH}} = 5M_{\odot}$, $r_{\text{BH}} \sim 10^{-6}/\text{yr}$). (For black holes, these authors provide only an off-the-cuff estimate; we have taken some liberty in interpreting it, choosing a mildly optimistic characteristic black hole mass.) Again using $T_L = 3$ yr and $\rho_g = 0.003/\text{Mpc}^3$,

we find

$$\left(\frac{S}{N}\right)_{\text{no,SR}} = 1.90 \left(\frac{\ln(2)}{\ln(1/P_{\text{no}})}\right)^{1/3} \times \left(\frac{S}{N}\right)_A (10M_{\odot}, 1\text{Gpc}).$$

In performing both calculations, we use Sigurdsson and Rees’s estimate that the density of galaxies with (spirals have low-mass holes; ellipticals and others tend to have more), so $\rho_g \sim 0.003/\text{Mpc}^3$. Also, we use a LISA lifetime $T_L = 3 \text{ yr}$.

In sum, we suspect that on astrophysical grounds we will have a 50% chance of seeing no source with S/N roughly greater than

$$\left(\frac{S}{N}\right)_{\text{no,guess}} \approx 2.5 \times \left(\frac{S}{N}\right)_A, \quad (2.34)$$

where the $(S/N)_A$ will chosen to be the most reasonable S/N over all orbital parameters (e) and black hole spins (a), given the fiducial parameters $d = 1\text{Gpc}$, $m = 10M_{\odot}$, and $M = 10^6M_{\odot}$.

2.4 Schwarzschild supermassive black hole (SMBH)

To illustrate this scheme in a case where all terms are algebraically tractable, we discuss the range of probable transition durations when the capturing hole has no angular momentum (Schwarzschild).

2.4.1 Choosing parameters

Rather than using E, L to characterize the orbit, when the orbit is confined in radius between two turning points (i.e., when it is bound), it is far simpler to characterize the potential $V = -(dr/d\tau)^2$ by the location of its 3 roots, r_{\pm}, \bar{r} , where r_{\pm} are the inner and outer turning points of the bound orbit and \bar{r} is the innermost root:

$$\begin{aligned} \left(\frac{dr}{d\tau}\right)^2 &= -V = \frac{1-E^2}{r^3} (r_+ - r)(r - r_-)(r - \bar{r}) \\ &= E^2 - \left(1 - \frac{2}{r}\right) \left(1 + \frac{L^2}{r^2}\right) \end{aligned} \quad (2.35)$$

Since we have only two free parameters, the three roots are not independent; they satisfy a self-consistency polynomial. For this reason, we introduce p, e —parameters analogous to semi-latus rectum and eccentricity in classical mechanics. Employing a consistency relation [generally Eq. (A.9) of Appendix A.2] to set \bar{r} , we find

$$r_{\pm} \equiv \frac{p}{1 \mp e}, \quad \bar{r} = \frac{2p}{p-4}, \quad (2.36)$$

$$E^2 = \frac{(p-2-2e)(p-2+2e)}{p(p-3-e^2)}, \quad L^2 = \frac{p^2 M^2}{p-3-e^2}. \quad (2.37)$$

These p, e parameters have the notable advantage that bound orbits (orbits that cannot escape to infinity) and non-plunging orbits (orbits which avoid the central black hole) are easy to describe: bound orbits have $e \in [0, 1]$, while non-plunging orbits have $0 < r_- - \bar{r} = p(p-6-2e)/[(1+e)(p-4)]$, or

$$z = p - 6 - 2e > 0. \quad (2.38)$$

As one approaches the transition, the maximum of the potential V decreases, r_- approaches \bar{r} , and z approaches 0. We can equivalently specify the location r of a transition by only one of p or e , with the other determined by $z = 0$. I will typically use e . For example, a transition of eccentricity e occurs at radius $r = r_- = \bar{r} = 2(3+e)/(1+e)$.

The parameters p, e used here are identical to those used the Teukolsky-equation-based inspiral literature [10, 11]. For example, the above discussion mirrors that in Cutler, Kennefick, and Poisson [10] between their equations (2.4) and (2.8), with the change of notation $r_1 \rightarrow r_-, r_2 \rightarrow r_+$ and $r_3 \rightarrow \bar{r}$.

2.4.2 Dependence of transition parameters on eccentricity

We know the potential [Eq. (2.35)]; hence we find that when $r_- = \bar{r}$ (=the transition radius) we have

$$V'' = -(1-E^2)2\frac{r_+ - r_-}{r_-^3} \quad (2.39)$$

and therefore, substituting r_{\pm} and E from Eqs. (2.36), (2.37) into $\tau_o = (V''/2)^{-1/2}$ we find

$$\tau_o = (3+e)\sqrt{\frac{2(9-e^2)}{e(1+e)^3}}. \quad (2.40)$$

Similarly, substituting $r = r_-$ into $\gamma = -g^{tt}E$ gives us the characteristic time required to make the transition:

$$\gamma\tau_o = \frac{2(3+e)^2}{\sqrt{e(1+e)^3}} \quad (2.41)$$

Since $d\phi/dt = g^{\phi\phi}L/(-g^{tt}E) = L/r^2(1-2/r)E$, we find

$$\left(\frac{d}{dr}\frac{d\phi}{dt}\right)_{r=r_{\max}} = \frac{2(r_{\max}-3)}{r_{\max}^{5/2}(r_{\max}-2)}, \quad (2.42)$$

when E, L are consistent with a circular orbit at $r = r_{\max}$. Using the above and Eq. (2.23), we

conclude that for moderate eccentricity the natural “transition extent” δr_{ref} is

$$\delta r_{\text{ref}} \approx \pi \frac{\sqrt{2e(3+e)}}{(1+e)(3-e)}. \quad (2.43)$$

This scale naturally varies with the scale of the potential (namely, $\delta r_{\text{ref}} \propto 1/\sqrt{V''}$) as $e \rightarrow 0$.

Further straightforward computations using the potential [Eq. (2.35)] reveal how r_{max} varies due to loss of E and L via wave emission when the particle is nearly in a circular orbit (so $dE \approx \Omega dL$):

$$\frac{dr_{\text{max}}}{dt} = \frac{dL}{dt} 2 \frac{(r_{\text{max}} - 3)^{3/2}}{r_{\text{max}} - 6}.$$

We therefore can express $I_{\text{ad,min}}$ in terms of (tabulated) known radiation-reaction angular momentum fluxes $\langle dL/dt \rangle$:

$$I_{\text{ad,min}} = \left(\frac{dL}{dt} \right)^2 \frac{(3-e)^2(3+e)}{2e^2(1+e)}. \quad (2.44)$$

To obtain a rough approximation of $I_{\text{ad,min}}$, rather than use the true dL/dt appropriate to circular orbits, we approximate dL/dt by the Peters-Mathews expression [Eq. (2.19)].

Our scheme ceases to apply when the eccentricity is below e_{min} defined by $\Delta I = I_{\text{max}}(e_{\text{min}})$ [Eq. (2.15)]. In the special case of $a = 0$, where $p_s = 6 + 2e$, the definition of I_{max} [Eq. (A.11)] reduces to

$$I_{\text{max}} = \frac{32}{27} \frac{e^3}{(9-e^2)(1+e)}. \quad (2.45)$$

2.4.3 Transition duration

With all the necessary elements assembled, we can apply our program [Eqs. (2.9), (2.11), (2.17)] to estimate the distribution in number of orbital cycles $N_c \equiv T_c \Omega / 2\pi$ we expect when a particle spirals into a nonspinning black hole at some fixed, known eccentricity e .

The results for $\eta = 10^{-5}$ are shown in Fig. 2.3. When our adiabatic approximation applies, we find that to a good approximation (within around 1 cycle) most transitions should have duration close to the shortest transition duration $N_{c,\text{ad}} = T_{c,\text{ad}} \Omega / 2\pi \in [3, 5]$. In particular, within the region $e > e_{\text{ad}}$ where our adiabatic approximation applies, almost all transitions will last for less than the Ori-Thorne (OT) circular duration; most will last for substantially less. At low eccentricity, most transitions seem to approach a result somewhat different than the OT circular estimate. Since OT use a different convention for δr ³ and since significant changes could still occur in the fundamentally nonadiabatic region between $e = 0$ and $e = e_{\text{ad}}$, we do not find the discrepancy troubling.

In the above, we show results for only $\eta = 10^{-5}$ (say, for $m = 10M_{\odot}$ and $M = 10^6$). As the

³Since all results depend (mildly) on the convention for transition extent, and since the Ori-Thorne prediction implicitly employs a characteristic length $\delta r_{\text{ref}} \approx (\text{few}) \times R_o \propto \eta^{2/5}$ with R_o given by Ori-Thorne Eq. (3.20), while the “standard” predictions [Eqs. (2.9),(2.11)] use Eq. (2.23), with $\delta r_{\text{ref}} \propto \eta^0$, we cannot guarantee that the results should be precisely compatible.

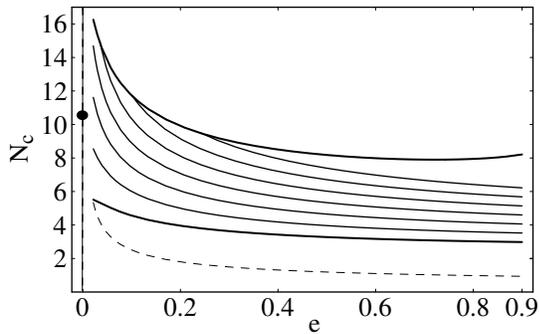


Figure 2.3: Plots of various predictions for the expected number of angular cycles (N_c) versus eccentricity (e) for a transition from inspiral to plunge with $\eta = 10^{-5}$. The top solid curve is the number of cycles when $T_c = T_{c,\text{ad}}$ [Eq. (2.17)], an estimate of the longest possible (adiabatic) transition duration. The bottom solid curve is the number of cycles when $T_c = T_{c,-}$ [Eq. (2.9)]. The 6 curves in between are the number of cycles when $T_c = T_{c,1}, \dots, T_{c,6}$ [Eq. (2.11)]; as only a fraction $\approx 10^{-1}, \dots, 10^{-6}$ of inspirals can have durations above these curves (respectively), they illustrate how few particles have durations significantly differing from $T_{c,-}$. The dot indicates the Ori-Thorne (OT) prediction for circular equatorial inspiral. The dashed curve is a characteristic-scale-based prediction based on $T_c = 4\gamma\tau_o$, used to illustrate the significance of the “logarithmic correction.” The plot starts at $e = e_{\text{ad}} \approx 0.0215$, at which point $\Delta I = I_{\text{max}}$.

variation of the duration with η is weak—we find $H(\eta_o = 10^{-5}) \in [0.1, 0.4)$ [Eq. (2.14)]—even substantially different test particle masses (e.g., $m \in [0.1, 30]M_\odot$ with $M = 10^6$) lead to results of the form above, scaled up or down by a factor $\lesssim 2$.

2.4.4 Prospects for LISA detecting a given transition

As discussed in Sec. 2.3, we can estimate the signal-to-noise ratio for a given transition using Eq. (2.27). For the standard case of a $10M_\odot$ particle falling into a $M = 10^6M_\odot$, application of that formula reveals no higher S/N than that predicted by Ori and Thorne; moreover, barring astrophysically unlikely masses, all transitions have too low a S/N to be detected. [See Fig. 2.7 below for details.] For example, if an inspiral of mass m into a 10^6M_\odot hole occurs at the fiducial distance 1 Gpc with $e = 1/3$, we have a 90% chance that $S/N \in [0.91, 1.01](m/10M_\odot)$.

The results for S/N can be well-approximated by way of Eq. (2.30) and a comparison with Ori and Thorne’s results for circular inspiral. (See Appendix A.3 for a summary of OT results). Specifically, using the fiducial case of $10M_\odot$ on 10^6M_\odot at 1 Gpc, for which we have $(S/N)_{\text{OT}} = 1.6$ and $N_{c,\text{OT}} = 10.5$, we find the general S/N for captures by a $M = 10^6M_\odot$ hole to be about

$$(S/N) \approx 1.6 \times \sqrt{\frac{N_c}{10.5}} \times \frac{m}{10M_\odot} \frac{1\text{Gpc}}{d}.$$

2.5 Kerr SMBH

The Kerr case follows similarly, save with an additional parameter (a).

2.5.1 Parameterizing orbits, potential

As before, it is simplest to characterize the potential by its three roots $r_{\pm} = p/(1 \mp e)$, \bar{r} :

$$\left(\frac{dr}{d\tau}\right)^2 = -V = \frac{1 - E^2}{r^3} (r_+ - r)(r - r_-)(r - \bar{r}). \quad (2.46)$$

And as before we can define $r_{\pm} = p/(1 \mp e)$; as before, we find a self-consistency relation $P(p, e, a, \bar{r})$ [Eq. (A.9)], permitting us to solve for $\bar{r}(p, e, a)$. As before, we can characterize the proximity to the last-stable surface by way of the separation between the two innermost roots ($r_- - \bar{r}$). Finally, as before, for each fixed black hole ($a = \text{const}$) and each particle exactly on the transition line from orbit to plunge, the particle can have $e \in [0, 1)$; its p will be constrained by the analogue of the Schwarzschild $p = 6 + 2e$: the self-consistency relation Eq. (A.10), which implicitly defines $p_s(e, a)$ such that $\bar{r}(p_s, e, a) = p_s/(1 + e)$.

2.5.2 Dependence of transition parameters on e, p

Since the potential has the same structure as before, the same general expression Eq. (2.39) applies, with E now determined by expressions in Appendix A.2. By explicitly differentiating the potential [Eq. (2.2)], inserting the definitions $r_{\pm} = p/(1 \mp e)$, and demanding the inner turning point is a maximum (so $\bar{r} = r_- = p/(1 + e)$), we find

$$V'' = -8e \frac{(1 + e)^3}{(3 - e)p_s^3} \quad (2.47)$$

and therefore know $\tau_o = (V''/2)^{-1/2}$ in terms of p, e at the transition.

The γ factor follows from inserting $r = p/(1 + e)$ into the usual expression for $dt/d\tau$:

$$\gamma \equiv \frac{dt}{d\tau} = -g^{tt}E + g^{t\phi}L, \quad (2.48)$$

where g^{tt} and $g^{t\phi}$ are known Kerr metric coefficients. Here, E and L are evaluated using the expressions (A.7) and (A.8) discussed in Appendix A.2, with $\bar{r} = p_s/(1 + e)$.

We obtain the transition extent δr_{ref} with the usual Eq. (2.23). This requires γ [Eq. (2.48), above], τ_o (also above), and $d(d\phi/dt)/dr$ [Eq. (2.20)].

Finally, as in the Schwarzschild case we estimate $I_{\text{ad, min}}$ [Eq. (2.16)] and thus $T_{c, \text{ad}}$ [Eq. (2.17)] via i) expressing $dr_{\text{max}}/d\tau$ in terms of dL/dt using explicit expressions for $r_{\text{max}}(E, L)$ and $dE = \Omega dL$,

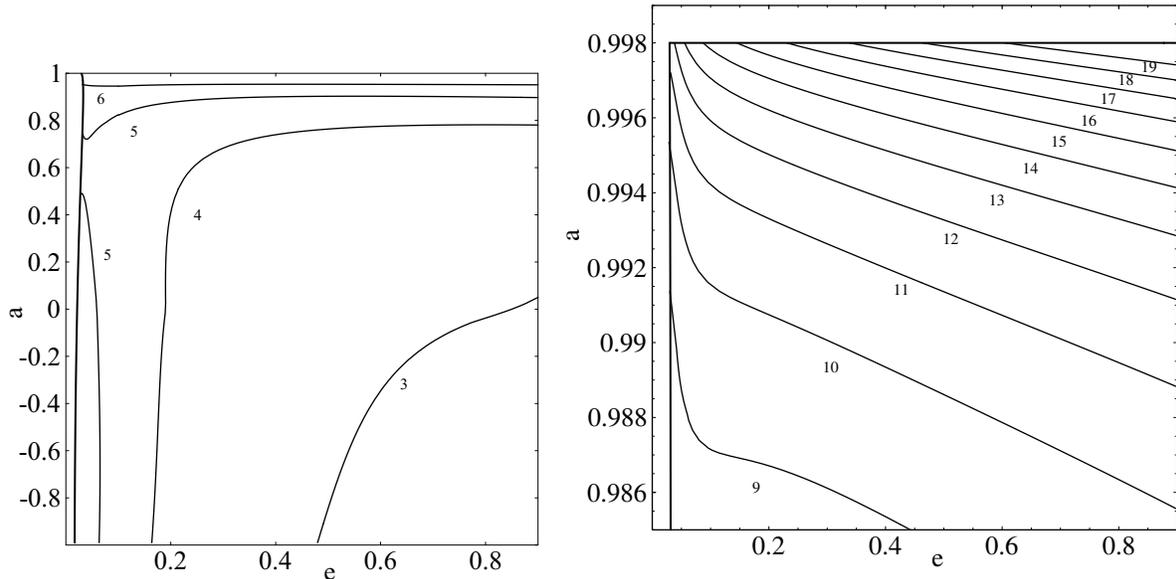


Figure 2.4: These two plots illustrate the shortest possible number of cycles $N_c = \Omega T_{c-}/2\pi$ [Eq. (2.9)] a transition could last, versus eccentricity (e) and black hole angular momentum (a) for the fiducial source ($10M_\odot$ into 10^6M_\odot). In both plots, contours are cut off, and bounding curves appear (shown heavy solid), when $\Delta I = I_{\text{ad},\text{min}}$ and when $a = 0.998$.

giving

$$\frac{dr_{\text{max}}}{dt} = \frac{dL}{dt} \frac{2r(2a\sqrt{r} + r^2 - 3r)^{3/2}}{(r^{3/2} + a)(r^2 - 6r + 8a\sqrt{r} - 3a^2)} \Big|_{r=r_{\text{max}}}, \quad (2.49)$$

(where we have used the orbital parameters E, L consistent with a circular orbit at $r = r_{\text{max}}$ [13]); then ii) using the Peters-Mathews expressions for dL/dt [Eq. (2.19)] to construct an approximate expression for dr_{max}/dt , which we then iii) insert in Eq. (2.16) to estimate the boundary between adiabatic and nonadiabatic transitions.

2.5.3 Transition duration

Combining these together, we can deduce the range of plausible transition durations for a test particle of eccentricity e falling into a hole of angular momentum a , measured as number of orbital cycles $N_c(e, a) = T_c\Omega/2\pi$. Plots of the number of cycles appropriate to $T_c = T_{c-}$ [Eq. (2.9)], to $T_c = T_{c,\text{ad}}$ [Eq. (2.17)], and to $T_c = T_{c,1}$ [Eq. (2.11)] appear in Figs. 2.4, 2.5, and 2.6, respectively.

These plots all assume a fiducial source ($10M_\odot$ on 10^6M_\odot). In these plots, we truncate the range of e, a allowed because, i) we need e larger than e_{ad} [Eq. (2.15)]; and ii) realistic astrophysical black holes have $a \leq 0.998$ [17]. Also, in these plots, we do not extend to $e \approx 1$ because we do not have data for I in this region, nor do we expect our estimate of ΔI [Fig. 2.2] to be reliable in this extreme.

At each a , we see behavior largely similar to the Schwarzschild results discussed in Sec. 2.4: i) almost all transitions take less time than the Ori-Thorne result for $e = 0$; ii) as we increase the

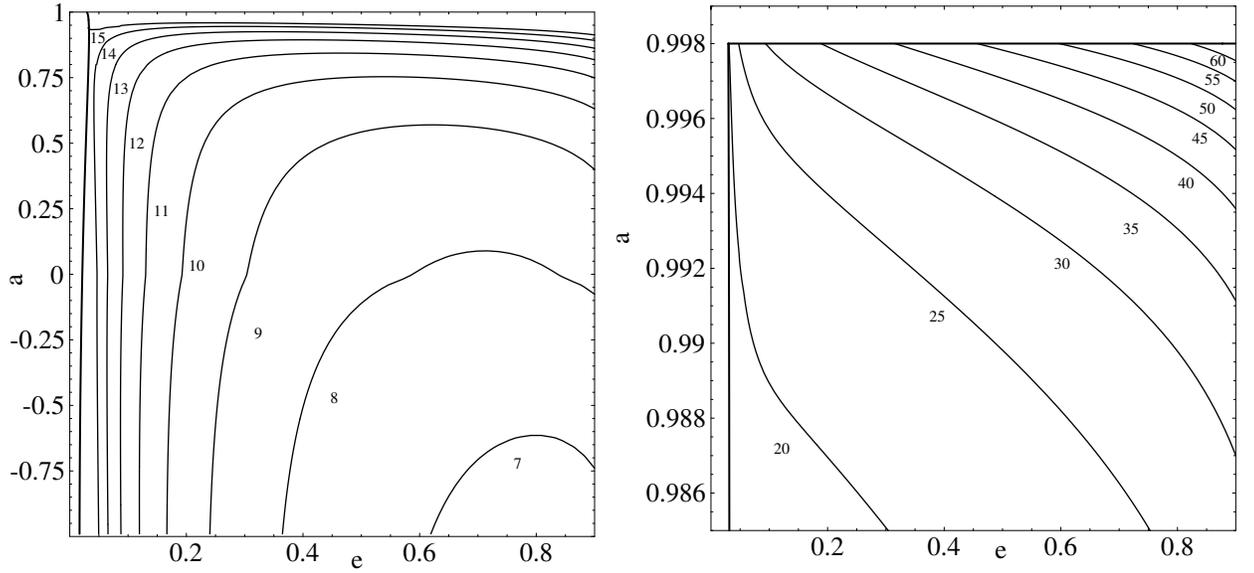


Figure 2.5: These two plots illustrate the longest possible number of cycles an (adiabatic) transition could last $N_c = \Omega T_{c,\text{ad}}/2\pi$ [Eq. (2.17)], versus eccentricity (e) and black hole angular momentum (a) for the fiducial source ($10M_\odot$ into 10^6M_\odot). Transitions of such long duration are extremely unlikely unless $e \approx e_{\text{ad}}$ [see Eq. (2.15), (2.10)]. In both plots, contours are cut off, and bounding curves appear (shown heavy solid), when $\Delta I = I_{\text{ad},\text{min}}$ and when $a = 0.998$.

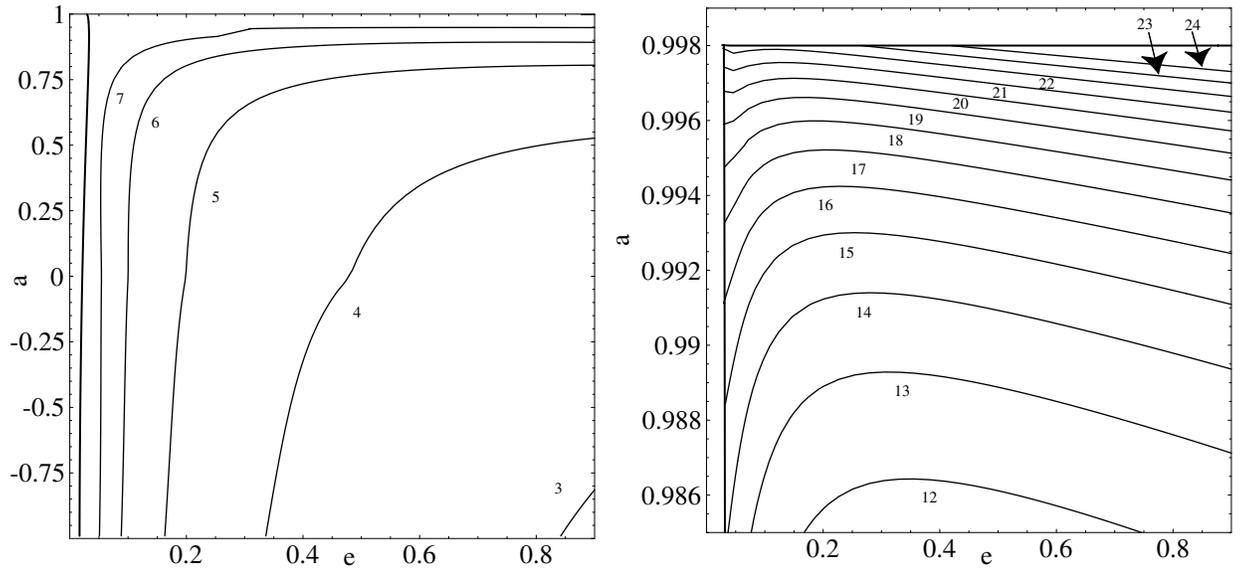


Figure 2.6: These two plots illustrate $N_c = \Omega T_{c,1}/2\pi$ versus e and a for the fiducial case. When the results of Fig. 2.5 are appreciably larger than these—that is, everywhere except near the left boundary $e = e_{\text{ad}} \sim 90\%$ of transitions with orbital parameters e, a will take less than the number of cycles shown in this plot to traverse the transition region.

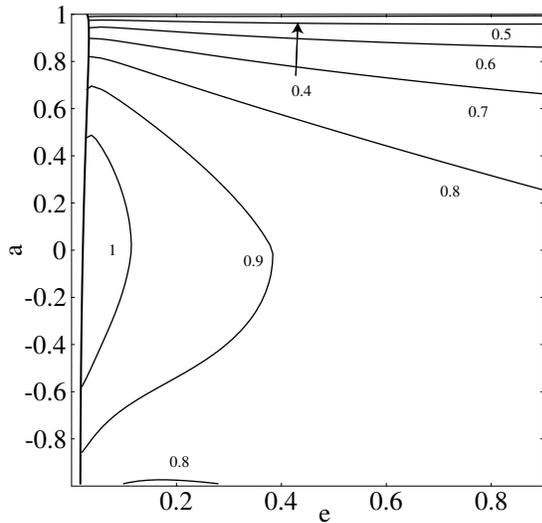


Figure 2.7: This plot illustrates S/N [Eq. (2.27)], given a fiducial source ($10M_{\odot}$ into 10^6M_{\odot} , at 1Gpc) given that the black hole has angular momentum a , the transition occurs at eccentricity e , and given the transition takes the shortest possible time (or $T_c = T_{c,-}$) [Eq. (2.9)]. As argued in the text, this time (and thus this S/N) will be to a good approximation characteristic of all transitions with those interaction parameters.

eccentricity, the transition duration decreases; and iii) since $T_{c,1} \approx T_{c,-}$ (compare Figs. 2.4 and 2.6), most transitions last close to the shortest-possible transition duration.

2.5.4 Prospects for LISA detecting a given transition

As in the Schwarzschild case, since eccentric usually transitions last for fewer angular cycles than their circular analogues, they are less detectable as well. Thus, in the fiducial case of captures of a $10M_{\odot}$ hole by a 10^6M_{\odot} hole, the data from Ori-Thorne Table II provides an upper bound on the S/N seen by LISA (shown in Fig. 2.7). Since this bound is small, we have little chance of seeing any given transition.

One should notice, however, that the distribution of S/N with orbital parameters is very flat and not much below 1. Therefore, only a modest improvement in LISA’s noise spectrum S_h could render most (measured by volume of parameter space) of the transitions detectable.

2.5.5 On probability of detection

Because LISA at present has so poor prospects for detecting the “fiducial” source ($m = 10M_{\odot}$ at 1Gpc), it has a poor chance of seeing any source at all. Even assuming all LISA sources had orbital parameters chosen to give the longest-plausible transition length (the OT circular inspiral duration, which has $S/N_A \lesssim 1.6$), by the estimate of Eq. (2.34) we expect we have a $\sim 50\%$ chance of no signal with $S/N \gtrsim 4$ being present in the datastream. With more realistic orbital parameters, we would

expect a $\sim 50\%$ chance of no signal $\gtrsim 2.3$. In other words, LISA has a good to excellent chance of *not* seeing any transitions from inspiral to plunge in its lifetime.

A modest improvement in LISA’s noise curve, however, would make a few circular (and to a lesser degree eccentric) transitions from inspiral to plunge detectable.

2.6 Summary

This paper has introduced a framework (depending on observational or other conventions) that extends the Ori-Thorne prediction for the transition duration from inspiral to plunge to include eccentric orbits. While the framework and applications contain many oversimplifications—most notably, the fit to $\Delta I(e, a)$ and the lack of a physically meaningful convention for δr_{ref} —the essential physics should be captured by Sec. 2.2.

This paper then applies that framework to probable LISA sources to suggest that, because an eccentric transition is generally only slightly briefer than a circular one, LISA should have only slightly worse prospects to resolve the transition from inspiral to plunge for eccentric orbits than for circular ones. While the prospects for detecting circular (and hence eccentric) transitions with LISA are not good, they are not necessarily bad: modest changes to the LISA noise floor could render a signal marginally detectable. Therefore, more detailed investigations could be of use.

Potentially, we could use other portions of orbits that pass close to the hole—for example, the previous few “bounces” off the inner portion of the radial potential—as probes of the strong-field metric. Analyzed separately (using the same ΔI framework) each of these “bounces” should provide in itself at best of order the same S/N as the transition. If the source has already been detected with good confidence, we should be able to coherently integrate over many such bounces and build up excellent S/N .

Finally, we could hope that eccentric inclined orbits might, by some happenstance of parameters, admit a regime of significantly longer transition times. The prospect seems unlikely, but the author may address it in a future paper.

Bibliography

- [1] F. D. Ryan, Phys. Rev. D. **56** (4) 1845
- [2] A. Buonanno, Y. Chen, M. Vallisneri, gr-qc/0205122
- [3] A. Buonanno and T. Damour, Phys. Rev. D **60**, 023517 (1999)
- [4] L.E. Kidder, C.M. Will, and A.G. Wiseman, Phys. Rev. D **47** 3281 (1993)
- [5] L. Lehner, CQG **18** (2001) R25
- [6] For LISA information, see <http://lisa.jpl.nasa.gov>
- [7] S. Detweiler and B. Whiting, gr-qc/0202086; Y. Mino, H. Nakano, and M Sasaki, gr-qc/0111074; and references therein.
- [8] A. Ori and K.S. Thorne, Phys. Rev. D. **62** 124022 (gr-qc/003032) Note that in their paper $N_{cyc} = 2N_c$ is the duration of the transition in quadrupolar gravitational wave cycles.
- [9] C. W. Misner, K. S. Thorne, and J. A. Wheeler, *Gravitation* (W. H. Freeman, San Francisco, 1973)
- [10] C. Cutler, D. Kennefick, and E. Poisson, Phys. Rev. D **50** 3816
- [11] K. Glampedakis and D. Kennefick, Phys. Rev. D **66** 044002 (2002)
- [12] P. C. Peters and J. Mathews, Phys. Rev. **131**, 435 (1963) The relevant equations, more transparently emphasized and written in terms of our parameters, may be found in K. Glampedakis, S. A. Hughes, and D. Kennefick, Phys. Rev. D **66** 064005 (2002)
- [13] J. M. Bardeen, W. H. Press, and S. A. Teukolsky, Astrophys. J. **178**, 347 (1972). Some of the relevant equations also appear in Ori and Thorne. Alternatively, one can use Eqs. (A.7) and (A.8) with $\bar{r} = p/(1 + e)$ and $p = p_s(e, a)$.
- [14] M. Freitag, astro-ph/0107193
- [15] S. Sigurdsson and M. J. Rees, MNRAS **284** 318

[16] J. Miralda-Escude and A. Gould, *Astrophys. J.* **545**, 847

[17] K. Thorne, *Astrophys. J.* **191**, 507 (1974)

Chapter 3

Reducing thermoelastic noise in advanced LIGO by flattening the light beams

[This chapter is closely adapted from a paper in preparation by E. d'Ambrosio, R. O'Shaughnessy, S. Strigin, K. Thorne, and S. Vyatchanin, *Reducing thermoelastic noise in gravitational-wave interferometers by flattening the light beams*. After LSC internal review, it will be submitted to Phys. Rev. D. Kip Thorne wrote the text of this paper.]

Abstract

In the baseline design for advanced LIGO interferometers, the most serious noise source is tiny, dynamically fluctuating bumps and valleys on the faces of the arm-cavity mirrors, caused by random flow of heat in the mirrors' sapphire substrates: so-called *thermoelastic noise*. We propose replacing the interferometers' baseline arm-cavity light beams, which have *Gaussian*-shaped intensity profiles that do not average very well over the dynamical bumps and valleys, by beams with *mesa*-shaped profiles that are flat in their central ~ 7 cm of radius, and that then fall toward zero as quickly as is allowed by diffraction in LIGO's 4 km arms; see Fig. 3.2. The mesa beams average the bumps and valleys much more effectively than the Gaussian beams. As a result, if the mirrors' substrate radii and thicknesses are held fixed at 15.7 cm and 13 cm, and the beam radii are adjusted so diffraction losses per bounce are about 10 ppm, replacing Gaussian beams by mesa beams reduces the thermoelastic noise power by about a factor 3. If other thermal noises are kept negligible, this reduction will permit advanced LIGO to beat the Standard Quantum Limit (circumvent the Heisenberg Uncertainty Principle for 40 kg mirrors) by about a factor 1.5 over a bandwidth about equal to frequency; optical (unified quantum) noise will become the dominant noise source; and the event rate for inspiraling neutron star binaries will increase by about a factor 2.5. The desired mesa beams can be produced from input, Gaussian-profile laser light, by changing the shapes of the arm

cavities’ mirror faces from their baseline *spherical* shapes (with radii of curvature of order 60 km) to “*Mexican-Hat*” (*MH*; sombrero-like) shapes that have a shallow bump in the center but are otherwise much flatter in the central 10 cm than the spherical mirrors, and then flare upward strongly in the outer 6 cm, like a sombrero; Fig. 3.3. In this paper we describe mesa beams and MH mirrors mathematically and we report the results of extensive modeling calculations, which show that the mesa-beam interferometers are *not* substantially more sensitive than the baseline Gaussian-beam interferometers to errors in the mirror figures, positions, and orientations. This has motivated the LIGO Scientific Community (LSC) to adopt MH mirrors and mesa beams as an option for advanced LIGO, to be studied further. The details of our modeling calculations are presented in companion papers.

3.1 Introduction and summary

The Laser Interferometer Gravitational-Wave Observatory (LIGO) is designed to support successive generations of interferometric gravitational-wave detectors. LIGO’s first interferometers are now in operation [9], and the (negative) results of its first gravitational-wave searches have recently been submitted for publication [10]. When they reach their design sensitivity (presumably next year), LIGO’s initial interferometers, together with their international partners, will reach out into the universe to distances where it is plausible, but not probable to detect gravitational waves [11]. After a planned upgrade to *advanced LIGO interferometers* (planned to begin in 2007), wave detection will be quite probable [11]. A baseline design for the advanced LIGO interferometers has recently been adopted [20], along with several options, not currently in the baseline, that merit further study and might be incorporated at a future date. This paper describes one of these options, which has been much discussed within the LIGO Scientific Community (LSC) but has not previously been presented in the published literature: the reshaping of the arm-cavity light beams so as to reduce thermoelastic noise.

3.1.1 The context: noise in advanced LIGO interferometers

For advanced LIGO’s baseline design [20], the dominant noise sources in the most interesting frequency range (above about 20 Hz) are thermoelastic noise and optical noise (also called “unified quantum noise”). Other thermal noises (most especially coating thermal noise) might, in the end, be important; but in this paper we shall assume them negligible and shall focus on the thermoelastic noise and optical noise.

In Fig. 3.1 we show the thermoelastic noise [5, 18], the optical noise [6, 7, 8], their sum (labeled total noise), and the standard quantum limit (SQL) for the advanced LIGO baseline design with sapphire mirrors [20]. This figure suggests (as is well known [6]) that, if the thermoelastic noise can

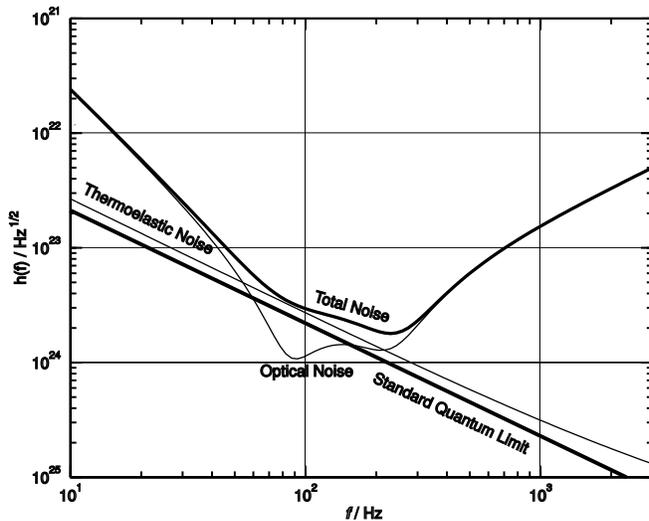


Figure 3.1: Noise curves for advanced LIGO.

be reduced significantly (and if other thermal noises can be kept negligible), then the advanced LIGO interferometers will be able to beat the SQL, and the interferometers' ranges (detectable distances) for astrophysical sources will be increased significantly.

In this paper we propose a method (“flattening the interferometers' light beams”) for reducing the thermoelastic noise, we evaluate the resulting increased range for neutron-star / neutron-star (NS/NS) binaries, and we explore practical issues related to our proposal. We have previously discussed our proposal, the increased NS/NS range, and the practical issues in presentations at meetings of the LIGO Scientific Collaboration [14, 2, 1, 4] and in an internal LIGO document [3].

3.1.2 The Physical Nature of Thermoelastic Noise; Motivation for Reshaping Beams

Our proposal is motivated by the physical nature of thermoelastic noise. This noise is created by the stochastic flow of heat (random motions of thermal phonons) within each test mass (mirror), which produces stochastically fluctuating hot spots and cold spots inside the test mass. The test-mass material (sapphire for the baseline design of advanced LIGO) expands in the hot spots and contracts in the cold spots, creating fluctuating bumps and valleys on the test-mass (mirror) faces. These face bumps influence the light beam's measurement of the test masses' positions: the interferometer's output phase shift is proportional to the difference of the test masses' average positions — with the average being the position of a mirrored test-mass face, weighted by the light's energy flux (its intensity distribution).

If the intensity distribution is “flat” (nearly constant) in most regions of high intensity, then the adjacent valleys and bumps (having been created by heat flow from one to the other) will average

out, giving low net thermoelastic noise. If, instead, the energy flux is sharply changing in most regions of high flux, then the adjacent valleys and bumps will not average well and the thermoelastic noise will be high. Also, the larger is the light beam, the better will be the averaging and thus the lower will be the noise.

These considerations suggest that large-radius, flat-topped beams with steep edges (e.g. the thick curve in Fig. 3.2 below) will lead to much smaller thermoelastic noise than small-radius, centrally peaked beams with gradually sloping sides (e.g., the thin, Gaussian curve in Fig. 3.2 below).

For Gaussian beams, the influence of beam radius r_o has been quantified by Braginsky, Gorodetsky and Vyatchanin [5] (who first pointed out the importance of thermoelastic noise for sapphire test masses): the thermoelastic noise power scales as $S_h^{\text{TE}} \propto 1/r_o^3$ (aside from small corrections due to the test masses' finite sizes [18]). This has motivated the baseline design for advanced LIGO interferometers with sapphire test masses: the beam radius r_o is chosen as large as possible, given the demand for small diffraction losses, $\mathcal{L}_0 \lesssim 10$ ppm per bounce in the interferometer's arm cavities.¹

The baseline design uses light beams with a Gaussian distribution of energy flux, since such beams are eigenfunctions of cavities with spherical mirrors, and spherical mirrors are a standard, well-developed technology. However, the Gaussian energy flux is far from flat: most of the energy is in regions where the flux is rapidly varying with radius (thin curve in Fig. 3.2 below), and correspondingly the thermoelastic noise is substantially larger than it would be with “flat-topped” beams (thick curve in Fig. 3.2). This has motivated a (previously unpublished) proposal by O’Shaughnessy and Thorne [14] to replace the Gaussian beams with flat-topped beams, while keeping the beam radius as large as is compatible with diffraction losses $\mathcal{L}_0 \lesssim 10$ ppm.

3.1.3 Summary of analysis and results

In Sec. 3.2 we construct an example of a flat-topped light beam — a flat-topped TEM00 mode of light that will resonate in an interferometer’s arm cavity, if the test-mass mirror faces are shaped appropriately. Because the intensity distribution of our flat-topped beam resembles a mesa in the southwest American desert, we call it a *mesa beam*² (a name suggested to us by Phil Willems). To produce this mesa beam as an eigenmode of a symmetric arm cavity one must give the mirror faces a shape, with a central bump and an upturned brim, that resembles a Mexican hat (or sombrero) (Fig. 3.3), so we we call the mirrors *Mexican-hat (MH) mirrors*.

We have not optimized our mesa beams’ intensity distribution so as to bring the thermoelastic noise to the lowest value possible, but in Sec. 3.2 we argue that our chosen mesa beams are likely to be close to optimal.

¹The 10 ppm is dictated by the following considerations: For the baseline design there is 125 W of input power to the interferometer and 830 kW of circulating power in each arm cavity. Ten ppm of diffraction loss per bounce results in a diffraction power loss in the arm cavities of $4 \times 10\text{ppm} \times 830\text{kW} = 33$ W, which is 25 per cent of the 125 W of input light, a reasonable value.

²It is also called a *flat-topped beam* and a *mexican-hat* or *MH* beam in the internal LIGO literature [14, 2, 1, 4, 3]

Figure 3.2 below shows the intensity distribution for our proposed mesa beam (thick curve) and compares it with the intensity distribution of the baseline Gaussian beam (thin curve), which has the same diffraction losses. Figure 3.3 below compares the mirror shapes that support these mesa and Gaussian beams as eigenmodes of a 4 km LIGO arm. In their inner 10 centimeters of radius, the MH mirrors that support mesa beams are much flatter than the spherical mirrors that support Gaussian beams, but in their outer 6 centimeters (the upturned brim region), the MH mirrors are far more curved.

As we shall see, this greater curvature at large radii compensates considerably for the flatter shape at small radii, enabling a mesa-beam interferometer to exhibit only modestly worse parasitic-mode behavior than a Gaussian-beam interferometer, and only modestly worse sensitivity to mirror tilts, displacements, and figure errors.

Three of us (O’Shaughnessy, Strigin and Vyatchanin; OSV [19]) have computed the substantial reductions in thermoelastic noise that can be achieved in advanced LIGO by replacing the baseline spherical mirrors and their Gaussian beams with MH mirrors and their mesa beams. The method of computation and the results are described in Sec. 3.3. Our principal conclusion is this:

1. By switching from the baseline (BL) spherical mirrors to MH mirrors with the same cylindrical test-mass diameters and thicknesses and the same 10-ppm-per-bounce diffraction losses as the BL, one can reduce the power spectral density of thermoelastic noise by a factor 0.34 and increase the event rate for compact-binary inspirals by a factor 2.6. Larger improvements could be achieved by using conical test masses with enlarged inner faces.

One might worry that the greater flatness of the MH mirrors, in the inner 10 cm where most of the light resides, will make mesa-beam interferometers much more sensitive to errors in the orientations, positions, and figures of the mirrors. We have explored this issue in great depth, with the conclusion that *mesa-beam interferometers are **not** substantially more sensitive to mirror errors than the BL Gaussian-beam interferometers*. Details of our explorations are given in companion papers by d’Ambrosio [12] and by O’Shaughnessy, Strigin and Vyatchanin [19], and our methods and conclusions are presented and discussed in Sec. 3.4 of this paper. Our quantitative conclusions, in brief, are these:

2. Among those parasitic optical modes of a perfect arm cavity, that are not strongly damped by diffraction losses, the parasite closest in frequency to the desired TEM00 mode is separated from it by 0.099 of the free spectral range in the BL Gaussian-beam case, and by 0.0404 of the free spectral range in the mesa-beam case; see Sec. 3.4.3. This factor ~ 2 smaller mode spacing leads to a modestly greater sensitivity of the mesa-beam interferometer to mirror errors.
3. The interferometer’s arm cavities are about four times more sensitive to mirror tilt when MH mirrors are used than for the BL spherical mirrors. When all four cavity mirrors are

tilted through angles θ about uncorrelated axes, the fractions of the carrier power driven into (dipolar) parasitic modes inside the arm cavities, and driven out the dark port, are about $0.001(\theta/0.01 \mu\text{rad})^2$ and $0.002(\theta/0.01 \mu\text{rad})^2$, respectively for MH mirrors; and $0.001(\theta/0.035 \mu\text{rad})^2$ and $0.002(\theta/0.035 \mu\text{rad})^2$ for BL mirrors; Sec. 3.4.4. This factor four greater sensitivity is not a serious issue, since it turns out that the strictest controls on mirror tilt come from the signal recycling cavity (and, if a heterodyne output were to be used, from the power recycling cavity), and not from the arm cavities; see below.

4. The sensitivity of the interferometer's arm cavities to transverse displacements of the end test-mass mirrors (ETMs) is nearly the same for MH mirrors as for the BL spherical mirrors. For uncorrelated displacements of the two ETMs through distances s , the fractions of the carrier power driven into (dipolar) parasitic modes inside the arm cavities, and driven out the interferometer's dark port, are about $100(s/1 \text{ mm})^2$ ppm and $200(s/1 \text{ mm})^2$ ppm, respectively, for MH mirrors; and $100(s/1.3 \text{ mm})^2$ ppm and $200(s/1.3 \text{ mm})^2$ ppm for BL spherical mirrors. For details, see Sec. 3.4.5.
5. For MH mirror figure errors with peak-to-valley height variations Δz in the innermost 10 cm by radius: after the control system has optimized the mirror tilts, the fractions of the carrier power driven into parasitic modes inside the arm cavities, and driven out the dark port, are about $0.0008(\Delta z/6 \text{ nm})^2$ and $0.0015(\Delta z/6 \text{ nm})^2$, respectively; Sec. 3.4.6. We do not know the corresponding constraints for BL spherical mirrors, but the measured mirror figure errors in the initial LIGO interferometers are of order 1 or 2 nm, which suggests that the MH arm cavities' required figure errors may be achievable. This is currently being explored.
6. The most serious constraints on mirror tilt and on mirror figure accuracy come not from the arm cavities but rather from the signal recycling (SR) cavity. The SR cavity and power recycling (PR) cavity operate approximately in the geometric optics regime and thus are nearly insensitive to whether one uses MH or spherical mirrors; Sec. 3.4.8. As a result, by switching from spherical to MH mirrors, one pays only a small penalty, in terms of mirror tilt constraints and figure-error constraints.
7. More specifically, the most severe constraints on tilt and figure error arise from the driving of signal power into parasitic modes when the signal light passes through the SR cavity. To keep the resulting increase in shot noise below one per cent in the standard wideband advanced LIGO interferometers, it is necessary to constrain the magnitude θ of the vectorial tilts of the input test-mass mirrors (ITM's) and signal recycling mirror (SRM) to $\theta_{\text{WB}}^{\text{BL}} \lesssim 0.024 \mu\text{rad}$ (for the baseline spherical mirrors) and $\theta_{\text{WB}}^{\text{MH}} \lesssim 0.016 \mu\text{rad}$ (for MH mirrors). For the third advanced interferometer, narrowbanded at $f \simeq 500 \text{ Hz}$ or $\simeq 1000 \text{ Hz}$, the constraint must be

tighter: $\theta_{\text{NB}}^{\text{BL}} \lesssim 0.011 \mu\text{rad}$, and $\theta_{\text{NB}}^{\text{MH}} \lesssim 0.007 \mu\text{rad}$. These are approximately the same as the constraints on LIGO-I tilt arising from the PR cavity, in the absence of an output mode cleaner. If there were no output mode cleaner in advanced LIGO and heterodyne readout were used in place of the baseline homodyne readout, then the constraint on tilts in the PR cavity (due to mode mixing for the RF sideband light used in the readout) would be about the same as that for wideband interferometers in the SR cavity. For the BL homodyne readout, no such PR constraint arises. The increase in shot noise scales as θ^2 ; and we estimate that our constraints are inaccurate by a factor $\lesssim 2$ due to ignoring correlations in the overlaps of certain parasitic modes, and for the narrowbanded interferometers, due to inaccuracy of the geometric optics approximation in the SR cavity. For details of all these issues, see Sec. 3.4.9.

8. We characterize the analogous constraints on mirror figure error by the *peak-to-valley* fluctuations in the mirror height in the central regions of the mirrors (regions enclosing 95 per cent of the light power; radius $\simeq 10$ cm for MH mirrors and $\simeq 8$ cm for baseline spherical mirrors), with the fluctuations averaged over ~ 3 cm (an averaging produced by breakdown of geometric optics in the SR cavity). Our estimated constraints for one per cent increase of shot noise are $\Delta z_{\text{WB}} \lesssim 2.0$ nm for wideband advanced LIGO interferometers and $\Delta z_{\text{NB}} \lesssim 1.0$ nm for narrowband, independently of whether the mirrors are MH or BL spherical—though the region over which the constraints must be applied is different, 10 cm radius for MH and 8 cm for BL. The increase in shot noise scales as Δz^2 , and our estimated constraints might be inaccurate by as much as a factor ~ 3 due to exploring only one representative shape for the figure errors, due to overlaps of certain parasitic modes, and for the narrowbanded interferometer due to inaccuracy of the geometric optics approximation in the SR cavity. These are approximately the same constraints as arise (in our calculations) from the PR cavity in LIGO-I, in the absence of an output mode cleaner. If there were no output mode cleaner in advanced LIGO and heterodyne readout were used, then the constraint on tilts in the PR cavity (due to mode mixing for the RF sideband light used in the readout) would be about the same as that for wideband interferometers in the SR cavity. For details of these conclusions, see Sec. 3.4.10.

Among all the constraints on mirror errors that arise from our modeling, the most serious are the last ones: SR-cavity-induced constraints on mirror figure errors to avoid a one per cent increase in shot noise. These constraints are nearly independent of whether the mirrors are spherical or MH. These constraints would be relaxed if the SR cavity were made less degenerate. This could be achieved by shaping the fronts of the ITMs as lenses that bring the light (Gaussian or mesa) to a focus somewhere near the SR mirror — and also near the PR mirror.

Because MH mirrors and their mesa beams produce such a great (factor 3) reduction of ther-

moelastic noise power, and they increase the sensitivity to mirror errors by only modest amounts, they have been adopted as options for advanced LIGO, and they may be of value for LIGO's future international-partner interferometers. In Sec. 3.5 we describe some of the future research that is needed in order to firm up our understanding of the pros and cons of MH mirrors and mesa beams.

3.1.4 Notation

We here summarize some of the notation used in the remainder of this paper. The numerical values are for advanced LIGO interferometers, including sapphire test-mass substrates, with the sapphire idealized as isotropic (its properties averaged over directions).

b: Diffraction lengthscale $b = \sqrt{\lambda L/2\pi} = 2.603$ cm for

light in the $L = 4$ km LIGO beam tubes; equal to a symmetric Gaussian beam's minimum possible radius at the end mirrors

(the radius at which the power flux has dropped by $1/e$ of its central value).

C_V : Specific heat of test-mass substrate per unit mass at constant volume [7.9×10^6 cm² s⁻² K⁻¹]

D: Mesa beam radius; Eq. (3.3)

E: Young's modulus of test-mass substrate [4×10^{12} g s⁻² cm⁻²]

ETM: End test mass of an arm cavity

H: Thickness of test mass

f: Gravitational-wave frequency at which noise is evaluated

\mathcal{F} : Finesse of an optical cavity

I: Noise integral for a test mass, Eq. (3.15)

ITM: Input test mass of an arm cavity

k: Wave number, equal to $2\pi/\lambda$

k_B : Boltzmann's constant [1.38×10^{-16} g cm s⁻² K⁻¹]

L: Interferometer arm length [4×10^5 cm]

\mathcal{L} : Diffraction loss in a single reflection off a mirror

M: Mass of test mass [4×10^4 g]

r Radius in transverse plane

- r_o : Radius of Gaussian light beam on test-mass face (radius at which the intensity has dropped by a factor $1/e$ from its central value) [$r_o = 4.23$ cm = $1.63b$ for baseline design]
- P_n : Fraction of interferometer's light power in mode n
- R_p : The physical radius of a test mass [15.7 cm]
- R : The radius of the mirror coated onto a test mass [equal to R_p or $R_p - 8$ mm]; also, the power reflectivity of a mirror
- s : Transverse displacement of an arm cavity's ETM
- $S_h(f)$: Spectral density of noise (thermoelastic or other) for detecting a gravitational wave h with optimal direction and polarization
- T : Temperature of test-mass substrate [300 K]
- u : Electric field of some light mode or superposition of modes (renormalized to unit norm, $\int |u|^2 d\text{Area} = 1$); usually evaluated at the transverse plane tangent to an ITM mirror face, with the light propagating away from the ITM. Subscripts identify the mode.
- U : Unnormalized electric field of some light mode.
- v : Same as u : Unit-normed electric field of some light mode or superposition of modes.
- Δz : The peak-to-valley mirror deformation (mirror figure error) in the central region
- α_l : Substrate's coefficient of linear thermal expansion
[5.0×10^{-6} K $^{-1}$]
- $\alpha_{1,2}$: Amplitude of excitation of an arm cavity's parasitic mode $u_{1,2}$ by a tilt of the cavity's ETM; Eq. (3.21)
- β_1 : Amplitude of excitation of the parasitic mode v_1 by mirror figure errors; Eq. (3.34)
- γ_0 : Overlap of arm cavity's fundamental mode u_0 with Gaussian mode u_d that drives it; Eq. (3.20)
- δ_ℓ : Fraction of the light power of some perturbed field u'_ℓ that is in parasitic modes; Eq. (3.47)
- κ : Thermal conductivity of test-mass substrate [33 W m $^{-1}$ K $^{-1}$]
- λ : Wavelength of laser light [1.064 μm]; also, in Sec. 3.4.9.3, a function appearing in the analysis of the signal recycling cavity.
- ρ : Density of test-mass substrate [4 g cm $^{-3}$]; also, amplitude reflectivity of signal recycling mirror

σ : Poisson ratio of test-mass substrate [0.23]

θ : angle of mirror tilt

Θ : expansion (fractional volume change) of substrate

$\omega = 2\pi f$: Angular frequency, corresponding to the frequency

f at which the noise S_h is evaluated

$\zeta_{1,2}$: Amplitude of excitation of an arm cavity's parasitic mode by transverse displacement of the ETM; Eq. (3.29)

3.2 Mexican-mat mirrors and the mesa modes they support

In this paper we study a specific variant of a mesa light beam and the MH mirrors that support it. We believe this variant to be near optimal for reduction of thermoelastic noise, but we have not carried out the (rather complex) analysis required to prove optimality.

3.2.1 Mesa fields

The flat-topped (mesa-shaped) eigenmode of an interferometer arm cavity, which we seek to construct, must have an intensity distribution that is nearly flat across most of the light beam, and that then falls as rapidly as possible (constrained by diffraction effects) at the beam's edges. Moreover, if (as in baseline advanced LIGO) the cavity's input test mass (ITM) and end test mass (ETM) have the same physical dimensions, then to minimize the thermoelastic noise at fixed net diffraction loss, the beam should be symmetric about the arm cavity's mid plane, so its beam radii D are the same on the two mirrors. Otherwise [since $S_h^{\text{TE}} \propto 1/D^3$ approximately, and diffraction losses increase exponentially rapidly with increasing D ; Eq. (B.8)], the mirror with the reduced beam radius and smaller diffraction loss will have its thermoelastic noise power increased, while that with the enlarged beam radius and larger diffraction loss will have its noise power decreased more modestly, leading to a net noise increase.

The fastest possible falloff, for light in an optical cavity of length L , is that on the edge of the *minimal Gaussian beam* — the Gaussian beam whose radius increases by a factor $\sqrt{2}$ in going from the beam waist (at the cavity's center plane) to the cavity's end mirrors. This minimal Gaussian, at the mirror planes, has the following (unnormalized) form

$$U_{\text{min Gauss}}(r) = \exp \left[\frac{-r^2(1+i)}{2b^2} \right], \quad (3.1)$$

where

$$b = \sqrt{L/k} = \sqrt{\lambda L/2\pi} = 2.603 \text{ cm} , \quad (3.2)$$

with $L = 4 \text{ km}$ the cavity length, $k = 2\pi/\lambda$ the wave number, and $\lambda = 1.064\mu\text{m}$ the wavelength of the light.

The minimal Gaussian's beam radius, $r_o = b = 2.603 \text{ cm}$, is somewhat smaller than the advanced LIGO baseline beam radius $r_{o \text{ BL}} = 4.23$. Correspondingly, the baseline Gaussian falls off much more gradually with radius than diffraction effects require, and produces much larger thermoelastic noise than necessary.

To produce a near-optimal flat-topped eigenmode, with near-minimal thermoelastic noise, we superpose minimal-Gaussian fields, with a field density that is constant out to a radius $r = D$ and then stops abruptly. More specifically, our chosen unnormalized eigenmode has the following form:

$$U(D, r) = \int_{\mathcal{C}_D} \exp \left[\frac{-[(x - x_o)^2 + (y - y_o)^2][1 + i]}{2b^2} \right] dx_o dy_o , \quad (3.3)$$

where $r \equiv \sqrt{x^2 + y^2}$ and the integration is over a circle \mathcal{C}_D of radius D : $\sqrt{x_o^2 + y_o^2} < D$.

By carrying out the y_o integral in Cartesian coordinates, with $y = 0$ and $x = r$, we obtain the following expression for $U(D, r)$, which we have used in much of our numerical work:

$$U(D, r) = b \sqrt{\frac{-2\pi}{1+i}} \int_{-D}^{+D} dx_o \exp \left[\frac{-(x_o - r)^2(1+i)}{2b^2} \right] \times \text{erfi} \left[\frac{\sqrt{D^2 - x_o^2}}{b} \sqrt{\frac{-(1+i)}{2}} \right] . \quad (3.4)$$

Here $\text{erfi}(z) = \text{erf}(iz)/i$ is the imaginary error function.

By converting to circular polar coordinates and performing the angular integral, we obtain the following simpler expression for our unnormalized eigenmode

$$U(D, r) = 2\pi \int_0^D \exp \left[\frac{-(r^2 + r_o^2)(1+i)}{2b^2} \right] \times I_0 \left[\frac{rr_o(1+i)}{b^2} \right] r_o dr_o . \quad (3.5)$$

Here I_0 is the modified Bessel function of order zero. Modes with other weightings of the minimal-radius Gaussians can be obtained by inserting a weighting function $f(r_o)$ into the integrands of Eqs. (3.3) and (3.5).

In Appendix B.1 we give some approximate formulae for $U(D, r)$ valid at large radii. These are

useful for quick, clipping-approximation computations of diffraction losses.

The squared norm of $U(D, r)$ (the area integral of its squared modulus) is given by the following approximate formula, which is accurate to within a fraction of a per cent for $3.0 \lesssim D/b \lesssim 6.0$ (the regime of interest to us):

$$N^2(D) \equiv \int_0^\infty |U(D, r)|^2 2\pi r dr = 4.66 - 50.58D + 62.10D^2. \quad (3.6)$$

We denote by u the normalized field on the mirror faces, and to distinguish it from a Gaussian field, we sometimes will use a subscript ‘‘mesa’’:

$$u_{\text{mesa}}(D, r) = u(D, r) = \frac{U(D, r)}{N(D)}. \quad (3.7)$$

3.2.2 Gaussian fields

The advanced LIGO baseline design uses arm cavities with spherical mirrors, which have Gaussian modes whose field at the mirror plane is (cf. [21])

$$u_G(r, r_o) = \frac{1}{\sqrt{\pi r_o^2}} \exp \left[-\frac{r^2}{2r_o^2} \left(1 - i \frac{b^2}{r_o^2 + \sqrt{r_o^4 - b^4}} \right) \right]. \quad (3.8)$$

Here r_o is the beam radius (at which the energy flux falls to $1/e$ of its central value). From the phase of this field one can read off the radius of curvature of the mirrors:

$$\mathcal{R}_c = L \left(\frac{r_o}{b} \right)^2 \left[\left(\frac{r_o}{b} \right)^2 + \sqrt{\left(\frac{r_o}{b} \right)^4 - 1} \right]. \quad (3.9)$$

3.2.3 Diffraction losses

In the baseline design of an advanced LIGO interferometer [20], the test masses are cylinders whose faces are coated with dielectric mirrors out to a radius R that is 0.8cm less than the cylinders’ physical radii

$$R = R_p - 0.8 \text{ cm}. \quad (3.10)$$

We shall explore MH mirrors that are coated in this same manner, $R = R_p - 0.8 \text{ cm}$ and also MH mirrors that are coated all the way out to the test-mass edges, $R = R_p$. The diffraction losses in each reflection of a cavity mode off a mirror are given, approximately, by the *clipping approximation*

$$\mathcal{L}_{\text{clip}} = \int_r^\infty |u(r)|^2 2\pi r dr. \quad (3.11)$$

Here $u(r)$ is the normalized field [$u_{\text{mesa}}(D, r)$ for a mesa mode and $u_G(r_o, r)$ for a Gaussian mode with infinite mirrors, $R = \infty$].

In actuality, the mirrors' edges at $r = R$ modify the field thereby causing the true diffraction losses to differ from this clipping formula. The true diffraction losses have been computed by OSV [19] via a numerical solution of the eigenequation for the cavity modes, and independently by d'Ambrosio [12] using an FFT code to propagate light in the cavity. The results are

$$\begin{aligned}\mathcal{L}_0 &\simeq 0.85\mathcal{L}_{\text{clip}} \text{ for mesa modes ,} \\ \mathcal{L}_0 &\simeq 2.5\mathcal{L}_{\text{clip}} \text{ for Gaussian modes ,}\end{aligned}\tag{3.12}$$

in the parameter regime of interest — though the numerical coefficients 0.85 and 2.5 can oscillate substantially as the beam radii and mirror radii are changed. When we need high-accuracy diffraction losses (e.g. in portions of Sec. 3.4), we compute them with care using the cavity eigenequation [19] or FFT code [12]).

3.2.4 Mirrors and normalized flux for mesa and gaussian modes

The baseline design for advanced LIGO interferometers has mirror radii $R = 14.9$ cm and Gaussian beam-spot radii $r_o = 1.63b = 4.23$ cm, corresponding to a diffraction loss of $\mathcal{L}_0 = 10$ ppm and a mirror radius of curvature $\mathcal{R}_c = 54$ km. The normalized energy flux $|u_G(r_o, r)|^2$ for this baseline Gaussian field is shown in Fig. 3.2, and the shape of the mirror (segment of a *sphere* with radius 54 km) is shown in Fig. 3.3.

A cavity made from MH mirrors with the baseline radius $R = 14.9$ cm and the baseline diffraction losses $\mathcal{L}_0 = 10$ ppm has a mesa beam radius $D = 3.43b = 8.92$ cm [computed from Eqs. (3.12) and (3.11)]. The normalized energy flux $|u_{\text{mesa}}(D, r)|^2$ for this mesa field is shown in Fig. 3.2. Notice how flat the top of this intensity profile is, and how much like a mesa the profile is shaped, and notice the contrast with the Gaussian profile.

The surfaces of the MH mirrors coincide with the mesa field's surfaces of constant phase; i.e., their height δz as a function of radius r is given by

$$k\delta z = \text{Arg}[u_{\text{mesa}}(D, r)] ,\tag{3.13}$$

where $k = 2\pi/\lambda$ is the light's wave number. This MH mirror shape is shown in Fig. 3.3. Notice the shallow bump in the middle and the flaring outer edges. This bump and flare resemble a Mexican hat (sombrero) and give the MH mirror its name.

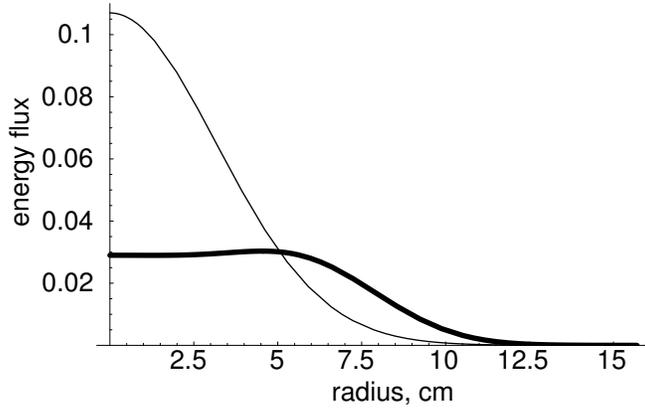


Figure 3.2: The power distributions of: (i) the Gaussian mode for the baseline mirrors with coated mirror radius $R = 15.7$ cm and beam radius $r_o = 1.73b = 4.50$ cm (thin curve), which has diffraction loss per bounce $\mathcal{L}_0 = 10$ ppm; and (ii) the mesa mode with $D = 3.73b = 9.71$ cm (thick curve) which, for this same coated mirror radius $R = 15.7$ cm, has the same diffraction loss per bounce $\mathcal{L}_0 = 10$ ppm.

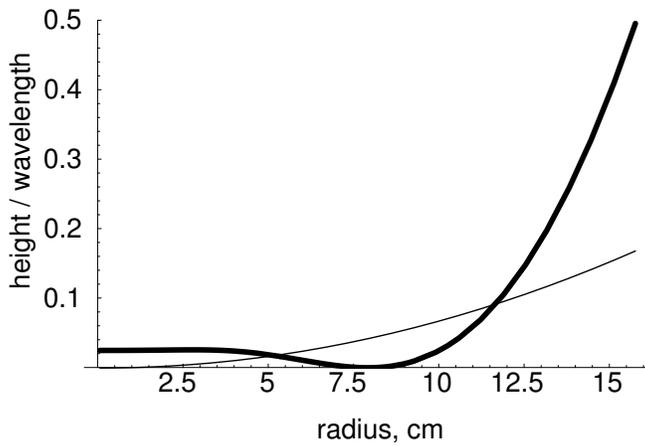


Figure 3.3: The shape of the spherical mirrors (thin curve) and MH mirrors (thick curve) that support the arm-cavity Gaussian mode (thin curve) and mesa mode (thick curve) of Fig. 3.2. The height is measured in units of the wavelength of the light, $\lambda = 1.064\mu\text{m}$.

3.3 Thermoelastic noise and neutron-star binary range for mesa-beam interferometers

3.3.1 Thermoelastic noise

3.3.1.1 Quantifying the thermoelastic noise: the noise integral and the NS/NS range

Building on the seminal work of Braginsky, Gorodetsky and Vyatchanin [5], Liu and Thorne [18] have used Levin's [17] direct method to derive the following formula for an interferometers' thermoelastic noise in terms of a noise integral I_A :

$$S_h(f) = 16\kappa k_B \left(\frac{\alpha_l ET}{C_V(1-2\sigma)\rho\omega L} \right)^2 \bar{I}; \quad (3.14)$$

where \bar{I} is the average, over the four test masses, of the thermoelastic noise integral,

$$\bar{I} = \frac{1}{4} \sum_{A=1}^4 I_A, \quad I_A = \frac{1}{F_o^2} \int_{V_A} (\vec{\nabla}\Theta)^2 d\text{volume}; \quad (3.15)$$

cf. Eqs. (3), (4) and (13) of [18]. Here the notation is as spelled out in Sec. 3.1.4, and we use numerical values (shown in Sec. 3.1.4) that assume the test-mass substrate is sapphire, idealized as an isotropic material. In Eq. (3.15), Θ is the expansion (fractional volume change) inside the test-mass substrate, produced by a static force with magnitude F_o and with profile identical to that of the light beam's intensity distribution over the test-mass face (e.g., Fig. 3.2), and the integral is over the volume V_A of test-mass A . Note that the dimensions of I_A and thence of \bar{I} are length/force² = s⁴g⁻²cm⁻¹.

Equation (3.14) shows that the frequency dependence of the thermoelastic noise is independent of the mirror shape and test-mass shape; it always has the same slope as the SQL (except in testbed systems with tiny mirrors and light beams [18], which are irrelevant in this paper). As a result, the thermoelastic noise produced by an advanced LIGO interferometer whose mirrors have some chosen shapes, divided by the thermoelastic noise of the baseline advanced LIGO interferometer, is equal to the ratio of the two interferometers' noise integrals

$$[S_h(f)/S_h^{\text{BL}}(f)]_{\text{TE}} = \bar{I}/\bar{I}_{\text{BL}} \quad (3.16)$$

This motivates our use of $\bar{I}/\bar{I}_{\text{BL}}$ as one measure of a candidate interferometer's thermoelastic noise.

We shall also use a second measure: The LIGO network's range for detection of NS/NS binaries, with network amplitude signal-to-noise ratio 8. In computing this range, we assume that (i) the three advanced LIGO interferometers (all with $L = 4$ km) are all being operated with signal recycling mirrors optimized for NS/NS inspiral (the operation mode tentatively planned for the first year of advanced LIGO observations), (ii) they all incorporate identical sapphire mirrors with the same

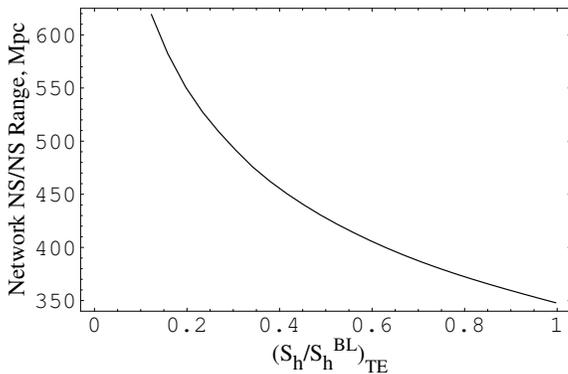


Figure 3.4: The LIGO network NS/NS range as a function of the thermoelastic noise power, in units of the baseline thermoelastic noise, $(S_h/S_h^{\text{BL}})_{\text{TE}}$. For each thermoelastic noise level, the advanced LIGO interferometer’s optical parameters (homodyne readout phase and signal-recycling mirror) are optimized to produce the greatest possible NS/NS range. The optimization has been performed for us by A. Buonanno and Y. Chen (private communication), assuming that the only significant noise sources are thermoelastic noise and optical (unified quantum) noise.

shapes and beam sizes, (iii) all thermal noises are negligible except thermoelastic noise, and (iv) the remaining interferometer parameters have their baseline advanced LIGO values [20] (e.g., the circulating power in each arm is 830 kW). This three-interferometer network range is larger by a factor $\sqrt{3} = 1.732$ than the single-4km-interferometer NS/NS range that is often used by the LIGO community and that is encoded into the “BENCH” LIGO software [15]. *For the baseline advanced LIGO design, the single-4km-interferometer NS/NS range is 200 Mpc, and the network NS/NS range (which we use) is 346 Mpc.*

Since the only noise source we change, in going from one candidate interferometer design to another, is the thermoelastic noise, the NS/NS range must be some function of $[S_h(f)/S_h^{\text{BL}}(f)]_{\text{TE}} = \bar{I}/\bar{I}_{\text{BL}}$.

Buonanno and Chen (private communication) have performed the optimization of the advanced LIGO optical-noise parameters (the homodyne detection phase and the position and reflectivity of the signal recycling mirror), as a function of the thermoelastic-noise level, to produce for us a curve of optimized NS/NS signal-to-noise ratio S/N as a function of $(S_h/S_h^{\text{BL}})_{\text{TE}}$. From that $S/N[(S_h/S_h^{\text{BL}})_{\text{TE}}]$, we have computed the corresponding network range, $(346 \text{ Mpc}) \times (S/N)(S/N)_{\text{BL}}^{-1}$ as a function of thermoelastic noise. We show that range in Fig. 3.4.

A third measure of a candidate interferometer’s performance is the ratio of its network event rate for NS/NS binaries to that of the baseline advanced LIGO network. Since the NS/NS binaries are very extragalactic, their event rate scales as the range cubed,

$$\text{Rate}/\text{Rate}_{\text{BL}} = (\text{Range}/346 \text{ Mpc})^3 . \quad (3.17)$$

In a companion paper [19], OSV evaluate the thermoelastic noise integral I_A numerically for a variety of test-mass shapes and beam radii. Here we summarize the most important conclusions of those computations and their implications for our three performance parameters: $[S_h(f)/S_h^{\text{BL}}(f)]_{\text{TE}}$ [Eq. (3.16)], NS/NS range [Fig. 3.4], and NS/NS event rate [Eq. (3.17)].

3.3.1.2 Baseline test masses with spherical mirrors and Gaussian beams

The baseline design of an advanced LIGO interferometer entails four identical test masses: sapphire cylinders with physical radii $R_p = 15.7$ cm, coated-mirror radii $R = R_p - 0.8$ cm, thickness $H = 13$ cm, density $\rho = 4$ g cm $^{-3}$ and mass $M = 40$ kg; and the baseline light beam at the test-mass face is Gaussian with beam radius $r_o = 4.23$ cm $= 1.63b$ so the diffraction loss per bounce is $\mathcal{L}_0 = 10$ ppm. For this baseline beam and test mass, OSV [19] find for the value of the noise integral

$$I_{\text{BL}} = 2.57 \times 10^{-28} \text{s}^4 \text{g}^{-2} \text{cm}^{-1} . \quad (3.18)$$

We advocate extending the mirror coating out to the test-mass edge so $R = R_p = 15.7$ cm, and increasing the Gaussian beam radius correspondingly, to $r_o = 4.49$ cm $= 4.49b$, so the diffraction losses are still 10 ppm. With this beam expansion, OSV find that the thermoelastic noise is reduced to $S_h/S_h^{\text{BL}} = I/I_{\text{BL}} = 0.856$ [19], from which we deduce (via Fig. 3.4) that the range for NS/NS binaries is increased from 346 Mpc to 364 Mpc, and the NS/NS event rate is increased by a factor $(364/346)^3 = 1.16$; see Table 3.1.

3.3.1.3 Cylindrical test masses with MH mirrors and mesa beams

OSV have computed the thermoelastic noise integral I for cylindrical test masses with mesa beams. The test masses' volumes were held fixed at the baseline value of 10^4 cm 3 (masses fixed at 40 kg), while their physical radii R_p and thicknesses H were varied. For each choice of R_p , two coated-mirror radii were chosen, $R = R_p - 8$ mm (the baseline choice) and $R = R_p$ (our proposed expansion of the coating). In all cases the mesa beam radius D was that value for which the diffraction losses are 10 ppm per bounce inside the cavity.

To within the accuracy of their computations, ~ 0.5 per cent, OSV [19] found that the thermoelastic noise integral I is minimized when the test-mass dimensions have their baseline values, $R = 15.7$ cm, $H = 13$ cm. In other words, the optimal test-mass shape is the same for mesa beams as for Gaussian beams. The optimized (10 ppm diffraction loss) radii D for the mesa beams, and the values of our three measures of interferometer performance are shown in Table 3.1, in two cases: for mirrors coated out to $R = R_p - 8$ mm (the baseline choice), and coated out to $R = R_p$.

As is shown in the table, *switching from Gaussian beams to mesa beams reduces the thermoelastic noise $S_h \propto \bar{I}$ by about a factor 3; it increases the NS/NS range from 346 Mpc to 465 Mpc if $R = R_p - 8$*

Table 3.1: Optimized light-beam configurations, their thermoelastic noise compared to the baseline, their neutron-star binary range, and their event rate for NS/NS inspiral divided by the baseline rate. All test masses are assumed to be cylinders with the baseline advanced LIGO dimensions: physical radius $R_p = 15.7$ cm and thickness $H = 13.0$ cm. The beam radii r_o and D are chosen so that the diffraction loss per bounce in the arm cavities is 10 ppm.

Coated Radius	Beam Shape and Radius	$\left(\frac{S_b}{S_{BL}}\right)_{TE}$	NS/NS Range	$\frac{\text{Rate}}{\text{Rate}_{BL}}$
$R = R_p - 8\text{mm}$	BL: Gaussian $r_o = 4.23\text{cm}$	1.000	346 Mpc	1.00
$R = R_p - 8\text{mm}$	mesa $D/b = 3.43$	0.364	465 Mpc	2.42
$R = R_p$	Gaussian $r_o = 4.49\text{cm}$	0.856	364 Mpc	1.16
$R = R_p$	mesa $D/b = 3.73$	0.290	497 Mpc	2.97

mm , and 497 Mpc if $R = R_p$; and it increases the NS/NS event rate by a factor $(465/346)^3 = 2.42$ if $R = R_p - 8$ mm, and to $(497/364)^3 = 2.55$ if $R = R_p$.

3.3.1.4 Conical test masses

By switching from cylindrical test masses to frustums of cones, with the same test-mass volume, one could further reduce, substantially, the thermoelastic noise and increase the NS/NS range and rate. For detailed explorations of this, see OSV [19].

We do not discuss this possibility in the present paper because the current technology for growing sapphire boules, from which to cut the advanced LIGO test masses, places a tight limit on the test-mass physical radius R_p . It cannot be much larger than the baseline $R_p = 15.7$ cm; and for that maximum radius, and test-mass volumes of order the baseline 10^4 cm, the optimal test-mass shape is cylindrical, with the baseline dimensions [19].

When it becomes possible, in the future, to grow larger sapphire boules, it might be worth considering test masses with frustum-of-cone shapes [19].

3.4 Sensitivity to mirror tilts, displacements and figure errors

The MH mirror figure (Fig. 3.3) is somewhat flatter than the baseline spherical mirror in its central 10 cm of radius where 95 per cent of the light power resides, but much more curved in its outer ~ 6 cm. One might worry that the central flatness will cause a mesa-beam interferometer to be unacceptably sensitive to mirror-tilt-induced, mirror-displacement-induced and figure-error-induced mixing of parasitic modes into the light beam's fundamental, mesa mode. We have investigated this

mode mixing and find that it is a modest problem, not a severe one. We describe our investigations and conclusions in this section. They have been described previously in our internal LIGO report [3], and a short summary of results was given in Sec. 3.1.3 above.

3.4.1 Foundations for investigation

3.4.1.1 Our tools of analysis

Our analysis of mode mixing and its consequences is based on three independent sets of tools. The first two sets are designed for studying the effects of mirror errors on the interferometer’s high-finesse arm cavities. The third set is for analyzing the highly degenerate power-recycling and signal-recycling cavities.

Our first tool set (developed by Richard O’Shaughnessy with confirming calculations by Sergey Strigin and Sergey Vyatchanin, and described in detail in OSV [19]) is an *integral eigenequation* for the modes of an arm cavity. In the limit of infinite mirror radii (i.e., neglecting diffraction losses), the cavity’s eigenmodes are orthonormal when integrated over the transverse plane; this is true for MH mirrors, just as for spherical mirrors. OSV have used their integral eigenequation to compute the modes with untilted, undisplaced and undeformed mirrors and with both infinite and finite radii. O’Shaughnessy has then tilted, displaced and deformed the ETM of one arm cavity and applied *first- and second-order perturbation theory* to its eigenequation to determine the tilt-induced and deformation-induced mode-mixing, the resulting fundamental eigenmode of the cavity with tilted and deformed ETM, its response to the driving beam, and the tilt-induced and deformation-induced power going out the interferometer’s dark port. The details of these calculations are given in OSV [19]. In the following subsections we describe the main results, we compare with computations via our second tool set, and we discuss the implications for the use of MH mirrors in advanced LIGO interferometers.

Our second tool set (developed by d’Ambrosio and described in detail in her companion paper [12]) is an adaptation of a standard LIGO research tool: an *FFT code for simulating interferometers*. D’Ambrosio has used this code to study the same arm-cavity phenomena that OSV have studied via the cavity eigenequation and perturbation theory. She presents the details of her computations and some associated perturbation theory analyses in Ref. [12]. In the following subsections we describe her main results, we compare them with the OSV eigenequation results, and we discuss their implications.

Our third tool set (developed by Thorne and described in Secs. 3.4.9 and 3.4.10 below) is designed to deal with the influence of mirror errors on the interferometer’s power-recycling and signal-recycling cavities. Because these cavities are highly degenerate, the light trapped in them and traversing them is moderately well described by the tools of geometric optics, augmented by approximate estimates

of the impact of diffraction; so these are the foundation for our third tool set.

3.4.1.2 Parameter values for comparing mesa and Gaussian configurations

When analyzing practical issues, we have carried out most of our computations for *fiducial MH mirrors* that have slightly larger coated radii, $R_M = 16\text{cm}$, than the *baseline spherical mirrors* (14.9 cm at present and 15.7 cm if our recommendation to coat the outer 8 mm is followed), and we have used mesa beams with modestly larger diffraction losses, $\mathcal{L}_0 = 18$ ppm, than the current baseline of about 10 ppm. In evaluating practical issues in this paper and its companions [12, 19] we will compare with baseline spherical mirrors and Gaussian beams that have this same enlarged coated radius $R_M = 16$ cm and diffraction loss $\mathcal{L}_0 = 18$ ppm.

Our mirror and beam parameters, then, are as follows:

- For the *baseline (BL) spherical mirrors*: mirror radius $R_M = 16$ cm (vs 15.7 cm for the true current baseline) and Gaussian beam radius $r_o = 4.70$ cm (vs. 4.23 cm for the true current baseline).
- For the *MH mirrors*: mirror radius $R_M = 16$ cm and mesa beam radius parameter $D = 4b = 10.4$ cm, where $b = \sqrt{\lambda L/2\pi} = 2.60$ cm, with $\lambda = 1.06$ μm and $L = 4$ km the light wavelength and arm length.

3.4.1.3 Reflection and transmission coefficients

For both configurations, fiducial MH and baseline spherical, we have assumed for the ITM mirror the advanced LIGO baseline power transmissivity $t_I^2 = 0.995$; we have chosen an idealized, perfectly reflecting end ETM; and for simplicity we have assumed that the only losses are those due to diffraction, which we treat as a separate mathematical entity from the reflection and transmission coefficients. Thus, our power reflection and transmission for the ITMs and ETMs are

$$\begin{aligned} r_I^2 &= 0.995, & t_I^2 &= 0.005; \\ r_E^2 &= 1.0, & t_E^2 &= 0. \end{aligned} \tag{3.19}$$

We have also assumed the baseline advanced LIGO transmissivities and reflectivities for the power-recycling (PR) and signal-recycling (SR) mirrors; see Secs. 3.4.9 and 3.4.10 below.

3.4.2 Driving a mesa-beam interferometer with a gaussian beam

One way to produce the desired mesa beams in the arm cavities is to drive the interferometer with Gaussian-beam light and let the arm cavities or a mode-cleaning cavity convert the light into mesa form. One of us, ED'A, has identified the Gaussian beam that has the greatest overlap with the

mesa beam of an MH-mirrored arm cavity. If one were to drive the MH arm cavity directly with a Gaussian beam, this would be the driving beam one would want to use. It has a beam radius $r_{od} = 6.92$ cm, compared to our baseline cavity's beam radius at the ITM of $r_o = 4.70$ cm. This Gaussian driving beam $u_d(r)$ has an overlap

$$\gamma_0^2 \equiv |\langle u_0 | u_d \rangle|^2 \equiv \left| \int u_0^* u_d d\text{Area} \right|^2 = 0.940 \quad (3.20)$$

with the cavity's fundamental mesa mode u_0 , which means that 94.0 per cent of the Gaussian driving-beam light will enter the MH-mirrored cavity, and 6.0 per cent will get rejected. See Ref. [12].

3.4.3 Parasitic modes in arm cavities

It is useful to think of the MH mirrors as having two regions: a central region with radius $\simeq 10$ cm, and an outer region with radius $\simeq 6$ cm. In its central region, the MH mirror is much flatter than the baseline spherical mirror; in its outer region, it is much more sharply curved; see Fig. 3.3. The flatness of the central region has led to concerns about degeneracies of modes and sensitivity to mirror tilts, displacements and figure errors.

O'Shaughnessy, Strigin and Vyatchanin have all independently solved the integral eigenequation for the modes of a LIGO arm cavity with MH mirrors. They have found (cf. slide 15 of [4]) that *among modes that are not strongly damped by diffraction losses, the one closest in frequency to the fundamental TEM00 mode u_0 is the lowest TEM01 mode (denoted u_1 below). Its frequency separation from the fundamental is 0.0404 of the cavity's free spectral range, which is 2.5 times smaller than for the baseline spherical-mirrored cavity, $0.099 \times$ (free spectral range).* Evidently, the sharp curvature of the MH mirrors' outer region compensates sufficiently for the flatness of their central region, to prevent the parasitic modes' frequencies from becoming near-degenerate with the fundamental.

3.4.4 Mirror tilt in arm cavities

Our modeling predicts that *mode mixing in the arm cavities of a mesa-beam interferometer produced by tilt of the ETM's or ITM's should be of no serious consequence, if the tilt angles are kept below about 10^{-8} rad.* In the following subsections we summarize the calculations that lead to this conclusion.

3.4.4.1 Parasitic mode mixing in arm cavities

Two of us have computed the influence of a tilt of the ETM on the fundamental mode of an arm cavity: ED'A has done this using an FFT code, and RO'S has done it by applying perturbation

theory to the arm cavity's integral eigenequation. The two computations agree on the following predictions:

When the ETM is tilted through an angle θ , the cavity's fundamental mode gets changed from $u_0(r)$ to

$$u'_0(\vec{r}) = (1 - \alpha_1^2/2)u_0(r) + \alpha_1 u_1(\vec{r}) + \alpha_2 u_2(\vec{r}) . \quad (3.21)$$

Here \vec{r} is position in the transverse plane, u_n are unit-norm superpositions of modes of the perfectly aligned cavity ($\int |u_n|^2 d\text{Area} = 1$), α_n are mode-mixing coefficients that scale as θ^n , and our computations have been carried out only up through quadratic order. The maximum tilt that can be allowed is of order 10^{-8} radian, so we shall express our predictions for the α_n in units of $\theta_8 \equiv \theta/10^{-8}$.

For our baseline spherical-mirrored cavities, u_0 is the (0,0) Hermite-Gaussian mode, u_1 is the (0,1) Hermite-Gaussian mode, u_2 is the (0,2) mode, and the dominant mixing coefficient α_1 is

$$\alpha_1^{\text{BL}} = \frac{1}{\sqrt{2}(1-g^2)^{3/4}} \left(\frac{\theta^{\text{BL}}}{b/L} \right) = 0.0064\theta_8 \quad (3.22)$$

(Eq. (56) of OSV [19]; cf. also ED'A [12]). Here $g = 0.952$ is our BL arm cavity's g-factor and $b = \sqrt{\lambda L/2\pi} = 2.603$ cm is its transverse diffraction scale.

For the MH cavity, u_0 is the mesa mode shown (for slightly different beam radius) in Fig. 3.2; and maps of u_1 , and u_2 , as computed by O'Shaughnessy [19] by solving the cavity eigenequation, are shown in Fig. 3.5. D'Ambrosio [12] has used her FFT code to compute maps that agree well with these. The MH mixing coefficients are

$$\alpha_1^{\text{MH}} = 0.0227\theta_8 , \quad \alpha_2^{\text{MH}} = 0.00018\theta_8^2 \quad (3.23)$$

(Eq. (57) of OSV [19]; cf. also ED'A [12]). Comparison of Eqs. (3.22) and (3.23) shows that, *to keep the dominant, dipolar mode-mixing coefficient α_1 at the same level in the MH cavity as in the baseline spherical-mirrored cavity requires controlling the MH mirrors' tilt angle θ four times more tightly.*

The fractional power in the dominant, dipolar parasitic mode is α_1^2 , which will be doubled to

$$\begin{aligned} P_1^{\text{BL arm total}} &\simeq 2(\alpha_1^{\text{BL}})^2 \simeq 0.001(\theta_8/3.5)^2 , \\ P_1^{\text{MH arm total}} &\simeq 2(\alpha_1^{\text{MH}})^2 \simeq 0.001\theta_8^2 , \end{aligned} \quad (3.24)$$

when the ITM and ETM are both tilted but about uncorrelated axes. This *suggests that, so far as the arm cavities are concerned (i.e., ignoring issues of tilt in the recycling cavities), the tilt of MH mirrors need not be controlled much better than $\simeq 10^{-8}$ rad.* We have verified this by examining the effects of the mode mixing on various cavity and interferometer parameters:

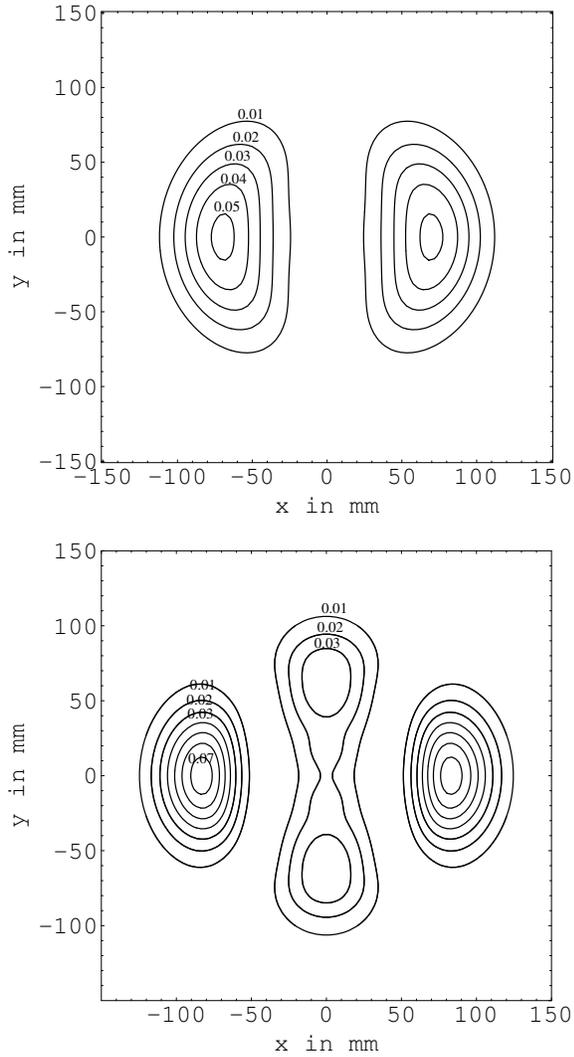


Figure 3.5: Contour diagram maps of the squared moduli $|u_1(\vec{r})|^2$ (top) and $|u_2(\vec{r})|^2$ (bottom) of the parasitic modes excited by tilt of the ETM of a MH arm cavity. These maps were generated by O'Shaughnessy [19] using the eigenequation for a cavity with infinite-radius MH mirrors, together with first- and second-order perturbation theory. The axis about which the ETM is tilted (rotated) is chosen to be vertical (the y axis). The units in which $|u_j(\vec{r})|^2$ is measured are $b^{-2} = (2.603\text{cm})^{-2}$; the peak values are 0.0526 for $|u_1|^2$, and 0.0726 for the peak of $|u_2|^2$ on the x -axis and 0.0398 for the peak on the y -axis. The contours are separated by 0.01 in units of b^{-2} ; the outermost contour in all cases is $0.01b^{-2}$.

3.4.4.2 Diffraction losses

One of us (ED'A), from her FFT simulations, has estimated the influence of ETM tilt on diffraction losses to be

$$\mathcal{L}'_0^{\text{MH}} = \mathcal{L}_0^{\text{MH}}(1 + 0.004\theta_8^2) \quad (3.25)$$

[cf. [12]]. This result been confirmed to a factor ~ 2 by RO'S by combining the clipping approximation with perturbation theory of the cavity's eigenequation; Eq. (59) of [19]. The influence of ITM tilt should be about the same, thus doubling the coefficient of θ_8^2 . *This tilt-induced increase in losses is so small that it can be ignored for tilt angles below $\sim 10^{-8}$ rad.*

3.4.4.3 Arm cavity gain

We have computed (ED'A via the FFT code and RO'S via perturbation theory) the following influence of ETM tilt on the arm cavity gain

$$G_{\text{cav}}^{\text{MH}} = 740(1 - 0.00057\theta_8^2) \quad (3.26)$$

(Eq. (62) of [19]; cf. also [12]). This result assumes the baseline values for the power transmissivities of the ITM and ETM, and assumes for simplicity that the only losses are diffraction losses; Sec. 3.4.1.3. The factor 740 assumes the cavity is driven by its best-fit Gaussian beam and thus is smaller by about $\gamma_0^2 = 0.940$ than the gain in the untilted, spherical-mirrored baseline cavity. The tilt of the ITM should produce about the same gain reduction as that of the ETM, thus doubling the coefficient of θ_8^2 to ~ 0.001 . This coefficient is small enough that *the tilt-induced decrease of MH arm-cavity gain will be negligible if θ is controlled to $\sim 10^{-8}$ rad.*

3.4.4.4 Dark port power

We have computed the influence of the tilt of one ETM on the dark-port output light (ED'A using the FFT code and RO'S using perturbation theory). Multiplying that result by four to deal with the case of all four cavity mirrors being tilted about uncorrelated axes, we find for the fraction of the interferometer's input power that winds up at the dark port in the fundamental mode u_0 and the parasitic modes u_1 and u_2 :

$$\begin{aligned} P_0^{\text{MH DP total}} &\simeq 1.0\theta_8^4 \text{ ppm} , \\ P_1^{\text{MH DP total}} &= 4\gamma_0^2(\alpha_1^{\text{MH}})^2 \simeq 2000\theta_8^2 \text{ ppm} , \\ P_2^{\text{MH DP total}} &= 4\gamma_0^2(\alpha_2^{\text{MH}})^2 \simeq 0.1\theta_8^4 \text{ ppm} . \end{aligned} \quad (3.27)$$

(Eqs. (64)–(66) of [19]; cf. also [12]). *Without an output mode cleaner, the dark-port power would primarily be in the dipolar mode u_1 , and for $\theta < 10^{-8}$ rad it would constitute < 0.2 per cent of the input light. The planned output mode cleaner will wipe out this u_1 power and the power in mode u_2 , leaving only the tiny fundamental-mode power, which should be totally negligible for θ below 10^{-8} rad.*

For comparison, the dark-port powers with the baseline Gaussian beams are

$$\begin{aligned} P_0^{\text{BL DP total}} &= 4(\alpha_1^{\text{BL}})^4 \simeq 1.0(\theta_8/3.5)^4 \text{ ppm} , \\ P_1^{\text{BL DP total}} &= 4(\alpha_1^{\text{BL}})^2 \simeq 2000(\theta_8/3.5)^2 \text{ ppm} , \end{aligned} \quad (3.28)$$

which shows once again that the BL spherical-mirrored arm cavities are ~ 4 times less sensitive to tilt than the MH arm cavities.

3.4.5 Transverse displacement of arm cavities' mirrors

When the ETM is displaced transversely through a distance s , the cavity's fundamental mode gets changed from $u_0(r)$ to

$$u'_0(\vec{r}) = (1 - \zeta_1^2/2)u_0(r) + \zeta_1 w_1(\vec{r}) + \zeta_2 w_2(\vec{r}) , \quad (3.29)$$

where the parasitic modes w_n , like the u_n 's, have unit norm, $\langle w_n | w_n \rangle = \int |w_n|^2 d\text{Area} = 1$, and have phase adjusted so the coupling coefficients ζ_n are real, and where $\zeta_n \propto s^n$.

R'OS has computed the coupling coefficients ζ_n for the BL spherical mirrors and for MH mirrors by applying perturbation theory to the cavity's eigenequation. In both cases ζ_2 is negligible compared to ζ_1 when the displacement is $s \ll b \equiv \sqrt{\lambda L/2\pi} = 2.60$ cm, so we shall ignore ζ_2 . R'OS finds that the MH and BL mirrors are roughly equally sensitive to transverse displacements; their coupling coefficients are (Eqs. (51) and (52) of [19]):

$$\begin{aligned} \zeta_1^{\text{BL}} &= \left(\frac{(1-g)^{1/4}}{\sqrt{2}(1+g)^{3/4}} \right) \frac{s}{b} = 0.008s_{\text{mm}} , \\ \zeta_1^{\text{MH}} &= 0.010s_{\text{mm}} . \end{aligned} \quad (3.30)$$

Here $g = 0.952$ is the BL arm cavity's g-factor, and s_{mm} is the ETM's transverse displacement in millimeters.

The corresponding fraction of the arm-cavity carrier power driven into the (dipolar) parasitic field w_1 is [Eq. (53) of [19]]

$$P_1^{\text{arm total}} = \zeta_1^2 \simeq \begin{cases} 100(s/1.3\text{mm})^2 \text{ppm} & \text{BL,} \\ 100(s/1.0\text{mm})^2 \text{ppm} & \text{MH.} \end{cases} \quad (3.31)$$

The fraction of the input carrier power driven out the dark port when the ETMs of both arm cavities are displaced through a distance s but in uncorrelated directions is about twice the above [Eq. (54) of [19]]:

$$P_1^{\text{DP total}} = 2\gamma_0^2\zeta_1^2 \simeq \begin{cases} 190(s/1.3\text{mm})^2\text{ppm} & \text{BL,} \\ 190(s/1.0\text{mm})^2\text{ppm} & \text{MH.} \end{cases} \quad (3.32)$$

These coupling coefficients and parasitic-mode powers are sufficiently small that transverse displacements are not a serious issue, and so shall ignore them in the rest of this paper. In any event, the low sensitivity to a change from BL spherical to MH mirrors makes displacement a non-issue in the any decision about whether to use MH mirrors.

3.4.6 Errors in the arm cavities' mirror figures

3.4.6.1 Billingsley's worst-case figure error

Garilynn Billingsley (of the LIGO Laboratory, Caltech) has provided us with a map of a worst-case figure error, $\delta z_{\text{wc}}(x, y)$ [height error as function of Cartesian coordinates in the transverse plane], produced by current technologies. Her map is based on the measured deviation of a LIGO-I beam-splitter substrate from flatness. The measured substrate had diameter 25 cm; she stretched its deviation from flatness (its “figure map”) to the baseline mirror diameter of 35.4 cm, fit Zernike polynomials to the stretched map, and smoothed the map by keeping only the lowest 36 Zernikes.

We show a contour diagram of the resulting figure map (figure “error”) in Fig. 3.6. In the central region (innermost 10 cm in radius), the peak to valley error Δz is about 30 nm, while in the outer region (10 cm to 16 cm in radius), it is about 110 nm. Billingsley thinks it likely that in the central region (which dominates our considerations), peak-to-valley errors of $\Delta z \sim 5$ nm may be achievable for MH mirrors — about 1/5 as large as in Fig. 3.6; and we have found that $\Delta z = 0.2 \times 30 = 6$ nm is small enough that the influences of the figure error scale, for $\delta z \lesssim 6$ nm, as Δz or Δz^2 with higher-order terms producing $\lesssim 10$ per cent corrections. Accordingly, in the analyses described below we shall use Billingsley's map, scaled down in height by a factor ε :

$$\delta z = \varepsilon \delta z_{\text{wc}}(x, y) , \quad (3.33)$$

and we shall use $\varepsilon = 0.2$ and $\Delta z = 6$ nm as our fiducial values for ε and Δz . Jean Marie Mackowski (an expert in coating mirrors) believes that $\Delta z \sim 2$ nm errors may be achievable for MH mirrors, if the mirror figure is produced by coating; this corresponds to $\varepsilon \sim 0.07$. For comparison, the figure errors for LIGO-I mirrors are $\Delta z \sim 1$ nm, but LIGO-I has spherical mirrors which presumably can be manufactured more accurately than MH mirrors.

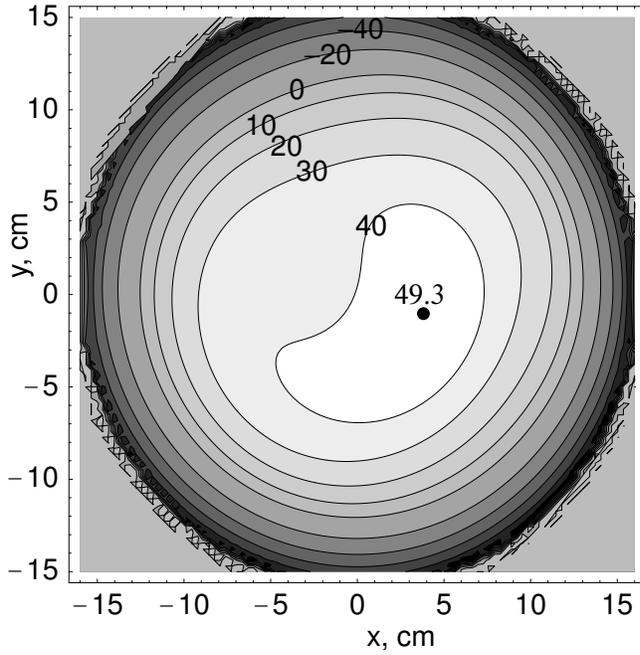


Figure 3.6: Contour diagram of Billingsley's worst-case figure error [height δz_{wc} in nanometers as a function of transverse Cartesian coordinates (x, y) in centimeters]. The hash at the outer edge of the mirror is an artifact of our numerical manipulation of Billingsley's map.

3.4.6.2 Mode mixing by figure errors without compensating tilt

We have computed the mode mixing in an MH cavity when Billingsley's worst-case figure error, reduced by $\varepsilon \lesssim 0.2$, is placed on the ETM. As in our tilt studies, the computation was done independently by ED'A using the FFT code and by RO'S using arm-cavity perturbation theory.

By analogy with Eq. (3.21), the fundamental mode with deformed ETM has the following form

$$u'_0 = (1 - \beta_1^2/2)u_0 + \beta_1 v_1, \quad (3.34)$$

where the parasitic mode v_1 , like the u_n 's, has unit norm $\langle v_1 | v_1 \rangle = \int |v_1|^2 d\text{Area} = 1$, and has its phase adjusted so that $\beta_1 \propto \varepsilon/0.2$ is real. By contrast with the tilt-induced mode mixing, where u_1 is dipolar (angular dependence $\cos \varphi$), the deformation parasite v_1 has a complicated shape that depends on the details of the deformation and that therefore contains a number of multipoles. A map of the power distribution $|\beta_1 v_1|^2$ of the admixed mode is shown in Fig. 3.7.

The fraction of the arm cavity power in the parasitic mode is

$$P_1^{\text{MH arm}} = \beta_1^2 = 0.0012(\varepsilon/0.2)^2 \quad (3.35)$$

(Eq. (70) of [19]; cf. also ED'A [12]), and the fraction of the interferometer's input power that goes

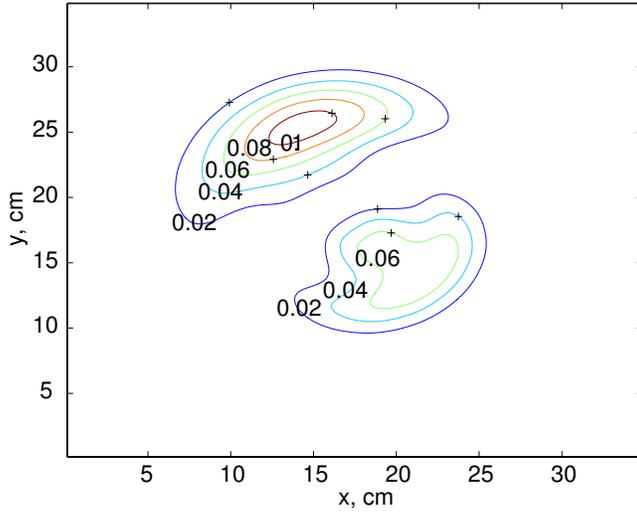


Figure 3.7: The power distribution $|u'_0 - u_0|^2 = |\beta_1 v_1|^2$ (in units $1/\text{m}^2$) of the deformation-induced parasitic mode when the deformation $\varepsilon \delta z_{\text{wc}}(x, y)$ with $\varepsilon = 0.2$ is applied to the ETM of an MH arm cavity. This map was computed by ED'A [12] using the FFT code; the map computed by RO'S using perturbation theory [19] is in reasonable agreement with this one (e.g., the heights of the two peaks are $\{0.111, 0.080\}$ in the FFT map, and $\{0.110, 0.088\}$ in the perturbation map). The parasitic power $|\beta_1 v_1|^2$ scales as $(\varepsilon/0.2)^2$.

out the dark port (if the interferometer is driven by the best-fit Gaussian mode u_d and if only one of the arm mirrors — one ETM — is deformed) is given by

$$P_1^{\text{MH DP}} = \gamma_0^2 \beta_1^2 = 0.0011(\varepsilon/0.2)^2 \quad (3.36)$$

(Eq. (71) of [19]; cf. also ED'A [12]). At the leading, ε^2 , order in the deformation (the order to which we have computed), this dark-port power is entirely in the parasitic mode v_1 . Our FFT and perturbation-theory calculations agree on the parasitic powers (3.35) and (3.36) to within about five per cent.

The dark-port power (3.36) and parasitic arm-cavity power (3.35) are influenced primarily by the figure error in the central (10 cm radius) region of the ETM, because about 96 per cent of the mesa-mode power is contained in that central region, and only about 4 per cent in the outer region — and of the outer 4 per cent, 3/4 (3 per cent) is in the annulus between 10 and 11 cm. The insensitivity to outer-region deformations is fortunate, because Billingsley tells us that it will be much easier to keep the figure errors small in the central region than in the outer region.

[RO'S has verified the insensitivity to the outer-region deformations by evaluating (via perturbation theory) the dark-port power for a mirror deformation

$$\delta z = \varepsilon_c \delta z_{\text{wc}}^{\text{central}} + \varepsilon_o \delta z_{\text{wc}}^{\text{outer}}. \quad (3.37)$$

Here $\delta z_{\text{wc}}^{\text{central}}$ is equal to δz_{wc} at $r < 9.6$ cm and is zero at $r > 12.2$ cm, and between 9.6 and 12.2 cm, $\delta z_{\text{wc}}^{\text{central}}/\delta z_{\text{wc}}$ falls linearly from 1 to zero; and similarly $\delta z_{\text{wc}}^{\text{outer}}$ is equal to δz_{wc} at $r > 12.2$ cm and is zero at $r < 9.6$ cm, and between 9.6 and 12.2 cm $\delta z_{\text{wc}}^{\text{outer}}/\delta z_{\text{wc}}$ grows linearly from 0 to 1. RO'S finds, as a function of the central-region and outer-region weightings,

$$P_1^{\text{MH DP}} = \gamma_0^2 \beta_1^2 = 0.0010(\epsilon_c/0.2)^2 + 1.5 \times 10^{-5}(\epsilon_o/0.2)^2 \quad (3.38)$$

(Eq. (72) of [19]); so that, for example, if Billingsley's worst-case perturbations are reduced by $\epsilon_c = 0.2$ in the central region (to $\Delta z = 6$ nm), but are kept at their full original strength $\epsilon_o = 1.0$ in the outer region (so $\Delta z = 110$ nm there), the outer region will contribute 40 per cent as much power to the dark port as the inner region.]

When all four arm-cavity mirrors are subjected to uncorrelated deformations, the arm-cavity parasitic power (3.35) will be increased by a factor 2 and the dark-port power (3.36) by a factor 4, to

$$\begin{aligned} P_1^{\text{MH arm total}} &= 2\beta_1^2 \simeq 0.0025(\Delta z/6\text{nm})^2, \\ P_1^{\text{MH DP total}} &= 4\gamma_0^2 \beta_1^2 \simeq 0.005(\Delta z/6\text{nm})^2 \end{aligned} \quad (3.39)$$

where Δz is the peak-to-valley mirror deformation in the central region. This suggests that, *so far as arm-cavity mode mixing is concerned, peak-to-valley figure errors of order 6 nm in the inner 10 cm are acceptable.*

3.4.6.3 Mode mixing by figure errors with compensating tilt

The parasitic mode v_1 (Fig. 3.7) contains a significant amount of dipolar field, as one sees from the asymmetry of the map. The advanced LIGO tilt control system, based on a quadrant-diode readout of asymmetry in the power distribution u'_0 , will tilt the mirror so as to remove the overlap between the deformed parasitic field v_1 and the dipolar-tilt parasitic field u_1 . ED'A and R'OS have independently computed that the optimal tilt is about $\theta_{\text{compensate}} = 1.3 \times 10^{-8}(\epsilon/0.2)$ radians about a line rotated 55 degrees from the x axis, and have computed the resulting field $u'_0 = u_0 + \beta_1 v_1 + \alpha_1 u_1$ with minimum parasitic-mode power. Figure 3.8 shows the mirror deformation after tilt, and Fig. 3.9 shows the parasitic power distribution $|\beta_1 v_1 + \alpha_1 u_1|^2$ for $\epsilon = 0.2$. Notice that the tilt has largely but not completely removed the dipolar asymmetry. Some residual dipolar field remains — that portion which cannot be compensated by a tilt.

We denote by v_{1c} the tilt-compensated parasitic mode and by β_{1c} its (real) amplitude, so $\beta_{1c} v_{1c} = \beta_1 v_1 + \alpha_1 u_1$. Then the cavity's eigenmode, with tilt compensation (including the second-order loss

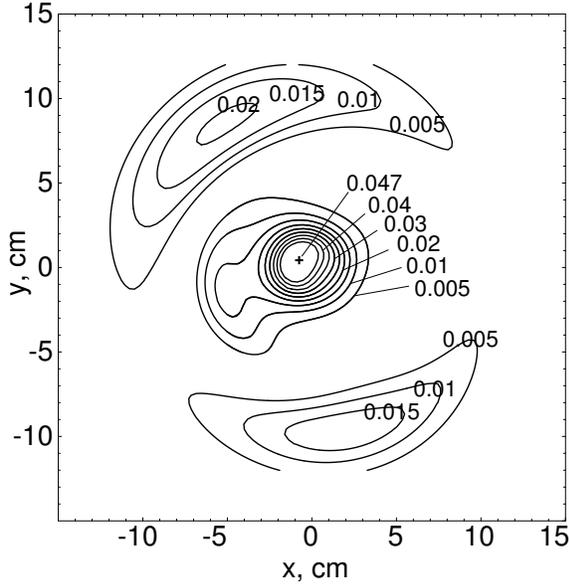


Figure 3.9: The power distribution $|u'_0 - u_0|^2 = |\beta_{1c}v_{1c}|^2$ (in units $1/\text{m}^2$) of the deformation-induced parasitic mode when the tilt-adjusted deformation $\delta z = 0.2\delta z_{\text{wc}} + \theta_{\text{compensate}} \sin(\varphi - 55^\circ)$ is applied to the ETM of an MH arm cavity. This map was computed by RO'S by applying perturbation theory to the cavity's eigenmode [19]. The map computed by ED'A using the FFT code [12] agrees reasonably well.

deformations as large as 10 nm. Recycling-cavity issues, discussed below, will place much tighter constraints on the mirror figures.

3.4.7 Influence of mirror tilt and figure errors on thermoelastic noise

When one MH mirror of an arm cavity is given the deformation $\varepsilon\delta z_{\text{wc}}$, the resulting deformation of the arm-cavity power distribution, $\delta|u'_0|^2 = |u'_0|^2 - |u_0|^2$ increases the thermoelastic noise. The following argument (due to RO'S [19]) shows that, at leading (linear) order in $\varepsilon/0.2$, *only the circularly symmetric portion of the parasitic mode $\beta_1v_1 = u'_0 - u_0$ contributes to the thermoelastic noise increase*: The fractional noise increase is equal to the fractional increase in the thermoelastic noise integral $I = \int (\vec{\nabla}\Theta)^2 d\text{vol}$ (Eq. (4.2) of [13]; slide 6 of [4]):

$$\frac{\delta S_h^{\text{TE MH}}}{S_h^{\text{TE MH}}} = \frac{\delta I}{I} \propto \int (\vec{\nabla}\Theta) \cdot (\vec{\nabla}\delta\Theta) d\text{vol}. \quad (3.44)$$

Here Θ is not a tilt angle but rather is the expansion (fractional increase of volume) of the substrate material when a static pressure $P \propto |u'_0|^2$ is applied to the mirror face. Since the unperturbed mesa beam is circularly symmetric, so will be the unperturbed expansion Θ , which means that only the circularly symmetric portion of the expansion perturbation $\delta\Theta$, and thence the circularly symmetric portion of the pressure perturbation $\delta P \propto \delta|u'_0|^2$ will contribute to the noise increase. At leading

(linear) order in the mirror deformation $\varepsilon/0.2$, the circularly symmetric portion of $\delta|u'_0|^2$ arises solely from the circularly symmetric portion of $\delta u'_0 = \beta_1 v_1$; thus, as claimed, only the circularly symmetric portion of $\beta_1 v_1$ can increase the thermoelastic noise.

This same argument shows that *the mesa-mode deformation produced by mirror tilt cannot influence the thermoelastic noise at first order in the tilt angle*; and therefore, *we need not be concerned about the influence of tilt on the thermoelastic noise — whether the tilt is unintended, or is being used in a controlled way to compensate the errors in the mirror figures.*

RO'S and SS have independently computed the fractional increase in the thermoelastic noise integral at linear order in $\varepsilon/0.2$; cf. Eq. (77) in OSV [19]. Their result, multiplied by four to account for four arm-cavity mirrors, is

$$\frac{\delta S_h^{\text{TE, MH total}}}{S_h^{\text{TE, MH}}} = 0.14(\varepsilon/0.2) \simeq 0.14(\Delta z/6\text{nm}) . \quad (3.45)$$

This 14 per cent increase of S_h , when all four mirrors are subjected to 6 nm figure errors in their central regions, is to be compared with the factor $1/0.34 = 295$ per cent decrease in thermoelastic noise achieved by switching from spherical mirrors to MH mirrors. There may also be a small increase in thermoelastic noise when a spherical mirror is deformed. *Assuming, conservatively, no deformation-induced noise increase for spherical mirrors, switching from spherical to 6nm-deformed MH mirrors will reduce the thermoelastic noise by a factor $0.34 \times 1.14 = 0.39$, which in turn will increase the distance for NS/NS binaries from 364 Mpc to about 490 Mpc (Fig. 3.4) and increase the event rate by about a factor 2.45*; cf. Table 3.1.

3.4.8 Mesa vs. Gaussian beams in recycling cavities

In LIGO-I interferometers, and in the baseline design for advanced LIGO interferometers, the greatest sensitivity to mirror tilts and figure errors occurs in the power recycling (PR) and signal recycling (SR) cavities. This sensitivity arises from the near degeneracy of the recycling cavities, which strongly enhances error-induced mode mixing. One might worry that for MH mirrors, with their greater central-region flatness, this severe mode mixing might be made substantially worse.

We have examined this question and conclude that *for the two wideband advanced LIGO interferometers, there is not much difference between the baseline Gaussian beams and the proposed mesa beams, with regard to their susceptibility to mode mixing in the recycling cavities.* The only significant difference arises from the fact that the mesa beams are larger and therefore sample, with significant power, larger-radii regions of the mirrors (the regions between, say, 8 cm radius and 10 cm radius), where the deformations may be worse.

The reason that the mode mixing is only marginally sensitive to the beam shape is quite simple: If heterodyne readout were used, then once RF modulated sideband light gets into the power recycling

cavity, it would make roughly $\mathcal{N}_{\text{PR}} \sim \frac{1}{2}(\text{cavity finesse } \mathcal{F}_{\text{PR}}) \simeq (\pi/2)/(1 - R_{\text{PR}}) \simeq 25$ round trips before losing 95 per cent of its power out the beam splitter's dark port. (Here $R_{\text{PR}} \simeq 0.94$ is the power reflectivity of the power recycling mirror). And once signal light gets into the signal recycling cavity, it makes roughly $\mathcal{N}_{\text{SR}} \sim \frac{1}{2}\pi/(1 - \rho\sqrt{R}) \sim 40$ round trips before losing 95 per cent of its power out the dark port or back into the arm cavities. (Here $R = 0.995$ is the ETM power reflectivity and $\rho = \sqrt{0.93}$ is the amplitude reflectivity of the SR mirror, in the notation of Buonanno and Chen [8].) The Fresnel length (transverse diffraction scale) for light that makes \mathcal{N} round trips in either recycling cavity with cavity length $\ell \simeq 10$ m is

$$r_F = \sqrt{\lambda_o 2\ell\mathcal{N}} \simeq 3\text{cm}\sqrt{\mathcal{N}/40}, \quad (3.46)$$

where $\lambda_o = 1\mu\text{m}$ is the light's wavelength. This Fresnel length is $\sim 1/2$ of the ~ 5 cm scales on which the ideal mirror shapes and the central-region worst-case mirror errors vary, and it is small compared to the ~ 15 to 20 cm diameter beams themselves. There thus is only modest diffractive coupling between light rays, and *the light bouncing back and forth in each recycling cavity is describable, to moderately good accuracy, by geometric optics*. Moreover, because the mirrors (whether MH or spherical) are nearly flat and nearly identical, the light's rays, to rather good accuracy, are all parallel to the optic axis and to each other and are thus decoupled from each other. If the mirrors and beam splitter were perfect and ideal in shape, the extreme length of their radii of curvature, $\gtrsim 50$ km, compared to the optical pathlength in the recycling cavities, $2\ell\mathcal{N} \lesssim 1$ km, would guarantee that the mesa beam would resonate equally well in the ideal MH-mirrored cavity or in the ideal spherical-mirrored cavity, or in a precisely flat-mirrored cavity; and the baseline Gaussian beam would also resonate, equally well, in all these cavities.

If the third interferometer is operated in narrow-band mode, then the number of round trips the signal light makes in the SR cavity will be much larger than 40, and the geometric optics approximation will begin to fail significantly. More specifically, for ETM reflectivity $R = 0.995$ and optimized narrow banding at $\{500 \text{ Hz}, 1000 \text{ Hz}\}$, the SR mirror's amplitude reflectivities are $\rho = \{0.994, 0.9985\}$ [see discussion following Eq. (3.53) below], corresponding to a number of round trips in the SR cavity $\mathcal{N}_{\text{SR}} \simeq \{180, 400\}$ and Fresnel lengths $r_F \simeq \{6\text{cm}, 9\text{cm}\}$. These Fresnel lengths (the transverse scale for diffractive light spreading) are about $1/3$ to $1/2$ the 95-percent-power diameter of the beam, 16 cm (BL Gaussian) and 20 cm (mesa). As we shall see below, this means that geometric optics can be used to get a rough upper limit on the fractional increase in shot noise due to tilt and irregularities of the SR cavity's mirrors, but not a reliable estimate of the shot-noise increase.

3.4.9 Increase in shot noise due to mirror tilts

3.4.9.1 Foundations

The mirror tilts produce a mismatch between various modes of the light, thereby increasing the shot noise. We shall focus on the shot noise increase at the minimum of the signal light's optical resonance in the arm cavity. This optical resonance is the one that is used to produce a noise minimum for a narrowbanded advanced LIGO interferometer, and it is the right-hand minimum (at $f \simeq 230$ Hz) of the optical noise for the standard wide-band advanced LIGO interferometer (Fig. 3.1), and approximately the minimum of the wide-band interferometer's total noise.

The (unit-norm) modes whose mismatch increases the shot noise are the following:

1. u_0 , the eigenmode of perfect arm cavities.
2. u'_0 , the carrier's eigenmode in an arm cavity with tilted ITM and ETM.
3. u'_s , the signal field's eigenmode in an arm cavity, at the center of its optical resonance, with tilted ITM, ETM, and SRM (signal recycling mirror).
4. u'_t , the field produced when u_0 is transmitted through the signal recycling cavity with tilted ITM and SRM.
5. u'_r , the reference-light field that is beat against the signal light to produce the input to the photodetector. For the baseline homodyne readout system, this will be carrier light u'_0 transmitted through the ITM and SR mirror to the photodiode. If heterodyne readout were used, it would be side-band light transmitted through the PR cavity and SR mirror to the photodetector.

For each primed field u'_ℓ , we denote by δ_ℓ the fraction of its light power that is in parasitic modes and thus has been lost from the fundamental mode u_0 due to mirror tilt:

$$\langle u_0, u'_\ell \rangle^2 = 1 - \delta_\ell. \quad (3.47)$$

The signal amplitude entering the photodetector is proportional to

$$S \propto \langle u'_r, \tilde{\tau}' u'_s \rangle \langle u'_s, u'_0 \rangle \langle u'_0, u_d \rangle. \quad (3.48)$$

The sequence of terms, from right to left, have the following meanings, and we approximate them as follows:

1. $\langle u'_0, u_d \rangle$ describes the influence of tilts on the driving of the arm cavity's eigenmode by the Gaussian driving field. For simplicity, we neglect the tiny coupling of the driving field to the second-order parasitic mode u_2 contained in u_0 and therefore approximate this coupling

amplitude by $\langle u'_0, u_d \rangle = \gamma_0 \langle u'_0, u_0 \rangle = \gamma_0(1 - \alpha_{1E}^2/2 - \alpha_{1I}^2/2) = \gamma_0(1 - \delta_0/2)$. Here the subscripts I and E denote the contributions from the tilts of the ITM and ETM.

2. $\langle u'_s, u'_0 \rangle$ describes the influence of tilts on the driving of the arm cavity's signal field by its carrier field (via the gravitational-wave-induced motion of the mirrors). For simplicity we neglect the (nonzero) overlap between the parasitic modes contained in u'_s and u'_0 , thereby obtaining $\langle u'_s, u'_0 \rangle = (1 - \delta_s/2 - \delta_0/2)$.
3. $\langle u'_r, \tilde{\tau}' u'_s \rangle$ describes the influence of tilts on (i) the passage of the signal u'_s , through the SR cavity (with cavity transmissivity $\tilde{\tau}'$ in the notation of Buonanno and Chen [8]), and on (ii) the overlap of the transmitted signal light with the reference light to produce the photodetector current. Again we neglect correlations between the parasitic components of the fields and therefore approximate the influence of the tilts by $\langle u'_r, \tilde{\tau}' u'_s \rangle \propto (1 - \delta_r/2 - \delta_s/2 - \delta_t/2)$. The δ_j terms represent the loss of overlap due to the parasitic-mode fields (assumed uncorrelated) contained in u'_r (the δ_r term), contained in u'_s (the δ_s term), and generated by the passage of the signal light through the SR cavity, whose mirror tilts deform the transmissivity $\tilde{\tau}'$, (the δ_t term).

If there is no mode cleaner on the interferometer output, then the rms amplitude of the shot noise is $N \propto \sqrt{\langle u'_r, u'_r \rangle} = 1$; i.e., the parasitic-mode components of u'_r contribute to the rms noise amplitude along with the fundamental-mode component. However, a mode cleaner will remove the parasitic components, so that $N \propto \langle u_0, u'_r \rangle = 1 - \delta_r/2$.

Combining the above approximations to the various terms, we find for the ratio of noise power to signal power (which is proportional to the spectral density of shot noise S_h^{shot} at the minimum of the optical resonance):

$$S_h^{\text{shot}} \propto \frac{N^2}{|S|^2} \propto 1 + 2\delta_0 + 2\delta_s + \delta_t + \begin{cases} \delta_r, & \text{no mode cleaner} \\ 0 & \text{with mode cleaner.} \end{cases} \quad (3.49)$$

(Note that we have ignored the increase in shot noise due to carrier-light parasitic fields going out the dark port, Eqs. (3.27) and (3.43), under the assumption that it is negligible, either because of an output mode cleaner or because the arm-cavity-mirror figures and tilts are adequately controlled.)

We shall now examine the various contributions to the shot noise one by one.

3.4.9.2 Carrier light in arm cavity

The fraction of the arm-cavity carrier power that is driven into parasitic modes by tilts of the ETM and ITM is $\delta_0 = \alpha_{1E}^2 + \alpha_{1I}^2$. The loss of this carrier power to parasites increases the shot noise by

$$\left(\frac{\delta S_h^{\text{shot}}}{S_h^{\text{shot}}} \right)_{\text{carrier}}^{\text{MH}} = 2\delta_0 = 4\alpha_1^2 = 0.01 \left(\frac{\theta}{2 \times 10^{-8}} \right)^2. \quad (3.50)$$

Here we have assumed that both mirrors are tilted through the same angle θ but about uncorrelated axes, we have assumed MH mirrors, and we have used Eq. (3.23) for α_1 . The baseline spherical mirrors are four times less sensitive to tilt, so to keep this contribution to the spectral density of shot noise below one per cent, we must control the ITM and ETM tilts to an accuracy

$$\theta_{1\%}^{\text{MH}} = 2 \times 10^{-8}, \quad \theta_{1\%}^{\text{BL}} = 8 \times 10^{-8}. \quad (3.51)$$

These are modest constraints on tilt.

3.4.9.3 Signal light in arm cavity

The signal recycling (SR) cavity presents a complex amplitude reflectivity $\tilde{\rho}' = e^{-\epsilon/F} e^{i\lambda/F}$ to the arm cavity's signal light (Eqs. (5) and (13) of Buonanno and Chen [8]). Here $F = c/2L$ is the interferometer's free spectral range, $\epsilon = \epsilon(R, \rho, \phi)$ and $\lambda = \lambda(R, \rho, \phi)$ are real functions of the ITM power reflectivity R , the SRM amplitude reflectivity ρ and the SR cavity's tuning phase $\phi = (k\ell)_{\text{mod } 2\pi}$, with ℓ the length of the cavity; and our notation is that of Buonanno and Chen [8]. Tilts of the ITM and SRM produce a spatially variable reflectivity $\tilde{\rho}'$. The spatial variations of the modulus $e^{-\epsilon/F}$ of $\tilde{\rho}'$ presumably will have much less influence on the arm cavity's signal eigenmode u'_s than the spatial variations of the phase. (This claim deserves to be checked.) Assuming this is so, then the dominant influence of an ITM or SRM tilt θ is to produce a spatially variable mirror displacement

$$\delta z = \theta r \sin \varphi \quad (3.52)$$

(where φ is azimuthal angle and r is radius), which in turn (*in the SR cavity's geometric optics limit*) produces a spatially variable phase of the cavity reflectivity, $\arg(\tilde{\rho}') = \delta\lambda/F = (d\lambda/d\phi)(k/F)\delta z$. If the cavity were replaced by a single mirror that is displaced through a distance δz_{eff} , then this phase change would be $2k\delta z_{\text{eff}}$. Correspondingly, the tilt of the ITM or SRM produces an effective mirror displacement $\delta z_{\text{eff}} = \mathcal{A}\delta z$, where the amplification factor \mathcal{A} is given by

$$\mathcal{A} = \frac{\delta z_{\text{eff}}}{\delta z} = \frac{d\lambda/d\phi}{2F} = (1 - R)\rho \frac{2\rho + (1 + \rho^2) \cos 2\phi}{(1 + \rho^2) + 2\rho \cos 2\phi}; \quad (3.53)$$

see Eq. (18) of Buonanno and Chen [8].

We shall focus on three configurations for the SR cavity: (i) The standard wideband advanced LIGO configuration (denoted “WB”), for which $R = 0.995$, $\rho = \sqrt{0.93}$, and $\phi = \pi/2 - 0.06$. (ii) An interferometer narrowbanded at a frequency $f = \lambda/2\pi \simeq 500$ Hz with bandwidth $\Delta f = \epsilon/2\pi \simeq 50$ Hz, which has a noise minimum of $\simeq 1 \times 10^{-24}/\sqrt{\text{Hz}}$; this configuration (which we shall denote “500”) is produced by $R = 0.995$, $\rho = 0.994$, and $\phi = 1.541$. (iii) An interferometer narrowbanded at $f = \lambda/2\pi \simeq 1000$ Hz with bandwidth $\Delta f = \epsilon/2\pi \simeq 50$ Hz (and so denoted “1000”), which has a

noise minimum of $\simeq 1 \times 10^{-24}/\sqrt{\text{Hz}}$ and parameters $R = 0.995$, $\rho = 0.9985$, $\phi = 1.556$. For these three configurations the amplification factor is

$$\mathcal{A}_{\text{WB}} = 0.27, \quad \mathcal{A}_{500} = 1.4, \quad \mathcal{A}_{1000} = 5.7. \quad (3.54)$$

We have chosen to compute the shot noise increase at the optical resonance so the signal field u'_s in the arm cavity will be on resonance, just as the carrier field is. This allows us to translate our carrier-field results over to the signal field with only one change: the influence of the tilts of the SRM and ITM must be multiplied by the amplification factor \mathcal{A} . Therefore, the fraction δ_s of the signal field's power that is in the tilt-induced parasitic modes is $\delta_s = \alpha_{1E}^2 + \mathcal{A}^2(\alpha_{1I}^2 + \alpha_{1SR}^2)$. The influence α_{1E}^2 of the ETM is the same as in the case of the carrier, which we have already dealt with, so we shall ignore it here and focus on the two mirrors that make up the SR cavity: the ITM and the SRM. If they both have the same tilt angles θ (but about uncorrelated axes) so $\alpha_{1I}^2 = \alpha_{1SR}^2 \equiv \alpha_1^2$, then these tilts produce a fractional increase in shot noise given by

$$\left(\frac{\delta S_h^{\text{shot}}}{S_h^{\text{shot}}} \right)_{\text{signal}} = 2\delta_s = 4\mathcal{A}^2\alpha_1^2. \quad (3.55)$$

This is greater by the factor \mathcal{A}^2 than the noise (3.50) due to loss of carrier light into parasitic modes, and correspondingly to keep this fractional increase of shot noise below one per cent requires controlling the ITM and SRM tilts to an accuracy $1/\mathcal{A}$ of that in Eq. (3.51):

$$\begin{aligned} \theta_{1\%}^{\text{MH WB}} &= 7 \times 10^{-8}, & \theta_{1\%}^{\text{BL WB}} &= 30 \times 10^{-8}, \\ \theta_{1\%}^{\text{MH } 500} &\gtrsim 1.4 \times 10^{-8}, & \theta_{1\%}^{\text{BL } 500} &\gtrsim 6 \times 10^{-8}, \\ \theta_{1\%}^{\text{MH } 1000} &\gtrsim 0.4 \times 10^{-8}, & \theta_{1\%}^{\text{BL } 1000} &\gtrsim 1.4 \times 10^{-8}. \end{aligned} \quad (3.56)$$

For the narrowbanded interferometers these limits are only lower bounds on $\theta_{1\%}$ because of the failure of the geometric optics limit. As we have seen, the Fresnel length for light trapped in the SR cavity is about 1/2 to 1/3 of the 95-percent-power beam diameter, so transverse spreading of the light will reduce somewhat the SR cavity's amplification factor \mathcal{A} and thence the influence of tilt on the beam asymmetry. We *guess* that this reduction might increase $\theta_{1\%}^{1000}$ by a factor of order 2 over the geometric-optics limit, (3.56); but since \mathcal{A} is only about 1 for $\theta_{1\%}^{1000}$, we *guess* that there is little

increase in $\theta_{1\%}^{500}$; so

$$\begin{aligned}\theta_{1\%}^{\text{MH } 500} &\simeq 1.4 \times 10^{-8}, & \theta_{1\%}^{\text{BL } 500} &\simeq 6 \times 10^{-8}, \\ \theta_{1\%}^{\text{MH } 1000} &\simeq 0.8 \times 10^{-8}, & \theta_{1\%}^{\text{BL } 1000} &\simeq 3 \times 10^{-8}.\end{aligned}\tag{3.57}$$

3.4.9.4 Transmission of signal light through SR cavity

When the ITM or SRM is tilted through an angle θ , producing a spatially dependent mirror displacement $\delta z = \theta r \sin \varphi$, it alters the SR cavity's transmissivity by a spatially dependent amount $\delta \tilde{\tau}' = (d\tilde{\tau}'/d\phi)k\delta z$, *in the geometric optics limit*. When an undistorted signal beam u_0 passes through this spatially variable transmissivity, a fraction

$$\langle u_0, |\delta \tilde{\tau}' / \tilde{\tau}'|^2 u_0 \rangle = \frac{1}{2} \mathcal{B}^2 k^2 \langle r^2 \rangle \theta^2.\tag{3.58}$$

gets transferred to parasitic modes. Here $k = 2\pi/\lambda_o$ is the wave number, $\langle r^2 \rangle = \langle u_0, r^2 u_0 \rangle$ is the beam's mean square radius, which has the values

$$\begin{aligned}\langle r^2 \rangle &= (6.95\text{cm})^2 && \text{for mesa beam} \\ \langle r^2 \rangle &= r_o^2 = (4.70\text{cm})^2 && \text{for BL Gaussian beam,}\end{aligned}\tag{3.59}$$

and

$$\mathcal{B}^2 = \left| \frac{d\tilde{\tau}'/d\phi}{\tilde{\tau}'} \right|^2 = \frac{4R\rho^2}{1 + R\rho^2 + 2\sqrt{R}\rho \cos 2\phi}.\tag{3.60}$$

Here we have used Eq. (11) of Buonanno and Chen [8] for $\tilde{\tau}'$, with the factor $e^{i\phi(\vec{r})}$ in the numerator removed, so as to obtain the transmissivity that carries the field from an input transverse plane to an output transverse plane in the presence of the mirror tilt (which gives ϕ its dependence on transverse position \vec{r}). For our three interferometer configurations, the values of \mathcal{B} are

$$\mathcal{B}_{\text{WB}} = 15, \quad \mathcal{B}_{500} = 33, \quad \mathcal{B}_{1000} = 66.\tag{3.61}$$

When both ITM and SRM are tilted through the same angle θ about uncorrelated axes, the total power transferred into parasitic modes is twice as large as Eq. (3.58) [i.e., δ_t is twice (3.58)], and correspondingly the fractional increase in shot noise is

$$\left(\frac{\delta S_h^{\text{shot}}}{S_h^{\text{shot}}} \right)_{\text{transmission}} = \delta_t = \mathcal{B}^2 k^2 \langle r^2 \rangle \theta^2.\tag{3.62}$$

Inserting the above values for B and $\langle r^2 \rangle$ and insisting that the shot noise not increase by more than

one per cent, we obtain the following constraints on the ITM and SRM tilt angles:

$$\begin{aligned}
\theta_{1\%}^{\text{MH WB}} &= 1.6 \times 10^{-8}, & \theta_{1\%}^{\text{BL WB}} &= 2.4 \times 10^{-8}, \\
\theta_{1\%}^{\text{MH } 500} &\gtrsim 0.7 \times 10^{-8}, & \theta_{1\%}^{\text{BL } 500} &\gtrsim 1.1 \times 10^{-8}, \\
\theta_{1\%}^{\text{MH } 1000} &\gtrsim 0.4 \times 10^{-8}; & \theta_{1\%}^{\text{BL } 1000} &\gtrsim 0.6 \times 10^{-8}.
\end{aligned}
\tag{3.63}$$

For the narrowbanded interferometers, the failure of the geometric optics limit dictates that these estimates of $\theta_{1\%}$ are lower limits; hence the “ \gtrsim ”. As in the case of signal light in an arm cavity reflecting off the SR cavity, so also here, we *guess* that these estimates are fairly good for narrow banding at 500 Hz and are roughly a factor 2 too severe at 1000 Hz, so

$$\begin{aligned}
\theta_{1\%}^{\text{MH } 500} &\simeq 0.7 \times 10^{-8}, & \theta_{1\%}^{\text{BL } 500} &\simeq 1.1 \times 10^{-8} \\
\theta_{1\%}^{\text{MH } 1000} &\simeq 0.7 \times 10^{-8}; & \theta_{1\%}^{\text{BL } 1000} &\simeq 1.1 \times 10^{-8}.
\end{aligned}
\tag{3.64}$$

Equations (3.63) for wideband interferometers and (3.64) for narrowband are the most severe of all our tilt constraints.

3.4.9.5 Reference light for baseline readout

The baseline design for the advanced LIGO interferometers includes an output mode cleaner and homodyne readout. As we have seen [Eq. (3.49)], the mode cleaner makes the shot noise insensitive to (first-order) losses of reference power into parasitic modes.

3.4.9.6 Transmission of RF reference light through a power recycling cavity without an output mode cleaner

In LIGO-I interferometers, by contrast with the baseline design of advanced LIGO, there is no output mode cleaner, and heterodyne readout is used in place of homodyne; i.e., the reference light is radio-frequency (RF) sidebands transmitted through the power recycling cavity to the output port. In this section we shall compare our approximate shot-noise analysis with more careful analyses of shot noise in LIGO-I, by analyzing this type of readout.

Suppose that the PRM or ITM is tilted through an angle θ and is thereby given the space-dependent displacement $\delta z = \theta r \sin \varphi$. Then, in the geometric optics limit, the RF reference light acquires, when passing through the PR cavity, a space-dependent phase shift $(\mathcal{F}/\pi)k\delta z$, where \mathcal{F} is the cavity finesse. The reference light emerging from the cavity therefore has the form $u'_r =$

$u_0 e^{i(\mathcal{F}/\pi)k\delta z}$, for which the fraction of light power in parasitic modes is

$$\left\langle u_0, \left(\frac{\mathcal{F}}{\pi}\right)^2 k^2 \delta z^2 u_0 \right\rangle = \frac{1}{2} \left(\frac{\mathcal{F}}{\pi} k\right)^2 \langle r^2 \rangle \theta^2. \quad (3.65)$$

When both the PRM and the ITM are tilted through the same angle θ but around uncorrelated axes, the parasitic mode power is twice as large [so δ_r is twice (3.65)], and the fractional increase in shot noise due to the loss of this reference-light power is then

$$\left(\frac{\delta S_h^{\text{shot}}}{S_h^{\text{shot}}}\right)_{\text{reference}} = \delta_r = \left(\frac{\mathcal{F}}{\pi} k\right)^2 \langle r^2 \rangle \theta^2. \quad (3.66)$$

For the LIGO-I interferometers \mathcal{F} is rather large, $\mathcal{F} \simeq 125$, which produces a strong sensitivity to mirror tilt. By contrast, for advanced LIGO \mathcal{F} is smaller, $\mathcal{F} \simeq 50$, which (as we shall see) compensates for the larger beam. The result would be about the same sensitivity to tilt as for LIGO-I, if there were no output mode cleaner in advanced LIGO and heterodyne readout were used.

Inserting the LIGO-I finesse $\mathcal{F} \simeq 125$ and mean square beam radius $\langle r^2 \rangle = b^2 = (2.6\text{cm})^2$, and constraining the shot noise increase to less than one per cent, we obtain the following constraint on the ITM and PRM tilts:

$$\theta_{1\%}^{\text{LIGO-I}} \simeq 1.6 \times 10^{-8}. \quad (3.67)$$

This is in remarkably good agreement with a much more careful computation by Fritschel et. al. [16], which gave³ $\theta_{1\%}^{\text{LIGO-I}} = 1.4 \times 10^{-8}$.

Inserting the advanced LIGO finesse $\mathcal{F} \simeq 50$ and mean-square beam radius $\langle r^2 \rangle = (6.95\text{cm})^2$ (mesa) and $(4.70\text{cm})^2$ (BL Gaussian), and constraining the increase in shot noise to no more than one per cent, we obtain

$$\theta_{1\%}^{\text{MH}} \simeq 1.5 \times 10^{-8}, \quad \theta_{1\%}^{\text{BL}} \simeq 2.5 \times 10^{-8},$$

As we have seen, an output mode cleaner will remove this shot noise increase, making these (rather modest) constraints no longer needed.

3.4.10 Increase in shot noise due to mirror figure errors

The increase in spectral density of shot noise due to mirror figure errors is given by the same equation $S_h^{\text{shot}} \propto 1 + 2\delta_0 + 2\delta_s + \delta_t + \{\delta_r \text{ or } 0\}$ as for mirror tilt [Eq. (3.49)], but now δ_ℓ is the fraction of power in the parasitic components of mode u'_ℓ due to figure errors rather than tilt.

³Their result (end of Sec. 2.A) is $\theta < 1.0 \times 10^{-8}$ for the tilts in pitch and in yaw, corresponding to a constraint $\theta < \sqrt{2} \times 10^{-8} \simeq 1.4 \times 10^{-8}$ on the magnitude of the vectorial tilt, for an 0.005 fractional decrease in amplitude signal to noise, which corresponds to our one per cent increase in S_h .

3.4.10.1 Carrier light in arm cavity

Deformations of the ITM and ETM, with optimized tilt compensation, drive a fraction $\delta_0 = \beta_{1cE}^2 + \beta_{1cI}^2$ into parasitic modes [Eqs. (3.40) and (3.41)]. Assuming the same peak-to-valley deformations Δz in the two mirrors' central regions, we obtain for the fractional increase in shot noise

$$\left(\frac{\delta S_h^{\text{shot}}}{S_h}\right)^{\text{MH}} = 2\delta_0 = 4\beta_{1c}^2 \simeq 0.01 \left(\frac{\Delta z}{15\text{nm}}\right)^2. \quad (3.68)$$

Correspondingly, to keep the shot noise increase below one per cent, we must constrain the ITM and ETM deformations to

$$\Delta z_{1\%}^{\text{MH}} \simeq 15\text{nm} \quad (3.69)$$

We have not carried out an analysis of the influence of the ITM and SRM mirror deformations on the baseline Gaussian arm cavity modes, and so cannot say what the analogous constraint is in the baseline case.

3.4.10.2 Signal light in arm cavity

As for tilt, so also for figure errors, the SR cavity amplifies the influence of the errors Δz by a factor \mathcal{A} [Eq. (3.53)], so the ITM and SRM deformations move a fraction $\delta_s = \mathcal{A}^2(\beta_{1cI}^2 + \beta_{1cSR}^2)$ of the arm cavity's signal light into parasitic modes. When the two figure errors have the same magnitude and we wish to keep the resulting shot noise increase below one per cent, this gives rise to constraints on the ITM and SRM figure errors that are $1/\mathcal{A}$ more severe than (3.69). Using the values (3.54) of \mathcal{A} for our three interferometers (wide-band, narrowbanded at 500 Hz and narrowbanded at 1000 Hz), and increasing the limit for the 1000-Hz narrowbanded case by a factor 2 due to failure of the geometric-optics limit [cf. Eq. (3.57) and associated discussion], we obtain the constraints

$$\begin{aligned} \Delta z_{1\%}^{\text{MH WB}} &\simeq 55\text{nm}, \\ \Delta z_{1\%}^{\text{MH } 500} &\gtrsim 10\text{nm}, \text{ with a guess of } \simeq 10\text{nm}, \\ \Delta z_{1\%}^{\text{MH } 1000} &\gtrsim 2.6\text{nm}, \text{ with a guess of } \simeq 5\text{nm}. \end{aligned} \quad (3.70)$$

3.4.10.3 Transmission of signal light through SR cavity

By the same analysis as for mirror tilt (Sec. 3.4.9.4), deformations $\delta z(x, y)$ of the ITM and SRM by the same peak-to-valley amounts Δz produce an increase in shot noise given by

$$\begin{aligned} \left(\frac{\delta S_h^{\text{shot}}}{S_h^{\text{shot}}} \right)_{\text{transmission}} &= \delta_t = 2\mathcal{B}^2 k^2 \langle (\delta z)^2 \rangle \\ &= \frac{1}{4} \mathcal{B}^2 k^2 (\Delta z)^2 \end{aligned} \quad (3.71)$$

[cf. Eq. (3.62) with $\langle (\delta z)^2 \rangle = \langle (\theta r \sin \varphi)^2 \rangle = \frac{1}{2} \langle r^2 \rangle \theta^2$]. Here $\langle (\delta z)^2 \rangle$ is the mean square deviation of ITM or SRM height from the desired figure, and we have approximated this by half the squared amplitude of mirror height fluctuations, which is 1/8 the square of the peak to valley height fluctuations, $(\Delta z)^2/8$. Inserting the values of \mathcal{B} for our three interferometers [Eq. (3.61)] and requiring that the shot noise increase by no more than one per cent, we obtain the following constraints on the ITM and SRM peak to valley deformations:

$$\begin{aligned} \Delta z_{1\%}^{\text{MH WB}} &= \Delta z_{1\%}^{\text{BL WB}} \simeq 2\text{nm}, \\ \Delta z_{1\%}^{\text{MH 500}} &= \Delta z_{1\%}^{\text{BL 500}} \lesssim 1\text{nm}, \text{ with a guess } \simeq 1\text{nm}, \\ \Delta z_{1\%}^{\text{MH 1000}} &= \Delta z_{1\%}^{\text{BL 1000}} \lesssim 0.5\text{nm}, \text{ with a guess } \simeq 1\text{nm}. \end{aligned} \quad (3.72)$$

Here as in Eqs. (3.70), (3.64) and (3.57), the breakdown of the geometric optics limit in the SR cavity has dictated a lower limit and a guess for the narrowbanded interferometers.

These are the most serious of our constraints on the mirror figures of advanced interferometers, and they are the same for mesa and baseline Gaussian beams, because transmission through the SR cavity is governed (at least roughly) by geometric optics. The one small difference is that the central region over which the peak-to-valley deformations are constrained (the region containing ~ 95 per cent of the light power) is larger for MH mirrors (central radius about 10 cm) than for the baseline Gaussian mirrors (central radius about 8 cm).

The mirror-figure constraint (3.72) is sufficiently severe, at least in the case of narrowbanded interferometers, that it might be worth considering reducing the degeneracy of the SR cavity by making the entrance faces of the ITM's into lenses that bring the signal light (and inevitably also the carrier light) to a focus somewhere near the SRM (and PRM). Since the constraint (3.72) is the same, whether the mirrors are MH or spherical, this recommendation is not dependent on whether mesa beams are implemented.

3.4.10.4 Reference light for baseline readout

As for mirror tilt, so for figure errors, the output mode cleaner in advanced LIGO makes its shot noise insensitive to the loss of reference light into parasitic modes.

3.4.10.5 Transmission of RF reference light through power recycling cavity without an output mode cleaner

In the LIGO-I-like case of no output mode cleaner and heterodyne readout with RF sideband light, the loss of reference light to parasites does increase the shot noise. By the same argument as for mirror tilt (Sec. 3.4.9.6), the net shot noise increase due to deformations $\delta z(x, y)$ of the ITM and PRM is

$$\begin{aligned} \left(\frac{\delta S_h^{\text{shot}}}{S_h^{\text{shot}}} \right)_{\text{reference}} &= \delta_r = 2 \left(\frac{\mathcal{F}}{\pi} k \right)^2 \langle (\delta z)^2 \rangle \\ &= \frac{1}{4} \left(\frac{\mathcal{F}}{\pi} k \right)^2 (\Delta z)^2 \end{aligned} \quad (3.73)$$

[cf. Eq. (3.66) with $\langle (\delta z)^2 \rangle = \langle (\theta r \sin \varphi)^2 \rangle = \frac{1}{2} \langle r^2 \rangle \theta^2$]. Inserting the finesse of LIGO-I ($\mathcal{F} \simeq 125$) and advanced LIGO ($\mathcal{F} \simeq 50$) and insisting that the shot noise not be increased by more than one per cent, we obtain the following constraints on the central-region peak-to-valley deformations of the PRM and SRM:

$$\begin{aligned} \Delta z_{1\%}^{\text{LIGO-I}} &= 0.8 \text{nm} , \\ \Delta z_{1\%}^{\text{MH}} &= \Delta z_{1\%}^{\text{BL}} = 2 \text{nm} , \end{aligned} \quad (3.74)$$

The LIGO-I constraint is rather severe, but is near the actual mirror errors (in the absence of thermal lensing). The advanced LIGO constraint is less severe, but is actually irrelevant because of advanced LIGO's output mode cleaner (and its use of homodyne readout rather than heterodyne).

3.5 Conclusions and Future Research

The thermoelastic benefits of MH mirrors are sufficiently great, and the tightened constraints that they place on mirror figures, positions and tilts are sufficiently modest, that MH mirrors have been adopted by the LSC as an option for advanced LIGO, and will be incorporated into future modeling along with spherical mirrors.

Two methods of manufacturing MH mirrors are being explored by the LIGO Laboratory: diamond cutting, and evaporative coating. The immediate goal is to determine the accuracy and reproducibility with which the desired MH mirror figures can be produced by each of these methods.

If MH mirrors are to be considered seriously for advanced LIGO, it is necessary to develop laboratory experience with them and with the mesa-beam optical cavities that they produce. An experimental effort in this direction has been initiated by Phil Willems of the LIGO Laboratory at Caltech.

As we have discussed in Secs. 3.4.8 and 3.4.10.3, the constraints on mirror figure in the recycling cavities are rather worrisome. More accurate studies of this are needed. As one aspect of these studies, it is important to check our geometric-optics-based claims that the MH recycling cavities are not much more sensitive to mirror figure errors than the baseline Gaussian-beam cavities.

If the recycling-cavity-induced constraints on mirror figure are found to be as serious as our estimates suggest, then it seems worthwhile to carry out studies of the option of converting the input faces of the ITM's into lenses that make the recycling cavities much less degenerate. PhD students in Thorne's research group may carry out these studies as part of the broadening experience that is a standard part of their education.

Bill Kells has suggested the possibility of operating the advanced LIGO interferometers initially with spherical mirrors and Gaussian beams, and later switching to mesa beams by altering only the ETM's. It seems to us that this option is worth detailed study. If the ITM input faces are turned into lenses that reduce the near degeneracy of the recycling cavities, then it might be possible to keep the lenses weak enough that the figures of the recycling mirrors can be the same for Gaussian and mesa beams, while still relaxing the mirror figure constraints to an acceptable level. This needs study.

Bibliography

- [1] V.B. Braginsky, E. d'Ambrosio, R. O'Shaughnessy, S. Strigin, K.S. Thorne, and S. Vyatchanin, *Presentation at august 2001 lsc meeting*, LIGO Report Numbers LIGO-G010333-00-D and LIGO-010297-00-D (15 August 2001).
- [2] ———, *Presentation at march 2001 lsc meeting*, LIGO Report Number LIGO-G010151-00-R (16 March 2001).
- [3] ———, *Status report on mexican-hat flat-topped beams for advanced ligo*, LIGO Report Number LIGOT030009-00-R (23 January 2003).
- [4] ———, *Presentation at september 2002 mit meeting on flat-topped beams for ligo-ii*, LIGO Report Number LIGO-G020543-00-R (6 September 2002).
- [5] V.B. Braginsky, M.L. Gorodetsky, and S.P. Vyatchanin, *Phys. Lett. A* **264** (1999), 1.
- [6] A. Buonanno and Y. Chen, *Phys. Rev. D* **64** (2001), 042006.
- [7] ———, *Phys. Rev. D* **65** (2002), 042001.
- [8] ———, *Phys. Rev. D* (2002), submitted; gr-qc/0208048.
- [9] The LIGO-I Collaboration, *Phys. Rev. D* (2003), description of state of IFOs during S1 searches, in preparation.
- [10] ———, *Phys. Rev. D* (2003), four papers on S1 searches in preparation.
- [11] C. Cutler and K.S. Thorne, *General Relativity and Gravitation: Proceedings of the 16th International Conference* (N. Bishop and S.D. Maharaj, eds.), World Scientific, 2002, p. 72.
- [12] E. D'Ambrosio, *Phys. Rev. D* **67** (2003), 102004.
- [13] E. d'Ambrosio, R. O'Shaughnessy, S. Strigin, K.S. Thorne, and S. Vyatchanin, *Reducing thermoelastic noise in gravitational-wave interferometers by flattening the light beams, draft of 3 september 2002*, *Phys. Rev. D* (2000), in preparation; available as the file beamshape020903.pdf at <http://www.cco.caltech.edu/~kip/ftp/>.

- [14] E. d'Ambrosio, R. O'Shaughnessy, and K.S. Thorne, *Presentation at august 2000 lsc meeting*, LIGO Report Number LIGO-G000223-00-D (16 August 2000).
- [15] L.S. Finn and the LIGO Scientific Collaboration, *Bench*, <http://fiji.nirvana.phys.psu.edu/~lsf/Benchmarks/>.
- [16] P. Fritschel, N. Mavalvala, D. Shoemaker, D. Sigg, M. Zucker, and G. Gonzalez, *Applied Optics* **37** (1998), 6734.
- [17] Yu. Levin, *Phys. Rev. D* **57** (1998), 659.
- [18] Y.-T. Liu and K.S. Thorne, *Phys. Rev. D* **62** (2000), 122002.
- [19] R. O'Shaughnessy, S. Strigin, and S. Vyatchanin, *Phys. Rev. D* (2003), submitted.
- [20] D. Shoemaker and the LIGO Scientific Collaboration, *Advanced ligo project book*, LIGO Report Number LIGO M030023-00M (2003).
- [21] A.E. Siegmann, *An introduction to lasers and masers*, McGraw-Hill, 1971.

Chapter 4

Calculations of thermoelastic noise and interferometer sensitivity to perturbation for the mesa-beam proposal for advanced LIGO

[This section consists of a draft version of a paper by R. O’Shaughnessy, S. Strigin, and S. Vyatchanin, to be submitted after LSC review to Phys. Rev. D. R. O’Shaughnessy wrote the text of this paper.]

Abstract

Thermoelastic noise will be the most significant noise source in advanced-LIGO interferometers with sapphire test masses. The standard plan for advanced-LIGO has optimized the optics, within the framework of conventional mirrors, to reduce thermoelastic noise. Recently, we and our collaborators have proposed going beyond the bounds of traditional optics to increase the effective beam spot size and thus lower thermoelastic noise. One particular proposal for mirror shapes (“Mexican-hat mirrors”) yields the class of “mesa” beams.

In this paper, we outline a general procedure for analyzing light propagating in individual arm cavities, and the associated thermoelastic noise, in the presence of arbitrary optics. We apply these procedures to study the Mexican-hat proposal. Results obtained by the techniques of this paper were presented elsewhere, to demonstrate that the Mexican-hat proposal for advanced-LIGO both significantly lowers thermoelastic noise and does not significantly complicate the function of the interferometer.

4.1 Introduction

Thermoelastic noise will be the dominant noise source in the advanced LIGO interferometers for frequencies at which they are most sensitive, should the advanced LIGO upgrade follow the baseline design [1, 2]. Other noise processes (e.g., conventional bulk thermal noise; coating thermal noise; unified optical noise) could be important, but are expected to be smaller in noise power, by a factor ~ 9 (at the minimum of the noise curve), than thermoelastic noise. As a result, the sensitivity of the advanced LIGO interferometers can be significantly increased.

In a recent paper—the mesa-beam interferometer paper (MBI) [4], summarized briefly in a short LIGO technical document (MBI-S) [5]—we and our collaborators have argued that a significant reduction in thermoelastic noise can be achieved by using modified optics that reshape the beam from a conventional (gaussian) profile into a *mesa-beam* shape. We furthermore performed calculations which imply (though do not comprehensively demonstrate) that this proposal would not place significant burdens on the experimental community (e.g., in the need for slightly improved tilt control systems and improved mirror figure accuracies).

The previous publications (MBI and MBI-S) survey the physically and practically relevant features of our proposal—that is, they demonstrate that the proposal should work, and sketch how to implement it. Those publications intentionally omitted many supporting details, to be provided in this paper and a companion paper by Erika d’Ambrosio [6]. In particular, this paper provides the computational details underlying our evaluations of thermoelastic noise in various interferometer configurations; both this paper and d’Ambrosio [6] provide computational details of our studies of the influence of mirror imperfections (tilts, displacements, figure errors) on interferometer performance).

Further, this paper considers a broader class of possible improvements than MBI. The MBI paper discussed using mesa beams reflecting off of the standard advanced LIGO cylindrical mirrors. In this paper, we analyze mesa beams reflecting off more generically shaped mirrors—not merely cylinders of different dimensions, but also frustum-shaped mirrors.¹

4.1.1 Outline of paper

In this paper, we provide details behind the calculations discussed in detail in MBI. For the reader not familiar with the principles behind thermoelastic noise, cavity optics, and cavity perturbation theory, the next four sections (Sections 4.2, 4.3, 4.4, and 4.5) provide a brief review of the equations and computational techniques we employ. To be more explicit, in Sec. 4.2 we describe how to construct an idealized interferometer based on mesa beams while remaining within LIGO design constraints (i.e. losses per bounce). [In the process, we also briefly review the theory of and introduce notation for optical cavity eigenfunctions and propagation operators.] We assume this interferometer, which

¹A frustum is a geometric shape arising between two parallel planes that intersect an axisymmetric cone perpendicular to the cone’s axis.

has mesa beams resonating within the arm cavities, is driven by conventional (i.e. Gaussian) lasers, as described in Sec. 4.3. In Sec. 4.4, we discuss how to analyze, via perturbation theory, the influence of mirror errors on interferometer performance. [In Section 4.7 we will apply this general formalism to mirror transverse position errors (displacement), orientation errors (tilt), and figure error (shapes).] Finally, in Sec. 4.5, we describe how we numerically implemented the calculations described in Sections 4.2 and 4.4. With those tools established, in Sections 4.6 and 4.7 we closely follow the arguments of MBI, carefully explaining the details behind those arguments.

Specifically, in Sec. 4.6, we use explicit forms for the mesa beam and its diffraction losses presented in Sections 4.2 and 4.5 to determine which mesa beams can operate in a given cavity, when the mesa beams are reflecting off of the standard advanced LIGO mirrors (i.e. cylinders of radius 15.7 cm and thickness 13 cm), as well as off mirrors of more generic shape (i.e. other cylinders and frustums). We then employ an expression for the thermoelastic noise associated with a given mirror and beam configuration (developed generally in Sec. 4.2 and discussed practically in Sec. 4.5) to determine the thermoelastic noise associated with a given configuration. Table 4.4 summarizes our results for designs which produce lower thermoelastic noise than the baseline advanced LIGO design.

Section 4.7 describes how sensitive an interferometer using Mexican-hat mirrors will be to perturbation. More explicitly, Sections 4.7.1, 4.7.2, and 4.7.3 demonstrate by way of explicit comparison that arm cavities using spherical and mesa-beam mirrors appropriate to advanced LIGO are roughly equally sensitive to perturbations. Section 4.7.3 describes how sensitive the overall interferometer (as measured by its cavity gain, its dark port power, and its thermoelastic noise) will be to tilt perturbations. Finally, Section 4.7.5 determines how sensitive (as measured by cavity gain, dark port power, and thermoelastic noise) a mesa-beam advanced LIGO design without signal recycling would be to mirror figure errors.

Augmented by additional results discussed in MBI, which demonstrate that the full interferometer including signal recycling will be only moderately more sensitive, we conclude mesa-beam interferometers could be used in advanced LIGO without serious concern.

Guide to reader

The MBI paper reports the *results* of calculations developed in this paper and in one by Erika d’Ambrosio [6], emphasizing their practical significance. On the other hand, this paper emphasizes the *method* by which those results were obtained. As a result, this paper has a dramatically different structure than MBI: two of the first sections (Sections 4.2 and 4.4) outline the equations that we have solved; the next section, Section 4.5, describes how we solved them; and two of the last sections, Sections 4.6 and 4.7, briefly summarize the specific technical results quoted in MBI and explain how those results were obtained.

The reader who wants only a cursory survey of the techniques used to obtain the MBI results

should skim Sections II-IV, emphasizing the summaries in Sections 4.2.5 and 4.4.5; this reader should completely skip over the implementation section (Sec. 4.5). Indeed, we recommend that a reader interested primarily in better understanding results presented in MBI should work *backwards*, first establishing a *cross-reference* between a result in MBI and a result in Sec. 4.6 or 4.7 of this paper, then reading in this paper the material surrounding the cross-reference.

On the other hand, a reader who wishes to verify or generalize our computations should read this paper from beginning to end.

4.1.2 Connection with other published work

The MBI paper [4] and the MBI-S LIGO technical document [5] survey our results, emphasizing their practical significance to advanced LIGO interferometer design. This paper provides details behind many relatively lengthy analytic or numerical calculations which those papers touched upon only briefly.

Another of our MBI collaborators, Erkia D’Ambrosio, has performed many of the same calculations (e.g., the sensitivity of the interferometer to tilt and to defects) using a sophisticated, standard, and trusted tool—the so-called “FFT code”—designed and developed specifically to study the behavior of general interferometers [6]. As noted in MBI and in her paper, our two independent approaches agree well.

4.1.3 Notation

4.1.3.1 Symbols

b The natural diffraction length scale (i.e. Fresnel length) associated with the previous two parameters: $b \equiv \sqrt{L/k} = 2.6$ cm.

C_V Specific heat of test mass, per unit mass, at constant volume

D Mesa beam characteristic scale, Eq. (4.5).

E Young’s modulus. An elastic parameter in the stress tensor for an isotropic material [Eq. (4.10b)]. (See Appendix B.2 for a discussion of the explicit value used.)

ETM End test mass of an arm cavity (i.e. the mass opposite the end light enters the cavity).

f Frequency, units s^{-1} . Typically, we are interested in values of f near the peak sensitivity of LIGO (i.e. $f \approx 100$ Hz).

g_{ab} Metric of 3-space. (See comments on tensor notation, below).

$h_{1,2}$ Height of mirror 1’s (or 2’s, respectively) surface.

I Thermoelastic noise integral, Eqs. (4.8) and (4.9).

ITM Input test mass of an arm cavity.

k The wavevector of light in the cavity, $k = 2\pi/\lambda$ with $\lambda = 1064\text{nm}$.

k_b Boltzmann's constant [$1.38 \times 10^{-16} \text{ g cm s}^{-1}\text{K}^{-1}$]

L The length of the LIGO arm cavity [$L = 3.99901\text{km}$].

\mathcal{L} Some number associated with diffraction losses. When used alone, denotes the total diffraction losses associated with one round trip through the cavity. When given certain subscripts, as with \mathcal{L}_1 , denotes the clipping-approximation estimate of diffraction losses associated with a single reflection off of mirror 1.

M Mass of the test mass [40 kg].

$P(r)$ The intensity distribution of laser light in a cavity, normalized to unity (i.e. $\int P(r)d(\text{area}) = 1$).

\vec{r} A vector transverse to the optic axis of the cavity (i.e. two-dimensional).

s Transverse displacement of an arm cavity's ETM.

T_{ab} Elastic stress tensor for an isotropic medium, Eq. (4.10b). (See comments on tensor notation, below).

u Electric field on some plane in the arm cavity (typically, the face of the ITM, propagating away from the ITM), renormalized to unit norm (i.e. $\int |u|^2 d(\text{area}) = 1$). Subscripts identify the specific system u refers to.

y^a The displacement vector field associated with an elastic distortion.

x, y The cartesian components of the 2-vector \vec{r} .

z A coordinate measured along the optic axis of the cavity (i.e. in a direction perpendicular to \vec{r}).

\hat{z}^a A unit vector pointing in the positive z direction.

α_l Coefficient of linear thermal expansion, assumed isotropic. For an isotropic material, the differential change in volume with temperature is $dV/dT/V = 3\alpha_l$. (See Appendix B.2.)

$\alpha_{1,2}$ Norm of the first- and second-order corrections to the optical state of the cavity when the mirror is tilted [i.e. in Eqs. (4.63) and (4.64)].

κ Thermal conductivity, assumed isotropic (see Appendix B.2).

η An eigenvalue of Eq. (4.13).

λ An eigenvalue of Eq. (4.29).

ρ Density of test mass. (See Appendix B.2).

σ Poisson ratio. An elastic parameter in the stress tensor for an isotropic material [Eq. (4.10b)]

θ An angle through which the mirror is tilted.

Θ The expansion associated with a displacement field $\Theta = \nabla_a y^a$ [e.g., in Eq. (4.9)]. Also denotes a unit step function [e.g., Eq. (4.11a)].

$\omega = 2\pi f$.

$\zeta_{1,2}$ Norm of the first- and second-order corrections to the optical state of the cavity when the mirror is displaced [i.e. in Eqs. (4.58) and (4.59)].

ζ_m Used *only* in Eq. (4.44), ζ_m denotes the m th zero of $J_1(x)$.

4.1.3.2 Tensor notation

The elasticity calculations (e.g., involving the elasticity tensor T_{ab}) involve tensors. While the authors prefer to interpret these expressions with an abstract tensor notation (i.e. Wald [7]), the reader may without loss of generality use global cartesian coordinates and interpret the latin indices a and b as running from $1 \dots 3$, indexing the coordinate directions. In this case, the reader should interpret ∇_a as the coordinate partial derivative $\partial/\partial x^a$ and g_{ab} as the 3×3 identity matrix.

4.1.3.3 State and operator notation

As explained briefly in Sec. 4.2.3.1, because propagation of the state of the field from one plane to another is merely a unitary transformation, we substantially simplify the many integral-equation operations by using a quantum-mechanics-style notation for the operators. For example, the integral operation

$$E(\vec{r}', z') = \int d^2 r U(\vec{r}', z'; \vec{r}, z) E(\vec{r}, z) \quad (4.1)$$

is represented as

$$E(z') = U(z', z) E(z) . \quad (4.2)$$

Further, typically the relevant planes (i.e. the z, z') are known, and omitted, in expressions like the above.

Under a similar philosophy, we also use quantum-mechanics-style notation for the values of the electric field on a specific plane; in other words, we represent states of the optical cavity using

quantum-mechanical state notation. For example, the optical state of the cavity may be equivalently denoted by the explicit functional form u or $u(\vec{r})$, or alternatively as the state vector $|u\rangle$. In a similar spirit, we index the potential solutions to the eigenequation [Eq. (4.29)] by an integer p , and denote the set of basis states as $|p\rangle$. More generally we can represent *any* function defined on a two-dimensional plane—not necessarily normalized; not necessarily a potential state—as a state vector.

We use this notation in particular to simplify notation for inner products between two states $|u\rangle$ and $|v\rangle$:

$$\langle v|u\rangle \equiv \int d^2r v^*(r)u(r) . \quad (4.3)$$

Finally, to simplify the frequent expressions that involve the norm of a state, we use the symbol $\|\cdot\|$ to denote L^2 norms:

$$\|u\|^2 \equiv \langle u|u\rangle = \int d^2r |u(\vec{r})|^2 \quad (4.4)$$

Many quantum mechanics textbooks written over the past half century employ this notation (cf., e.g., [8, 9, 10]).

4.2 Theory of an idealized Mexican-hat arm cavity

In this section, we review the equations needed to understand the design of an arm cavity which (i) uses mesa beams, (ii) reduces thermoelastic noise in an advanced LIGO interferometer, and (iii) satisfies advanced LIGO design constraints (i.e. the diffraction losses per bounce are less than the advanced LIGO design threshold of 10ppm, and the mirrors have mass 40 kg).

Specifically, we introduce the mesa beam, defined for the purposes of this paper by Eq. (4.5). We define the mirror surfaces which confine this beam by Eq. (4.6). From Eq. (4.5) we can determine the beam intensity profile. We can relate the beam intensity profile to the power spectrum of thermoelastic noise [Eq. (4.7)] via Eqs. (4.9) and (4.10). Finally, we describe how to limit attention to only those mesa beams with low diffraction losses, using an approximate approach [the clipping approximation, Eq. (4.16)] which we will test against an exact expression [Eq. (4.15)].

4.2.1 Mexican-hat mirrors and mesa beams

The MBI paper presented explicit forms for certain specific mesa beams (“canonical mesa beams”) produced inside a cavity bounded by two *identical* mirrors. [In MBI, only the values of the beam fields at the surface of the ITM (specifically, values propagating away from the ITM surface) were given; Appendix B.3 provides a more comprehensive discussion.] We summarize Eqs. (3.3), (3.6), and (3.7) to provide a compact summary of how to compute the values u_{mesa} of a mesa beam’s

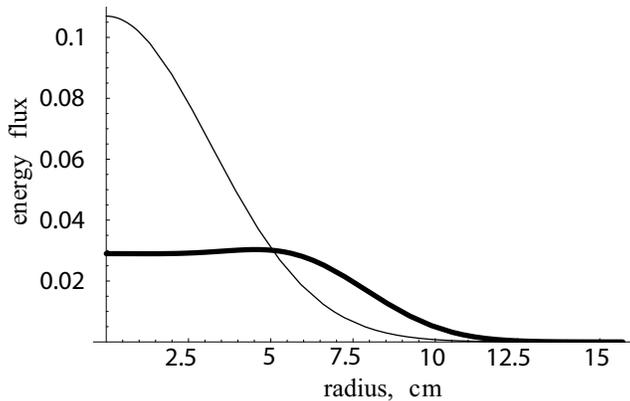


Figure 4.1: The heavy solid curve is an example of a mesa beam’s intensity profile; the lighter curve is the intensity profile of a Gaussian beam with similar losses. More precisely, this curve shows the energy fluxes per unit area (normalized to unity) for the mesa beam (thick curve) and Gaussian beam (thin curve) that have 10ppm of diffraction losses on a mirror with coated radius 15.7 cm. [cf. Fig. 3.3]

normalized electric field on the mirror face, at radius r from the optic axis:

$$U(\vec{r}, D) = \int_{r' < D} d^2r' \exp \left[-\frac{(\vec{r} - \vec{r}')^2(1+i)}{2b^2} \right] \quad (4.5a)$$

$$N^2(D) \equiv \int_0^\infty |U(D, r)|^2 2\pi r dr \quad (4.5b)$$

$$u_{\text{mesa}}(D, r) \equiv U(D, r)/N(D) \quad (4.5c)$$

where in the first integral the integration is over all points with $r' \equiv |\vec{r}'| < D$. [For readers who wish to numerically explore mesa fields themselves, note that MBI Eq. (2.5) provides a more efficient means of calculating the mesa amplitude function $u(D, r)$.] Figure 4.1 provides an example of a mesa-beam intensity distribution $|u(D, r)|^2$.

As discussed in Appendix B.3 and in MBI Sec. II D [cf. Eq. (3.13)], the mirrors that reflect mesa beams back into themselves (denoted *Mexican-hat mirrors*) necessarily have a continuous height function h_{MH} given by

$$kh_{\text{MH}} = \text{Arg}[u_{\text{mesa}}(D, r)] . \quad (4.6)$$

where $k = 2\pi/\lambda$ is the wavenumber of light in the arm cavity. The solid curve in Figure 4.2 provides an example of a Mexican-hat mirror.

4.2.2 Thermoelastic noise

Using the expressions in the previous section, one can design idealized interferometers which operate with mesa beams and Mexican-hat mirrors. The thermoelastic noise associated with the resulting interferometer design can be discovered with the techniques of Liu and Thorne (LT) [12], who employ

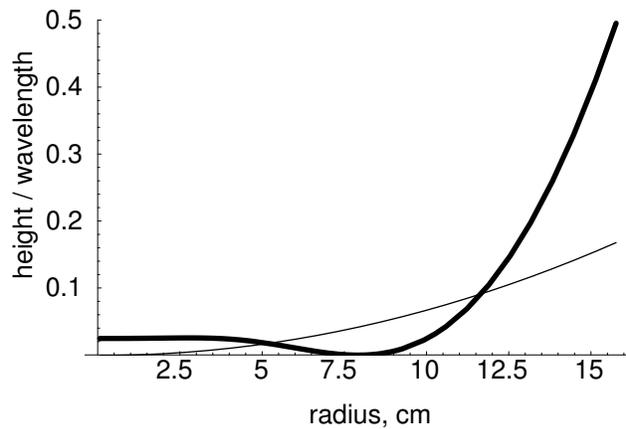


Figure 4.2: The heavy solid curve is an example of a Mexican-hat mirror; the lighter curve is an example of a spherical (i.e. parabolic) mirror. More precisely, this curve shows the spherical and Mexican-hat mirrors which produce the beams described in Fig. 4.1 [cf. Fig. 3.3].

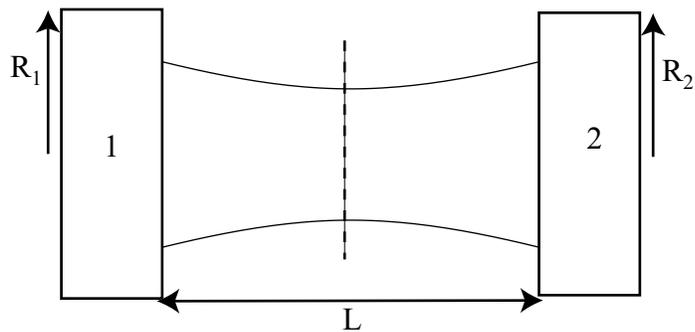


Figure 4.3: This figure outlines our notation for an optical cavity: L is the separation between mirrors, and $R_{1,2}$ are the physical radii of the two bounding mirrors. The optical cavities studied in detail in this paper are bounded by two *identical* mirrors (e.g., $R_1 = R_2$).

Levin's direct method [13] to derive the following formula for an interferometers' thermoelastic noise in terms of a noise integral I_A :

$$S_h(f) = 4\kappa k_B \left(\frac{\alpha_l ET}{C_V(1-2\sigma)\rho\omega L} \right)^2 \sum_{A=1}^4 I_A ; \quad (4.7)$$

cf. LT Eqs. (3), (4) and (13); the notation used here has been described in Sec. 4.1.3.

The sum on $A = 1, 2, 3, 4$ is over the four test masses, and the quantity summed is the noise integral

$$I_A = \frac{2}{F_o^2} \int_{V_A} \langle (\vec{\nabla}\Theta)^2 \rangle d\text{volume} . \quad (4.8)$$

In Eq. (4.8), Θ is the expansion (i.e. fractional volume change) resulting from a pressure—sinusoidally oscillating at frequency $f = \omega/2\pi$, proportional to the light beam intensity, with net force F_o —imposed on the face of the mirror; the integral is over V_A , the volume of test mass A ; and the average $\langle \dots \rangle$ is over one oscillation period $1/f$. [Note that the dimensions of I_A are length/force² = s⁴g⁻²cm⁻¹.]

In the analysis that leads to the above expressions [Eqs. (4.7) and (4.9)], LT [12] have made certain idealizations and approximations, some better than others.² We shall employ this expression as it stands, using the precise parameter values provided in Appendix B.2

4.2.2.1 Effective static elasticity model for I_A

For the frequencies at which advanced LIGO will be most sensitive ($f \sim 100\text{Hz}$), the interior of the test mass will respond effectively instantaneously to the imposed pressure profile in the model problem described above.³ As a result, by going to the accelerating frame of the test mass and following that accelerating frame for one period to simplify the average, we can approximate I_A by

$$I_A \approx \frac{1}{F_o^2} \int_{V_A} (\vec{\nabla}\Theta)^2 d\text{volume} . \quad (4.9)$$

where the expansion $\Theta = \nabla_a y^a$ arises due to the response (i.e. local displacement field) y^a of the test-mass substrate to the following *static* conditions:

1. *Static (accelerating-frame) force*: Some net force $F_o \hat{z}$ acts uniformly throughout the test mass (i.e. provides a force density F_o/V_A). To be specific, using the equations for an isotropic elastic

²For example, they idealize the test-mass material as an isotropic material; however, the proposed test mass material, sapphire, is only moderately well approximated as isotropic. They also assume the elastic response of the solid is *quasistatic* (introducing errors of order the sound crossing time over the gravitational wave period, $\sim 10^{-3}$) and that heat diffuses slowly so the elastic response of the solid to the imposed pressure profile is *adiabatic* (introducing errors of order the diffusion length over the beam size, or $\sim 0.3 \text{ mm}/10 \text{ cm} \sim 3 \times 10^{-3}$).

³The mirrors considered in this paper have characteristic dimension of order $H \sim 10 \text{ cm}$. Thus, the sound crossing time is of order $H/c_s \sim H/\sqrt{E/\rho} \sim 10^{-5} \text{ sec}$. Since the body responds elastically to imposed forces on times of order a few sound-crossing times, the elastic response to a force which is imposed at frequency $\sim 100\text{Hz}$ is effectively instantaneous. Equivalently, the same conclusion follows because the gravitational wave frequency is far below the frequency of any resonance of the test mass.

solid (cf. Blandford and Thorne [14]), the response field y^a satisfies

$$\nabla^a T_{ab} = \hat{z}_b F_o / V_A, \quad (4.10a)$$

$$T_{ab} \equiv -\frac{E}{(1+\sigma)(1-2\sigma)} \times \left[\frac{(1-2\sigma)}{2} (\nabla_a y_b + \nabla_b y_a) + \sigma \Theta g_{ab} \right]. \quad (4.10b)$$

2. *Static pressure on mirror face:* An equal and opposite net force $-F_o \hat{z}$ acts on the mirror face, with a distribution $P(r)$ proportional to the beam intensity profile. As a result, the displacement field y^a must satisfy

$$T_{ab}(r, z=0) \hat{z}^b = -F_o P(r) \hat{z}_a, \quad (4.10c)$$

where the mirror's front (reflecting) surface is at $z=0$ and where the mirror lies below $z=0$.

3. *Break translation symmetry:* Finally, translation symmetry must be broken (i.e. so numerical simulations converge to a unique solution). To break translation symmetry, we fix “the location” of the mirror, where that location is determined as the average location of points in a set \mathcal{R} . Therefore, to specify a unique solution, we require the displacement field y^a satisfy

$$\langle y^a \rangle_{\mathcal{R}} = 0 \quad (4.10d)$$

where \mathcal{R} is some arbitrary nonempty set and $\langle \cdot \rangle_{\mathcal{R}}$ denotes the average over \mathcal{R} .

[This quasistatic approximation to I_A and the overall thermoelastic power spectrum [Eq. (4.7)] involves errors only of order 10^{-3} relative to the general Liu-Thorne expression (4.9) for I_A [12]; cf. footnotes 2 and 3, or Appendix B.4.]

4.2.3 Propagation, the eigencondition, and diffraction losses

In the previous two sections, we described mesa beams and the technique by which one can determine, given a beam and four mirrors that support that beam (two identical ITMs and two identical ETMs), the thermoelastic noise an interferometer using those mirrors and beams would experience. But not all combinations of mesa beams and mirrors satisfy advanced LIGO design constraints; the advanced LIGO design specifications require each arm cavity have low diffraction losses per bounce (10ppm).

In this section, we describe how the diffraction losses can be precisely computed [Eq. (4.15)] and estimated [the clipping approximation, Eq. (4.16)]. Using the diffraction losses, one can then find combinations of mesa beams and mirrors which produce low diffraction losses per bounce.

Our analysis proceeds by describing in general terms the propagation of light in an arm cavity,

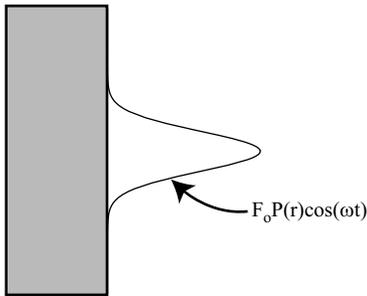


Figure 4.4: Using the fluctuation-dissipation theorem, the effect of thermal fluctuations on the phase fronts of a beam reflecting off the surface of a mirror (and hence the thermoelastic noise induced by the mirror in the phase) can be computed by applying a sinusoidally-oscillating force proportional to the beam intensity profile. [For the purposes of illustration, a Gaussian beam profile is shown, not necessarily to scale.]

focusing in particular on eigenmodes and their loss per round trip. In the process of justifying our simple estimate of diffraction loss, the clipping approximation [Eq. (4.16)], we will review the definitions and properties of eigenmodes of the arm cavity. This review establishes notation and conventions for our presentation of optical perturbation theory, as discussed in Sec. 4.4.

4.2.3.1 Principles of paraxial optics

The light in LIGO arm cavities is well-described by the paraxial approximation. In this section, primarily to establish notation conventions, we briefly review the propagation of light under the paraxial approximation. A detailed description of the relevant physics can be found in standard references.⁴ Briefly, in this approximation, we can completely describe the state of the wave by the wavelength λ of the light used and the values of the wave amplitude on some fixed plane $z = \text{constant}$ (i.e. transverse to the optic axis) and at some time t (cf. [14]). The values at any causally-related later combination t', z' —which must be separated from the transverse plane z at time t by a light ray (i.e. $t' - t = |z - z'|$, mod reflections)—can be deduced by applying the appropriate linear functional to these states.

For example, we could characterize the optical state by the value of the electric field on some plane z .⁵ We would then relate the field at any other plane and at any other time to our initial state via the linear operation

$$E(\vec{r}', z', t') = \int d^2r U(\vec{r}', z'; \vec{r}, z) E(\vec{r}, z, t),$$

cf. Eq. (4.1). For compactness and clarity, we employ a quantum-mechanics-motivated notation,⁶ in

⁴ For a more pedagogical and yet brief presentation, we recommend Blandford and Thorne's treatment [14]. Other pedagogical (and technical) treatments can be found in many laser physics books, e.g., [15, 16, 17, 18]. We also recommend the collection of original research articles on laser and resonator physics compiled by Barnes [19].

⁵ We limit attention to a single polarization, e.g., the polarization excited by LIGO.

⁶ Most quantum mechanics textbooks written over the past half-century (e.g. [8, 9, 10]) have adopted a similar

which the integrals are suppressed and the operation above is denoted by [cf. Eq. (4.2)]

$$E(z', t') = U(z', z)E(z, t).$$

By way of example, the following are kernels of integral operators which describe free propagation down a length L ; reflection off a mirror of height $h_{1,2}$; and a “window” that cuts out all light outside a radius $r = R_{1,2}$, respectively:

$$G_L(\vec{r}, \vec{r}') \equiv -i \frac{k}{2\pi L} \exp i \left[\frac{(\vec{r} - \vec{r}')^2}{2L/k} + kL \right] \quad (4.11a)$$

$$G_{1,2}(\vec{r}, \vec{r}') \equiv -\delta(\vec{r} - \vec{r}') \exp [-2ikh_{1,2}(r)] \quad (4.11b)$$

$$T_{1,2}(\vec{r}, \vec{r}') \equiv \delta(\vec{r} - \vec{r}') \Theta(r - R_{1,2}) \quad (4.11c)$$

where $\Theta(x)$ is a step function equal to 1 when $x > 0$ and 0 otherwise, and where the negative sign in G_1 arises because of boundary conditions on the electric field at a perfectly reflecting surface; cf. [14, 15, 16], and other previously noted references on paraxial optics for further details.

These propagation operations can be combined to generate more complicated processes. For example, while G_1 describes reflection off an *infinite* mirror of height h_1 , $G_1 T_1$ describes the reflection of light off a *finite* mirror of height h_1 and radius R_1 .

4.2.3.2 Describing paraxial propagation through an arm cavity

When we have an arm cavity bounded by two cylindrical mirrors separated by a length L , of cylinder radius $R_{1,2}$ and surface heights $h_{1,2}$ respectively (see Fig. 4.3), we can use the above propagators [Eq. (4.11)] to describe the free, undriven propagation of light within the arm cavity⁷. For example, as we demonstrate graphically in Fig. 4.5, we can express the field at the surface of mirror 1—specifically, the field for light heading *away* from the mirror surface—at time $t + 2L/c$ in terms of the field at that surface at time t as follows:

$$u(t + 2L/c) = G_1 T_1 G_L G_2 T_2 G_L u(t) \quad (4.12)$$

for the operators on the right side as given in Eq. (4.11) and for $|u\rangle$ denoting some polarization of the electric field.

We are in particular interested in eigenmodes—that is, states u so the beam returns proportional to itself:

$$\eta e^{2ikL} u = G_1 T_1 G_L G_2 T_2 G_L u \quad (4.13)$$

operator notation.

⁷The description of an arm cavity driven by a source laser adds only a straightforward inhomogeneous term to our discussion; the eigenmodes of the homogeneous term are as always in such problems of paramount importance.

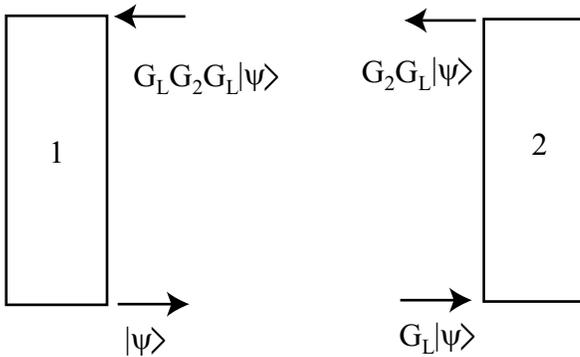


Figure 4.5: This figure demonstrates how propagation operators described in Eq. (4.11) are used to evolve the field from one plane of constant z to another. For brevity, the operators T_1 and T_2 , which allow for the finite extent of the mirrors' surfaces, have been omitted in this figure.

for some η . [We factor the phase e^{2ikL} out of η to eliminate the effect of the phase factor e^{ikL} present in the operator G_L .] These eigenmodes are *resonant* when ηe^{2ikL} is *real* and positive. The arm cavity length L and frequency k are tuned so one state, the *ground state* u_o , is resonant; in this case, η_o , k , and L satisfy

$$\text{Arg} [\eta_o e^{2ikL}] = 0 . \quad (4.14)$$

The mesa beam [Eq. (4.5)] is designed specifically to be one of the eigenmodes of the Mexican-hat mirrors [Eq. (4.6)] when the mirrors are infinite (i.e. $T_1 = T_2 = \mathbf{1}$), and when the cavity length is properly tuned to admit it. When the mirrors are finite, the mesa beam shape will no longer be precisely supported by the arm cavity: the true eigenmodes will have a slightly different shape than the ideal mesa-beam form; too, the eigenvalue η will no longer be of unit magnitude.

4.2.3.3 Diffraction losses: Exactly

From Eqs. (4.12) and (4.13), we know that when the initial optical state is an eigenmode, the optical state after n round trips decays as measured by its norm, which evolves as $\|u(t + n2L/c)\| = |\eta|^n \|u(t)\|$. Thus, if $|\eta| < 1$ the magnitude of u (i.e. the L^2 norm of u) decreases with each round trip.

Since our model above only permits losses from diffraction, and since the power in the cavity is proportional to $\int d^2r |u|^2$ (i.e. the L^2 norm of u , squared), the following quantity is the diffraction loss per round trip:

$$\mathcal{L}_{\text{net}} = 1 - |\eta|^2 . \quad (4.15)$$

When the two mirrors are identical, as is our case, we can subdivide the loss in two, to obtain a meaningful diffraction loss “per bounce”.

We can therefore extract from numerical solutions to Eq. (4.13) (cf. Sec. 4.5.2) the magnitude of

η and, by the above expression [Eq. (4.15)], deduce the mesa-beam diffraction losses in the presence of finite mirrors. We do *not*, however, use this method in this paper.

4.2.3.4 Diffraction losses: The clipping approximation

In practice the exact expression for the diffraction losses given above [Eq. (4.15)] is slow to evaluate, because we must solve for fine details of the eigenequation for each mirror size $R_{1,2}$ of interest. To obtain a rough estimate of the diffraction losses, usually accurate to a factor of order unity, we will often use instead the clipping approximation.

The clipping approximation estimates diffraction losses by assuming the beam is not significantly changed by those losses; the beam has to a good approximation the same shape as it would have in the presence of infinite mirrors. Assuming the beam profile is known, we then directly compute the losses by determining the power lost off the edges of finite mirrors on each reflection. Specifically, we assert $\mathcal{L}_{\text{net}} \approx \mathcal{L}_{\text{clip}}$, with

$$\mathcal{L}_{\text{clip}} = \mathcal{L}_1 + \mathcal{L}_2 \quad (4.16a)$$

$$\mathcal{L}_1 = |(1 - T_1)u_o|^2 = \int_{r>R_1} |u_o|^2 d^2r \quad (4.16b)$$

$$\mathcal{L}_2 = |(1 - T_2)G_L u_o|^2 = \int_{r>R_1} |G_L u_o|^2 d^2r \quad (4.16c)$$

where u_o is an eigensolution to the propagation equation [Eq. (4.13)] computed assuming infinite mirrors.

4.2.3.5 Diffraction losses of the ideal mesa beam, computed via the clipping approximation and corrected to estimate the true diffraction losses

In this paper, we approximate the diffraction losses of the mesa beam resonating between two identical Mexican-hat mirrors by the clipping approximation [Eq. (4.16)] applied to the mesa beam [Eq. (4.5)]. These losses are shown in Fig. (4.6).

The true diffraction losses per bounce generally differ from the estimates of clipping approximation. For example, for a $D = 4b$ mesa beam resonating between $R = 16$ cm cylindrical mirrors, the true diffraction losses are about 18.5 pmm, while the clipping approximation estimates diffraction losses of 21 pmm. Based on this good agreement, in this paper we assumed the true diffraction losses were well-approximated by the clipping-approximation estimate, with relative errors of order a factor 1.2.

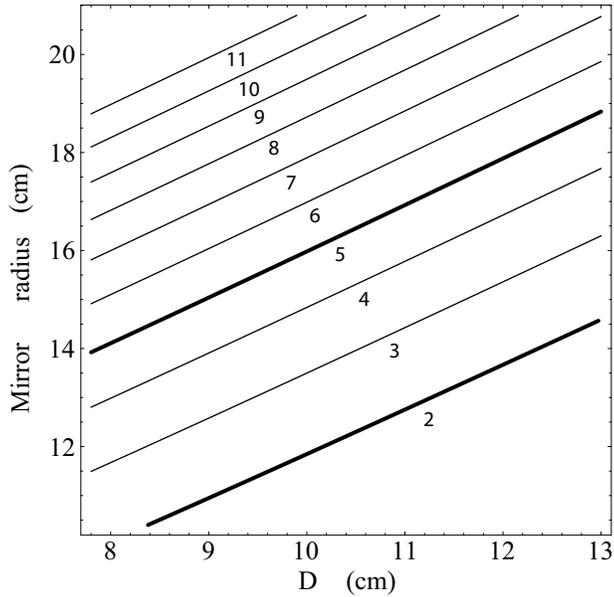


Figure 4.6: A contour plot of the (\log_{10} of the) diffraction losses of the mesa beam [Eqs. (B.15) with $\bar{z} = L/2$, i.e. (B.17)] according to the clipping approximation [Eq. (4.16)], as a function of mirror radius R and averaging scale D . Solid curves indicate losses 10^{-n} for $n = 2, 3, \dots, 11$. The two heavy solid curves correspond to $n = 2$ and $n = 5$; these two curves have special significance to LIGO design (they determine the maximum value of losses off the back and front faces of an ITM, per bounce; cf. Sec. 4.2.4).

4.2.4 Advanced LIGO design constraints

The advanced LIGO interferometers are composites of many interrelated systems; each component of that system has been designed making assumptions about the other components. In particular, the advanced LIGO interferometers require (roughly) the following constraints on mirror and beam designs:

1. *Mirror mass of 40 kg*: We consider only 40 kg test-mass mirrors, because the present advanced LIGO design for the suspension and seismic isolation system requires 40kg mirrors. [Also, the advanced LIGO suspension system cannot support heavier mirrors (Phil Willems, private communication)].⁸
2. *Mirror front radius limited by arm-cavity diffraction losses*: We consider only combinations of mirror front face radius R and beam size (e.g. if mesa beams are used, D) which have diffraction losses equal⁹ to 10 parts per million (ppm).¹⁰

⁸The support system could hold up lighter mirrors. However, thermoelastic noise decreases as the mirrors and beams are increased in size (in equal proportions). We therefore limit attention to the most massive possible mirrors.

⁹The advanced LIGO design could accept mirrors with lower losses. However, since thermoelastic noise decreases (as a general rule) when the beam radius increases, we limited attention to the largest possible beams.

¹⁰This constraint is required to keep the diffraction losses bounded by a reasonable portion of the total loss. Specifically, for the baseline design there is 125 W of input power to the interferometer and 830 kW of circulating power in each arm cavity. The 10 ppm of diffraction loss per bounce results in a diffraction power loss in the arm

3. *Mirror back radius limited by power-recycling-cavity diffraction losses:* Finally, we consider test-mass mirrors whose back face is sufficiently large that light entering the arm cavity through the back face loses at most 1 percent (i.e. 10,000 ppm) of the input power.¹¹ When evaluating this constraint, we assume the input light is in the same state as the arm cavity light (i.e. in a cavity eigenmode).

[This constraint only matters for noncylindrical mirrors (i.e. frustums) with a relatively small back face size; for cylindrical mirrors, it holds automatically.]

4.2.5 Summary: Exploring mesa-beam arm cavity configurations

To summarize, we design and evaluate new mesa-beam interferometer configurations by the following process:

1. *Pick a configuration:* We (i) select two axisymmetric mirrors (i.e. cylinders or frustums) which have the *same* front face radius and (ii) select some value D . We shape the mirror faces [Eq. (4.6)] and tune the length of the arm cavity [Eq. (4.14)] so the ground state eigenmode of that cavity—the eigenmode which is approximately a mesa beam with length scale D —is resonant.
2. *Check that the configuration satisfies the advanced LIGO design constraints:* If (i) the mirrors masses are not equal to 40 kg, if (ii) the physical diffraction losses in the arm cavity are greater than 10 pmm per bounce [Eq. (4.15)], or if (iii) the diffraction losses for input light are greater than one percent, then we stop and try again: the configuration does not satisfy the advanced LIGO design constraints described in Sec. 4.2.4. [In this paper, we use only the clipping approximation, Eq. (4.16), to estimate diffraction losses.]
3. *Evaluate the thermoelastic noise of this configuration:* We finally evaluate the noise integrals I_1 and I_2 for each of the two arm-cavity mirrors [Eq. (4.10)], using as beam pressure profile the (normalized) beam intensity profile [Eq. (4.5)]:¹²

$$P(r) = |u_{\text{mesa}}(r, D)|^2 . \quad (4.17)$$

[Section 4.5.1 will describe in detail precisely how we evaluated the two thermoelastic noise integrals.]

cavities of $4 \times 10 \text{ ppm} \times 830 \text{ kW} = 33 \text{ W}$, which is 25 per cent of the 125 W of input light, a reasonable value.

¹¹ As with the constraint involving the front face size (cf. footnote 10), this constraint is required to keep the diffraction losses bounded by a reasonable proportion of the total loss. Specifically, for the baseline design the power impinging on each ITM is 1.05 kW, so one per cent diffraction loss (i.e. 10,000 ppm) corresponds to losing $2 \times 0.01 \times 1.05 \text{ kW} = 21 \text{ W}$ at the ITM input, which is 17 per cent of the 125 W total laser power.

¹²Ideally, we *should* use as beam intensity profile the true eigenstate appropriate to two *finite* mirrors, as obtained by solving Eq. (4.13). However, because diffraction losses are low, the eigenstate for finite mirrors is very well approximated by the eigenstate for infinite mirrors; and the latter is the mesa-beam amplitude u_{mesa} .

These thermoelastic noise integrals then determine [via Eq. (4.7)] the overall thermoelastic noise of an interferometer using two identical arm cavities, where each cavity consists of this particular mirror and beam configuration.

4.3 Driving a mesa-beam arm cavity with a Gaussian beam

Ideally, a mesa-beam arm cavity would be driven by light already in the mesa-beam state. But if the optics required to generate mesa-beam inputs are unavailable or too inconvenient, Gaussian beams – generated by conventional optics—can be used to drive the arm cavities. With the proper choice of Gaussian, the arm cavity behaves almost as if the ideal mesa beam was used.

In the remainder of this paper, as in MBI [4], we shall assume properly-chosen Gaussian beams are used to drive the interferometer.

In her companion paper [6], Erika d’Ambrosio provides a comprehensive analytic and numerical discussion of Gaussian beams driving mesa-beam cavities [cf., e.g., ED’A Eq. (3)]. In this section, we briefly summarize some of her results on mesa-beam cavities driven by gaussian beams.

4.3.1 Beam and cavity used in this section

4.3.1.1 Fiducial mesa-beam cavity

Rather than study all possible interferometer configurations (i.e. all possible mirror and beam sizes), we select a single *fiducial* mesa-beam arm cavity: the cylindrical mirrors have front face size 16 cm and the mesa beams have $D = 4b = 10.4$ cm, where $b = \sqrt{\lambda L/2\pi} = 2.60$ cm (and where $\lambda = 1.064\mu\text{m}$ and $L = 4\text{km}$ are the light wavelength and arm length) [cf. MBI Sec. IV A 2]. This fiducial arm cavity has diffraction loss 18 pmm.¹³

We also assume the mirrors bounding the arm cavity —both the input test mass (ITM) and end test mass (ETM) —have no intrinsic losses, except for diffraction losses (which we treat separately). Specifically, we assume the ETM is perfectly reflecting, and assume the ITM has power transmissivity $t_I = 0.995$ [1]. Thus, our power reflection and transmission coefficients for the ITM and ETM are

$$\begin{aligned} r_I^2 &= 0.995, & t_I^2 &= 0.005; \\ r_E^2 &= 1.0, & t_E^2 &= 0. \end{aligned} \tag{4.18}$$

[This same fiducial cavity will appear frequently elsewhere; for example, we explore perturbations of this particular fiducial cavity when exploring the sensitivity of the interferometer to tilt.]

¹³While not *quite* appropriate—the diffraction losses are slightly larger than the advanced LIGO design cutoff of 10 pmm [1]—this fiducial design is sufficiently similar to acceptable advanced LIGO designs.

4.3.1.2 Gaussian beam used to drive the fiducial arm cavity

As described in MBI Sec. IV B, this fiducial arm cavity will be driven by a Gaussian beam of form given in MBI Eq. (2.8), with Gaussian beam radius $r_{od} = 6.92\text{cm}$. As in MBI, we denote this state by $|u_d\rangle$. This state has overlap

$$\gamma_o^2 \equiv |\langle u_d | u_o \rangle|^2 = 0.940 \quad (4.19)$$

with the arm cavity's ground state u_o [cf. MBI Eq. (4.2)]. [Erika d'Ambrosio has found that this particular Gaussian beam gives the largest possible overlap with our fiducial beam; cf. ED'A Eq. (3).]

Strictly speaking, we use a Gaussian state $|u_d\rangle$ which is shifted by precisely the phase offset necessary to make γ_o real and positive; therefore,

$$\gamma_o = \sqrt{0.940} = 0.970 . \quad (4.20)$$

4.3.2 Operation of fiducial cavity when driven by a Gaussian beam

Eq. (4.19) implies a fraction 0.94 of the (gaussian) driving-beam light will enter the Mexican-hat-mirrored cavity; the remaining 6 percent will be reflected. The fraction of the light that enters the cavity is amplified (due to resonant interaction with the bounding cavity mirrors), but eventually leaks back out the same way it entered (minus that fraction lost through diffraction losses). The combination of the 6 percent reflected light and 94 percent transmitted light (modulo losses) then returns back towards the beamsplitter.

This small section provides a more quantitative view of this entire process.

4.3.2.1 Decomposing light along the cavity eigenstates

The transmission and reflection processes are best understood in terms of the resonant eigenstate of the cavity. We therefore rewrite the driving beam in terms of a projection along and perpendicular to the ground state:

$$\begin{aligned} |u_d\rangle &= \mathcal{P}_o |u_o\rangle + (1 - \mathcal{P}_o) |u_o\rangle \\ &= \gamma_o |u_o\rangle + (1 - \mathcal{P}_o) |u_o\rangle \end{aligned} \quad (4.21)$$

where the projection operator \mathcal{P}_o is defined so $\mathcal{P}_o v = |u_o\rangle \langle u_o | v \rangle$ for any state v , i.e.

$$\mathcal{P}_o \equiv |u_o\rangle \langle u_o| , \quad (4.22)$$

and where γ_o is defined by (cf. Sec. 4.3.1.2)

$$\gamma_o \equiv \langle u_o | u_d \rangle ; \quad (4.23)$$

the phase of u_d has been designed so the unperturbed γ_o is real. Since the input state u_d is normalized,

$$|\mathcal{P}_o u_d|^2 + |(1 - \mathcal{P}_o)u_d|^2 = \|u_d\|^2 = 1. \quad (4.24)$$

The portion $\gamma_o u_o$ of the driving field excites the mesa-beam cavity, while the portion $(1 - \mathcal{P}_o)u_d$ cannot resonate and thus gets fully reflected off the input test mass (ITM).

4.3.2.2 Amplification of light entering cavity

A fraction γ_o^2 of the light power enters the arm cavity, resonates, and is amplified by a factor

$$\frac{1 + r_I}{1 - r_I} \left(1 - \frac{2\mathcal{L}_o}{1 - r_I} \right) \quad (4.25)$$

within the arm cavity,¹⁴ where \mathcal{L}_o denotes the diffraction losses per bounce for the ground state of the arm cavity. Therefore, relative to the overall input power, the light in the arm cavity is amplified by a factor

$$\mathcal{G}_{\text{anal}} = \gamma_o^2 \frac{1 + r_I}{1 - r_I} \left(1 - \frac{2\mathcal{L}_o}{1 - r_I} \right) \approx 737. \quad (4.26)$$

[While the term $\mathcal{L}/(1 - r_I)$ is very small in this case and can be ignored, we provide the entire expression now for use in Sections 4.4 and 4.7, when we study the effect of mirror perturbations on the cavity gain.]

4.3.2.3 Output: Transmission of light out of cavity, and recombination with reflected light

The cavity's output light is the combination of light reflected at the ITM (which does not enter the cavity) and light exiting the arm cavity through the ITM. After some algebra (cf. [6]),¹⁵ we find that

¹⁴The second (unnumbered) equation in d'Ambrosio [6] gives the transmitted field through the arm cavity. Dividing this by the second mirror transmissivity, squaring it to get an expression for beam power, and then specializing to $r_2 = 1$ and $r_1 = r_I$, we obtain the first factor in Eq. (4.25). The second factor arises when losses are properly taken into account on each bounce.

¹⁵To understand the propagation of light that enters the cavity, evaluate the first equation in d'Ambrosio [6] with $r_2 = 1$, $t_1 = t_I$ and $r_1 = r_I$, and both exponential factors unity. Therefore, for input light in the ground state, the reflected light is also precisely in the ground state, cf. the first term in Eq. (4.27). On the other hand, light that is in other states does not enter the cavity; therefore, the second term in that equation does not occur (or, more properly, the two phase factors take some random value with absolute value unity, rather than unity). Therefore, for light not in the cavity ground state, the reflected light is the same shape, but picks up a phase factor -1 , cf. the second term in Eq. (4.27).

given a driving beam u_d decomposed according to Eq. (4.21) the reflected light u_r is in the state

$$|u_r\rangle = \mathcal{P}_o |u_d\rangle - (1 - \mathcal{P}_o) |u_d\rangle \quad (4.27)$$

$$= (2\mathcal{P}_o - 1) |u_d\rangle . \quad (4.28)$$

where for simplicity we have ignored the (small) effects of diffraction losses on the light that resonates in the cavity.

4.4 Effect of mirror perturbations on arm cavities and the interferometer

In Section 4.2, we described how to design an idealized mesa-beam arm cavity and to evaluate the thermoelastic noise associated with such a design.

In this section, we shall describe in extremely general terms how one can use perturbation theory to study the influence of defects on a single ideal (i.e. lossless) arm cavity. We make no reference to the specific details of the cavity used, save that—for technical convenience—we limit attention in our perturbation theory to a system with two identical, infinite mirrors. [Recall the canonical mesa beams, as defined in Sec. 4.2.1, are designed for symmetric cavities.] Naturally the methods for perturbing arbitrary mirrors bear considerable similarity to methods for perturbing purely spherical mirrors. Many explicit tools and concepts carry over with little change to our more generic approach.

Specifically, in Sec. 4.4.1 we introduce the orthonormal basis set of solutions we shall use to construct the perturbation theory expansion. Then, in Sec. 4.4.2, we list the explicit second-order perturbation expansions we use to deduce the effect of mirror defects upon the optical state of the cavity. Finally, in Sec. 4.4.3, we describe how changes in the optical state of the cavity influence other quantities that can be deduced from that state, such as the cavity diffraction losses [Eq. (4.16)] and the thermoelastic noise integrals I_A [Eq. (4.9)].

4.4.1 Basis states for perturbation theory

We will use perturbation theory to analyze the effect of small changes in the arm cavity mirrors on the solutions to the eigenequation (4.13). To construct perturbation theory expansions by methods similar to those used in quantum mechanics (cf., e.g., [8, 9, 10]), we *prefer* to use as a basis for the perturbation expansion the eigenmodes of the initial equation [Eq. (4.13)].¹⁶ Fortunately, as shown in Appendix B.5, so long as the two mirrors are identical and infinite, the set of eigenmodes of Eq. (4.13) are a complete, orthogonal set, independent of the mirror shape.

¹⁶Actually, as Erika d’Ambrosio has frequently reminded us, we can construct perturbation theory just as well in the more general case, by using a dual basis. See, for example, Appendix A of d’Ambrosio [6].

At this point, we could introduce basis states (and notation for basis states) which are defined as eigensolutions to Eq. (4.13). For technical reasons—our methods simplify practical evaluation of the perturbation equation¹⁷—we prefer instead to use a basis designed to simplify “half” (i.e. a square root of) the round-trip operator, with a certain phase factor removed. In other words, we use as basis states the eigensolutions $|p\rangle$, λ_p to the equation

$$\lambda_p |p\rangle = e^{-ikL} G_1 G_L |p\rangle \quad (4.29)$$

where $|p\rangle$ is some state and p denotes some index over all basis states. [By an argument following that given in Appendix B.5, this equation admits a complete set of orthogonal solutions $|p\rangle$.] These states correspond directly to solutions to the full eigenequation. Explicitly, if we insert a solution $|p\rangle$ to Eq. (4.29) into the eigenequation [Eq. (4.13)], we immediately conclude that $|p\rangle$ is also an eigensolution of Eq. (4.13), with eigenvalue

$$\eta_p = \lambda_p^2 . \quad (4.30)$$

4.4.2 Effect of perturbations on light in the arm cavities

When the mirror shapes $h_{1,2}$ are deformed, the light propagating in the cavity changes. Given a basis of states and a specific problem to perturb [Eq. (4.13)], we employ conventional techniques from quantum mechanics (cf., e.g., [8, 9, 10]) to compute that change.

4.4.2.1 Results of perturbation expansion

Using conventional quantum-mechanics-style techniques (cf., e.g., [8, 9, 10], but note the operator we perturb is unitary rather than hermitian), we can devise a perturbative expansion for the cavity ground state eigenvalue η'_o and state u'_o in powers of the height perturbation δ_2 . The derivations of the sometimes-long expressions noted here are provided in Appendix B.6.

For simplicity, we provide perturbation theory expansions only to second order, and only to describe the effects due to changes in the height of one mirror (i.e. h_2 of mirror 2) by an amount δh . To express these changes in height in dimensionless form, we introduce the variable δ_2 :

$$\delta_2 \equiv 2kh_2 . \quad (4.31)$$

¹⁷On the one hand, by removing the common factor $\exp(ikL)$ present in G_L , we produce a naturally dimensionless eigenvalue problem (as k and L enter only through the single length parameter b). On the other hand, by focusing on “half” of the propagator rather than a round-trip operator, we can easily deduce the relationship between the fields at the two ends of the cavity. The latter proves helpful, because we will study perturbations of the mirror at one end of the cavity (end 2), yet represent the state of the cavity field by its values on the other end (end 1).

Change in eigenphase: The eigenphase η_o of the ground state $|u_o\rangle = |0\rangle$ changes as

$$\eta'_o = \eta_o \left(1 - i \langle 0 | \delta_2 | 0 \rangle - \frac{\langle 0 | \delta_2^2 | 0 \rangle}{2} - \sum_{k \neq 0} \frac{\eta_k |\langle 0 | \delta_2 | k \rangle|^2}{\eta_o - \eta_k} \right) \quad (4.32)$$

Change in eigenstate: When we construct the perturbation theory expansion to second order, we find the ground state changes according to an expression of the form

$$u'_{o,\text{pt}} \approx |0\rangle + |\psi^{(1)}\rangle + |\psi^{(2)}\rangle. \quad (4.33a)$$

In this expression, $\psi^{(1)}$ and $\psi^{(2)}$ denote those terms first- and second-order in δ_2 in the perturbation expansion, respectively. To be explicit, when we perform the perturbation theory expansion [details of which are provided in Appendix B.6], we find

$$|\psi^{(1)}\rangle = \sum_{k \neq 0} -\frac{i\lambda_o\lambda_k}{\eta_o - \eta_k} |k\rangle \langle k | \delta_2 | 0 \rangle \quad (4.33b)$$

$$|\psi^{(2)}\rangle = \sum_{k \neq 0} |k\rangle \frac{\lambda_o\lambda_k}{\eta_o - \eta_k} \left[-\frac{1}{2} \langle k | \delta_2^2 | 0 \rangle + \frac{\eta_o}{\eta_o - \eta_k} \langle k | \delta_2 | 0 \rangle \langle 0 | \delta_2 | 0 \rangle - \sum_{p \neq 0} \frac{\eta_p}{\eta_o - \eta_p} \langle k | \delta_2 | p \rangle \langle p | \delta_2 | 0 \rangle \right]. \quad (4.33c)$$

By construction $|\psi^{(1)}\rangle$ and $|\psi^{(2)}\rangle$ are orthogonal to the unperturbed ground state $|0\rangle$. As a result, this expression (4.33) for $|\psi'\rangle$ is *not* normalized: we find, working to second order, that the norm of $|u'_{\text{pt}}\rangle$ is

$$\langle u'_{o,\text{pt}} | u'_{o,\text{pt}} \rangle \approx 1 + \|\psi^{(1)}\|^2 + \mathcal{O}(\delta_2^3). \quad (4.34)$$

where we use the shorthand $\|\psi\|^2 \equiv \langle \psi | \psi \rangle$. Therefore, the physically appropriate normalized perturbed state $|u'_o\rangle$ is given by the expression

$$\begin{aligned} |u'_o\rangle &\approx \frac{1}{\sqrt{1 + \|\psi^{(1)}\|^2}} \left(|0\rangle + |\psi^{(1)}\rangle + |\psi^{(2)}\rangle \right) \\ &= |0\rangle + |\psi^{(1)}\rangle + \left(|\psi^{(2)}\rangle - \frac{\|\psi^{(1)}\|^2}{2} |0\rangle \right) + \mathcal{O}(\delta_2^3). \end{aligned} \quad (4.35)$$

4.4.2.2 Estimating convergence of perturbation expansion

Perturbation theory is only effective when higher order terms provide only a small correction to lower order terms. To test the convergence of the series, we compare the magnitudes of the first two perturbative corrections. When $\|\psi^{(2)}\|/\|\psi^{(1)}\| \ll 1$, we believe the series converges and our expressions should be effective.

4.4.3 Implications of change in optical state for other quantities

In the previous section, we described the effect of perturbations $\delta h_{1,2}$ on light propagating in a single Fabry-Perot arm. These perturbations cause the beam shape incident on the two mirrors to change, generally in a different way at each mirror. Therefore, quantities that depend on the state of the beam at each mirror, such as the diffraction losses and thermoelastic noise integrals I_A , also change. In this section, we loosely describe how the lowest order effect of these changes can be characterized.

4.4.3.1 Change in diffraction losses under perturbation of one mirror

Unfortunately, a systematic treatment of diffraction losses within the context of perturbation theory proves very tricky, not the least because we must represent both the height change and the effect of finite mirror size as perturbations, then use many states to insure the perturbation theory expansion converges properly for the effects of diffraction.

Rather than perform a truly accurate, well-motivated computation, we will in this paper loosely estimate diffraction losses by the clipping approximation applied to the perturbed beam state. By way of example, we can estimate the diffraction losses produced during a single reflection off of mirror 1 or off mirror 2 in the presence of a perturbation of mirror 2 by expanding the appropriate clipping approximation estimate \mathcal{L}_1 or \mathcal{L}_2 [cf. Eq. (4.16)]:

$$\mathcal{L}_1 \approx \langle u|O_1|u \rangle , \quad (4.36a)$$

$$\mathcal{L}_2 \approx \langle G_L u|O_2|G_L u \rangle , \quad (4.36b)$$

$$O_1 \equiv 1 - T_1 . \quad (4.36c)$$

To be explicit, we can expand the clipping-approximation estimate for \mathcal{L}_1 to second order in δ_2 as

$$\begin{aligned} \mathcal{L}_1 &\approx \langle 0|O_1|0 \rangle + 2\text{Re} \left(\langle \psi^{(1)}|O_1|0 \rangle \right) \\ &+ \langle \psi^{(1)}|O_1|\psi^{(1)} \rangle - \|\psi^{(1)}\|^2 \langle 0|O_1|0 \rangle \\ &+ 2\text{Re} \left(\langle \psi^{(2)}|O_1|0 \rangle \right) + \mathcal{O}(\delta_2^3) . \end{aligned} \quad (4.37)$$

A similar expression is found for \mathcal{L}_2 , with O_1 replaced by $G_L^\dagger O_2 G_L$. [Because the beam profile will generally change in different ways at the two ends due to a perturbation localized at only one end, generally $\mathcal{L}_1 \neq \mathcal{L}_2$.]

4.4.3.2 Change in cavity gain due to perturbations of one mirror

The amount of power resonating in the cavity also changes, in part through modified diffraction losses [Eq. (4.37)] but also through a change in the resonant ground state of the cavity from u_o (i.e.

$|0\rangle\rangle$ to u'_o [Eq. (4.35)]. Because the resonant state of the cavity changes, the overlap γ_o changes to γ'_o , which to second order can be approximated by

$$\gamma'_o \equiv \langle u'_o | u_d \rangle = \left(1 - \frac{\|\psi^{(1)}\|^2}{2} \right) \gamma_o + \gamma_1 \|\psi^{(1)}\| + \gamma_2 \|\psi^{(2)}\|, \quad (4.38a)$$

where we define

$$\gamma_1 \equiv \langle \psi^{(1)} | u_d \rangle / \|\psi^{(1)}\|, \quad (4.38b)$$

$$\gamma_2 \equiv \langle \psi^{(2)} | u_d \rangle / \|\psi^{(2)}\|. \quad (4.38c)$$

To be very specific, we can find the perturbed cavity gain $\mathcal{G}'_{\text{anal}}$ by evaluating Eq. (4.26) to second order in δh_2 ,

$$\mathcal{G}'_{\text{anal}} = \gamma_o'^2 \frac{1+r_I}{1-r_I} \left(1 - \frac{\mathcal{L}'_1 + \mathcal{L}'_2}{1-r_I} \right), \quad (4.39)$$

using Eq. (4.37) for the perturbed diffraction losses \mathcal{L}_1 and \mathcal{L}_2 and using Eq. (4.38) for the change in overlap between the perturbed cavity ground state and the driving beam.

[Though we do not write out the resulting second-order expansion in full detail, we will apply these methods to compute to second order the effects of mirror tilt on cavity gain in Sec. 4.7.]

4.4.3.3 Change in dark port power due to perturbations of one mirror

An interferometer consists of two arms. When the beamsplitter recombines the two fields leaving the two arm cavities, the light going out the dark port of the interferometer (u_{dp}) is the interference between the light reflected off the two arms (denoted I and II):

$$u_{\text{dp}} = \frac{1}{\sqrt{2}}(u_{r,\text{II}} - u_{r,\text{I}}). \quad (4.40)$$

If both cavities are identical, then $u_{r,\text{II}} = u_{r,\text{I}} = u_r$ for u_r given by Eq. (4.28) and no light exits the dark port. If the cavities are perturbed, however, power will generically go out the dark port.

For example, if only cavity II is perturbed, then the dark port beam state is

$$\begin{aligned} u_{\text{dp}} &= \frac{1}{\sqrt{2}}(u'_r - u_r) \\ &= \sqrt{2}(\mathcal{P}'_o - \mathcal{P}_o) u_d \\ &= \sqrt{2}(\gamma'_o u'_o - \gamma_o u_o) \\ &= \sqrt{2} \left[\left(-\gamma_o \|\psi^{(1)}\|^2 + \gamma_1 \|\psi^{(1)}\| + \gamma_2 \|\psi^{(2)}\| \right) |u_o\rangle \right. \\ &\quad \left. + \left(\gamma_o + \gamma_1 \|\psi^{(1)}\| \right) |\psi^{(1)}\rangle \right. \\ &\quad \left. + \gamma_o |\psi^{(2)}\rangle \right] + O(\delta_2^3). \end{aligned} \quad (4.41a)$$

where going from the first line to the second we use Eq. (4.28) [for u_r]; from the second to the third we use Eqs. (4.22) [for \mathcal{P}_o] and (4.23) [for γ_o]; and from the third to the fourth we use the perturbation expansions (4.33) [for u'_o] and (4.38) [for γ'_o].

The dark port power P_{dp} (as a relative fraction of interferometer input power, which is twice the input power to each arm cavity) can be expressed as

$$P_{\text{dp}} = \|u_{\text{dp}}\|^2/2, \quad (4.41\text{b})$$

an expression which we shall not expand here.

4.4.3.4 Influence of perturbations of one mirror on thermoelastic noise

The thermoelastic noise integrals associated with each mirror (1 and 2) in a given arm cavity both change because the beam profile at each mirror changes. To evaluate those changes, we insert the modified state [Eq. (4.35)] into the thermoelastic noise integral I_A [Eq. (4.9)].

For example, to evaluate the linear-order change in thermoelastic noise at mirror 1, we first take the new normalized intensity profile $P(r)$ at mirror 1, given by

$$\begin{aligned} P'_1(r) &\equiv |u'(r)|^2 \\ &\approx |u(r)|^2 + 2\text{Re}\left(u_0^*(r)\psi^{(1)}(r)\right) + O(\delta_2^2). \end{aligned} \quad (4.42)$$

Using this new intensity profile, we solve the thermoelastic noise model problem [Eq. (4.10)] for the expansion Θ' and in particular the first-order change in expansion $\delta\Theta \equiv \Theta' - \Theta$. Finally, we insert this first order change $\delta\Theta$ into the expression for I_A with $A = 1$, linearized about the background intensity profile:

$$\delta I_1 = \frac{2}{F_o^2} \int_{V_1} (\nabla_a \Theta)(\nabla^a \delta\Theta) d\text{volume} \quad (4.43)$$

where Θ is the expansion produced by the unperturbed pressure profile P .

The beam shape at mirror 2 also changes; by a similar construction, we can find its effect on the thermoelastic noise integral I_2 .

4.4.4 Observations which simplify our computations

4.4.4.1 Symmetry and the influence of first-order changes in state

The ground-state beam of physical interest is isotropic. As a result, when we perform the computations outlined above, we find only the *axisymmetric* (about the optic axis) part of $|\psi^{(1)}\rangle$ contributes to first-order changes in the three integral quantities of physical interest (the overlap γ_o ; diffraction losses; and thermoelastic noise).

By way of example, consider the first-order change in thermoelastic noise [Eq. (4.43)]. Since the unperturbed pressure profile and thus unperturbed expansion (Θ) are isotropic, the above integral couples only to the axisymmetric part of $\delta\Theta$ and therefore δP and $\psi^{(1)}$ [cf. Eq. (4.42)].

Similarly, to lowest order displacement and tilt perturbations have odd parity; therefore, symmetry insures that the first-order change γ_1 [Eq. (4.38)] is zero.

4.4.4.2 Assuming (roughly) equal contributions from the two mirrors to changes in I_A and diffraction losses

The change in mirror shape at mirror 2 leads to roughly comparable changes in state in the beam profile at both ends. To a rough approximation, then, we can assume that the changes in thermoelastic noise and diffraction loss at mirror 2 will be comparable to their changes at mirror 1.

4.4.5 Summary: Using perturbation theory to explore the sensitivity of the arm cavity and interferometer to perturbations

To summarize, we evaluate how sensitive mesa-beam arm cavities and interferometers are to mirror perturbations as follows:

1. *Assume the arm cavity is fiducial:* The perturbation expansions above can be applied to any arbitrary mirror configuration (i.e. to any specific mesa-beam D).¹⁸ In principle, we *could* apply perturbation theory to all possible mirror configurations generated in Sec. 4.2. Instead, to avoid repeating computations which should yield nearly identical results, we apply perturbation theory only to the *fiducial* advanced LIGO arm cavity presented in Section 4.3.1.1.
2. *Consider only perturbations of one ETM:* Similarly, we change only one end test mass in one arm cavity, rather than apply perturbations to each test-mass mirror.
3. *Compute the eigenmodes of the fiducial arm cavity:* We find the natural eigenmodes of the fiducial arm cavity—and therefore the basis states $|p\rangle$ in our perturbation expansion—by numerically solving the basis-state eigenequation (4.29). [The numerical code which solved this integral eigenequation is described in Sec. 4.5.2.]
4. *Find how the ground state of arm cavity changes:* Next, we apply the perturbation expansion [Eq. (4.33)] to find how the ground state of the arm cavity changes when each perturbation of physical interest is applied: tilt, displacement, and mirror figure error. In other words, for each of these perturbations δh_2 , we use Eq. (4.33) to find the first- and second- order corrections $|\psi^{(1)}\rangle$ and $|\psi^{(2)}\rangle$ to the state of the cavity.

¹⁸Our technique assumes the cavity eigenmodes of an arm cavity with finite mirrors are well-approximated by the modes of a cavity with infinite mirrors. Therefore, the only relevant parameter remaining is the mesa-beam averaging scale D . [This approximation ignores quantities of order the diffraction losses, $10 \text{ pmm} \sim 10^{-5}$.]

5. *Determine how the thermoelastic noise integral I_A for each mirror changes:* Given the changed beam state, we can recompute the thermoelastic noise integral I_A for each mirror (cf. Sections 4.2 and 4.5.1), using the perturbed beam state $u'_o(r)$ and the perturbed beam intensity profile $P'(r) = |u'_o(r)|^2$. More directly, the series expansion of thermoelastic noise can be discovered by a series expansion of the thermoelastic noise integral, as sketched (to first order) in Sec. 4.4.3.4.

6. *Describe how the cavity gain and dark port power change:* Finally, given the changed beam state, we can also recompute the arm cavity gain [Eq. (4.39)] and interferometer dark port power [Eq. (4.41)] when the ETM of one cavity is perturbed. These two expressions both depend on how the perturbed arm cavity interacts with the driving beam [i.e. on γ_1 and γ_2 ; cf. Eq. (4.38)] and on the perturbed arm cavity state itself [i.e. on $||\psi^{(1)}||$ and $||\psi^{(2)}||$; cf. Eq. (4.33)]. Also, the cavity gain depends on how the diffraction losses of the arm cavity change [Eq. (4.37)].

4.5 Numerical implementations of the thermoelastic noise integral, the optical eigenequation, and optical perturbation theory

In the previous two sections, we described the abstract expressions we must evaluate to design and analyze an arm cavity bounded by Mexican-hat mirrors. In this section, we describe how we implemented and solved those equations.

Where independent methods were used to perform a particular computation, we indicate the different techniques used.

4.5.1 Thermoelastic noise for perfect (undeformed) mirrors

To evaluate the thermoelastic noise power spectrum [Eq. (4.7)], we need to perform two tasks. First, we must solve the elasticity problem described in Eq. (4.10); then, using the result, we can evaluate the integral I_A , which enters directly in Eq. (4.7).

We employed three methods to address these two tasks. The first two methods—a numerical (finite-element) solution and an exact analytic solution—applied only to special circumstances: axisymmetric beam profiles, and cylindrical mirrors. The third method was an approximation based on assuming the mirror to be half-infinite; we used it only as a quick, easy-to-evaluate check on the qualitative behavior of the previous two methods.

4.5.1.1 General approach

We usually solved the elasticity equations (4.10b)-(4.10d) using a commercial two-dimensional finite-element code [22]. We chose the region \mathcal{R} in Eq. (4.10d) to be one of the mirror faces.¹⁹ [To use a two dimensional code, we limited attention to *axisymmetric* mirrors and pressure profiles.]

The commercial finite-element code we employed gave us the displacement vector y^a . We then used postprocessing code discussed in Appendix B.8 to evaluate the derivatives and integrals needed in Eq. (4.9).

4.5.1.2 Special case: cylindrical mirror

Independently, we employed the analytic elasticity solution Liu and Thorne (LT) [12] developed for cylindrical mirrors with axisymmetric pressure profiles imposed on them.

Liu and Thorne constructed a solution to the elasticity equations for a cylinder [Eqs. (4.10b)-(4.10d)] which they apply to Eq. (4.9) to find an explicit expression for the thermoelastic noise produced by gaussian beams on cylindrical mirrors. By replacing a single equation in their expressions, we can convert their solution to one appropriate to *arbitrary* axisymmetric beam intensity profiles.

To be explicit, LT Eq. (44) gives an expression that is precisely 1/2 of I (1/2 because of averaging which we have factored out but which LT retain), in terms of quantities defined in LT Eqs. (35) and (36). To generalize to a generic axisymmetric pressure profile, one need only change LT Eq. (36)—the only equation which involves the specific pressure profile—so that it involves the intensity function $P(r)$ defined in Eq. (4.10c) for $m > 0$:²⁰

$$p_m = \frac{2}{a^2 J_0^2(\zeta_m)} \int_0^a P(r) J_0(\zeta_m r/a) r dr . \quad (4.44)$$

Here a is the radius of the cylindrical mirror (denoted R elsewhere in this paper); ζ_m is the m th zero of $J_1(x)$; and the functions J_0 and J_1 are the zeroth and first order cylindrical Bessel functions. The sum converges rapidly: typically, only a handful of terms in the infinite sum [LT Eq. (45)] are required.

4.5.1.3 Approximate technique: half-infinite mirrors

If the mirror is sufficiently large compared to the imposed pressure profile $P(r)$, the elasticity problem we must solve [Eq. (4.10)] can be well approximated by a solution to a similar problem with the mirror boundary taken to infinity. In this case, as the force density term [Eq. (4.10c)] goes to zero,

¹⁹The code works faster if the region \mathcal{R} is a point or set of points. However, the code is significantly more susceptible (on physical grounds—the points act like “nails” in the mirror) to small errors in the neighborhood of points, errors that contribute significant erroneous expansion. Therefore, as a practical compromise we chose \mathcal{R} to be a surface.

²⁰The LT coefficient p_o is independent of the pressure profile shape; it is always $1/\pi a^2$.

we need only solve for the response of a half-infinite (i.e. filling the region $z < 0$) elastic medium to an imposed surface pressure. As described in greater detail in Appendix B.9, analytic expressions exist for the response, permitting us to find an compact expression for the thermoelastic noise integral. We find

$$I = \left(\frac{(1 + \sigma)(1 - 2\sigma)}{2\pi E} \right)^2 \int d^2 \vec{K} \left| \vec{K} \right| \left| \tilde{P}(\vec{K}) \right|^2 \quad (4.45)$$

where $\tilde{P}(K)$ is the two-dimensional Fourier transform of $P(\vec{r})$:

$$\tilde{P}(\vec{K}) \equiv \int d^2 \vec{r} e^{-i\vec{K} \cdot \vec{r}} P(\vec{r}) . \quad (4.46)$$

Recall that $P(\vec{r})$ is normalized to unity (cf. Sec. 4.2.2).

4.5.2 Numerically solving for the resonant optical eigenmodes of a cavity bounded by arbitrary axisymmetric mirrors

To test the validity of the clipping approximation [Eq. (4.16)] against the exact diffraction losses [Eq. (4.15)] and to generate the set of basis solutions needed to construct perturbation theory [Eq. (4.29)], we must numerically solve for the cavity eigenmodes. In this section, we describe numerically how we converted the cavity eigenmode integral equation into a numerical eigenproblem, which we then solved with standard numerical tools (e.g., Mathematica).

4.5.2.1 Setting up the problem to be solved; preliminary analytic simplifications

Rather than solve the eigenequation for a full round trip through a symmetric cavity [Eq. (4.13) with $R_1 = R_2$], we exploit symmetry and instead study the closely related problem of *half* a round trip through a symmetric cavity. In other words, we plan to diagonalize the basis state eigenequation [Eq. (4.29)] when that eigenequation is restricted to the space of functions defined on $R_1 = R_2 \equiv b X_{\max}$:

$$\eta |u\rangle = T_1 G_1 G_L T_1 \exp(-ikL) |u\rangle \quad (4.47)$$

If we write out the appropriate eigenequation for this “half-a-round-trip” operator, making use of the definitions in Eq. (4.11) and using the dimensionless spatial units $\vec{X} = \vec{r}/b$, we find the following integral equation:

$$\eta \psi(\vec{X}) = \int d^2 X' \mathcal{G}(X; X') \psi(\vec{X}') , \quad (4.48a)$$

$$\mathcal{G}(\vec{X}; \vec{X}') \equiv \frac{-i}{2\pi} \Theta(|\vec{X}'| - X_{\max}) e^{i \left[\frac{(\vec{X} - \vec{X}')^2}{2} - \bar{h}(X) \right]} , \quad (4.48b)$$

$$\bar{h} \equiv 2kh . \quad (4.48c)$$

Here, $\psi(\vec{X}) = u(\vec{r})/b$ denotes a dimensionless representation of the state u .

Since the mirror surfaces are (ideally) axisymmetric, the operators above all commute with rotation around to the optic axis. We can therefore mutually diagonalize this operator and the generator of rotations. Therefore, we require ψ be proportional to $\exp(im\varphi)$ for some integer m :

$$\psi(X, \varphi) \equiv \Phi(X)e^{im\varphi}. \quad (4.49)$$

Substituting in this form for ψ , we reduce the problem to a series of one-dimensional integral equations, one for each $|m|$:

$$\eta\Phi(X) = \int_0^{X_{\max}} \bar{X}d\bar{X}\mathcal{G}_m(X, \bar{X})\Phi(\bar{X}), \quad (4.50a)$$

$$\mathcal{G}_m(X, \bar{X}) = -i^{m+1}J_m(X\bar{X})e^{i\left[\frac{X^2+\bar{X}^2}{2}-\bar{h}(X)\right]}. \quad (4.50b)$$

4.5.2.2 Method of numerical solution

For each m of interest, we represent the integral operator on the right side of Eq. (4.50) as a matrix. For simplicity, we discretize space in a uniform grid $X_A = AX_{\max}/(N-1)$ for $A = 0 \dots N$, define $\Phi_A = \Phi(X_A)$ and similarly, and approximate the integral by a simple quadrature rule.²¹ Thus, we approximate the integral equation of Eq. (4.50) by the matrix eigenvalue problem

$$\eta\Phi_A = \sum_{B=0}^N \frac{X_{\max}^2 B}{(N-1)^2} \mathcal{G}_m(X_A, X_B)\Phi_B \quad (4.51)$$

This equation can be solved for η and Φ by any standard eigensolution package.

4.5.2.3 Interpreting and applying the numerical solution

To summarize, to find numerical approximations to eigensolutions ψ of Eq. (4.47)—an eigenequation similar to the basis-state eigenequation [Eq. (4.29)] but for mirrors of finite radius $R = bX_{\max}$ —we construct and solve the matrix eigenvalue problem (4.51).

The approximate eigensolutions η_p and Φ so obtained provide numerical approximations to the two cavity-eigenproblems of technical interest:

- *Solutions for true cavity eigenmodes:* This method provides us with precisely the numerical cavity modes needed to understand mesa beams in the presence of finite mirrors. [Specifically, eigenmodes and eigenvalues of Eq. (4.47) are also eigenmodes of Eq. (4.13).] In particular, we can use the norm $|\eta|$ to determine the round-trip diffraction losses according to Eq. (4.15).

²¹We in fact used the equation Eq. (4.51) as stated. More sophisticated quadrature techniques, such as Gaussian quadrature, offer greater accuracy with fewer points (and hence significantly less computation time).

- *Solutions for the basis states:* Further, by making the mirrors sufficiently large that diffraction effects can be ignored²², solutions to Eq. (4.47) provide good approximations to solutions to the basis eigenvalue problem [Eq. (4.29)].

4.5.3 Numerical implementation of perturbation theory

To use perturbation theory exactly [i.e. Eqs. (4.32) and (4.33)], one needs an infinite collection of states. In practice, we limited attention to a handful: the lowest-lying three to five modes for each $|m| = 0, \dots, 7$; more modes were used when the rate of convergence of the perturbation theory expansion suggested more were needed.

4.5.4 Numerical exploration of changing diffraction losses

Given the δh_2 -induced perturbed beam state $|\psi\rangle$ at mirror 1, we can compute the (clipping approximation estimate of the) diffraction losses at mirror 1 for a given height perturbation δh_2 at mirror 2. Largely, we simply evaluated the integrals (i.e. inner products and norms) required to construct Eq. (4.37). However, to provide an independent numerical check (i.e. to insure we had no typographical or structural errors), we also evaluated the diffraction losses directly, using the definition Eq. (4.16), for a sequence of height perturbations $\delta h(\varepsilon) \equiv \varepsilon \delta h_2$ with $\varepsilon = 0, 0.1, 0.2, \dots, 1$; then fitted a second-degree polynomial to the resulting data points to extract the series coefficients in Eq. (4.37).

4.5.5 Numerical investigation of changes in thermoelastic noise when one mirror shape is perturbed

Finally, we can use the known form of $\psi^{(1)}(r, \varphi)$ [from Eq. (4.33b)] in the procedure outlined earlier [cf. Eqs. (4.42) and (4.43)] to evaluate the first-order change in thermoelastic noise associated with mirror 1 due to changes in the shape of mirror 2.

Specifically, given $\psi^{(1)}(r, \theta)$, we use Eq. (4.42) to find how the pressure profile changes. We then use the first-order change δP to the pressure profile in the LT expression for thermoelastic noise (see Sec. 4.5.1.2), linearized about the response to the unperturbed cavity beam intensity P . [We will not provide the rather lengthy but straightforward linearization of the LT expressions here; any computer-algebra system can easily reproduce the desired expansion.]

²²As a practical matter, we test the quality of our solutions by observing the convergence of the norm $|\eta|$ as we increase the number of points used to represent the state.

4.6 Dependence of thermoelastic noise on mirror and beam shape

In this section, we explore the dependence of the thermoelastic noise integral I_A [Eq. (4.9)] on the various arm cavity parameters available to us: (i) the mirror’s dimensions (i.e. cylinder height and radius) and shape (i.e. frustum or cylinder), (ii) the beam size (e.g., the mesa-beam size D), and (iii) the type of beam resonating in the arm cavity (i.e. Gaussian or mesa). The thermoelastic noise integral provides a simple way to characterize how the thermoelastic noise power spectrum S_h [Eq. (4.7)] of these configurations compares to S_h for the baseline advanced LIGO configuration (S_h^{BL}):

$$S_h/S_h^{\text{BL}} = I/I_{\text{BL}} , \quad (4.52)$$

where I_{BL} (evaluated below) is the value of the thermoelastic noise integral for the baseline advanced LIGO configuration [Eq. (4.53)], and where I is the value of the thermoelastic noise integral for all four (identical) mirrors in the interferometer.

Specifically, in this section we compare the following types of mirror and beam configurations, all of which satisfy the advanced LIGO design constraints [Section 4.2.4]: (i) the baseline advanced LIGO configuration (which uses Gaussian beams and cylindrical mirrors); (ii) an improved baseline configuration (which also uses Gaussian beams and cylindrical mirrors); (iii) configurations with mesa-beam light resonating between identical cylindrical mirrors; and (iv) configurations with mesa-beam light resonating between identical frustum mirrors. [In this section, we restrict attention to arm cavities with identical mirrors; in Appendix B.10 we discuss generalizations to arm cavities bounded by *nonidentical* mirrors.] In Table 4.4 we summarize the optimal (i.e. lowest value of I/I_{BL}) configurations we found for each class. The results of items tabulated in this section are applied, in MBI Sec. III (cf. MBI Table I), to produce advanced LIGO designs with lower thermoelastic noise than the baseline design.

Our evaluations have been performed independently by all three co-authors (RO’S, SS, and SV), using multiple methods (both finite-element solutions and infinite-sum analytic solutions) when appropriate.

4.6.1 Baseline advanced LIGO configuration

The baseline design of an advanced LIGO interferometer [1] has four identical cylindrical sapphire test masses (i.e. physical radius $R_p = 15.7$ cm; thickness $H = 13$ cm; mass 40 kg) whose surfaces are coated over most of their surface (i.e. out to a radius $R = R_p - 8$ mm). These mirrors’ surfaces are designed so the largest possible Gaussian consistent with the 10 ppm diffraction loss constraint (Sec. 4.2.4) resonates in the arm cavity [i.e. a Gaussian beam with radius $r_o = 4.23$ cm = $1.63b$; cf.

Table 4.1: The thermoelastic integral I for a cylindrical test mass and a Gaussian beam, in units of the value $I_{\text{BL}} = 2.57 \times 10^{-28} \text{s}^4 \text{g}^{-2} \text{cm}^{-1}$ for the advanced LIGO baseline design. The values of I/I_{BL} are estimated to be accurate to within one per cent. Also shown are the diffraction losses \mathcal{L}_0 off each bounce of the test mass (in parts per million, or ppm), computed using the clipping approximation. The first row shows the baseline design (denoted BL) for the advanced LIGO interferometers.

	R	R_p [cm]	H [cm]	r_o	r_o/b	I/I_{BL}	\mathcal{L}_0 [ppm]
BL	$R_p - 8\text{mm}$	15.7	13.	4.23	1.63	1.000	10
	R_p	15.7	13.	4.49	1.73	0.856	10

Appendix B.7 and MBI Eq. (2.8)] for a discussion of Gaussian beams].

The thermoelastic noise integral for an arm cavity bounded by cylindrical mirrors and using a Gaussian beam is found by (i) constructing the Gaussian amplitude function $u_G(r, r_o)$ [MBI Eq. (2.8)] and its associated beam intensity profile $P(r) = |u_G(r, r_o)|^2$; (ii) solving for the elastic expansion Θ that arises due to $P(r)$ in the elastic model problem of Eq. (4.10) described in Sec. 4.2.2 (cf. Sec. 4.5.1 for numerical methods); and (iii) inserting the resulting expansion into the definition of the thermoelastic noise integral I [Eq. (4.9)]. For the baseline beam and test mass, the resulting value of the noise integral I is

$$I_{\text{BL}} = 2.57 \times 10^{-28} \text{s}^4 \text{g}^{-2} \text{cm}^{-1} . \quad (4.53)$$

[The advanced LIGO cylindrical mirrors and Gaussian beams are optimal: these beams produce very nearly the lowest thermoelastic noise possible using Gaussian beams reflecting off identical 40 kg cylindrical mirrors, where those cylinders are coated out to a radius $R = R_p - 8\text{mm}$ and where the diffraction losses per bounce are restricted to less than 10 ppm.]

4.6.2 Improved baseline advanced LIGO configuration

The conventional baseline described above wastes the last 8mm of mirror face size. We can improve upon the thermoelastic noise merely by eliminating the uncoated ring in the last 8mm, i.e. by coating the mirror out to the edge and adjusting the beamspot size to fill in the extra space. If the mirror coating extends out to the test-mass edge so $R = R_p = 15.7$ cm, and the Gaussian beam radius is correspondingly increased to $r_o = 4.49$ cm so the diffraction losses are still 10 ppm, then the thermoelastic noise is reduced to $I/I_{\text{BL}} = 0.856$; see Table 4.1.

[Again, the *same* advanced LIGO cylindrical mirrors (i.e. with unchanged physical radius and thickness) produce very nearly the lowest possible thermoelastic noise, among all 40 kg cylindrical mirrors coated out to their physical radius.]

4.6.3 Mesa beams reflecting off identical 40 kg cylindrical mirrors

Arm cavities with *mesa beams* reflecting off cylindrical mirrors admit configurations with even lower thermoelastic noise than the improved baseline. To explore the advantages of mesa beams, we

Table 4.2: The thermoelastic integral I for a cylindrical test mass and a mesa beam, in units of the value $I_{\text{BL}} = 2.57 \times 10^{-28} \text{s}^4 \text{g}^{-2} \text{cm}^{-1}$ for the Advanced LIGO baseline design. The values of I/I_{BL} are estimated to be accurate to within one per cent. The first four test masses, like the baseline, are mirror coated only out to $R = R_p - 0.8 \text{cm}$; the last five are coated all the way out to the test-mass edge, $R = R_p$.

$R[\text{cm}]$	$R_p[\text{cm}]$	$H[\text{cm}]$	D/b	I/I_{BL}	$\mathcal{L}_0 [\text{ppm}]^1$
$R_p - 8 \text{mm}$	14.67	14.79	3.00	0.414	10
$R_p - 8 \text{mm}$	15.70	13.00	3.43	0.364	10
$R_p - 8 \text{mm}$	17.11	10.87	4.00	0.442	10
$R_p - 8 \text{mm}$	19.58	8.30	5.00	1.000	10
R_p	13.94	16.38	3.00	0.373	10
R_p	15.70	13.00	3.73	0.290	10
R_p	16.37	11.88	4.00	0.313	10
R_p	18.85	8.96	5.00	0.628	10
R_p	21.36	6.98	6.00	1.69	10

evaluated the thermoelastic noise integral (via the method described in Sec. 4.2.5) for two one-parameter²³ families of cylindrical mirrors and mesa beams: 40 kg mirrors with the largest possible mesa beams resonating off their front faces (i.e. set by 10ppm diffraction losses; cf. Sec. 4.2.3.5),²⁴ coated either out to (i) their physical radius R_p or (ii) only out to $R_p - 8 \text{ mm}$. Table 4.2 summarizes our results for each one-parameter family.

Whether the mirror is coated out to the full physical mirror radius or not, in both cases quadratic fits to $I/I_{\text{BL}}(R_p)$ give minima rather near the baseline physical radii $R_p = 15.7 \text{ cm}$. Indeed, to within our accuracy of computation, the same mirror shape used as the baseline advanced LIGO design ($R_p = 15.70 \text{cm}$ and $H = 13.00 \text{cm}$) gives the optimal thermoelastic noise for mesa beams with 10 ppm diffraction loss. The beam radii D and thermoelastic noise I/I_{BL} for these near-optimal mexican-hat test masses are shown in Table 4.4 below.

4.6.4 Mesa beams reflecting off identical 40 kg frustum mirrors

Roughly speaking, two mirrors with a larger front face radius permit a wider beam to resonate in the arm cavity and yield even lower thermoelastic noise. Therefore, we explore arm cavities bounded by frustum-shaped mirrors (cf. footnote 1), which expand the front face of the mirror at the expense of the back face.

More specifically, we considered arm cavities resonating with mesa beams of beam radius parameter D bounded by 40 k test-mass mirrors whose front and back faces were as small as diffraction losses would permit (i.e. the front face produces precisely 10 pmm diffraction losses with arm-cavity

²³The arm cavity has three free parameters (mirror radius; mirror thickness; and mesa-beam radius parameter D) and two constraints (mirror mass and diffraction losses per bounce).

²⁴As expected, thermoelastic noise decreased with increasing mesa beam radius; we obtain the lowest value of thermoelastic noise integral I when the mesa beam radius is as large as possible, consistent with the diffraction constraint.

Table 4.3: The thermoelastic integral I for a frustum input test mass (ITM) and a Mexican-hat beam, in units of $I_{\text{BL}} = 2.57 \times 10^{-28} \text{s}^4 \text{g}^{-2} \text{cm}^{-1}$. The values of I/I_{BL} are estimated to be accurate to within one per cent. The last two columns show the diffraction losses for each bounce off the test mass front and back faces, in ppm (parts per million). The last column is for the light impinging from the beam splitter onto face 2 of the test mass (radius R_2). The second-to-last column is for the light inside the arm cavity, on face 1 of the test mass (radius R_1).

R	$R_{p1}[\text{cm}]$	$R_{p2}[\text{cm}]$	$H[\text{cm}]$	D/b	I/I_{BL}	$\mathcal{L}_0[\text{ppm}]\text{front}$	$\mathcal{L}_0[\text{ppm}]\text{back}$
$R_p - 8\text{mm}$	14.67	10.57	19.81	3.00	0.355	10	10^4
$R_p - 8\text{mm}$	15.70	11.56	17.00	3.43	0.253	10	10^4
$R_p - 8\text{mm}$	17.11	12.88	14.06	4.00	0.207	10	10^4
$R_p - 8\text{mm}$	17.45	13.22	13.45	4.13	0.208	10	10^4
$R_p - 8\text{mm}$	19.58	15.27	10.43	5.00	0.285	10	10^4
R_p	13.94	9.88	22.24	3.00	0.345	10	10^4
R_p	15.70	11.56	17.00	3.73	0.198	10	10^4
R_p	16.37	12.19	15.49	4.00	0.175	10	10^4
R_p	17.29	13.04	13.75	4.39	0.162	10	10^4
R_p	18.85	14.58	11.33	5.00	0.193	10	10^4
R_p	21.36	17.00	8.62	6.00	0.398	10	10^4

mesa beam light of scale D ; the back face produces precisely one percent diffraction losses for the same mesa-beam input light; cf. Sec. 4.2.4).²⁵ These arm cavities satisfy advanced LIGO design constraints (cf. Sec. 4.2.4).

Table 4.3 summarizes the thermoelastic integrals I/I_{BL} for two one-parameter family of designs: (i) the mirror is coated out to its physical radius $R = R_p$ and (ii) the mirror is only coated out to $R = R_p - 8$ mm. By fitting a quadratic to $(D/b, I/I_{\text{BL}})$, we estimate the optimal mirror dimensions and associated beam radii D . Our optimal results appear in Table 4.4.

4.7 Interferometer sensitivity to mirror perturbations

In Sec. 4.6, we found mirror and beam configurations for the advanced LIGO arm cavity which indeed possess lower thermoelastic noise than the baseline advanced LIGO design. But our primary modification—the change to Mexican-hat optics for the cavity arms—involves employing mirrors which have never before been used in an interferometer. Naturally, then, we must make every effort to demonstrate that this radical proposal will not introduce new problems.

For example, the Mexican-hat mirror has a very flat central region (cf. Fig. 4.1). In our early presentations of this proposal, it was suggested to us that such a mirror design might make the resulting interferometer substantially more susceptible to errors, be they from static tilts and displacements or mirror figure error.

²⁵While in principle we could consider any combination of front and back face radii R_1 and R_2 and any mesa beam radius D such that all three satisfy the LIGO design constraints presented in Sec. 4.2.4, we found that for R_2 greater than or equal to the minimum radius allowed by diffraction losses the thermoelastic noise integral increases with R_2 (for fixed mirror mass and front face size). Therefore, we limited attention to R_2 as small as possible. Also, as usual, we limited attention to mesa beams as large as diffraction losses on the mirror front face permit.

Table 4.4: Optimized test-mass and light beam configurations, their thermoelastic noise compared to the baseline. [A subset of this table appears as Table I in MBI [4].]

Test Masses $\{R_{p1}, R_{p2}; H\}$	Beam Shape	$\left(\frac{S_h}{S_{BL}^h}\right)_{TE}$
BL: cylinders, $R = R_p - 8\text{mm}$ $\{15.7, 15.7; 13.0\}$	BL: Gaussian $r_o = 4.23\text{cm}$	1.000
BL: cylinders, $R = R_p - 8\text{mm}$ $\{15.7, 15.7; 13.0\}$	mesa $D/b = 3.73$	0.364
identical frustums, $R = R_p - 8\text{mm}$ $\{17.11, 12.88, 14.06\}$	mesa $D/b = 4.00$	0.207
BL: cylinders, $R = R_p$ $\{15.7, 15.7; 13.0\}$	Gaussian $r_o = 4.49\text{cm}$	0.856
BL: cylinders, $R = R_p$ $\{15.7, 15.7; 13.0\}$	mesa $D/b = 3.73$	0.290
identical frustums, $R = R_p$ $\{17.29, 13.04, 13.75\}$	mesa $D/b = 4.39$	0.162

In this section, we examine this concern by examining the effect of perturbations on both gaussian and mesa-beam arm cavities. More explicitly, we by applying the tools described in Sec. 4.4 to two fiducial beams (cf. MBI Sec. IV A 1):

- *Fiducial mesa beam*: The fiducial mesa-beam arm cavity, described in Sec. 4.3.1.1, which has mesa beams with $D = 4b = 10.4$ cm.
- *Fiducial Gaussian beam*: A fiducial Gaussian beam arm cavity, which has beam radius $r_o = 4.70$ cm (i.e. a g -value $g = 0.952$).²⁶ [This fiducial Gaussian-beam arm cavity *differs* from the baseline (cf. MBI Sec. IV A 1).] The fiducial Gaussian beam cavity is chosen to have the same diffraction losses, on a mirror of the same coated radius, as the mesa-beam arm cavity (i.e. so the two fiducial cavities we compare are similar).

Using these two fiducial beams, we demonstrate that mesa-beam and Gaussian-beam interferometer designs for advanced LIGO will have broadly (i.e. within a factor \sim a few) similar sensitivity to perturbations. In short, we demonstrate that the mesa-beam proposal will not introduce undue sensitivity of the arm cavities and interferometers to mirror errors.

4.7.1 Frequency distribution of parasitic modes as a measure of arm cavity sensitivity to perturbations

Given the denominators present in Eq. (4.32) and (4.33), a system will generically be more unstable to perturbations if the eigenphases of excited modes [i.e. $\arg(\eta_k)$] are close to the eigenphase of the resonant state. We can therefore crudely characterize the influence of generic perturbations by the

²⁶The value $g = 1 - L/\mathcal{R}_1$, for \mathcal{R}_1 the radius of curvature of spherical mirrors, can be related to the Gaussian beam radius r_o using formulae presented in Appendix B.7.

distribution of eigenmodes nearby the ground state, also called the parasitic mode distribution. This discussion provides the basis for MBI Sec. IV B.

The resonant frequencies $\omega = kc$ are determined when, after one round trip through the arm cavity, light in a given state interferes constructively with itself. Therefore, the light must be in an eigenstate $|p\rangle$ of the cavity [Eq. (4.13)] and, moreover, the frequency of the light must be chosen so the eigenvalue of the round-trip eigenequation is *real*, or chosen so [cf. Eq. (4.14)]

$$2\pi n = 2L\omega/c + \text{Arg}(\eta_p) \quad (4.54)$$

for n some integer. The same eigenmode $|p\rangle$ resonates at a specific frequency, and every other frequency separated from that frequency by the free spectral range $\omega_{FSR} = \pi c/L$. Within each free spectral range, different eigenmodes (i.e. $|p\rangle, |q\rangle$) are occur at different frequencies, separated by an amount uniquely determined by their eigenvalues η :

$$\Delta\omega_{pq} \equiv \omega_p - \omega_q = \text{Arg}(\eta_p) - \text{Arg}(\eta_q) . \quad (4.55)$$

For Gaussians beams, the eigenmodes are distributed within each free spectral range regularly; each nearest neighbor is separated by frequency [cf. Eq. (B.38)]

$$\begin{aligned} \Delta\omega &= \frac{c}{L} \times \cos^{-1} g = \omega_{FSR} \times \frac{\cos^{-1} g}{\pi} \\ &= 0.099 \times \omega_{FSR} , \end{aligned} \quad (4.56)$$

using the g -value $g = 0.952$ for our Gaussian baseline beams.

For the mesa beams, the frequency distribution of parasitic modes must be obtained numerically, by (i) solving the eigenequation for the eigenvalues η_p [Eq. (4.13), or equivalently Eq. (4.29), using the numerical methods of Sec. 4.5.2] and then by (ii) using those eigenvalues in Eq. (4.55) to deduce $\Delta\omega_{p0}$ and therefore the distribution of parasitic modes. Table 4.5 lists the values for $\Delta\omega_{p0}/\omega_{FSR}$ for a few states. Among states with low diffraction losses (i.e. with $\mathcal{L} = 1 - |\eta_p|^2 < 10^{-2}$), most modes are very well separated from the ground state; the nearest parasitic mode of a mesa beam cavity is only a factor ~ 2.5 closer to the ground state frequency than the nearest parasitic mode of a Gaussian-beam arm cavity. Therefore, as discussed in MBI Sec. IV C, we crudely expect the mesa-beam cavity to be only marginally more sensitive to perturbations than a Gaussian-beam arm cavity.

Table 4.5: For a LIGO arm cavity with fiducial Mexican hat mirrors ($D/b = 4$, $R = 16$ cm), this table gives the separation $\Delta\omega/\omega_{\text{FSR}}$ of the eigenfrequencies of parasitic modes from the eigenfrequency of the fundamental mesa-beam mode.

	$l = 0$	$l = 1$	$l = 2$	$l = 3$
$p = 0$	0.0	0.0404	0.1068	0.1943
$p = 1$	0.1614	0.2816	0.4077	-0.4581
$p = 2$	0.4303	-0.4140	-0.2570 ¹	-0.0812 ¹
$p = 3$	-0.2330 ¹	-0.0488 ¹	0.1406 ¹	— ¹

4.7.2 Effect of displacement on cavities bounded by spherical and Mexican-hat mirrors

The discussion of the previous section gives us good reason to believe that mesa-beam and gaussian-beam arm cavities will display the same sensitivity (i.e. within a factor ~ 2.5) to mirror perturbations. Here, we test this proposition when the ETM is displaced through a distance $\vec{s} = s\hat{x}$ – or, more explicitly, when mirror 2 is perturbed by $\delta h_2 = \delta h_{\text{disp}}$, given by

$$\begin{aligned}
 \delta h_{\text{disp}} &= (\vec{s} \cdot \vec{r}) \frac{1}{r} \frac{dh_2}{dr} \\
 &+ \frac{1}{2} \frac{|\vec{s} \times \vec{r}|^2}{r^3} \frac{dh_2}{dr} + \frac{1}{2} \frac{(\vec{s} \cdot \vec{r})^2}{r^2} \frac{d^2 h_2}{dr^2} \\
 &+ O(s^3) \\
 &= (xs) \frac{1}{r} \frac{dh_2}{dr} + \frac{1}{2} \frac{(xs)^2}{r^2} \frac{d^2 h_2}{dr^2} \\
 &+ \frac{1}{2} \frac{(ys)^2}{r^3} \frac{dh_2}{dr} + O(s^3).
 \end{aligned} \tag{4.57}$$

Section 4.4.5 summarizes the step-by-step process we follow to explore the influence of perturbations (here, displacement). However, because we expect—and our calculations below confirm—that whatever the precise mirror shapes, the properties of the interferometer will depend only weakly on displacement,²⁷ we do not complete all the steps that procedure includes (e.g., we do not compute the change in thermoelastic noise or cavity gain with displacement).

4.7.2.1 Displacement of spherical mirrors

To evaluate the perturbation expansion [Eqs. (4.33) and (4.35)] for a cavity bounded by two identical spherical mirrors subjected to a displacement of its ETM through a distance s [Eq. (4.57)], we use analytic techniques specialized to spherical mirrors (e.g., Hermite-Gauss basis functions; cf. Appendix B.7). After some algebra (described in detail in Appendix B.7), we find the ground state

²⁷The natural length parameter for the problem is $b = \sqrt{\lambda L/2\pi} = 2.6$ cm, the diffraction length. When we perform perturbation theory, we find results which vary in powers of s/b . Since the LIGO control system will control displacements to much smaller than 1 cm, displacements have a relatively small effect on the LIGO interferometer.

of the perturbed cavity, to first order, to be [cf. MBI Eq. (4.12)]

$$|u\rangle = |0\rangle + \zeta_1^{\text{BL}} |(1,0)\rangle + \mathcal{O}(s^2), \quad (4.58a)$$

$$\begin{aligned} \zeta_1^{\text{BL}} &= \frac{(1-g)^{1/4}}{\sqrt{2}(1+g)^{3/4}} (s/b) \\ &= 0.008(s/1 \text{ mm}). \end{aligned} \quad (4.58b)$$

[For clarity and for consistency with other work, we have absorbed a phase into the definition of the $|(1,0)\rangle$ Hermite-Gauss state.] Here, $g \equiv 1 - L/R_c = 0.952$ is the g -value for the two fiducial spherical mirrors.

4.7.2.2 Displacement of Mexican-hat mirrors

Similarly, if we repeat the above calculation for Mexican-hat mirrors using the methods summarized in Sec. 4.4.5, we find the ground state of the cavity changes to [cf. MBI Eq. (4.12)]

$$|\psi\rangle = |0\rangle + \zeta_1^{\text{MH}} |w_1\rangle + \mathcal{O}(d^2) \quad (4.59a)$$

$$|w_1\rangle \equiv |\psi^{(1)}\rangle / \zeta_1^{\text{MH}} \quad (4.59b)$$

$$\zeta_1^{\text{MH}} \equiv 0.262 (s/b) = 0.010 (s/1 \text{ mm}). \quad (4.59c)$$

Note that unlike the Gaussian-beam case, $|w_1\rangle$ is a unit-norm *superposition* of eigenmodes of the unperturbed (mesa-beam) cavity, rather than an eigenmode of that cavity itself.

4.7.2.3 Power in parasitic modes

Perturbations cause changes in the resonant ground state. Equivalently, perturbations couple the ground state to the parasitic modes, causing power to bleed from the resonant ground state of the unperturbed arm cavity into these other arm cavity eigenmodes.

In the case of displacement, the power in the parasitic modes is easily distinguished from the carrier light by symmetry: while the carrier light is axisymmetric, the lowest-order changes in state are dipolar (i.e. $|\psi^{(1)}\rangle \propto |(1,0)\rangle$ for Gaussian beams). The fraction of the total arm cavity light power in the dipolar parasitic modes is [cf. MBI Eq. (4.13)]

$$P_1 = \zeta_1^2 \simeq \begin{cases} 100(s/1.3\text{mm})^2 \text{ppm} & \text{BL,} \\ 100(s/1.0\text{mm})^2 \text{ppm} & \text{MH.} \end{cases} \quad (4.60)$$

4.7.2.4 Power out the dark port

If one arm cavity's ETM is displaced, then the light leaving the two arm cavities will not interfere destructively at the dark port. The precise amount of power P_{dp} out the dark port depends on the driving beam [cf. Eq. (4.41), which depends on γ_0 , γ_1 , and γ_2]. If the interferometer is driven by the

optimal Gaussian beam (cf. Sec. 4.3.1.2), then the power out the dark port is approximately entirely in a dipolar mode, with net power [i.e. Eq. (4.41) to lowest order, with $\gamma_1 = 0$; cf. MBI Eq. (4.14)]

$$P_{\text{dp}}^{\text{total}} = 2\gamma_0^2 \zeta_1^2 \simeq \begin{cases} 190(s/1.3\text{mm})^2 \text{ppm} & \text{BL,} \\ 190(s/1.0\text{mm})^2 \text{ppm} & \text{MH.} \end{cases} \quad (4.61)$$

4.7.3 Effect of tilt on the resonant eigenstate of an arm cavity bounded by spherical and Mexican-hat mirrors

Because both the spherical and Mexican-hat mirrors planned for advanced LIGO are very flat, the planned advanced LIGO arm cavities will necessarily be fairly sensitive to tilt.²⁸ This section demonstrates that, though Mexican-hat mirrors have a very flat central region—much more so than their spherical counterpart (cf. Fig. 4.2)—an arm cavity bounded by Mexican-hat mirrors will only be somewhat (i.e. a factor \sim few) more sensitive to tilt.

If the ETM of an arm cavity is tilted through an angle θ about its y axis, the mirror surface is effectively perturbed by $\delta h_2 = \delta h_{\text{tilt}}$:

$$\delta h_{\text{tilt}} = \theta x + O(\theta^3). \quad (4.62)$$

4.7.3.1 Tilt-induced changes in the arm cavity ground state

When this perturbation is inserted into the perturbation expansion (4.33) and the terms in that expansion are evaluated in the case of *spherical mirrors* (using special properties of spherical mirrors and Hermite-Gauss basis states; cf. Appendix B.7), we find first- and second-order corrections to the state [i.e terms in Eqs. (4.33) and (4.35)] to be given by [cf. MBI Eqs. (4.3) and (4.4)]

$$|\psi_1\rangle = \alpha_1^{\text{BL}} |(1, 0)\rangle, \quad (4.63a)$$

$$|\psi_2\rangle = \alpha_2^{\text{BL}} |(2, 0)\rangle, \quad (4.63b)$$

$$\begin{aligned} \alpha_1^{\text{BL}} &\equiv \frac{1}{\sqrt{2}(1-g^2)^{3/4}} (\theta_y \times \sqrt{kL}) \\ &= 0.0064(\theta/10^{-8}), \end{aligned} \quad (4.63c)$$

$$\begin{aligned} \alpha_2^{\text{BL}} &\equiv \frac{4}{\sqrt{2}(1-g^2)^{1/2}(1-g)} (\theta_y \times \sqrt{kL})^2 \\ &= 0.00046(\theta/10^{-8})^2. \end{aligned} \quad (4.63d)$$

Here $g = 0.952$ is the g -value for the fiducial cavity, the states $|(m, n)\rangle$ denote states in the Hermite-Gauss basis, and these states have been adjusted in phase to make α_1 and α_2 real.

If this same expansion is evaluated using *Mexican-hat mirrors* via the numerical procedure out-

²⁸For example, to order of magnitude, a tilt angle $\theta \sim \sqrt{\lambda/2\pi L} = 6 \times 10^{-6}$ should produce a perturbation of order unity in the optical state of the cavity. Because an order-unity change in state implies fairly dramatic change in the interferometer, tilts much smaller still (of order $\text{few} \times 10^{-8}$) can cause serious difficulty with the interferometer.

lined in Sec. 4.4.5, we find first- and second-order corrections to the state [i.e terms in Eq. (4.35)]. We will not provide an explicit form for these states here. These corrections have norms given by the Mexican-hat analogues of Eq. (4.63) for α_1 and α_2 [cf. MBI Eq. (4.5)]:

$$\begin{aligned}\alpha_1^{\text{MH}} &\equiv \|\psi^{(1)}\| = 14.78 (\theta\sqrt{kL}) \\ &= 0.0227(\theta/10^{-8})\end{aligned}\tag{4.64a}$$

$$\begin{aligned}\alpha_2^{\text{MH}} &\equiv \|\psi^{(2)}\| = 74.97 (\theta\sqrt{kL})^2 \\ &= 0.00018(\theta/10^{-8})^2.\end{aligned}\tag{4.64b}$$

We can use these expressions to define normalized representations of the first- and second-order corrections:

$$u_1 \equiv \left| \psi^{(1)} \right\rangle / \alpha_1^{\text{MH}},\tag{4.64c}$$

$$u_2 \equiv \left| \psi^{(2)} \right\rangle / \alpha_2^{\text{MH}}.\tag{4.64d}$$

4.7.3.2 Parasitic mode power excited by tilt

The largest correction to the perturbed arm cavity ground state is a *dipolar* perturbation (i.e. $|1, 0\rangle$ for Gaussians; $|u_1\rangle$ for mesa beams). Therefore, interpreting this change as an excitation of dipolar parasitic modes, the fractional power in the dipolar parasitic modes is [cf. MBI Eq. (4.6)]

$$P_1 = \alpha_1^2 \simeq \begin{cases} 0.0005(\theta/3.5 \times 10^{-8})^2 & \text{BL,} \\ 0.0005(\theta/1.0 \times 10^{-8})^2 & \text{MH,} \end{cases}\tag{4.65}$$

when the ETM is tilted through an angle θ .

4.7.3.3 Tilt-induced changes in the diffraction losses of the ground state of the resonant arm cavity (Mexican-hat only)

The diffraction losses associated with the ground state—which we approximate by the clipping approximation losses \mathcal{L}_1 and \mathcal{L}_2 [Eq. (4.16)]—also change when the beam state changes. As described in Sec. 4.4.3.1, we find the perturbed value for, say, \mathcal{L}_1 merely by expanding the expression for \mathcal{L}_1 , obtaining the general expansion Eq. (4.37).

If we evaluate Eq. (4.37) for the case of a cavity bounded by two Mexican hat mirrors with one mirror (the ETM, i.e. mirror 2) tilted through an angle θ , we find we can rewrite Eq. (4.37) for the losses at mirror 1 (the ITM) in terms of the expansion mentioned above (i.e. in terms of α_1 , α_2 , u_1 ,

and u_2):

$$\begin{aligned}\mathcal{L}'_1 &= \mathcal{L}_1 + \alpha_1^2 (\mathcal{L}_A^{(2)} - \mathcal{L}_1) + \alpha_2 \mathcal{L}_B^{(2)} \\ &\approx 18 \text{ pmm} [1 + 0.0025(\theta/10^{-8})^2]\end{aligned}\tag{4.66}$$

[cf. MBI Eq. (4.7)] where

$$\begin{aligned}\mathcal{L}_1 &= \langle u_o | O_1 | u_o \rangle = 18 \text{ pmm} , \\ \mathcal{L}_A^{(2)} &= \langle u_1 | O_1 | u_1 \rangle = 96 \text{ pmm} , \\ \mathcal{L}_B^{(2)} &= 2\text{Re} \langle u_o | O_1 | u_2 \rangle = 29 \text{ pmm} .\end{aligned}\tag{4.67}$$

and where u_1 and u_2 are defined by Eq. (4.64).

Using a similar technique we can also evaluate the change in the clipping approximation diffraction losses at mirror 2. [As discussed Sec. 4.4.4.2, we for simplicity assume the diffraction losses at both mirrors remained the same as the mirror tilted, or $\mathcal{L}_1 = \mathcal{L}_2$. Since the diffraction losses influence physical quantities like the cavity gain fairly little, we require only a rough estimate of their sensitivity to perturbations.]

4.7.4 Effect of tilt on arm cavities and interferometers using Mexican-hat mirrors but driven by Gaussian beams

Tilt of one ETM causes changes in several important properties of an interferometer that uses mesa beams (and is driven by the optimal Gaussian): (i) the amount of thermoelastic noise present in the dark port signal; (ii) the amount of power present in the perturbed arm cavity (i.e. the arm cavity gain); and (iii) the amount of input light power leaving the dark port. In this section we compute the changes in these three quantities by applying the general techniques presented in Sec. 4.4.3 to the case of tilt.

4.7.4.1 Tilt and thermoelastic noise

To lowest order in $\theta/(b/L)$, tilt has *no effect* on thermoelastic noise: at first-order, tilt only excites odd-parity modes, which produce nonaxisymmetric intensity perturbations, and nonaxisymmetric intensity perturbations do not contribute to lowest-order changes in the thermoelastic noise integral (cf. Sections 4.4.3.4 and 4.4.4.1).

4.7.4.2 Tilt and cavity gain, when driven by a Gaussian beam

Equation (4.39) provides a general expression for the cavity gain for a perturbed cavity; this expression depends on (i) the amount of light power γ'_o entering the perturbed arm cavity [Eq. (4.38), evaluated using the perturbed mesa-beam state from Eq. (4.64)], which we can express using the perturbation parameters γ_1 and γ_2 ,

$$\gamma_1 = 0 \quad \gamma_2 = -0.070 + 0.013i ; \quad (4.68)$$

and (ii) on the diffraction losses of the resonant state of the tilted arm cavity [Eq. (4.66) above].²⁹ Combining these expressions, we find the arm cavity power, relative to the input beam power, to be given by [cf. MBI Eq. (4.8)]

$$\begin{aligned} \mathcal{G}'_{\text{anal}} &= \mathcal{G}_{\text{anal}} \left[1 - \alpha_1^2 + \frac{\alpha_2}{\gamma_0} (\gamma_2 + \gamma_2^*) - 2 \frac{\mathcal{L}'_1 - \mathcal{L}_1}{1 - r_I} \right] \\ &= 737 \left[1 - 5.5 \times 10^{-4} (\theta/10^{-8})^2 \right] . \end{aligned} \quad (4.69)$$

4.7.4.3 Tilt and dark port power, when driven by a Gaussian beam

Finally, when one ETM mirror in a mesa-beam interferometer is tilted, the beamsplitter sends light to the dark port. Equation (4.41) provides a general expression for the light u_{dp} leaving the interferometer through the dark port; for tilt, this expression evaluates to

$$u_{\text{dp}} = \sqrt{2} \left[(-\gamma_o \alpha_1^2 + \gamma_2 \alpha_2) |u_o\rangle + \alpha_1 \gamma_o |u_1\rangle + \alpha_2 \gamma_o |u_2\rangle \right] \quad (4.70)$$

where we use (i) the definitions $\alpha_1 = \|\psi^{(1)}\|$ and similarly [cf. Eq. (4.64)], and (ii) $\gamma_1 = 0$, to simplify the general equation (4.41).

The corresponding fraction of the interferometer's power that exits the interferometer through the dark port in the fundamental mode u_o and in the parasitic modes u_1 and u_2 is [cf. MBI Eq. (4.9); note the MBI expression will be larger by a factor 4]³⁰

$$P_{\text{dp},0} = |-\gamma_o \alpha_1^2 + \gamma_2 \alpha_2|^2 \simeq 0.256 (\theta/10^{-8})^4 \text{ pmm} \quad (4.71)$$

$$P_{\text{dp},1} = \gamma_0^2 (\alpha_1^{\text{MH}})^2 \simeq 478 (\theta/10^{-8})^2 \text{ pmm} , \quad (4.72)$$

$$P_{\text{dp},2} = \gamma_0^2 (\alpha_2^{\text{MH}})^2 \simeq 0.024 (\theta/10^{-8})^4 \text{ pmm} . \quad (4.73)$$

²⁹When computing the diffraction losses for a tilted cavity, we assume the same diffraction losses associated with a bounce off mirror 1 and mirror 2.

³⁰The MBI expression estimates the effect when all four mirrors are tilted about uncorrelated axes; it therefore is larger than our result, which describes the effect of tilting only one mirror.

4.7.5 Effect of mirror figure error on mesa-beam interferometers

In this section, we explore the sensitivity of individual arm cavities and the overall interferometer to mirror figure error, when the beam resonating in the arm cavities is a fiducial mesa beam. More specifically, in this section we (i) distort the ETM by a physically plausible amount (i.e. an amount estimated from actual mirror figure error measurements of LIGO mirrors); (ii) apply perturbation theory (cf. Sec. 4.4.5) to deduce the change in resonant ground state of the arm cavities; and then (iii) use the resulting modified beam state to deduce how the power spectrum of thermoelastic noise will change due to the perturbed beam cross section.

MBI [4] uses the computations performed in this section to place constraints on the accuracy of machining required of the mirrors used in a mesa-beam interferometer; cf. MBI Sections IV F and IV G.

4.7.5.1 GLB’s worst-case figure error

GariLynn Billingsley has provided us with a map of a worst-case figure error, $\delta z_{wc}(x, y)$ [height error as function of Cartesian coordinates in the transverse plane], produced by current technologies. Her map is based on the measured deviation of a LIGO-I beam-splitter substrate from flatness. The measured substrate had diameter 25 cm; she stretched its deviation from flatness (its “figure map”) to the baseline advanced LIGO mirror diameter of 35.4 cm, fit Zernike polynomials to the stretched map, and smoothed the map by keeping only the lowest 36 Zernikes. MBI Fig. 6 provides a contour diagram of the resulting figure map (figure “error”). In the central region (innermost 10 cm in radius), the peak to valley error Δz is about 30 nm, while in the outer region (10 cm to 16 cm in radius), it is about 110 nm.

Billingsley (private communications) thinks it likely that in the central region (which dominates our considerations), peak-to-valley errors of $\Delta z \sim 5$ nm (about 1/5 as large as in MBI Fig. 6) may be achievable; Jean Marie Mackowski believes even smaller errors can be obtained with coating methods (cf. MBI Sec. IV F 1). Accordingly, in the analyses described below we shall use Billingsley’s map, scaled down in height by a factor ε [cf. MBI Eq. (4.15)]:

$$\delta z = \varepsilon \delta z_{wc}(x, y) , \tag{4.74}$$

and we shall use $\varepsilon = 0.2$ and $\Delta z = 6$ nm as our fiducial values for ε and Δz .

Also, Billingsley thinks it likely that the outer regions of the mirror will be of significantly lesser quality than the inner regions. To study the sensitivity of the interferometer and thermoelastic noise to errors in the exterior, we divided GLB’s perturbations into two regions by a weight function [cf.

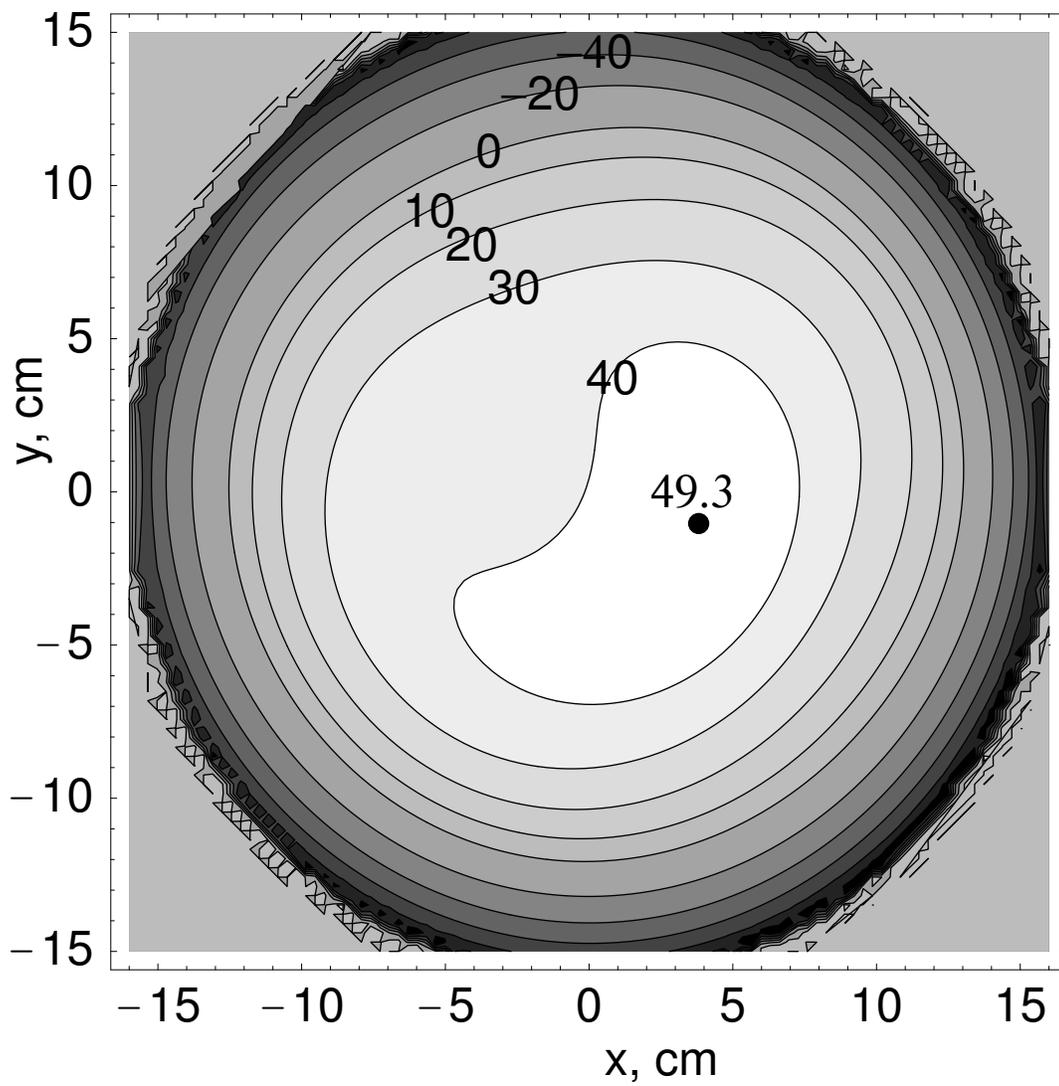


Figure 4.7: Contour diagram of GariLynn Billingsley's worst-case mirror figure error [height δz_{wc}] in nanometers, cf. MBI Fig. 6.

MBI Eq. (4.19)]:

$$\begin{aligned} \delta z &= \varepsilon_c W(r/b - 3.7) \delta z_{wc} \\ &+ \varepsilon_o [1 - W(r/b - 3.7)] \delta z_{wc} \end{aligned} \quad (4.75)$$

where $W(x) = 1$ if $x > 1$, $W(x) = x$ if $x \in [0, 1]$, and $W(x) = 0$ if $x < 0$. This tapering function permits us to give the central ~ 10 cm in radius one level of figure error (determined by ε_o) and the outer ~ 6 cm a different level (determined by ε_c).

4.7.5.2 Effect of mirror figure errors, without compensating tilt

Perturbation theory (implemented according to Sec. 4.4.5) permits us to represent the change in the mesa-beam ground state when mirror 2 is perturbed by GLB's worst case perturbation, reduced by a factor ε [Eq. (4.74)]. Using the numerical prescription from Sec. 4.5.3 to evaluate Eq. (4.33), the first- and second-order corrections to the state are found to be [cf. MBI Eqs. (4.16) and (4.17)]

$$|\psi_1\rangle = \beta_1 |v_1\rangle, \quad (4.76a)$$

$$|\psi_2\rangle = \beta_2 |v_2\rangle, \quad (4.76b)$$

$$\beta_1 \equiv |\psi_1| = 0.167\varepsilon \quad (4.76c)$$

$$\beta_2 \equiv |\psi_2| = 0.123\varepsilon^2. \quad (4.76d)$$

[The scale factor ε must be fairly small (~ 0.2) for the second order correction to be significantly smaller than the first order correction (i.e. for the series to converge, and for perturbation theory to be useful.)] The series expansion can be used, as in the case of displacement and tilt, to determine physical quantities such as

- *The power in parasitic modes in the arm cavity:* MBI Fig. 7 shows a map of the power distribution $|\psi^{(1)}(r)|^2 = |\beta_1 v_1(r)|^2$ in the most-strongly-excited parasitic mode combination $|v_1\rangle$. The net power in parasitic modes is well approximated by the integral of this quantity [cf. MBI Eq. (4.17)]:

$$P_1 = |\beta_1|^2 = 0.0012(\varepsilon/0.2)^2. \quad (4.77)$$

- *The power leaving the dark port:* Equation (4.41) tells us that the fraction of the interferometer's overall laser light input power going out the dark port is approximately [i.e. Eq. (4.41) to lowest order; cf. MBI Eq. (4.18)]

$$\begin{aligned} P_{dp} &\approx (|\gamma_o|^2 + |\gamma_1|^2) |\beta_1|^2 \\ &\simeq \gamma_o^2 |\beta_1|^2 = 0.0011(\varepsilon/0.2)^2 \end{aligned} \quad (4.78)$$

where we neglect γ_1 and compute only the lowest order term in the dark port power.³¹

Naturally, the interferometer is more sensitive to errors in the inner ~ 10 cm, where the beam power is large, than to the outer ~ 6 cm, where the beam power is small. To investigate this effect, we considered height perturbations of the form discussed in Eq. (4.75); to characterize the sensitivity of the beam shape to these defects, we examined the norm $\beta_1^2 = \|\psi^{(1)}\|^2$ [Eq. (4.76)], which is necessarily a quadratic form in ε_c and ε_o [cf. MBI Eq. (4.20)]. Using the same methods as for Eq. (4.76), we found the norm β_1^2 to be

$$\begin{aligned}\beta_1^2 &= 0.027\varepsilon_c^2 + 4 * 10^{-4}\varepsilon_o^2 + 3.6 * 10^{-6}\varepsilon_c\varepsilon_o \\ &\approx 0.00108[(\varepsilon_c/0.2)^2 + 0.37\varepsilon_o^2].\end{aligned}\tag{4.79}$$

If β_1 is used as a characteristic example of the sensitivity of physical quantities (e.g., the power out the dark port) to height perturbations, then a mesa-beam interferometer is around 8 times more sensitive to mirror figure errors in the interior region of each test-mass mirror (i.e. its inner ~ 10 cm) than it is to perturbations outside that region.

4.7.5.3 Mirror figure errors, with compensating tilt

The tilt control system automatically and dynamically reorients the mirrors in response to what it interprets as tilt. Specifically, the mirror tilt control system (i) *measures signals containing information about beam asymmetry*, such as the output of a quadrant photodiode; (ii) *computes the mirror tilt that would generate these asymmetries*, using perturbation theory expansions [i.e. Eq. (4.33), to first order]; and then (iii) *tilts all four LIGO mirrors to eliminate the apparent tilts* that the system computed in the previous step. Therefore, since mirror figure error *also* produces beam asymmetries, the tilt control system of the interferometer will act to partially compensate (the dipolar part of) the static mirror defect.

The precise quantity the tilt control system measures to deduce the tilt angle is not important: different approaches to tilt control interpret the optical state of our arm cavity [i.e. interpret the state $|u'_o\rangle$ defined by the corrections in Eq. (4.76)] in a manner fairly independent of the method used (i.e. the compensating tilt we calculate depends little on the method we use to calculate it). In this section, we assume the tilt control system acts to minimize the dipolar component of the arm cavity beam power.

Therefore, we find the tilt-compensated state by (i) adding together the results of a tilt perturbation [i.e. a generalization of Eq. (4.63) that depends on two tilt angles, θ_x and θ_y] and the mirror figure error perturbation [Eq. (4.76), which depends on ε]; (ii) evaluating the norm $\|\psi^{(1)}\|$ of the

³¹Unlike displacement and tilt, for mirror figure error γ_1 is generically nonzero, since the perturbation admits an axisymmetric part.

first-order perturbation; and then (iii) finding the tilt angles θ_x and θ_y which minimize that norm. The resulting optimal tilt angle is [cf. MBI Sec. IV H 3]

$$\theta_x = +0.98 \times 10^{-8}(\varepsilon/0.2), \quad (4.80a)$$

$$\theta_y = +0.69 \times 10^{-8}(\varepsilon/0.2), \quad (4.80b)$$

$$\theta = \sqrt{\theta_x^2 + \theta_y^2} = 1.2 \times 10^{-8}(\varepsilon/0.2) \quad (4.80c)$$

(in the limit of small ε , so linear theory applies). MBI Figure 8 shows the height of the surface of the mirror after the compensating tilt is applied.

After the tilt is applied, the first-order correction to the resonant mesa-beam state is a combination of tilt and mirror-figure perturbations. For brevity, we denote the net first-order correction to the mesa-beam state when tilt is applied by

$$|\psi^{(1)}\rangle = \beta_{1,c} |v_{1,c}\rangle, \quad (4.81a)$$

$$\beta_{1,c} \equiv \|\psi^{(1)}\| 0.02(\varepsilon/0.2). \quad (4.81b)$$

The square of the norm (i.e., $\beta_{1,c}^2$), as usual, is the fraction of arm cavity power which is in parasitic modes [cf. MBI (4.23)]:

$$P_{1c}^{\text{arm}} = |\beta_{1,c}|^2 = 0.0004(\varepsilon/0.2)^2. \quad (4.82)$$

MBI Figures 9 shows a map of the power in parasitic modes, $|\beta_{1,c} v_{1,c}(r)|^2$. Also, the norm of this first-order correction provides an estimate of the power leaving the dark port of this interferometer (i.e. an interferometer with one tilted, defective mirror) when the interferometer is driven with Gaussian beams [i.e. Eq. (4.41 to lowest order; cf. MBI (4.24)]:³²

$$\begin{aligned} P_{dp} &\approx (|\gamma_o|^2 + |\gamma_1|^2) |\beta_{1,c}|^2 \\ &\simeq \gamma_o^2 |\beta_1|^2 = 0.00038(\varepsilon/0.2)^2 \end{aligned} \quad (4.83)$$

4.7.5.4 Influence of mirror figure errors on thermoelastic noise

Because the ETM's figure error distorts the beam resonating in the arm cavity, the thermoelastic noise produced by each mirror bounding that arm cavity changes by some small amount. Given the change in state deduced above, we know how the beam profile at the ITM changes. We can therefore evaluate, using the discussion of Sec. 4.4.3, the change in thermoelastic noise associated with the beam reflecting off the ITM's face.

To be explicit, to compute the first-order effects of a perturbation δP to the thermoelastic

³²As in the untilted case, we neglect the (nonzero) term $|\gamma_1|^2$.

noise integral I_1 of mirror 1 in this fiducial case, we first select the axisymmetric portion $\delta P_o(r) = \int d\varphi P(\vec{r})/2\pi$ of δP ; we then linearize the analytic expressions derived by Liu and Thorne for cylindrical mirrors (cf. Sec. 4.5.1.2), using δP_o as the magnitude of the perturbation³³; and finally we extract from this linearization the first-order change in I_1 .

When we apply this technique to the fiducial beam reflecting off a cylindrical mirror of radius 16 cm and thickness 13 cm which is deformed by GariLynn Billingsley’s mirror distortion [Eq. (4.74)], we find the thermoelastic noise integral for that mirror changes to

$$I_1/(2 \times 10^{-28}) = 0.632 [1 + 0.035(\varepsilon/0.2)] \quad (4.84)$$

where I_1 denotes the total thermoelastic noise integral for mirror 1 when mirror 2 is deformed by our scaled height perturbation (4.74). [Roughly speaking, we expect the beam and hence the thermoelastic noise to change in a qualitatively similar fashion at mirror 2 (cf. Sec. 4.4.4.2)].

MBI Section IV G applies this result to deduce how sensitive the power spectrum of thermoelastic noise is to uncorrelated mirror figure errors on all four mirrors [cf. MBI Eq. (4.27)].

4.8 Conclusions

In this paper, we have described both the theory and practice needed to obtain the results summarized in MBI. Specifically, we have developed analytic and numerical tools to evaluate the following:

- *Thermoelastic noise integrals*: We developed practical techniques for finding the thermoelastic noise for nonstandard optical systems (i.e. noncylindrical finite mirrors and unusual beam shapes) [Sections 4.2.2 and 4.5.1]. We tabulated the thermoelastic noise, relative to the current advanced LIGO baseline, for many alternative mirror and beam configurations [Tables 4.1 - 4.4 in Section 4.6]. We found many configurations with lower thermoelastic noise than the baseline advanced LIGO configuration.
- *Eigenmodes of an arm cavity bounded by Mexican-hat mirrors*: We also wrote numerical code to find the eigenmodes of an optical cavity [Sections 4.2.3 and 4.5.2]. We computed and tabulated many of the eigenmodes of an arm cavity bounded by Mexican-hat mirrors (none of which appear explicitly in this paper).
- *Second-order optical perturbation theory*: Finally, we developed expressions for second-order optical perturbation theory [Sec. 4.4]. We applied perturbation theory extensively to study the sensitivity of mesa-beam interferometers to perturbations (i.e. to mirror figure error, mirror tilt, and mirror displacement).

³³When linearizing the LT equations, the natural relationship between δP and δp_m follows from Eq. (4.44). Note, however, that p_o is independent of $P(r)$ [because $P(r)$ are all normalized] and thus $\delta p_o = 0$.

The results found here are used in MBI to conclude that mesa-beams interferometer designs offer clear advantages over the baseline advanced LIGO design, without being substantially more sensitive to mirror figure error, tilt, or displacement perturbations.

In MBI [4], in Erika d’Ambrosio’s paper [6], and this paper, we and our collaborators have only taken the first steps towards the design of a practical mesa-beam advanced LIGO proposal. For example, more perturbative calculations—this time, applied to the final design, rather than to a fiducial case—are needed, for the design of the control system (i.e. so the relationship between light on various photodiodes and the correcting tilt applied to the LIGO mirrors can be established). Further, in this paper we have only begun to explore the space of all possible mirrors. For simplicity, we chose to fix the mirror mass to 40 kg. In practice, however, with sapphire, mirror designs are limited by fabrication limits (i.e. the radius of the mirror is limited by the radius of the sapphire boule one can grow) rather than by weight limits. Therefore, before a final design is chosen, more mirror designs (including cylinders with $m > 40$ kg) should be examined. Finally, in this paper, our Mexican-hat designs were limited to symmetric cavities (i.e. using *identical* mirrors). In practice, asymmetric designs offer the possibility of lower thermoelastic noise and greater practical convenience³⁴. Further work is necessary on the design, construction, and operation of asymmetric MH-like cavities.

³⁴For example, rather than design all four mirrors to be identical MH, one may want to operate with some spherical mirrors first, then replace a few mirrors (i.e. the ETMs) with MH-like mirrors later. Bill Kells has proposed using a flat ETM and first a spherical and then an MH-like ITM.

Bibliography

- [1] LIGO II Conceptual Project Book, available at <http://www.ligo.caltech.edu/docs/M/M990288-A1.pdf>
- [2] The LIGO Scientific Collaboration (LSC) maintains a website describing plans for the first-generation upgrade: <http://www.ligo.caltech.edu/advLIGO/>. This site was created for (and contains) the advanced LIGO NSF review. An older advanced LIGO website is at <http://www.ligo.caltech.edu/~ligo2/>.
- [3] Peter Fritschel maintained a website which provided a succinct summary of an older advanced LIGO design, including material parameters: <http://www.ligo.caltech.edu/~ligo2/scripts/l2refdes.htm> (quoted 28-June-2003)
- [4] E. Ambrosio, R. O'Shaughnessy, S. Strigin, K. S. Thorne, and S. Vyatchanin, to be submitted to Phys. Rev. D [i.e. Chapter 3 of this paper]
- [5] E. Ambrosio, R. O'Shaughnessy, S. Strigin, K. S. Thorne, and S. Vyatchanin, LIGO Report Number LIGO T030009, available at <http://ligo.caltech.edu>
- [6] E. D'Ambrosio, Phys. Rev. D **67**, 102004
- [7] R. Wald, *General Relativity*, University of Chicago Press, 1984.
- [8] P. A. M. Dirac, *The principles of quantum mechanics*, 4th edition (Oxford University Press, 1958)
- [9] L. D. Landau and E. M. Lifshitz, *Quantum Mechanics*, 3rd edition (Pergamon Press) 1977
- [10] J. S. Townsend, *A modern approach to quantum mechanics*, (McGraw-Hill, New York) 1992
- [11] V. Braginsky, M. Gorodetsky, and S. Vyatchanin, Phys. Lett. A **264**, 1 (1999).
- [12] Y. T. Liu and K. Thorne, Phys. Rev. D **62**, 1222002.
- [13] Y. Levin, Phys. rev. D **57**, 659 (1998)

- [14] R. D. Blandford and K. S. Thorne, *Applications of classical physics*, unpublished. Draft versions are available at <http://www.pma.caltech.edu/Courses/ph136/yr2002/index.html>
- [15] A. Siegman, *An Introduction to Lasers and Masers* (McGraw-Hill, 1971).
- [16] A. Siegman, *Lasers* (University Science, Mill Valley, 1986); errata URL: http://www-ee.stanford.edu/~siegman/lasers_book_errata.txt
- [17] A. Yariv, *Quantum electronics*, 3rd edition (John Wiley & Sons, 1989).
- [18] L. Ronchi, *Optical resonators*, in *Laser physics handbook*, vol. 1, F. T. Arecchi and E. O. Schulz-Dubois, eds. (North-Holland Publishing, 1972)
- [19] F. M. Barnes, ed., *Laser Theory*, (IEEE Press, 1972)
- [20] R. G. Beausoleil and D. Sigg, *J. Opt. Soc. Am. B* **107**, 1597 (1997). Also available as LIGO Report Number P990004, available at <http://ligo.caltech.edu>
- [21] R. G. Beausoleil et al. LIGO Report Number P020026, available at <http://ligo.caltech.edu>
- [22] FEMLAB, described at <http://www.comsol.com>
- [23] L. Landau and E. M. Lifshitz, *Theory of elasticity*, 2nd edition (Pergammon Press, 1970)
- [24] N. Nakagawa et al., *Phys Rev. D* **65**, 102001.
- [25] K. Numata, CaJAGWR talk, pdf slides available at <http://cajagwr.caltech.edu/pdf/numata.pdf> and video available at <http://cajagwr.caltech.edu/scripts/numata.ram>
- [26] R. P. Feynman and A. H. Hibbs, *Quantum Mechanics and Path Integrals*, (McGraw-Hill, 1965).

Chapter 5

Numerical relativity

[This section, along with the associated Appendix C, is precisely the text of R. O’Shaughnessy, gr-qc/0307001, submitted to Phys. Rev. D]

Abstract

Many alternative formulations of Einstein’s evolution have lately been examined, in an effort to discover one which yields slow growth of constraint-violating errors. In this paper, rather than directly search for well-behaved formulations, we instead develop analytic tools to discover which formulations are particularly ill-behaved. Specifically, we examine the growth of approximate (geometric-optics) solutions, studied only in the future domain of dependence of the initial data slice (e.g., we study transients). By evaluating the amplification of transients a given formulation will produce, we may therefore eliminate from consideration the most pathological formulations (e.g., those with numerically-unacceptable amplification). This technique has the potential to provide surprisingly tight constraints on the set of formulations one can safely apply. To illustrate the application of these techniques to practical examples, we apply our technique to the 2-parameter family of evolution equations proposed by Kidder, Scheel, and Teukolsky, focusing in particular on flat space (in Rindler coordinates) and Schwarzschild (in Painleve-Gullstrand coordinates).

5.1 Introduction

Recently developed numerical codes offer the possibility of extremely accurate and computationally efficient evolutions of Einstein’s evolution equations in vacuum [1]. To take full advantage of these new techniques to perform an unconstrained evolution of initial data and boundary conditions, we must address an unpleasant fact: many choices for evolution equations and boundary conditions permit ill-behaved, unphysical solutions (e.g., growing, constraint-violating solutions) near physical solutions.

By way of example, when Kidder, Scheel, and Teukolsky (KST) evolved a single static Schwarzschild hole as a test case, they found evidence suggesting that their evolution equations and boundary conditions, when linearized about a Schwarzschild background, admitted growing, constraint-violating *eigenmodes* [1] [2]. These eigenmodes were excited by generic initial data (i.e. roundoff error); grew to significant magnitude; and were directly correlated with the time their code crashed. As this example demonstrates, the existence and growth of ill-behaved solutions limit the length of time a given numerical simulation can be trusted—or even run.

For this reason, some researchers have explored the analytic properties of various formulations of Einstein’s equations [2, 3, 4, 5, 6, 7, 8] and boundary conditions [9, 10, 11, 12] used in numerical relativity, searching for ways to understand and control these undesirable perturbations.

In this paper, we discuss one particular type of undesirable perturbation: short-wavelength, transient wave packets. [For the purposes of this paper, a transient will be any solution defined in the future domain of dependence of the initial data slice. Depending on the boundary conditions, the solution may or may not extend farther in time, outside the future domain of dependence. Inside the future domain of dependence, however, “transient solutions” are manifestly independent of boundary conditions.] Depending on the evolution equations and background spacetime used, these transients can potentially grow significantly (i.e. by a factor of more than 10^{16} in amplitude). Under these conditions, even roundoff-level errors in initial data should produce transients that amplify to unit magnitude. Once errors reach unit magnitude, then guided by the KST results discussed above, we expect nonlinear terms in the equations to generically cause these errors to grow even more rapidly, followed shortly thereafter by complete failure of a numerical simulation. In other words, if the formulation and background spacetime permit transients to amplify by 10^{16} , we expect numerical simulations of these spacetimes to quickly fail.

In this paper we develop conditions which tell us when such dramatic amplification is *assured*. Specifically, we describe how to compute the amplification of certain transients for a broad class of partial differential equations (first-order symmetric hyperbolic PDEs) that includes many formulations of Einstein’s equations. If this amplification is larger than 10^{16} , then we know we should not evolve this formulation numerically.

5.1.1 Outline of remainder of paper

In this paper, we analyze the growth of transients. [Remember, in this paper a transient is any solution defined in the future domain of dependence of the initial data slice.] Rather than study all possible formulations, we limit attention to a class of partial differential equations we can analyze in a coherent, systematic fashion: first-order symmetric hyperbolic systems. Furthermore, because we concern ourselves only with stability and the growth of small errors, we limit attention to linear perturbations upon some background. Finally, to be able to produce concrete predictions, we restrict

attention to those transients which satisfy the geometric optics approximation.

In Sec. 5.2 we introduce an explicit ray-optics-limit solution to first-order symmetric hyperbolic linear systems—a class which includes, among its other elements, linearizations of certain formulations of Einstein’s equations. We provide explicit ODEs which determine the path (i.e. ray) and amplitude of a geometric-optics solution, in terms of initial data at the starting point of the ray. Then, in Sec. 5.3, we introduce wave packets as solutions which are confined to a small neighborhood of a particular ray. We further define two special classes of wave packet—coherent wave packets and prototypical coherent wave packets—which, because of their simple, special structure, are much easier to analyze. Finally, in Sec. 5.4, we introduce and discuss the technique (energy norms) we will use to characterize the amplitude of wave packets. In particular, we provide an explicit expression [Eq. (5.20)] for the growth rate of energy of a prototypical coherent wave packet.

To demonstrate explicitly how the techniques of the previous sections can be applied to produce the growth rate of transients, in Sections 5.5 and 5.6 (as well as Appendix C.5) we describe by way of example how our methods can broadly be applied to the two-parameter formulations that Kidder, Scheel, and Teukolsky (KST) have proposed [1]. Specifically, Sections 5.5 and 5.6 will respectively describe wave packets on flat-space (written in Rindler coordinates) and radially propagating transients on a Schwarzschild-black-hole background (expressed in Painleve-Gullstrand coordinates).

Finally, to demonstrate explicitly how expressions for the growth rate of transients can be used to filter out particularly pathological formulations, in Sections 5.7 and 5.8 we use the results for the growth rates of transients obtained in Sections 5.5 and 5.6 to determine what pairs of KST parameters (γ and \hat{z}) *guarantee* significant amplification of some transient propagating on a Rindler and Painleve-Gullstrand background, respectively.

Guide to the reader

While the fundamental ideas behind this paper—the study of wave packets and the use of their growth rates to discover ill-behaved formulations—remains simple, when we attempted to perform practical, accurate computations, we quickly found the simplicity of this idea masked behind large amounts of novel (but necessary) notation. We therefore found it difficult to simultaneously satisfy the casual reader—who wants only a summary of the essential results, and who is still evaluating whether the results and the methods used to obtain them are worthy of further attention—and the critical reader—who needs comprehensive understanding of our methods in order to evaluate, duplicate, and (potentially) extend them. We have chosen to slant the paper towards the towards the critical reader; thus this paper is a *comprehensive* and *pedagogical* introduction to our techniques.

While this paper can be consumed in a single reading, for the reader interested in a brief summary of the essential ideas and results, or for anyone making a first reading of this paper, the author recommends only reading the most essential details. First and foremost, the reader should

understand the scope and significance of this paper (i.e. read the abstract and Sec. 5.1). Next, the reader should follow the general description of the techniques in Sections 5.2, 5.3, and 5.4 in detail. Subsequently, the reader should examine our demonstration that our techniques indeed give correct results for growth rates (cf. the introduction to Section 5.5 and the summary of that section’s results in Section 5.5.4). Finally, to understand how these techniques can be used to discover ill-behaved formulations, the reader should examine Sections 5.7 and 5.8.

The more critical reader may wish to test and verify our computations. This reader should then review Sections 5.2, 5.3, and 5.4 again, then work through Sections 5.5 and 5.6 in detail (returning to the earlier sections for reference as necessary). This reader will also benefit from the general approach to KST 2-parameter formulations discussed in Appendix C.5.

Finally, the most skeptical readers will want to examine the conceptual underpinnings of and justifications for our every computation. This reader should simply follow the text as presented, but carefully read every footnote and appendix as they are mentioned in the text. In particular, this reader will want to review our Appendices C.2 (for a justification of our ray-optics techniques) and C.1 (for many useful identities used in the previous appendix and elsewhere in the paper) as well as Appendix C.3 (for a more detailed discussion of prototypical coherent wave packets, a key element in our computational method).

5.1.2 Connection with prior work

5.1.2.1 Study a short-time, rather than long-time, instability mechanism

First and foremost, we should emphasize that our work differs substantially from all previous work on this subject: we very explicitly restrict attention to amplification over only a short time (i.e. a light-crossing time). On the one hand, unlike other work, because of this restriction, our claims—being independent of boundary conditions—apply to *all* boundary conditions. On the other, because we forbid ourselves from studying our solutions outside the future domain of dependence of the initial data slice—even though, in practice, we could draw some elementary conclusions¹—in this paper we choose not to make any claims about how a formulation of Einstein’s equations will behave at late times (i.e. its late-time stability properties).

5.1.2.2 Study an instability mechanism, not necessarily the dominant one

In other papers which attempt to address the stability properties of various formulations of Einstein’s equations—for example, Lindblom and Scheel (LS) [2]—the authors try (somewhat naturally) to

¹In fact, because these solutions are high-frequency solutions, we can quite easily determine their interaction with most boundary conditions. For example, maximally-dissipative boundary conditions (i.e. the time derivatives of all ingoing characteristic fields are set to zero) imply, in the geometric-optics limit, that all solutions on ingoing rays will be zero. In particular, that implies that, when wave packets reach the boundary, they leave without reflecting. Other boundary conditions may also be easily analyzed.

understand the *dominant* instability mechanism. Unfortunately, we do not fully understand all the dominant instability mechanisms which can occur in generic combinations of evolution equations, boundary conditions, and background spacetimes. Indeed, while some theoretical progress has been made towards estimating the dominant instability mechanisms (i.e. LS), for generic “reasonable” formulations (i.e. those which we have not excluded based on other known pathologies, such as being weakly hyperbolic), we currently can only reliably determine how effective simulations will be by running those simulations. And simulations are slow.

In this paper, instead of studying the *dominant* instability, we study *an* instability (transients) which we can easily understand and rigorously describe. We use this instability to discover particularly troublesome formulations of Einstein’s equations: those which have trouble with transients.

5.1.2.3 Short-wavelength approximations

This paper makes extensive use of geometric optics, a special class of short-wavelength approximation. Several authors have applied short wavelength techniques to study the stability of various formulations of Einstein’s equations [3, 4, 6]. These techniques, however, have generally been applied to systems whose coefficients do not vary in space, limiting their validity either to very small neighborhoods of generic spacetimes, or to flat space. Previous analyzes have thus obtained only a description of local plane wave propagation: in other words, local dispersion relations. In this paper, with the geometric optics approximation, we describe how to glue these local solutions together. Such gluing is essential if we are to obtain a good approximation to a global solution of the PDE and hence a concrete, reliable estimate of the amplification of a transient. In this sense, the present paper is the logical extension of work by Shinkai and Yoneda (see, e.g., [4]), an attempt at converting their analysis to precise, specific conditions one can impose which insure that transients do not amplify.

5.1.2.4 Energy norms

This paper also employs the energy-norm techniques introduced by Lindblom and Scheel (LS) [2]. Energy norms provide a completely generic approach to determining the growth rate given a known solution and, moreover, can be used to *bound* the growth of generic solutions. While LS choose to apply these techniques to study a different class of solution—large-scale solutions whose growth presently limits their numerical simulations—these techniques remain generally applicable. We use them to characterize the growth of wave packets.

5.2 Ray optics limit of first-order symmetric hyperbolic systems

In classical electromagnetism, certain short-wavelength solutions to Maxwell's equations can be approximated by a set of ordinary differential equations for independently-propagating rays: a set of equations for the path a ray follows, and a set of equations which determine how the solution evolves along a given ray [13]. This limit is known as the ray optics (or geometric-optics) limit. In this section, we construct an analogous limit for arbitrary first-order symmetric hyperbolic linear systems.

5.2.1 Definitions

We study a specific region of four-dimensional coordinate space (t, \vec{x}) , on which at each point we have a N -dimensional (real) vector space V of “fields” $u \in V$.

Inner products: On the space of fields, an *inner product* is a map from two vectors u, v to a real number with certain properties (bilinear, symmetric, and positive-definite). The inner product is assumed to be smooth relative to the underlying four-manifold. The canonical inner product on R^N (i.e. the N -dimensional dot-product, relative to some basis of fields which is defined everywhere throughout space) is denoted (\cdot, \cdot) , and does not vary with space. We can represent any other inner product in terms of the canonical inner product and a map $S : V \rightarrow V$ as (u, Sv) , where $(u, Sv) = (Su, v)$.

An operator Q is said to be *symmetric* relative to the inner product generated by S if $(u, SQv) = (Qu, Sv)$ for all u, v . In other words, an operator Q is symmetric if it is equal to its own conjugate relative to S , denoted Q^\dagger and defined by $(u, SQv) = (Q^\dagger u, Sv)$ for all u, v . Equivalently, the conjugate Q^\dagger relative to S may be defined in terms of the transpose Q^T (i.e. the conjugate relative to $S = 1$):

$$Q^\dagger \equiv S^{-1}Q^T S . \quad (5.1)$$

First-order symmetric hyperbolic linear systems (FOSHLS): A first-order symmetric hyperbolic linear system has the form

$$[\partial_t + A^a(x, t)\partial_a - F(x, t)]u(\vec{x}, t) = 0 \quad (5.2)$$

for $u(x, t)$ a smooth function from the underlying four-manifold into the N -dimensional space of fields, for A^a and F some (generally space and time dependent²) linear operators on that space, and

²As a practical matter, we will limit attention in this paper to A^a and F varying slowly (or not at all) in time; therefore, all time dependence in the operators A^a , F , and S may usually be neglected. For completeness, however,

for A^a a symmetric operator relative to some inner product.

If more than one inner product makes A symmetric, henceforth, when talking about a specific FOSHLS, we shall fix one specific (arbitrary) inner product throughout the discussion, and therefore some specific S .

Characteristic fields and speeds: For all 3-vectors p_a , $A^a p_a$ is symmetric relative to the inner product generated by S . It has a set of eigenvalues, eigenspaces, and (for each eigenspace) basis eigenvectors,³ denoted as follows:

- $\omega_j(t, \vec{x}, \vec{p})$ are the eigenvalues of $A^a p_a$;
- $B_j(t, x, p)$, where j runs from 1 to the number of distinct eigenvalues of $A^a p_a$, are the eigenspaces of $A^a p_a$; and
- $v_{j,\alpha}(t, \vec{x}, \vec{p})$ are some orthonormal basis of eigenvectors for the space $B_j(t, x, p)$, where α runs from 1 to the dimension of B_j .

Because $A^a p_a$ is symmetric relative to the inner product induced by S , the eigenspaces are orthogonal relative to the inner product, and the eigenspaces are complete. Finally, at each point (x, p) and for each eigenspace, there is a unique projection operator $P_j(t, x, p)$ which satisfies $P_j v = v$ if $v \in B_j$, $P_j v = 0$ if $v \in B_k$ with $k \neq j$.

We require $A^a p_a$ and its eigenvalues, eigenspaces, and projection operators to vary smoothly over all x^a and p_b in the domain. [We do not demand the eigenvectors themselves to be smooth save in the neighborhood of each point (x^a, p_b) : topological constraints may prevent one from defining an eigenvector everywhere (i.e. for all p_a given x^a)⁴.]

Group velocity and acceleration: We define the group velocity V_j^a and group acceleration $a_{j,a}$ via

$$V_j^a(t, \vec{x}, \vec{p}) \equiv \frac{\partial}{\partial p_a} \omega_j(t, \vec{x}, \vec{p}) , \quad (5.3)$$

$$a_{j,a}(t, \vec{x}, \vec{p}) \equiv -\frac{\partial}{\partial x^a} \omega_j(t, \vec{x}, \vec{p}) . \quad (5.4)$$

We shall make frequent use of an alternative expression for the group velocity, Eq. (C.2), which is discussed in Appendix C.1. Among other things, Eq. (C.2) implies

$$\omega_j(x, p) = V_j^a p_a .$$

we retain time dependence for readers who may wish to apply these techniques to more generic systems.

³Any real symmetric matrix can be diagonalized, with orthogonal eigenspaces.

⁴For example, in the first-order representation of the scalar wave equation, two of the eigenvectors at each point (x, p) are essentially vectors transverse to the surface $|p|$. These cannot be extended over the sphere.

5.2.2 Form of ray-optics solution

We now construct a solution which approximately satisfies Eq. (5.2). Our method works by constructing a set of characteristics (i.e. rays), then integrating some amplitude equations along each characteristic (as an ODE) to find the amplitudes farther along the ray.

In this section, we only introduce the results of our analysis. In Appendix C.2, we provide a more comprehensive justification of our ray-optics approach.

Ray-optics solution

Rather than express our solution in terms of the original N -dimensional variable u , we introduce $N + 1$ new variables $d_{j,\alpha}$ and ϕ and parametrize the original state by

$$u = \bar{u}e^{i\phi} \quad (5.5)$$

where we further expand \bar{u} in terms of the eigenvectors $v_{l,\alpha}$ of $A^a\partial_a\phi$ at each point (t, \vec{x}) :

$$\bar{u} = \sum_l \sum_\alpha d_{l,\alpha}(t, x, \partial\phi)v_{l,\alpha}(t, x, \partial\phi) . \quad (5.6)$$

[For notational clarity, the arguments t , \vec{x} , and $\partial_a\phi$ to the functions ϕ , $v_{l,\alpha}$, and $d_{l,\alpha}$ will in the following be usually omitted.]

In terms of these new variables, a ray-optics solution is a solution to the following equations, for some fixed j :

$$0 = [\partial_t + V_j^a(x, \partial\phi)\partial_a] \phi , \quad (5.7a)$$

$$0 = d_{l,\beta} \quad \text{for } l \neq j , \text{ and} \quad (5.7b)$$

$$0 = [\partial_t + V_j^a\partial_a] d_{j,\alpha} + \sum_\beta d_{j,\beta} \left(v_{j,\alpha}, S(\partial_t + A^a\partial_a - F)v_{j,\beta} \right) . \quad (5.7c)$$

When we substitute solutions to the ray-optics equations [Eq. (5.7)] back into the original FOSHLS [Eq. (5.2)], as described in detail in Appendix C.2, we find the geometric-optics solutions are excellent approximate solutions to the original PDE, so long as certain mild conditions continue to hold [e.g., the oscillations in ϕ remain rapid compared to any other length or time scale].

5.2.3 Interpreting the geometric-optics equations

We introduce the geometric-optics solution precisely because it simplifies the PDE—in particular, because it converts the problem of solving a general PDE [Eq. (5.2)] into the problem of solving

uncoupled ODEs [Eq. (5.7)]. Specifically, these ODEs consist of the the phase equation [Eq. (5.7a)]—which determines the path of the ray leaving a point \vec{x} consistent with initial data for ϕ with gradient $\partial_a\phi(\vec{x})$ —and the polarization equations [Eqs. (5.7b) and (5.7c)]—which allow us to propagate the $d_{l,\alpha}$ along each ray.

But while these equations are now ODEs, their structure is not particularly transparent. In this section, we rewrite the phase equation [Eq. (5.7a)] and the polarization equation [Eq. (5.7c)] to better emphasize their properties and physical interpretation.

5.2.3.1 Path of the ray

The physical significance of the phase equation [Eq. (5.7a)] becomes much easier to appreciate when it is rewritten in first-order form. When we differentiate that expression and re-express the result as an equation for $k_a \equiv \partial_a\phi$, we find

$$\begin{aligned} 0 &= \partial_t k_a + V_j^b \partial_b k_a + [\partial_a V_j^b(x, k)] k_b \\ &= \partial_t k_a + V_j^b \partial_b k_a - a_{j,a}(x, k) . \end{aligned} \quad (5.8)$$

[While k does depend on x , because $(\partial_{k_c} V_j^b) k_b = 0$ the last term in the first line does indeed simplify into $-a_{j,a}$, as stated]. Solutions to this PDE may be constructed by gluing together solutions to the following pair of coupled ODEs for $\vec{x}(t)$ and $\vec{k}(t)$:

$$\frac{dx^a}{dt} = V_j^a(\vec{x}, \vec{k}) \quad (5.9a)$$

$$\frac{dk_a}{dt} = a_{a,j}(\vec{x}, \vec{k}) . \quad (5.9b)$$

By using the definitions of V_j^a and $a_{a,j}$, we find these are precisely Hamilton's equations, using $\omega_j(t, x, k)$ as the Hamiltonian.

These two equations define the rays (i.e. characteristics). Given initial data for k_a which has $k_a = \partial_a\phi$ in a 3-dimensional neighborhood of a point, we have a unique ray emanating from each point in that neighborhood. Solutions to Eq. (5.8) follow from joining the resulting rays emanating from each point in the neighborhood together; and solutions for ϕ [i.e. Eq. (5.7a)] follow by integrating the phase out along each ray.

5.2.3.2 Propagating polarization along ray

In practice, the polarization equation [Eq. (5.7c)] is difficult to interpret: since it involves spatial derivatives of basis vectors, and since we have freedom to choose our basis vectors $v_{j,\alpha}$ arbitrarily within each subspace B_j , we cannot transparently disentangle meaningful terms from convention-induced effects.

To constrain the basis and simplify the equation, we sometimes choose a basis in the neighborhood of the ray of interest which satisfies the *no-rotation condition* [discussed at greater length in Appendix C.1.2]:

$$(v_{j,[\alpha}, S(\partial_t + A^a \partial_a) v_{j,\beta]}) = 0. \quad (5.10)$$

where the square brackets denote antisymmetrization over α and β [i.e. $X_{[\alpha\beta]} = (X_{\alpha\beta} - X_{\beta\alpha})/2$]. The no-rotation condition completely constrains the antisymmetric part of an operator [i.e. the left side of Eq. (5.10)]; the condition that the basis vectors $v_{j,\alpha}$ remain orthogonal constrains that operator's symmetric part; and therefore the basis $v_{j,\alpha}$ is necessarily completely specified at any point along a ray in terms of initial data for the basis.

Using the no-rotation condition, we find the polarization equation becomes the less-arbitrary expression [Appendix C.1.3]:

$$\begin{aligned} 0 = & \left(\partial_t + V_j^a \partial_a + \frac{1}{2} \partial_a V_j^a(x, \partial\phi) \right) d_{j,\alpha} \\ & - \sum_{\beta} \left(v_{j,\alpha}, \left[SF + \frac{1}{2} \partial_t S + \frac{1}{2} \partial_a (SA^a) \right] v_{j,\beta} \right) d_{j,\beta} \end{aligned} \quad (5.11)$$

where in the above $v_{j,\alpha}$ is a no-rotation basis. In Sec. 5.3 we will use this expression to motivate the definition of prototypical coherent wave packets, which have an exceedingly simple growth rate.

5.2.4 When do geometric-optics solutions exist?

Given initial data (say, for k_a and $d_{j,\alpha}$ on some initial compact region), we can in practice always find a solution to the geometric-optics equations [Eq. (5.7)] valid for some small interval δt (i.e. by using general PDE existence theorems, like Cauchy-Kowaleski). However, for general initial data we cannot solve the phase equation [Eq. (5.7a)] for an arbitrary time T . By way of example, even if we find each individual ray [i.e. each solution to Eq. (5.9) emanating from each initial-data point] emanating from our initial data region out to time T , these rays may cross before time T , rendering the geometric-optics solution for $d_{j,\alpha}$ both singular and inconsistent at the ray-crossing point. (A similar problem arises in classical geometric optics.) Furthermore, depending on the structure of A^a , certain rays may not even admit extension to time T (i.e. certain rays may be future-inextendible, precisely like rays striking singularities in GR).

A proper treatment of these technical complications is considerably beyond the scope of this paper. In practice, we will assume we have chosen initial data so that our geometric-optics solution can be evolved to any time T , unless it involves transport into a manifest singularity (i.e. a singularity of the spacetime used to generate the FOSHLS) before time T . Furthermore, we will assume the solution is *well-behaved*—that is, the congruence has finite values for k_a , V_j^a , $a_{a,j}$ and their first derivatives. With a well-behaved solution to the phase equation [Eq. (5.7a)], we may always find a

finite, consistent solution to the polarization equation [Eq. (5.7c)] in terms of the initial data.⁵

5.3 Defining wave packets

In Sec. 5.2, we have constructed approximate solutions to linearized first-order symmetric hyperbolic PDEs in the geometric optics limit. These solutions are constructed by integrating ODEs for (and along) rays [Eq. 5.7]. Since each ray evolves independently, we are naturally led to consider *wave packets*—that is, ray-optics solutions which are nonzero only in a (four-dimensional) neighborhood of some (four-space) ray.

In this section, we outline how wave packets may be generally constructed. We also describe the two special classes of wave packets, coherent wave packets and prototypical coherent wave packets (PCWP), which will be the focus of discussion henceforth.

5.3.1 Constructing wave packets

A wave packet that persists for a time T is some solution to the geometric-optics equations [Eq. (5.7)] which is nonzero only in some small neighborhood of a ray (i.e. nonzero only within some coordinate length δ from the central ray).

From a constructive standpoint, while we can easily construct solutions from initial data for k_a and $d_{j\alpha}$, we have no transparent way, besides solving the equations themselves, to determine whether a particular set of initial data for k_a even generates a congruence which exists and remains well-behaved (e.g., $\partial^a V_b$ and $\partial_a k_b$ both finite) for time T , let alone whether the specific combination of initial data for k_a and $d_{j\alpha}$ yields a geometric-optics solution with support only within a given distance δ from a ray.

Still, physically we *expect* we can avoid these technical complications. For example, we *expect* that, for all rays of physical interest, we can extend the central ray of interest to time T (i.e. characteristics of physical interest can be extended as long as physically necessary). We *expect* that singular congruences k_a can be avoided by proper choice of initial k_a data (e.g., the ray equations do not require all congruences near the ray of interest to diverge or come to a focus). And given a well-behaved congruence, we expect we can always choose initial data for $d_{j\alpha}$ in a sufficiently small neighborhood so the solution for $d_{j,\alpha}$ is nonzero only within some fixed distance δ from the central ray.

Thus, as a proper treatment of these technical complications is considerably beyond the scope of this paper, we shall henceforth simply assume that a wave packet solution can always be constructed about any ray of physical interest.

⁵Since the polarization equation [expressed as Eq. (5.7c) or as Eq. (5.11)] is linear in the polarization fields $d_{j\alpha}$, it therefore admits well-behaved solutions for the evolution of $d_{j,\alpha}$ along a well-behaved ray so long as the linear operators present in that equation are well-behaved.

5.3.2 Specialized wave packets I: Coherent wave packets

Since rays propagate independently, one can choose arbitrary initial data, and in particular arbitrary polarization directions w , and still obtain a wave-packet solution. Here, w is defined by

$$w \equiv \bar{u}/|\bar{u}| \quad |\bar{u}| \equiv [(\bar{u}, S\bar{u})]^{1/2}. \quad (5.12)$$

We prefer to further restrict attention to those wave packets which have a single, dominant polarization direction w present initially (and therefore for all time). In other words, we require w vary slowly across the wave packet's spatial extent. Wave packets with this property we denote *coherent wave packets*.

5.3.3 Specialized wave packets II: Prototypical coherent wave packets (PCWPs)

While coherent wave packets have a simple polarization structure, characterized by some polarization direction w , this polarization structure need not necessarily have a transparent relationship to the terms present in the polarization equation [Eq. (5.7c); or equivalently Eq. (5.11) if we use a no-rotation basis]. Therefore, we define *prototypical coherent wave packets* (PCWPs) as wave packets which have at each time their polarization direction w equal to one of the eigenvectors $f_j^{(\mu)}$ of the operator O_j :

$$O_j \equiv P_j \left\{ F + \frac{1}{2} S^{-1} [\partial_t S + \partial_a (S A^a)] \right\} P_j \quad (5.13)$$

$$O_j f_j^{(\mu)} \equiv o_{j\mu} f_j^{(\mu)} \quad (5.14)$$

where μ , running from 1 to the dimension of B_j , indexes the eigenvectors of O_j . For simplicity, we assume O_j has a complete set of eigenvectors.⁶

If PCWPs exist, we expect—because of their relationship to the terms of the polarization equation [Eq. (5.11)]—the propagation of their polarization will be much easier to understand. Most notably, as we will show in the next section [Sec. 5.4], prototypical coherent wave packets have particularly simple expressions for their growth rates [i.e. Eq. (5.20)].

PCWPs will exist as exact solutions to the polarization equation [Eq. (5.11)] only in certain special circumstances; for example, most of the polarizations to be discussed in Sections 5.5 and 5.6 admit exact PCWP solutions. However, as demonstrated in more detail in Appendix C.3, we do not

⁶The behavior of the polarization equation when O_j has Jordan blocks is straightforward (i.e. we converge to some specific eigenvector in the Jordan block; we obtain no change to the final predictions for exponential growth rates; we only add at most a polynomial in t to the amplitude functions) but tedious to describe in detail. Moreover, in all physically interesting cases we have examined, Jordan blocks have not appeared in O_j ; we have been able to choose a complete set of basis eigenvectors.

expect the polarization equation to generically admit PCWP solutions.

Nonetheless, as discussed in greater detail in Appendix C.3, a PCWP with $w = f_j^{(\nu)}$ is a good approximate solution to the polarization equation when the eigenvalue $o_{j\nu}$ of O_j is sufficiently large. Indeed, by rewriting the polarization equation in the basis $f_j^{(\mu)}$, we can show *generic* coherent wave packets will rapidly converge to a PCWP with $w = f_s^{(\nu_o)}$ for ν_o indexing the eigenvalue of O_j with largest real part. In other words, based on Eq. (5.20), when coherent wave packets grow quickly, they can always be well-described by a PCWP.

5.4 Describing and bounding the growth rate of wave packets

Since a wave packet is narrow and we care little about its precise spatial extent, we commonly characterize the wave packet by a single number (e.g., a peak amplitude) rather than a generic distribution of polarization over space. Unfortunately, the maximum value of the amplitudes $d_{j,\alpha}$ depend on the spatial extent of the wave packet—in other words, it depends on our choice of congruence, rather than the central ray itself.

Because the amplitude function is subject to focusing effects (through the term $\partial_a V^a$), we choose to describe the magnitude of the wave packet by the magnitude of its energy norm. Introduced by Lindblom and Scheel (LS), the energy norm is an integral quantity analogous to energy [2]; and, like the energy of a wave packet solution to Maxwell’s equations, the energy norm will not be susceptible to focusing effects.

In this section, we describe how energy norms can be used to characterize the magnitude of wave packets. We also obtain special expressions for the growth rates of coherent wave packets [Eq. (5.19)] generally and prototypical coherent wave packets [Eq. (5.20)] in particular.

Also, for completeness, in Appendix C.4 we provide an explicit, rigorous bound for the growth rate of energy which will not be otherwise used in the paper.

5.4.1 Energy norms and the magnitude of geometric-optics solutions

Lindblom and Scheel define the energy norm by way of two quadratic functionals of a solution u [LS Eqs. (2.3) and (2.8)]. When expressed in terms of our notation, these functionals are

$$\epsilon \equiv (u^*, Su) \quad E \equiv \int \mu d^3x \epsilon. \quad (5.15)$$

Unlike LS, we do not generically have a preferred spatial metric; we therefore replace the factor \sqrt{g} present in LS Eq. (2.8) by the more generic μ .⁷

⁷ Unlike LS, we are not necessarily working with a metric space; therefore, we have no preferred measure on the coordinate space and therefore allow for an arbitrary, as-yet-undetermined measure factor μ .

We may substitute in the expressions appropriate to a ray-optics solution to obtain excellent approximations to the energy. By way of example, the energy E_j of a geometric-optics solution propagating in the j th polarization may be expressed as

$$\begin{aligned} E_j &\approx \int \mu d^3x \sum_{\alpha,\beta} d_{j,\alpha}^* d_{j,\beta} (v_{j,\alpha}^*, S v_{j,\beta}) \\ &= \int \mu d^3x \sum_{\alpha} |d_{j,\alpha}|^2 \end{aligned} \quad (5.16)$$

where the terms neglected are small in the geometric optics limit and where the second line holds because by construction the basis $v_{j,\alpha}$ is orthonormal.

5.4.2 Energy norms and the growth rate of wave packets

Following the techniques of Lindblom and Scheel, we can use energy norms and conservation-law techniques to obtain a general expression for the growth rate of a wave packet.

To follow their program, we must generate a conservation law. Define, therefore, an energy current j^a [i.e. LS Eq. (2.4)]

$$j^a \equiv (u^*, SA^a u).$$

The quantities ϵ and j^a obey the conservation-law-form equation

$$\begin{aligned} \partial_t \epsilon + \mu^{-1} \partial_a (\mu j^a) &= (u^*, SFu) + (Fu^*, Su) \\ &+ (u, [\partial_t S + \mu^{-1} \partial_a (\mu SA^a)] u) \end{aligned}$$

[i.e. the analogue of LS Eqs. (2.5) and (2.6)].

For a wave-packet solution, which is concentrated at each time to a small spatial region, the current j^a drops to zero rapidly, and is in particular zero at the manifold boundary. As a result, when we integrate the conservation law, we find the energy obeys the equation

$$\frac{dE}{dt} = \int \mu d^3x (u^*, SQ u) \quad (5.17a)$$

$$Q \equiv F + S^{-1} [F^T S + \partial_t S + \mu^{-1} \partial_a (\mu SA^a)] \quad (5.17b)$$

where F^T is defined so $(u, Fv) = (F^T u, v)$ for all u, v (i.e. F^T is the transpose). [In LS, the analogous equations are (2.7) and (2.9); in our case, however, we have no surface term involving j_a because the solution falls off rapidly away from the wave packet.]

We can show Q is symmetric relative to S .⁸ We can also show that Q is closely related to the

⁸ Because S and SA^a are symmetric relative to the canonical inner product, so are their derivatives. And if T is

symmetric part of the operator O_j [Eq. (5.13)]:

$$P_j Q P_j = O_j + O_j^\dagger + \frac{\partial_a \mu}{\mu} P_j V_j^a . \quad (5.18)$$

5.4.3 Energy norms and the growth rate of coherent wave packets

Since coherent wave packets are both localized and possess a well-defined polarization direction w , we find Eq. (5.17) becomes, for coherent wave packets,

$$\frac{1}{E} \frac{dE}{dt} \approx (w^*, S Q w) \quad (5.19)$$

where the right side is evaluated at the location of the wave packet at the current instant.

Because we still need the appropriate polarization direction w to make use of the above expression—a direction we can only obtain from the polarization equation [Eq. (5.11)]—Eq. (5.19) provides only an alternate perspective on the growth of wave packets, not an entirely independent approach to the evolution of the amplitude.

5.4.4 Energy norms and the growth rate of PCWPs

In the special case of a PCWP, however, we do know the polarization direction w : it is one of the normalized eigenvectors $f_j^{(\mu)}$ of the operator O_j [see Sec. 5.3.3]. In this case, we find the energy growth rate for a PCWP with $w = f_j^{(\mu)}$ to be

$$\frac{1}{E_{j\mu}} \frac{dE_{j\mu}}{dt} = o_{j\mu} + o_{j\mu}^* + \frac{\partial_a \mu}{\mu} V_j^a . \quad (5.20)$$

[Here, we used Eq. (5.18) in Eq. (5.19).]

5.5 Geometric optics limit of KST: Rindler

In the previous sections (Sections 5.2, 5.3, and 5.4), we have developed a procedure for computing the evolution and amplification of ray-optics solutions in general and prototypical coherent wave packet solutions in particular. To provide a specific demonstration of these methods, we demonstrate how to construct the geometric optics limit (as described in Section 5.2) and compute the growth rate of wave packets (as described in Sections 5.3 and 5.4) when the first-order hyperbolic system is the 2-parameter first-order symmetric hyperbolic system Kidder, Scheel, and Teukolsky introduced (see their Section II J), linearized about a flat-space background in Rindler coordinates.

symmetric relative to the canonical inner product, then $S^{-1}T$ is symmetric relative to the inner product generated by S .

Our computations in this section proceed as follows. First, we review Rindler coordinates and the effects of using Rindler coordinates as the background in the linearized KST equations. We then describe the limited set of rays we will study (i.e. rays that propagate only in the x direction). Subsequently, we construct the explicit form of the polarization equation [Eq. (5.7c)] for packets that propagate only in x . [The analysis simplifies substantially because the basis vectors used do not vary with x ; therefore, the derivatives present in Eq. (5.7c) disappear.] The analysis of the polarization equation leads us directly to an explicit expression for the growth of energy of a coherent wave packet [Eq. (5.17)] in general and a prototypical coherent wave packet in particular [Eq. (5.20)].

Finally, to verify our expressions give an accurate description of the growth of PCWPs, we compare them against the results of numerical simulations.

5.5.1 Generating the FOSHLS using the background Rindler space

Flat space in Rindler coordinates is characterized by the metric

$$ds^2 = -x^2 dt^2 + (dx^2 + dy^2 + dz^2), \quad (5.21)$$

for $x > 0$. Using this spacetime as a background, we can linearize the KST 2-parameter formulation to generate a FOSHLS of the form of Eq. (5.2)—and in particular find explicit forms for the operators A^a and F . For example, we find that the principal part has the form [KST Eq. (2.59), along with the definition of $\hat{\partial}_o$ in KST Eq. (2.10)]:

$$\partial_t \delta g_{ij} \simeq 0 \quad (5.22a)$$

$$\partial_t \delta P_{ij} + x g^{ab} \partial_a \delta M_{bij} \simeq 0 \quad (5.22b)$$

$$\partial_t \delta M_{kij} + x \partial_k \delta P_{ij} \simeq 0 \quad (5.22c)$$

As the right-hand sides of these equations are very long, we shall not provide them, or an explicit form for F , in this paper. The right hand side depends on the two continuous KST parameters, \hat{z} and γ [1].

Using the FOSHLS obtained by linearizing, we can proceed generally with any linear analysis, including a construction of the geometric-optics limit.

5.5.2 Describing local plane waves by diagonalizing $A^a \hat{x}_a$

The geometric-optics limit is a short-wavelength limit. Naturally, then, the first step towards the geometric-optics limit is understanding the plane-wave solutions in the neighborhood of a point. We find these solutions by substituting into Eq. (5.2) the form $u \propto u_o \exp i(k \cdot x - \omega t)$; assuming k and ω are large, so we may disregard the right side; assuming both u_o and A^a are locally constant; and then

solving for u_o and the relationship between k_a and ω . In other words, we find those local-plane-wave solutions by diagonalizing $A^a k_a$, as discussed in Sec. 5.2, to find eigenvalues ω_j and eigenvectors $v_{j,\alpha}$, where j indexes the resulting eigenvalues and α indexes the degenerate eigenvectors for each j .

Because the principal part is both simple and independent of the two KST parameters (\hat{z} and γ), we can diagonalize it by inspection. For every propagation direction, the eigenvalues are precisely $\omega_s(x, k) = s|k|$ for $s = \pm 1, 0$. For our purposes, we study only propagation in the x direction. Thus, we need only the eigenfields of $A^a \hat{x}_a$, which are [see KST Eq. (2.61) and also Appendix C.5.1.3]

$$U_{ab}^g = g_{ab} \quad (5.23a)$$

$$U_{y,ab}^0 = M_{yab} \quad (5.23b)$$

$$U_{z,ab}^0 = M_{zab} \quad (5.23c)$$

$$U_{ab}^\pm = \frac{1}{\sqrt{2}} (P_{ab} \pm M_{xab}) \quad (5.23d)$$

These expressions may be interpreted as equivalent to the basis vectors $v_{j,\alpha}$, as discussed in Appendix C.5 [see Appendix C.5.1.3, and in particular Eq. (C.25)].

5.5.3 Deriving the polarization and energy equations, for propagation in the x direction on the light cone

In this section, we describe how to construct and analyze the polarization equation [Eq. (5.7c)] and energy equation [Eq. (5.17)] for wave packets propagating in the x direction. For technical convenience, we limit attention to rays which propagate on the light cone—in other words, which travel on one of the two null curves of the metric:

$$dx/dt = sx$$

for $s = \pm 1$. [In terms of the above representation of the eigenspaces of $A^a \hat{x}_a$ discussed above, only the fields U^\pm , given in Eq. (??), propagate on the light cone.]

5.5.3.1 Essential tool: Diagonalizing $P_s F P_s$

We have the polarization equation [Eq. (5.7c)] and a basis [Eq. (5.23), or equivalently Eq. (C.25)]; the application is straightforward. We can, however, substantially simplify our expression by changing the basis used to expand \bar{u} from $v_{j,\alpha}$ to the basis of eigenvectors $f_s^{(\mu)}$ of $P_s F P_s$, defined by the normalized solutions to

$$F f_s^{(\mu)} = \zeta_{s,\mu} f_s^{(\mu)} .$$

[Equivalently, we may define these eigenvectors in component fashion. For each s , the matrix $(v_{s,\alpha}, Fv_{s,\beta})$ admits a complete set of normalized eigenvectors $f_{s,\alpha}^{(\mu)}$:

$$\sum_{\beta} (v_{s,\alpha}, Fv_{s,\beta}) f_{s,\beta}^{(\mu)} = \zeta_{s,\mu} f_{s,\alpha}^{(\mu)} .$$

Using these eigenvectors, we regenerate $f_s^{(\mu)} = \sum_{\alpha} f_{s,\alpha}^{(\mu)} v_{s,\alpha}$, which are eigenvalues of $P_s F P_s$.]

These eigenvectors may be classified according to their symmetry properties under rotations about the propagation axis x :

- *Symmetric-traceless-transverse 2-tensor* [basis vectors correspond to the fields U_{yz}^s and $(U_{yy}^s - U_{zz}^s)/\sqrt{2}$] One subspace corresponds to the 2-dimensional space of symmetric-traceless 2-dimensional tensors transverse to the propagation direction. The operator $P_s F P_s$ is degenerate in this subspace; the single eigenvalue associated with this subspace is given by $\zeta_{s,t}$, defined by

$$\zeta_{s,t} = -s \tag{5.24a}$$

- *Transverse 2-vector* [basis vectors correspond to the fields U_{xz}^s and U_{xy}^s] Another subspace corresponds to the 2-dimensional space of 2-dimensional vectors transverse to the propagation direction. Again, the operator is degenerate on this space. The eigenvalue of F in this subspace is given by $\zeta_{s,v}$ for

$$\zeta_{s,v} = -s \frac{1 + \gamma}{-1 + 2\gamma} \tag{5.24b}$$

- *2-scalars* [spanned by vectors corresponding to the fields U_{xx}^s and $(U_{yy}^s + U_{zz}^s)$] Finally, the 2-dimensional space of rotational 2-scalars has its degeneracy broken by F . For each s , we find two eigenvalues, denoted $\zeta_{s,s1}$ and $\zeta_{s,s2}$, with values

$$\zeta_{s,s1} = -s \tag{5.24c}$$

$$\zeta_{s,s2} = -s \frac{1 + 2\gamma^2}{-1 + 2\gamma} \tag{5.24d}$$

These eigenvectors f are linearly independent. Indeed, symmetry guarantees that—with the exception of the two 2-scalar eigenvectors - most of the eigenvectors are mutually orthogonal.

5.5.3.2 Polarization equation for general geometric-optics solutions

We can apply these eigenvectors to rewrite the polarization equation [Eq. (5.7c)] using the basis $f_s^{(\mu)}$. Specifically, we define $D_{j\mu}$ by the expansion $d_{s,\alpha} = \sum_{\mu} D_{s\mu} f_{s,\alpha}^{(\mu)}$. Noting our basis vectors $f_s^{(\mu)}$

are independent of space and time, we find a set of independent equations for the $D_{s\mu}$ of the form

$$(\partial_t + sx\partial_x)D_{s\mu} = \zeta_{s\mu}D_{s\mu} . \quad (5.25)$$

This equation, along with the explicit forms for the basis vectors $f_s^{(\mu)}$, tells us how to evolve arbitrary polarization initial data along our congruence.

5.5.3.3 Energy equation for general geometric-optics solutions

Similarly, we may rewrite expressions for the energy E [Eq. (5.15), or Eq. (5.16)] and growth rate $E^{-1}dE/dt$ [Eq. (5.17)] using the basis $f_s^{(\mu)}$. For example, we define energy of the wave packet by Eq. (5.16), using a measure $\mu = \sqrt{g} = 1$ consistent with the flat spatial metric of the background. We find, using symmetry properties of the eigenvectors to simplify the sum,

$$\begin{aligned} E &= \int d^3x \sum_{\mu \in \{t,v\}} |D_{s\mu}|^2 \\ &+ \int d^3x \, 2\text{Re} \left[D_{s,s1}^* D_{s,s2} \left(f_s^{(s1)*}, S f_s^{(s2)} \right) \right] \end{aligned} \quad (5.26)$$

The growth rate of energy $E^{-1}dE/dt$ can be obtained in two ways:

1. First, we can explicitly differentiate Eq. (5.26), using Eq. (5.25) to simplify as necessary.
2. Alternatively, we can employ the general expression for the growth rate of geometric-optics solutions [Eq. (5.19)]. [To do so, we express Q in terms of O_s via Eq. (5.18). Then we find the following explicit expression for O_j by using Eq. (C.28) from Appendix C.5, which in this case tells us

$$P_s S^{-1} [\partial_t S + \partial_a (S A^a)] P_s = s P_s \quad (5.27)$$

when we rewrite the results of that expression in an operator, rather than component, notation. Finally, we employ the basis $f_s^{(\mu)}$. Because of Eq. (5.27), we know the eigenvectors $f_s^{(\mu)}$ of $P_s F P_s$ are equivalently eigenvectors of O_s .]

In either case, one concludes

$$\begin{aligned} \frac{dE}{dt} &= \int d^3x \sum_{\mu} |D_{s\mu}|^2 [2\text{Re}(\zeta_{s\mu}) + s] \\ &+ \int d^3x \, 2\text{Re} \left[D_{s,s1}^* D_{s,s2} \left(f_s^{(s1)*}, S f_s^{(s2)} \right) \right. \\ &\quad \left. \times (\zeta_{s,s1}^* + \zeta_{s,s2} + s) \right] . \end{aligned} \quad (5.28)$$

The above equations remain completely generic and apply to all ray-optics solutions that propagate along the congruence $dx/dt = sx$.

5.5.3.4 Energy equation in a special case: PCWPs

As Eq. (5.25) demonstrates, the polarizations do not change direction as they propagate. In other words, if a wave packet initially has only $D_{s\mu} \neq 0$ for some specific pair of (s, μ) , then the wave packet will always have $D_{s\mu} \neq 0$ only for that s and μ . Moreover, as noted in the discussion surrounding Eq. (5.27), the basis vectors $f_s^{(\mu)}$ used to define the $D_{s\mu}$ are eigenvectors of O_s . Following the discussion of Sec. 5.3.3, we call such a solution a prototypical coherent wave packet.

For a wave packet solution which is confined to the $(s\mu)$ polarization, we need only one term in each sum to find the energy $E_{s\mu}$ and growth rate $E_{s\mu}^{-1} dE_{s\mu}/dt$:

$$E_{s\mu} = \int d^3x |D_{s\mu}|^2 \quad (5.29a)$$

$$\frac{1}{E_{s\mu}} \frac{dE_{s\mu}}{dt} = 2\text{Re}(\zeta_{s\mu}) + s. \quad (5.29b)$$

[The above expression was obtained directly from Eq. (5.28). Equivalently, we can obtain the same result using Eq. (5.20) by way of Eq. (5.27).]

To be very explicit, we find using Eq. (5.24) the growth rates of the tensor (t) and one of the scalar ($s1$) polarizations to be constant, independent of γ but depending on which direction the packet propagates ($s = \pm 1$):

$$\frac{1}{E_{s,t}} \frac{dE_{s,t}}{dt} = \frac{1}{E_{s,s1}} \frac{dE_{s,s1}}{dt} = -s \quad (5.30a)$$

We also find the vector (v) and remaining scalar ($s2$) polarizations have a growth rate which varies with γ , according to

$$\frac{1}{E_{s,v}} \frac{dE_{s,v}}{dt} = -s \left(2 \frac{1+\gamma}{-1+2\gamma} - 1 \right) \quad (5.30b)$$

$$\frac{1}{E_{s,s2}} \frac{dE_{s,s2}}{dt} = -s \left(2 \frac{1+2\gamma^2}{-1+2\gamma} - 1 \right) \quad (5.30c)$$

5.5.4 Comparing growth rate expressions to simulations of prototypical coherent wave pulses

In Eq. (5.30) we tabulated the expected growth rates of energy for each possible coherent wave packet. To demonstrate that these expressions are indeed correct, we compare these predicted growth rates with the results of numerical simulations of wave packets propagating on a Rindler background.

5.5.4.1 Specific simulations we ran

To test the validity of our expressions, we used a 1D variant of the KST pseudospectral code kindly provided by Mark Scheel. He developed this code to study the linearized KST equations on a Rindler background (e.g., to produce the results shown in Lindblom and Scheel Sec. IV A [2]).

We ran this code at a fixed, high resolution (512 collocation points in the x direction) on a computational domain $x \in [0.01, 1]$ with various wave-packet initial data. Specifically, we used a wave packet profile proportional to

$$W(x) = A \cos(2\pi x/\lambda) \exp[-(x - x_c)^2/\sigma^2] \quad (5.31)$$

with $A = 10^{-5}$, $x_c = 0.55$, $\sigma = 0.1$, and $\lambda = 0.01$. The precise initial data used depended on the polarization we wanted:

- *Tensor* When we wanted a tensor polarization, we used initial data for a single left-propagating 2-tensor component: $U_{xy}^- = W$, with all other characteristic fields zero. In other words, we used initial data $P_{yz} = M_{xyz} = W(x)/2$ with all other fields zero.
- *Vector* When we wanted a vector polarization, we used initial data for a single left-propagating 2-vector component: $U_{xz}^- = W$, with all other characteristic fields zero. In other words, we used initial data $P_{xz} = M_{xxz} = W/2$ with all other components zero.
- *Scalar 1 (s1)* When we wanted to excite the left-propagating s1 polarization, we used initial data $P_{xx} = M_{xxx} = W/2$.
- *Scalar 2 (s2)* After some algebra, one can demonstrate that to excite the s2 polarization, we should use initial data $P_{yy} = M_{xyy} = W/4$ and $P_{yy} = M_{xyy} = -W/4$.

To avoid the influences of boundaries, we only studied the results of the simulations out to a time $t \sim 0.1$.

5.5.4.2 Results

For each polarization (t , v , $s1$, and $s2$), we found that wave packets remained in the initial polarization, with little contamination from other fields. For example, when exciting the tensor polarization, we found all fields other than U_{xy} remained small.

The wave packets' energy grew exponentially, with growth rates that agreed excellently with Eq. (5.30). For example, the polarizations $s1$ and t both had growth rates consistent with unity to a part in a thousand. Our expressions for the growth rates for $s2$ and v also agreed well with the results of numerical simulations, as shown in Fig. 5.1 for left-propagating pulses ($s = -1$).

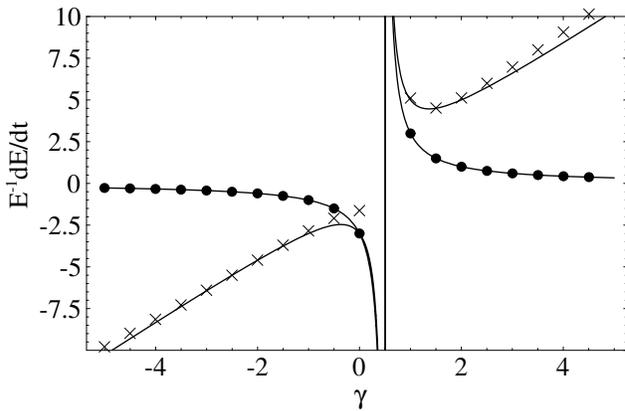


Figure 5.1: The two solid curves show the theoretically predicted growth rates for the “vector” [v , Eq. (5.30b)] and one of the scalar [$s2$, Eq. (5.30c)] polarizations, when those polarizations propagate to the left ($s = -1$). The circles show the results for numerical simulations of the vector wave packet; the crosses show the results for wave packets in the $s2$ polarization. Both predictions agree very well with simulations.

5.6 Geometric optics limit of KST: PG

In this section, we study another example of the geometric optics formalism: the propagation of radially-propagating wave packets evolving according to the KST 2-parameter formulation of evolution equations, linearized about a Painleve-Gullstrand background.

Our analysis follows the same course as the Rindler case addressed in Sec. 5.5. We first review Painleve-Gullstrand coordinates and the effects of using these coordinates as the background in the linearized KST equations. Subsequently, we construct the explicit form of the polarization and energy equations [Eqs. (5.11) and (5.17)] for packets that propagate radially on the light cone. Finally, in a departure from the Rindler pattern, we also add an analysis of the “zero-speed” modes that propagate against the shift vector.

5.6.1 Generating the FOSHLS using a background Painleve-Gullstrand space

A Schwarzschild hole in Painleve-Gullstrand coordinates is characterized by the metric

$$ds^2 = -dt^2 + \left(dr + \sqrt{\frac{2}{r}} dt \right)^2 + r^2 d\Omega^2 . \quad (5.32)$$

We shall use this metric in cartesian spatial coordinates [i.e. $z = r \cos \theta$, $x = r \sin \theta \cos \phi$, $y = r \sin \theta \sin \phi$] as the background spacetime in the KST equations. Linearizing about this background, we obtain the explicit FOSHLS we study in the remainder of this section.

As before, we shall not provide the very complicated derivative-free terms (i.e. F) explicitly in

this paper. The principal part, however, remains simple by design; in this case, we have [KST Eq. (2.59), along with the definition of $\hat{\partial}_o$ in KST Eq. (2.10)]:

$$(\partial_t - \beta^a \partial_a) g_{ij} \simeq 0 \quad (5.33a)$$

$$(\partial_t - \beta^a \partial_a) P_{ij} + g^{ab} \partial_a M_{bij} \simeq 0 \quad (5.33b)$$

$$(\partial_t - \beta^a \partial_a) M_{kij} + \partial_k P_{ij} \simeq 0 \quad (5.33c)$$

with $\beta^a = \sqrt{2/r} \hat{r}^a$.

5.6.2 Local plane waves and diagonalizing $A^a \hat{r}_a$

As discussed generally in Sec. 5.2 and by way of a Rindler example in Sec. 5.5.2, to understand how wave packets propagate radially we must first understand how local plane waves propagate radially, which in turn requires we diagonalize $A^a \hat{r}_a$. The basis vectors and eigenvalues are addressed in detail and in a more general setting in Appendix C.5.1.3. In brief, the eigenvalues are $\omega_s(x, k) = s|k| - \beta^a k_a$ with $s = \pm 1, 0$ and the eigenvectors correspond directly to the Rindler results [i.e. Eq. (5.23), with $x \rightarrow r$; the similarity exists because we can use symmetry without loss of generality to demand the ray propagate radially in the x direction, along $\hat{r} = \hat{x}$].

5.6.3 Deriving the polarization and energy equations, for radial propagation on the light cone

Almost half (12 of the 30 characteristic fields) naturally are associated with wave packets that propagate at the speed of light of the background spacetime (i.e. $s = \pm 1$). In other words, they propagate on characteristics that correspond to null curves of the Painleve-Gullstrand metric [Eq. (5.32)]. For radially propagating characteristics, that means

$$dr/dt = V_s^r \quad (5.34)$$

$$V_s^r \equiv s - \sqrt{2/r} \quad (5.35)$$

with $s = \pm 1$. The resulting null curve structure is shown in Fig. 5.2.

Because both this case and the Rindler case discussed in Sec. 5.5.3 possess rotational symmetry about the propagation axis, the equations governing these two cases prove exceedingly similar. The analysis follows the same course.

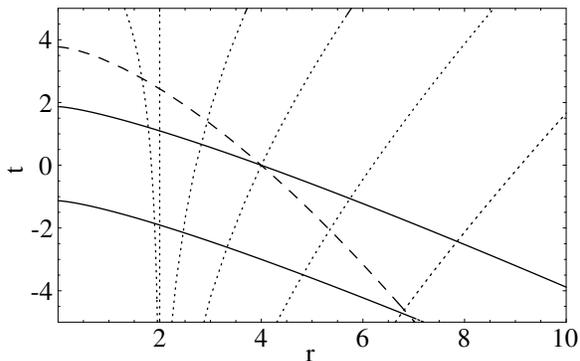


Figure 5.2: Examples of the three types of radially propagating rays of the KST 2-parameter system linearized about a Painleve-Gullstrand background. The solid lines show rays propagating inward at the speed of light (V_r^-). The dotted lines show rays propagating “outward” at the speed of light (V_r^+). Finally, the dashed curve shows the rays that propagate inside the light cone (at speed V_r^0). The quantities V_r^s are defined in Eq. (5.35).

5.6.3.1 Essential tool: Diagonalizing $P_s F P_s$ with $s = \pm 1$

As in the Rindler case, we will rewrite the polarization and energy equations by using eigenvectors $f_s^{(\mu)}$ of $P_s F P_s$. Because we again have rotational symmetry about the propagation direction, we can again decompose the eigenvectors into a set of two scalars ($s1$ and $s2$), a 2-vector v , and a symmetric-traceless-2-tensor t . The eigenvalues may be expressed using

$$\zeta_{s,\mu} \equiv \bar{\zeta}_{s,\mu} / \sqrt{2} r^{3/2} \quad (5.36)$$

where the $\bar{\zeta}_{s\mu}$ are defined by

$$\bar{\zeta}_{s,s1} = -3 \quad (5.37a)$$

$$\bar{\zeta}_{s,s2} = \left[\frac{7}{2} + 3\gamma - \frac{(33 + 91\hat{z} + 24\hat{z}^2)}{4(1 + 3\hat{z})(1 - 2\gamma)} \right] \quad (5.37b)$$

$$\bar{\zeta}_{s,v} = \frac{3 - 3\hat{z} - 5\gamma}{1 - 2\gamma} \quad (5.37c)$$

$$\bar{\zeta}_{s,t} = 1 \quad (5.37d)$$

The eigenspaces are, by symmetry, spanned by precisely the same fields as in the Rindler case. In particular, as in the Rindler case the eigenvectors do not change as we move along a ray.

5.6.3.2 Polarization equation for $s = \pm 1$

For polarizations which propagate radially on the light cone (i.e. $s = \pm 1$), the polarization equation [Eq. (5.7c)] can be written as

$$0 = [\partial_t + V_s^r \partial_r] d_{s,\alpha} + s \frac{d_{s,\alpha}}{r} - \sum_{\beta} d_{s,\beta} \left(v_{s,\alpha}, SF v_{s,\beta} \right). \quad (5.38)$$

where we make use of Eqs. (C.26) and (C.27) to simplify the right side, and where we observe $\partial_a \hat{r}^a = 2/r$.

As in the Rindler case, we may expand the amplitude $\bar{u} = \sum_{\mu} D_{s\mu} f_s^{(\mu)}$ in terms of the basis $f_s^{(\mu)}$, and thereby arrive at polarization propagation equation precisely analogous to the Rindler result [compare with Eq. (5.25)]:

$$\left[\partial_t + \left(s - \sqrt{\frac{2}{r}} \right) \partial_r \right] D_{s\mu} = \left(\zeta_{s\mu} - \frac{s}{r} \right) D_{s\mu}. \quad (5.39)$$

These equations may be integrated to describe the evolution of polarization along any individual radial ray.

5.6.3.3 Energy equation for $s = \pm 1$

Because symmetry guarantees a close similarity between this Painleve-Gullstrand case and the Rindler case, we find the energy E of a geometric-optics-limit solution propagating on the light cone radially inward ($s = -1$) or outward ($s = +1$) can be expressed with precisely the same expression we used in the Rindler case: Eq. (5.26). [In this case, we again use a measure $\mu = 1$ compatible with the background flat spatial cartesian-coordinate metric.]

The rate of change of this energy, dE/dt , can be obtained in two ways. On the one hand, we can directly form E , convert to spherical coordinates, differentiate the resulting expression for dE/dt , and use Eq. (5.39). On the other hand, we can find dE/dt using the general expression of Eq. (5.17), an expression we simplify by using i) the relation between Q and O_j given in Eq. (5.18), ii) the basis $f_s^{(\mu)}$ of eigenvectors of $P_s F P_s$, and iii) the expression [obtained from Eq. (C.28) and converted from a component to an operator expression]

$$P_s [\partial_t S + \partial_a (S A^a)] P_s = -\frac{3}{\sqrt{2} r^{3/2}} P_s. \quad (5.40)$$

In either case, we conclude

$$\begin{aligned} \frac{dE}{dt} &= \int d^3x \sum_{\mu} |D_{s\mu}|^2 \frac{2\text{Re}(\bar{\zeta}_{s\mu}) - 3}{\sqrt{2}r^{3/2}} \\ &+ \int d^3x 2\text{Re} \left[D_{s,s1}^* D_{s,s2} \left(f_s^{(s1)*}, S f_s^{(s2)} \right) \right. \\ &\quad \left. \times \frac{\bar{\zeta}_{s,s1}^* + \bar{\zeta}_{s,s2} - 3}{\sqrt{2}r^{3/2}} \right]. \end{aligned} \quad (5.41)$$

In particular, for prototypical coherent wave packets—that is, wave packets where s and μ are the same everywhere in the packet—we can express the growth rate of the energy $E_{s\mu}$ of the wave packet as

$$\frac{1}{E_{s\mu}} \frac{dE_{s\mu}}{dt} = \frac{2\text{Re}(\bar{\zeta}_{s\mu}) - 3}{\sqrt{2}r^{3/2}} \quad (5.42)$$

where r is the current location of the packet.

5.6.4 Deriving the polarization and energy equations, for radial propagation against the shift vector

The remaining 18 fields propagate inward against the shift vector, at speed $V_o = -\sqrt{2/r}$.

We shall not follow the same pattern we used to address propagation on the light cone [on a Rindler background in Sec. 5.5.3 and on a Painleve-Gullstrand background in Sec. 5.6.3]. In those sections, we provided extensive discussion and background—the explicit form of the polarization equation; a modified form of the polarization equation in an alternative basis; explicit expressions for the growth rate of energy general geometric-optics solutions; explicit demonstration that PCWP solutions existed—before finally recovering the growth rate of PCWPs. Instead, for pedagogical and other reasons [see Sec. 5.6.4.3], we shall take a briefer, more practical approach better suited to extracting precisely the information needed to decide when some coherent wave packet can amplify a significant amount within the future domain of dependence.

Specifically, following the arguments at the end of Sec. 5.3.3, we expect that—whether or not PCWPs exist as exact solutions to the polarization equation – when the largest eigenvalue $o_{o\nu}$ of O_o is particularly large, a generic coherent wave packet will rapidly converge to a PCWP with $w = f_o^{(\nu)}$. In other words, we expect that when the growth rates are large, the growth rate of generic coherent wave packets can be obtained by finding the largest value of $dE/dt/E$ for PCWPs [i.e. the maximum of Eq. (5.20) over μ].

In short, we continue to evaluate Eq. (5.20) to get growth rates, though now we trust the results only when the growth rates are large.

5.6.4.1 Growth rate of PCWPs

To evaluate the growth rate of PCWPs, we must diagonalize O_o :

$$O_o = P_o \left\{ F + \frac{1}{2} S^{-1} [\partial_t S + \partial_a (S A^a)] \right\} P_o$$

However, from Eq. (5.40) we know the term in square brackets is diagonal. Therefore, diagonalizing O_o to obtain eigenvalues $o_{o\mu}$ and eigenvectors $f_j^{(\mu)}$ is equivalent to diagonalizing $P_o F P_o$ for eigenvalues $\zeta_{o\mu}$ and eigenvectors $f_o^{(\mu)}$. The eigenvalues of the two operators are related by

$$o_{o\mu} = \zeta_{o\mu} - \frac{3}{2\sqrt{2}r^{3/2}} \quad (5.43)$$

We shall express the eigenvalues $\zeta_{o\mu}$ of $P_o F P_o$ in terms of the dimensionless rescaled quantities L_μ , defined implicitly by

$$\zeta_{o\mu} = L_\mu \times \sqrt{2}/r^{3/2} \quad (5.44)$$

Substituting Eq. (5.43) into the general expression for the growth rate of PCWPs [Eq. (5.20)], we find that a PCWP in the polarization μ will have energy grow at rate

$$\frac{1}{E_{o\mu}} \frac{dE_{o\mu}}{dt} = \left[2\text{Re}(L_\mu) - \frac{3}{2} \right] \frac{\sqrt{2}}{r^{3/2}} \quad (5.45)$$

where r is the instantaneous location of the packet.

5.6.4.2 Essential tool: Diagonalizing $P_o F P_o$

To obtain explicit growth rate expressions using Eq. (5.45), we need the eigenvalues of $P_o F P_o$, expressed according to Eq. (5.44).

As in the previous two cases, the eigenspaces of $P_o F P_o$ may be decomposed into distinct classes, depending on their symmetry properties of rotation about the propagation axis. These spaces are as follows:

- *Helicity-0* a four-dimensional space of rotational scalars (“helicity-0” states), with eigenvalues given by Eq. (5.44) with

$$L_{s1,s2} = \frac{-1 + 3\hat{z} + 18\hat{z}^2 \pm \sqrt{3}\sqrt{Y_1}}{4(1 + 3\hat{z})} \quad (5.46a)$$

$$L_{s3} = \frac{1}{4}(-30 + 19\eta + 12\hat{z} - 6\eta\hat{z}) \quad (5.46b)$$

$$L_{s4} = \frac{3}{2}(1 + 2\hat{z}) \quad (5.46c)$$

Here, we use $\eta \equiv -2/(\gamma - 1/2)$ and Y_1 given according to an expression listed below [Eq.

(5.47)].

- *Helicity-1* an 8-dimensional space of rotational 2-vectors (“helicity-1” states), with doubly-degenerate eigenvalues given by

$$L_{v1} = -2 \quad (5.46d)$$

$$L_{v2} = \frac{1}{2}(1 + 6\hat{z}) \quad (5.46e)$$

$$L_{v3,v4} = \frac{3}{8}(5 + 8\hat{z}) + \frac{\eta(13 + 83\hat{z} + 84\hat{z}^2)}{32(1 + 3\hat{z})} \pm \frac{\sqrt{Y_2}}{32(1 + 3\hat{z})} \quad (5.46f)$$

Again, we use $\eta \equiv -2/(\gamma - 1/2)$. The expression for Y_2 is given below [Eq. (5.48)].

- *Helicity-2* a four-dimensional space of symmetric-traceless-2-tensors (“helicity-2” states), with doubly-degenerate eigenvalues given by

$$L_{t1} = \frac{3}{2}(1 + 2\hat{z}) \quad (5.46g)$$

$$L_{t2} = 2 + 3\hat{z} \quad (5.46h)$$

- *Helicity-3* and finally a 2-dimensional space of helicity-3 states, with eigenvalue

$$L_3 = 3(1 + \hat{z}) \quad (5.46i)$$

In the above discussion, $Y_{1,2}$ are defined by

$$Y_1 = (1 + 3\hat{z})(-5 + 5\hat{z} + 24\hat{z}^2 + 36\hat{z}^3) \quad (5.47)$$

$$Y_2 = 1296(1 + 3\hat{z})^2 + \eta^2(13 + 83\hat{z} + 84\hat{z}^2)^2 - 24\eta(1 + 3\hat{z})(89 + 199\hat{z} + 132\hat{z}^2) \quad (5.48)$$

and we use the shorthand $\eta \equiv -2/(\gamma - 1/2)$.

5.6.4.3 Aside: Why can't we follow the previous pattern?

Unlike all cases previously discussed, a handful of the eigenvectors depend weakly on position. As a result, the use of a basis which diagonalizes O_o does not offer as dramatic a simplification as it did in our earlier analyzes of the polarization equation [Sections 5.5.3 and 5.6.3]. To be explicit, if we rewrite the polarization equation in the basis $f_o^{(\mu)}$ in the fashion of those earlier analyzes, we obtain

[see Eq. (C.17)]

$$\sum_{\nu} D_{o\nu} M_{\mu\nu} = \left(\partial_t + V_o^a \partial_a + \frac{1}{2} \partial_a V_o^a - o_{o\mu} \right) D_{o\mu} \quad (5.49)$$

for $M_{\mu\nu}$ some *nonzero*, position-dependent matrix coupling the various $D_{o\mu}$.

5.7 Transients and limitations on numerical simulations: Rindler

In earlier sections, we developed—in general [Sections 5.2, 5.3, and 5.4] and for specific examples [e.g., Sec. 5.5 analyzes propagation of transients according to the KST 2-parameter formulation of Einstein’s equations, linearized about a Rindler background] – tools to analyze the growth of special (i.e. prototypical coherent wave packet) geometric-optics-limit transient solutions. In this section, we demonstrate how these tools can be used to discover when a particular formulation of Einstein’s equations [here, some specific member of the KST 2-parameter system] which is linearized about a specific background [here, flat space in Rindler coordinates] admits some massively-amplified transient solution.

Specifically, in this section we apply the general tools developed in an earlier section [Sec. 5.5] to determine the largest possible amplification of a prototypical coherent wave packet while it remains within the future domain of dependence of some initial data slice. In Sec. 5.7.1 we describe the initial data slice we chose and the subset of transient solutions we studied. In Sec. 5.7.2, we apply the tools developed in an earlier section [Sec. 5.5] to determine the amplification of each transient. We also find an expression for the largest possible amount a transient can amplify. Finally, in Sec. 5.7.2, we invert our expression to determine which pairs of KST parameters (\hat{z}, γ) admit transients that amplify in energy by more than 10^{32} (i.e. in amplitude by more than 10^{16}).

5.7.1 Transients studied

We limit attention to the future domain of dependence of the initial-data slice $x \in [0.01, 1]$ at $t = 0$. Since the KST 2-parameter formulation has fields which propagate at (but no faster than) the speed of light, the future domain of dependence of this slice is precisely what we would obtain using Einstein’s equations: a region bounded by the two curves $x_- \equiv 0.01 \exp t$ and $x_+ \equiv \exp(-t)$. This region is shown in Fig. 5.3. The future domain of dependence extends to time

$$T_{\max} \equiv \ln 10, \quad (5.50)$$

at which point the two bounding curves intersect.

Geometric-optics solutions are defined on rays [i.e. solutions to Eq. (5.8)]. While three classes of rays exist in this region—those ingoing at the speed of light ($dx/dt = -x$); those outgoing at

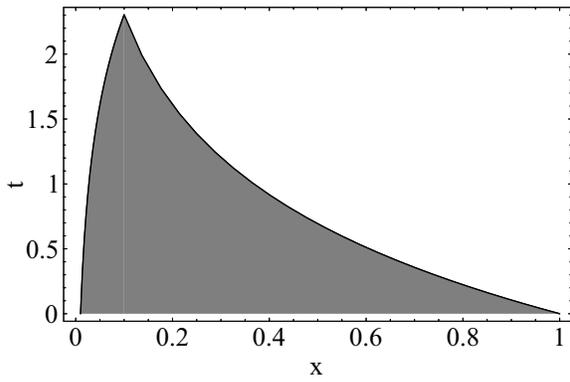


Figure 5.3: The shaded region is the future domain of dependence of the region $x \in [0.01, 1]$ for the KST 2-parameter formulation of evolution equations linearized about a Rindler background. Transients are any solutions which are defined in this region. We study all the prototypical coherent wave packets which propagate on the light cone (i.e. according to $dx/dt = \pm x$).

the speed of light ($dx/dt = +x$); and those which have fixed coordinate position—we for simplicity chose to study only the amplification of transients that propagate on the light cone.

5.7.2 Amplification expected

For each ray that propagates on the light cone ($dx/dt = \pm x$) within the future domain of dependence, and for each polarization on that ray, we can compute the amplification in energy. If $R_{s,\mu} \equiv E_{s\mu}^{-1} dE_{s\mu}/dt$ [see Eq. (5.30)], we can express the ratio of energy of the wave packet when it exits the future domain of dependence at time t_{out} to the initial energy at time $t = 0$ as

$$\mathcal{A}_{s\mu}(x_o) = E_{s\mu}(t_{\text{out}})/E_{s\mu}(0) = \exp(t_{\text{out}} R_{s\mu}) .$$

We have explicit expressions for $R_{s\mu}$; we can compute $t_{\text{out}}(x_o, s)$ for each initial point x_o and for each propagation orientation (i.e. for each s); and we therefore can maximize $\mathcal{A}_{s\mu}(x_o)$ over all possible choices of initial location (x_o), propagation direction (s), and polarization (μ) to find the largest possible ratio \mathcal{A} of initial to final prototypical coherent wave packet energy.

In fact, because for each polarization of prototypical coherent wave packet, the growth rates of energy is independent of time and space, the largest amplifications possible always occur along the longest-lived rays—in other words, along the two bounding rays x_+ and x_- , which both extend to $t_{\text{out}} = T_{\text{max}}$. Therefore, we conclude that, while within the future domain of dependence of the slice $x \in [0.01, 1]$, the largest amount the energy of any prototypical coherent wave packet can amplify is given by the factor

$$\mathcal{A} = \exp(T_{\text{max}} R_{\text{Rind}}) \tag{5.51}$$

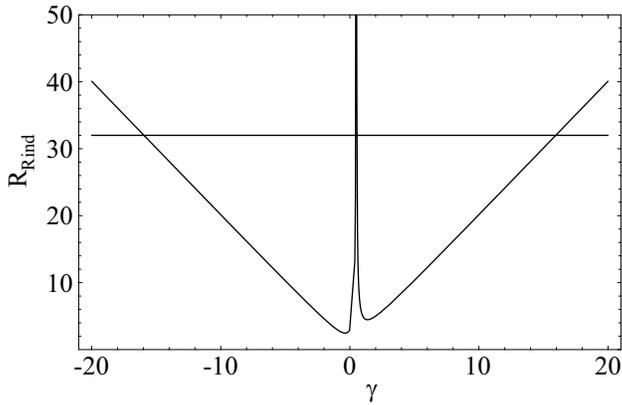


Figure 5.4: The solid curve is the theoretical prediction for the largest growth rate of wave packets that propagate on the light cone [Eq. (5.52)]. The horizontal line is the value 32. According to arguments made in Sec. 5.7.2, those γ which have $R_{\text{Rind}} > 32$ have some prototypical coherent wave packet which, in the future domain of dependence, amplifies in energy by more than 10^{32} .

where R_{Rind} is given by

$$\begin{aligned} R_{\text{Rind}} &\equiv \max_{\mu, s} R_{s\mu} = \max_{\mu, s} [2\text{Re}(\zeta_{s\mu}) + s] \\ &= \max \left(1, \left| 2 \frac{1+\gamma}{-1+2\gamma} - 1 \right|, \left| 2 \frac{1+2\gamma^2}{-1+2\gamma} - 1 \right| \right). \end{aligned} \quad (5.52)$$

KST formulations which definitely possess some ill-behaved transient solution when linearized about Rindler

Finally, we can invert Eq. (5.51) to find those combinations of KST parameters (\hat{z}, γ) which permit some transient (in particular, some prototypical coherent wave packet) to increase in energy by more than a factor 10^{32} (i.e. 10^{16} in amplitude). The condition may be expressed as either $\mathcal{A} > 10^{32}$ or, equivalently, as $R_{\text{Rind}} > 32$. The function R_{Rind} is shown in Fig. 5.4, along with the line $R_{\text{Rind}} = 32$.

Therefore, we know that some transient can amplify in energy by more than 10^{32} if i) $\gamma > (33 + \sqrt{949})/4$, ii) $\gamma < -(31 + \sqrt{1077})/4$, or iii) $\gamma > 29/64$ and $\gamma < (33 - \sqrt{949})/4$.

5.7.3 Relevance of our computation to numerical simulations

We have demonstrated that the KST 2-parameter formulation of Einstein's equations always admits, at any instant, prototypical coherent wave packet solutions which grow exponentially in time. Generically, we expect that at each instant (including in the initial data) these solutions are excited by errors in the numerical simulation (e.g., truncation and roundoff). They then propagate and grow; eventually, they reach the computational boundary.

Our calculations above describes the largest amount any such wave packet solution could possibly

grow by the time it reaches the computational boundary. If that amplification factor is sufficiently large that the wave packets reach “unit” amplitude (i.e., whatever magnitude is needed to couple to nonlinear terms strongly), here conservatively assumed to be 10^{16} , then we expect any simulation using that particular combination of KST parameters will quickly crash.

Aside: What happens to PCWPs at late times?

Eventually, the wave packets excited by numerical errors will reach the computational boundary. What happens afterward depends strongly on the precise details of the boundary conditions.

For example, maximally-dissipative boundary conditions (i.e., the time derivatives of all ingoing characteristic fields are set to zero) will allow the wave packet to leave the computational domain entirely (with some small amount of reflection that goes to zero in the geometric-optics limit). In this case, at late times no transient will ever amplify by more than the amount described above (in Sec. 5.7.2).

On the other hand, other choices for boundary conditions could cause wave packets to reflect back in to the computational domain. In these circumstances, the outcomes are far more varied —at late times, the wave packet could potentially grow, could decay to zero, or could enter a repetitive cycle where on average its amplitude is constant.⁹

Therefore, without some more specific proposal for boundary conditions, we cannot make useful statements regarding the late-time development of this instability process—or, in other words, we cannot study the growth of coherent wave packets for more than a light crossing time.

5.8 Transients and limitations on numerical simulations: PG

In this section, we provide another example of how tools developed earlier for the analysis of prototypical coherent wave packets – in general [Sections 5.2, 5.3, and 5.4] and for specific examples [e.g., Sec. 5.6 analyzes propagation of transients according to the the KST 2-parameter formulation of Einstein’s equations, linearized about a Painleve-Gullstrand background]—can be applied to discover which formulations of Einstein’s equations permit ill-behaved transients.

Specifically, in this section we study the propagation of coherent wave packets in the 2-parameter KST form of Einstein’s evolution equations, linearized about Schwarzschild written in Painleve-Gullstrand (PG) coordinates. The theory needed to understand the propagation and growth of radially-propagating coherent wave packets has been developed in an earlier section [Sec. 5.6]. We apply our techniques to a handful of coherent wave packet transient solutions, to discover conditions

⁹In fact, in this particular case, we expect that if a wave packet with growth rate $1/\tau$ reflects, then symmetry and the structure of the Rindler growth rates [i.e., Eq. (5.30)] insures that the reflected ray has growth rate $-1/\tau$. Therefore, on average, the wave packet has a zero growth rate.

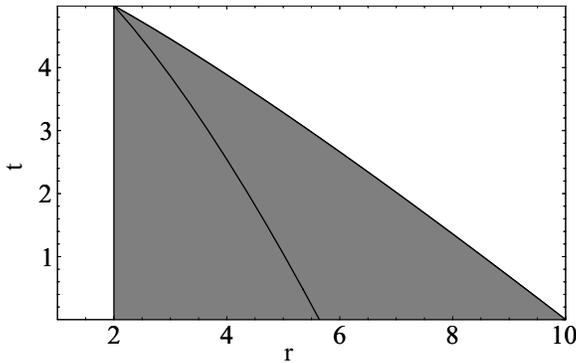


Figure 5.5: The shaded region is the future domain of dependence of the region $r \in [2, 10]$ for the KST 2-parameter formulation of evolution equations. Transients are any solutions which are defined in this region. For reasons emphasized in the text, the rays that last for the longest coordinate time prove particularly helpful. These rays are the left and right boundaries (i.e., the horizon and a ray that propagates inward at the speed of light from $r = 10$) and one ray propagating against the shift vector which emanates from their intersection.

on the two KST parameters (\hat{z}, γ) which permit amplification of those transients' energy by a factor $1/\epsilon_e^2$ for $\epsilon_e = 10^{-16}$.

To provide concrete examples of estimates, we assume the initial data slice contains the region $r \in [2, 10]$. So any influence from boundary conditions cannot muddle our computations, we limit attention to coherent wave packets which are defined in the future domain of dependence of that slice.

5.8.1 Transients studied

We limit attention to the future domain of dependence of the region $r \in [2, 10]$ at $t = 0$. Since the KST 2-parameter formulation has fields which propagate at (but no faster than) the speed of light, the future domain of dependence of this slice is precisely what we would obtain using Einstein's equations: the region shown in Fig. 5.5. In particular, the future domain of dependence is bounded on the left by the generators of the horizon (trapped at $r = 2$) and on the right by rays travelling inward at the speed of light. This ingoing ray reaches $r = 2$ at the endpoint of the future domain of dependence, at time $t = t_{\max}$ defined by

$$\begin{aligned} t_{\max} &\equiv \int_2^{10} \frac{dr}{1 + \sqrt{2/r}} \\ &= 4[3 - \sqrt{5} + \operatorname{csch}^{-1}(2)] \approx 4.98 \end{aligned} \tag{5.53}$$

In our future domain of dependence, we have three classes of solutions to the ray-propagation equation [Eq. (5.8)]: those ingoing at the speed of light ($V_- = -1 - \sqrt{2/r}$); those ingoing with the shift ($V_o = -\sqrt{2/r}$); and those outgoing ($V_+ = 1 - \sqrt{2/r}$) [Eq. (5.35)].

5.8.2 Amplification conditions

For each of the three classes of rays ($s = \pm 1, 0$) propagating radially in the future domain of dependence [Fig. 5.5] and for each polarization on that ray, we can compute the amplification in energy using $R_{s\mu} \equiv E_{s\mu}^{-1} dE_{s\mu}/dt$ [see Eqs. (5.42) and (5.45)]. Specifically, for a wave packet starting at $r = r_o$ at time $t = 0$, propagating in the s -type congruence and in the polarization μ , the energy at the time $t_{\text{out}}(r_o, s)$ the ray exits the future domain of dependence is larger than the initial energy by a factor

$$\mathcal{A}_{s\mu}(r_o) \equiv E_{s\mu}(t_{\text{out}})/E_{s\mu}(0) \quad (5.54a)$$

$$\ln \mathcal{A}_{s\mu}(r_o) = \int_0^{t_{\text{out}}} dt R_{s\mu} = \int_{r_o}^{r_{\text{out}}} \frac{dr}{V_r^s} R_{s\mu}. \quad (5.54b)$$

We then search over all r_o , over all propagation directions s , and over all polarizations μ to find the largest amplification factor \mathcal{A} .

In fact, as in the Rindler case, we immediately know which rays produce the largest possible amplification, so we can perform the maximization by inspection.

- *Outgoing at light speed:* Since the amplification of energy increases as r gets smaller ($dE/dt/E \propto 1/r^{3/2}$) and with the duration of the ray in time, manifestly the generator of the horizon— with both the longest duration and the smallest r of all outgoing rays —will provide the largest possible amplification.

Since the ray of interest has fixed radial location $r = 2$, we find $\zeta_{+,\mu}$ is constant for all polarizations. Thus, the energy of a prototypical coherent wave packet in polarization μ increases by a factor $A_{+\mu}$, for $A_{+\mu} = \exp(t_{\text{max}}\zeta_{+\mu})$. In other words,

$$\begin{aligned} \ln \mathcal{A}_\mu &= [2\text{Re}(\bar{\zeta}_{+,\mu}) - 3] [3 - \sqrt{5} + \text{csch}^{-1}(2)] \\ &\approx 1.245 [2\text{Re}(\bar{\zeta}_{+,\mu}) - 3]. \end{aligned} \quad (5.55)$$

[The values for each $\bar{\zeta}_{+\mu}$ are given in Eq. (5.37).]

- *Ingoing at light speed:* The longest ray —namely, the right boundary of the future domain of dependence—permits the greatest possible amplifications. Thus, among all possible ingoing rays, the largest amplification factor for the polarization μ is given by $\mathcal{A}_{-\mu}$:

$$\begin{aligned} \ln \mathcal{A}_{-\mu} &= [2\text{Re}(\bar{\zeta}_{-,\mu}) - 3] \cdot \frac{\ln 5 - 2\text{csch}^{-1}(2)}{2} \\ &\approx 0.323 [2\text{Re}(\bar{\zeta}_{-,\mu}) - 3]. \end{aligned} \quad (5.56)$$

[The values for each $\bar{\zeta}_{-\mu}$ are given in Eq. (5.37).]

Note that since $\zeta_{-\mu} = \zeta_{+\mu}$, the outgoing transients trapped on the horizon grow *more* than the ingoing ones over the same time interval¹⁰.

- *Ingoing with lapse*: The amplification of energy increases both with ray length and with proximity to $r = 0$ (since growth rates go as $1/r^{3/2}$). Thus, the longest ray propagating at this speed contained in the future domain of dependence gives the best chances. That ray starts with $r = r_L$, with r_L defined so the ray terminates at the horizon at $t = t_{\max}$:

$$r_L \equiv \left[\frac{(4 + 3t_{\max})^2}{2} \right]^{1/3}. \quad (5.57)$$

Thus, we find the largest possible amplification among those polarizations that have $s = 0$ to be given by $\mathcal{A}_{0\mu}$:

$$\begin{aligned} \ln \mathcal{A}_{0\mu} &= \left[2\text{Re}(L_\mu) - \frac{3}{2} \right] \times \int_2^{r_L} \frac{dr}{\sqrt{2/r}} \frac{\sqrt{2}}{r^{3/2}} \\ &= \left[2\text{Re}(L_\mu) - \frac{3}{2} \right] \ln(r_L/2). \end{aligned} \quad (5.58)$$

[The values for each L_μ are given in Eq. (5.46).]

5.8.3 Results: Some KST parameters which have transients which amplify by 10^{32}

Under the proper choice of KST parameters, shown shaded in Fig. 5.6, one of the three types of ray ($s = \pm 1, 0$) may admit some prototypical coherent wave packet of polarization μ whose energy amplifies by 10^{32} [i.e., $\mathcal{A}_{s\mu} \geq 10^{32}$]. The clear region in Fig. 5.6 indicates KST parameters for which we have not yet found a transient which amplifies by 10^{32} .

5.8.4 Generalizations of our method which could generate stronger constraints on KST parameters

With our *extremely* conservative approach—eliminating those formulations with wave-packet solutions which amplify by 10^{32} in the future domain of dependence—we have already eliminated a broad region of parameter space. By relaxing some of our very restrictive assumptions, we expect we could discard still more KST parameters:

1. *Lower amplification cutoff*: Currently, we require an *enormous* amplification before we eliminate a formulation; relaxing the requirement on amplification excludes more systems.

¹⁰This should be expected: the ingoing and outgoing wave packets have similar growth rates at any given radius; we limit attention to rays which persist for a fixed time; and the outgoing modes we study remain closer to the horizon, where the growth rate is larger.

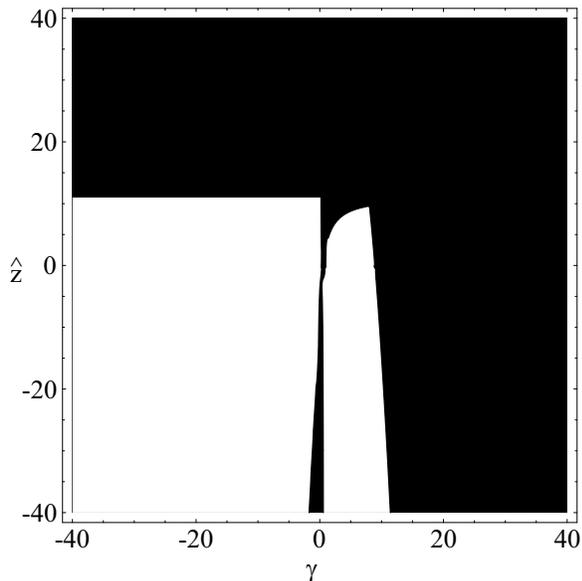


Figure 5.6: The shaded region indicates those KST parameters which produce some radially-propagating prototypical coherent wave packets which amplify their energy by greater than 10^{32} within the future domain of dependence of the slice $r \in [2, 10]$. Note a large proportion of parameter space has been excluded.

2. *Consider more transients* Currently, we compute the amplification of only a handful of transients; a consideration of other transients (for example, in the neighborhood of circular photon orbits in PG) may allow us to exclude additional parameters.
3. *Consider a larger region* Currently, we limit attention only to the future domain of dependence of the initial data slice. Certain rays, however, remain within the computational domain for far longer. For example, in the PG case, rays near the horizon remain in the domain for arbitrarily long;¹¹ even the slowly-infalling rays last substantially longer than the domain of dependence. Therefore, by considering the amplification of transients over a longer interval, we will discover significantly greater amplification and thus exclude a significantly broader class of formulations of Einstein's equations.
4. *Combine with boundary conditions* Finally, if we determine how geometric-optics solutions interact with boundary conditions, we can generalize our approach and address the *late-time stability* properties of the evolution equations—or, in other words, address the stability properties of the full initial-plus-boundary value problem.

¹¹One must take care to use the rays near the horizon in a sensible fashion. While analytically the rays remain within the computational domain for arbitrarily long times, one cannot expect wave packet solutions to be resolved and present in a numerical solution for arbitrarily long: the code has a finite smallest resolved scale. In practice, one must remember that whatever amplification one computes must be realistically attainable by some numerical simulation of fixed (though perhaps high) resolution in the coordinates of interest.

5.9 Conclusions

In this paper, we have demonstrated that certain transients (prototypical coherent wave packets) can be used to veto a significant range of proposed formulations of Einstein’s equations. We have described in considerable pedagogical detail precisely how to construct expressions for (or estimates of) the growth rate of prototypical coherent wave packets [i.e., Eq. (5.20)], verify those estimates, and employ them to veto proposed formulations of Einstein’s equations. These expressions employ no free parameters or knowledge of the solution, aside from a choice of plausible rays to examine. Moreover, despite the sometimes exhaustive details provided in Sections 5.5 and 5.6, the key tool—the growth rate of prototypical coherent wave packets [Eq. (5.20)] —is easy to apply, with little conceptual, notational, or computational overhead (see, for example, the brief Sec. 5.6.4.1 and its application in Sec. 5.8). Whether they are used conservatively, as in this paper, or generalized along the lines suggested in Sec. 5.8.4 (i.e., using more rays and larger fragments of spacetime), we believe these techniques will provide a useful way to bound the number of proposed formulations before further tests are conducted (for example, by the more ambitious Lindblom-Scheel energy-norm method) to decide whether a given formulation can produce effective simulations.

While our the specific examples of analyses in this paper have employed linearizations of the field equations themselves, we could just as well linearize a FOSH system representing evolution equations for the constraint fields [see, for example, KST Eqs. (2.40-2.43)]. The evolution equations for the constraints have been emphasized by many other authors as a probe of unphysical behavior. Since the general arguments of Sections 5.2, 5.3, and 5.4 do not depend on the precise FOSHLS used, we can perform a calculation following the same patterns as (for example) Sec. 5.8 to discover ill-behaved formulations.¹²

In this paper, we have also discovered curious properties of modes trapped on the horizon of a Schwarzschild hole in PG coordinates (Sec. 5.8). Analytically, we would expect that, if any growth rate for modes trapped on the horizon were positive, then these modes should grow without bound and be present in the evolution at late times. Numerically, however, we know that no *resolved* wave packets can appear at late times: such solutions would have to initiate arbitrarily close to the horizon, inconsistent with resolved, finite-resolution initial data. Still, marginally-resolved solutions of similar character could potentially behave in an implementation-dependent fashion, seeding outgoing modes which then propagate and amplify into the domain for all time. We shall explore this possibility in a future paper.

¹²The author expects no new information can be obtained by such an analysis. Moreover, because the constraint equations, when written in first-order form, involve many more variables than the field equations themselves, such an analysis should prove substantially more challenging.

Bibliography

- [1] L. E. Kidder, M. A. Scheel, and S. A. Teukolsky, Phys. Rev. D **64** (2001) 064017
- [2] M. A. Scheel and L. Lindblom, Phys. Rev. D **66** (2002) 084014
- [3] M. Alcubierre, G. Allen, B. Brügmann, et al. Phys.Rev. D **62** (2000) 124011
- [4] G. Yoneda and H. Shinkai, gr-qc/0204002 (brief description). Shinkai also provides a long but readable review of his program in gr-qc/0209111.
- [5] G. Calabrese, L. Lehner, and M. Tiglio, Phys. Rev. D **65** (2002) 104031
- [6] O. Brodbeck, S. Frittelli, P. Hubner, O. Ruela, J. Mat. Phys. **40** 909 [gr-qc/9809023]
- [7] O. Sarbach et al, Phys. Rev. D **66** (2002) 064002
- [8] B. Kelly et al, Phys. Rev. D **64** (2001) 084013
- [9] G. Calabrese and O. Sarbach, gr-qc/0303040.
- [10] G. Calabrese et al, gr-qc/0209017
- [11] B. Szilagyi and J. Winicour, gr-qc/0205044
- [12] S. Frittelli and R. Gomez, gr-qc/0302017 and gr-qc/0302032
- [13] M. Born and E. Wolf *Principles of Optics*, 6th edition, Pergamon Press (1980)

Appendix A

Appendix for compact object inspiral

A.1 Evolution of the maximum

As the conserved constants E , L evolve, the height of any local maximum in the potential V [Eq. (2.1)] will similarly evolve. Because we are at a local maximum, we can find a simple expression for the rate of change of the value of the potential at that local maximum:

$$V_{\max}(E, L) = V(r_{\max}(E, L), E, L) \quad (\text{A.1})$$

(where r_{\max} , the location of the potential's local maximum, is a solution to $dV/dr = 0$). In general, to order of magnitude, we expect it goes as $\sim V/\tau_{\text{gw}}$ for τ_{gw} the gravitational wave timescale $\tau_{\text{gw}} \sim E/\dot{E} \sim L/\dot{L}$. But when the particle is nearly on a circular orbit, then the source of radiation nearly satisfies helical symmetry and therefore $dE \approx \Omega dL$ for Ω the angular frequency of the circular orbit. And in these special conditions V changes even more slowly than we would normally expect.

To be explicit, we evaluate dV_{\max}/dt , which (because we are at a maximum) we can generally express as follows:

$$\frac{dV_{\max}}{dt} = \frac{\partial V}{\partial L} \left[\frac{dL}{dt} + \frac{dE}{dt} \frac{\partial V / \partial E}{\partial V / \partial L} \right]_{r=r_{\max}}. \quad (\text{A.2})$$

We can most transparently prove the necessary result by rewriting the potential V more abstractly than the standard form of Eq. (2.2). Recall one derives the radial potential from the constancy of the test-particle's rest mass [e.g., $g^{ab}p_a p_b = -m^2$]. Since the Kerr metric in Boyer-Lindquist coordinates has form $g_{ab} = g_{tt}dt^2 + 2g_{t\phi}dtd\phi + g_{\phi\phi}d\phi^2 + g_{rr}dr^2 + g_{\theta\theta}d\theta^2$, by employing the definitions of E and L (e.g., $-p_t \equiv \text{energy}$) and the existence of an equatorial orbit, one obtains the (first integral of the) radial geodesic equation Eq. (2.1) with $V = (1 + E^2 g^{tt} + L^2 g^{\phi\phi} - 2ELg^{t\phi})/g_{rr}$. One can similarly show, by employing the definitions of the ‘‘raised’’ components (e.g., $p^t \equiv mdt/d\tau$ for τ proper time), that $d\phi/dt = (g^{\phi\phi}L - g^{t\phi}E)/(-g^{tt}E + g^{t\phi}L)$ and $dt/d\tau = -g_{rr}\partial V/\partial E$. Using

these expressions in Eq. (A.2), we conclude that

$$\frac{dV_{\max}}{dt} = -\frac{dt/d\tau}{g_{rr}} \left[\frac{dL}{dt} - \frac{dE}{dt} \frac{1}{d\phi/dt} \right]_{r=r_{\max}}. \quad (\text{A.3})$$

When the potential admits a nearly-circular orbit (angular frequency $\Omega = d\phi/dt$) near the local maximum, we have $dE \approx \Omega dL$ in the emitted radiation and therefore dV_{\max}/dt is smaller than normal. [Similar arguments apply to the minimum, and prove that stable circular orbits evolve to stable circular orbits.]

A.2 Kerr Parameters and Constraints

In the Teukolsky-equation-based inspiral literature [1, 2], when the orbit is *bound* (=does not fall into hole or escape to infinity) it is characterized not by physical parameters (E, L, a) but by the location of its radial turning points (r_{\pm}) and the remaining root of its potential (\bar{r}):

$$V = \frac{E^2 - 1}{r^3} (r_+ - r) (r - r_-) (r - \bar{r}) \quad (\text{A.4})$$

$$= -(E^2 - 1) - \frac{2}{r} + \frac{[L^2 - a^2 (E^2 - 1)]}{r^2} - \frac{2(L - aE)^2}{r^3}. \quad (\text{A.5})$$

To further simplify the algebra involved, one replaces r_{\pm} by a parameterization analogous to classical mechanics (semi-latus rectum and eccentricity):

$$r_{\pm} = \frac{p}{1 \mp e} \quad (\text{A.6})$$

(with p, e both real, positive). After replacing r_{\pm} by p, e using $r_{\pm} = p/(1 \mp e)$, we find the following explicit forms for E, L in general: (specify retrograde orbits by a negative a sign)

$$E = \sqrt{1 - \frac{(1 - e^2)}{2p + (1 - e^2)\bar{r}}}, \quad (\text{A.7})$$

$$L = \left[\text{sign}(a) \frac{p\sqrt{\bar{r}}}{\sqrt{2p + (1 - e^2)\bar{r}}} + aE \right]. \quad (\text{A.8})$$

The parameters p, e, \bar{r} , however, are not fully independent. The coefficients of $V(r)$ satisfy a polynomial equation when the coefficients are expressed in terms of L, E, a ; the coefficients must satisfy the same polynomial when the coefficients are written using r_{\pm}, \bar{r} . Re-expressing that polynomial in terms of p, e , we find

$$0 = p^2 [p(\bar{r} - 2) - 4\bar{r}]^2 + a^4 [p + (1 - e^2)\bar{r}]^2 - 2a^2 p \{ 4(1 - e^2)\bar{r}^2 + 2p^2(2 + \bar{r}) + p\bar{r} [8 + (\bar{r} - 2)(1 - e^2)] \}. \quad (\text{A.9})$$

As a practical matter, we usually specify a particle by (p, e, a) and then solve for \bar{r} and hence all other orbital parameters.

A.2.1 Separatrix

The transition from stable to unstable occurs when $\bar{r} = r_- = p/(1+e)$. We can attempt to express this relation in terms of the usual orbital parameters (p, e, a) . Substituting in the above polynomial [Eq. (A.9)], we get

$$0 = p_s^2 (p_s - 6 - 2e)^2 + a^4 (e - 3)^2 (1 + e)^2 - 2a^2 (1 + e) p_s [14 + 2e^2 + p_s (3 - e)] \quad (\text{A.10})$$

where $p = p_s(e, a)$, defined by (appropriate solutions to) the above relation, determines the separatrix between stable and unstable geodesic orbits.

A.2.2 Minimum of potential

At the transition from stable to unstable orbits, the maximum of the potential is at $r_- = \bar{r} = p_s/(1+e)$. Therefore, by symbolically differentiating Eq. (A.4), we find $r_{\min} = 3p_s/(3-e)$. Inserting this expression into V and noting $V(r_{\max}) = 0$, we find $I_{\max} \equiv V(r_{\max}) - V(r_{\min})$:

$$I_{\max} = \left(\frac{4e}{3}\right)^3 \frac{1}{(3-e)(1+e)p_s} \quad (\text{A.11})$$

A.3 Results from Ori and Thorne

So that the reader can more easily compare our estimates against those of Ori and Thorne, we provide approximations to their results.

In Ori and Thorne's Table II, they tabulate S/N and N_c for $10M_\odot$ on 10^6M_\odot , using fiducial distance 1 Gpc to set the amplitude scale. [In this table, OT list the number of quadrupolar gravitational wave cycles $N_{\text{cyc}} = 2N_c$.] We can approximate their results by the two functions

$$\begin{aligned} \log_{10} N_{c,\text{OT}} &\approx 0.92 + 0.110x + 0.0478x^2, \quad a \in [0.9, 0.99] \\ &\approx 1.03 + 0.073x + 0.138x^2, \quad a < 0.9 \end{aligned} \quad (\text{A.12})$$

and

$$\begin{aligned} \log_{10} \left(\frac{S}{N}\right)_{\text{OT}} &\times 10^{-22} \\ &\approx 0.3156 + 0.3193x + 0.0451x^2, \quad a \in [0.8, 0.99] \\ &\approx 0.191 - 0.226x - 0.448x^2, \quad a < 0.8 \end{aligned} \quad (\text{A.13})$$

where $x \equiv \log_{10}(1 - a)$.

Bibliography

[1] C. Cutler, D. Kennefick, and E. Poisson, *Phys. Rev. D* **50** 3816

[2] K. Glampedakis and D. Kennefick, *Phys. Rev. D* **66** 044002 (2002)

Appendix B

Appendix for thermoelastic noise

B.1 Approximate formulae for mesa modes

By inserting into Eq. (3.5) the asymptotic expansion (at large argument) for the Bessel function

$$I_0(z) = \frac{1}{\sqrt{2\pi z}} e^z, \quad (\text{B.1})$$

setting $b = 1$, $r_o = D - \xi$, and

$$\alpha = \sqrt{(1+i)/2}, \quad (\text{B.2})$$

we bring Eq. (3.5) into the approximate form

$$U(D, r) = \sqrt{\frac{\pi}{\alpha^2 r}} \int_0^D \sqrt{D - \xi} \exp[-\alpha^2 (r - D + \xi)^2] d\xi. \quad (\text{B.3})$$

By expanding $\sqrt{D - \xi}$ as a power series in ξ/D up to some order n , and then performing the integral in Eq. (B.3) analytically, we obtain expressions for $U(D, r)$ with various accuracies. The least accurate expression, $n = 0$ (obtained by setting $\sqrt{D - \xi} = \sqrt{D}$, integrating, and discarding a term proportional to $\text{erfc}(\alpha r)$ that is negligible compared to $\text{erfc}[\alpha(r - D)]$ at the relevant radii, $r \sim D$ or larger) is

$$U_0(D, R) = \frac{\pi}{2\alpha^2} \sqrt{\frac{D}{r}} \text{erfc}[\alpha(r - D)]. \quad (\text{B.4})$$

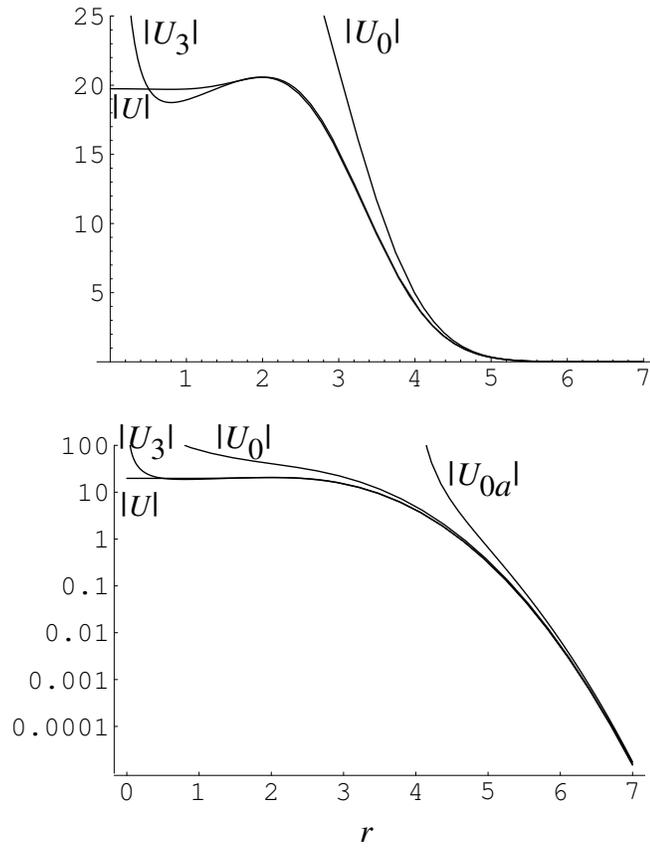


Figure B.1: Comparison of the moduli of the exact flat-topped mode $U(4, r)$, the zero-order approximation $U_0(4, r)$ [Eq. B.4)], the third-order approximation $U_3(4, r)$ [Eq. (B.5)], and the asymptotic approximation to the zero-order approximation, $U_{0a}(4, r)$ [Eq. (B.7)]. Top: linear plot; bottom: logarithmic plot.

Here $\text{erfc}(z)$ is the complementary error function, $1 - \text{erf}(z)$. At order $n = 3$ we get a much more accurate expression:

$$\begin{aligned}
 U_3(D, r) = & \frac{1}{64\alpha^5 D^2 \sqrt{Dr}} \\
 & \times \left\{ \pi \left[\alpha^3 (2r^3 - 10Dr^2 + 30D^2r + 10D^3) + \alpha(3r - 5D) \right] \right. \\
 & \quad \times \left[\text{erfc}[\alpha(r - D)] - \text{erfc}[\alpha r] \right] \\
 & + 2\sqrt{\pi} \left[e^{-\alpha^2 r^2} [\alpha^2 (15D^2 - 5Dr + r^2) + 1] \right. \\
 & \quad \left. \left. - e^{-\alpha^2 (r-D)^2} [\alpha^2 (11D^2 - 4Dr + r^2) + 1] \right] \right\}. \tag{B.5}
 \end{aligned}$$

Here erfc is the complementary error function, $\text{erfc}(z) = 1 - \text{erf}(z)$.

The two approximations $U_0(4, r)$ and $U_3(4, r)$ are compared with the exact mode $U(4, r)$ in Fig. B.1. At $r > 2$, U_3 is highly accurate; at $r > 5$, U_0 is highly accurate. The analytic formula (B.4)

for U_0 shows that the diffraction-induced tail of this mode falls off very rapidly outside $r = D$ — slightly more rapidly than would the tail of a Gaussian centered on $r = D$. Of course, this is to be expected since $U(D, r)$ is constructed from a superposition of Gaussians that are centered on radii $\leq D$.

A useful but cruder approximation to $U(D, r)$ can be obtained by inserting the asymptotic expansion of $\operatorname{erfc}[\alpha(r - D)]$ into Eq. (B.4) for U_0 . The result is

$$U_{0a}(D, r) = \frac{\sqrt{\pi D/r}}{2\alpha^3(r - D)} e^{-\alpha^2(r-D)^2}, \quad (\text{B.6})$$

which has a modulus

$$|U_{0a}(D, r)| = \sqrt{\frac{\pi D}{\sqrt{2}r}} \frac{e^{-(r-D)^2/2}}{r - D}. \quad (\text{B.7})$$

This modulus is accurate to within several tens of per cent in the regime of interest for diffraction, $D \sim 4$ and $R \gtrsim 6$ (in units of b). By inserting this approximation into Eqs. (3.11) and (3.7) for diffraction losses in the clipping approximation, we deduce that the diffraction losses for mesa beams scale with beam-spot radius D as

$$\mathcal{L}_{\text{clip}}^{\text{mesa}} \propto \exp(2RD/b^2). \quad (\text{B.8})$$

Here R is the mirror radius and we have restored the diffraction lengthscale b , which was set to zero throughout this appendix.

B.2 Sapphire material parameters: Notation and values

While sapphire is not an isotropic material, it can be reasonably approximated by isotropic elastic and thermodynamic properties. In this paper, we therefore treat sapphire as an isotropic medium, with the following specific values for physical parameters:¹

$$E = 4 \times 10^{12} \text{erg/cm}^3 \quad (\text{B.9})$$

$$\rho = 4.0 \text{ g/cm}^3 \quad (\text{B.10})$$

$$\sigma = 0.28 \quad (\text{B.11})$$

$$\alpha = 5.5 \times 10^{-6} \text{ K}^{-1} \quad (\text{B.12})$$

$$\kappa = 3.3 \times 10^6 \text{ ergs cm}^{-2} \text{ K}^{-1} \text{ s}^{-1} \quad (\text{B.13})$$

$$C_V = 7.7 \times 10^6 \text{ cm s}^{-2} \text{ K}^{-1} \quad (\text{B.14})$$

¹These numbers are taken directly from the advanced LIGO project book [1], and the advanced LIGO summary web page [3].

B.3 General description of mesa beams

In Section 4.2.1 we only briefly describe the form of mesa beams evaluated at the surfaces of an arm cavity's two mirrors. In this appendix, we provide a much more thorough treatment of mesa beams.

B.3.1 Constructing mesa beams

Mesa beams are constructed by averaging some gaussian beams (with waist size A_o and waist location z_w) over some disc of size D and then normalizing the result. Specifically, we construct mesa beams that propagate towards positive z by first averaging a (non-normalized) gaussian beam over a disc of radius D :

$$U(\vec{r}, z, D) = e^{ikz} \times \int_{r' < D} d^2r' U_g(\vec{r} - \vec{r}', z; \bar{z}, z_w) \quad (\text{B.15a})$$

$$U_g(\vec{r}, z; \bar{z}, z_w) = \exp \left[-\frac{r^2 \left[1 - i \frac{(z - z_w)}{\bar{z}} \right]}{2A_o \sqrt{1 + \frac{(z - z_w)^2}{\bar{z}^2}}} \right] \times e^{-i \tan^{-1}(z/\bar{z})} \quad (\text{B.15b})$$

[where A_o is a function of \bar{z} , given by $\bar{z} \equiv kA_o^2$, and where z_w is the location of the beam waist]. The result of the average is necessarily axisymmetric. We then normalize the result:

$$N^2(D, z) \equiv \int_0^\infty |U(r, z, D)|^2 2\pi r dr \quad (\text{B.15c})$$

$$u_{\text{mesa}}(\vec{r}, z, D) \equiv U(r, z, D)/N(D, z) e^{ikz} \quad (\text{B.15d})$$

To construct the associated Mexican-hat mirror which will reflect this mesa beam back into itself (propagating in the opposite direction) at location z_r , we require the Mexican-hat mirror height function h_{MH} to be continuous and to satisfy

$$u_{\text{mesa}}(r, z_r, D) \exp(-2ikh_{\text{MH}}) = u_{\text{mesa}}(r, z_r, D)^* \quad (\text{B.16})$$

for $u(\vec{r}, z_r, D)$ the values of a mesa beam at the mirror plane $z = z_r$. This expression is equivalent to MBI Eq. (2.13) [cf. Eq. (4.6) in this paper]. These requirement uniquely specify the Mexican-hat mirror shape h_{MH} .

B.3.1.1 Canonical mesa beams: mesa beams for symmetric cavities

The mesa beams presented in MBI and discussed henceforth in this paper — denoted *canonical mesa beams* — are assumed to have particularly special form. First, it is assumed that the cavity confining

the mesa beam is symmetric, with mirror 1 placed at $z = -L/2$ and mirror 2 placed at $z = L/2$; also it is assumed that the waist location, z_w , is placed precisely between them (i.e. $z_w = 0$). Second, it is assumed that the gaussian used to construct the mesa beam is the *minimal-diffraction* gaussian for the symmetric cavity. The minimal diffraction gaussian has $\bar{z} = L/2$ and $z_w = 0$, implying that the beam width is $b = \sqrt{L/k}$ at the mirror surfaces and $b/\sqrt{2}$ at the beam waist at $z = 0$.

B.3.2 Canonical mesa beams at the mirror surfaces

In this paper and in MBI, we only need to know the form of canonical mesa beams at the mirror surfaces. In this case, since the beam is canonical (i.e. has symmetric cavities and minimal diffraction gaussians, so $z_w = 0$ and $\bar{z} = L/2$, implying $A_o = b/\sqrt{2}$) and since we evaluate it at the mirror surfaces $z = \mp L/2$, we find that Eq.(B.15b) simplifies to

$$U_{g,\min}(\vec{r}, z = \mp L/2) = \exp\left[-\frac{r^2 [1 \pm i]}{2b^2}\right] e^{\pm i\pi/4}. \quad (\text{B.17})$$

Inserting this expression into the definition (B.15) of general mesa beams, we find that the mesa beam form at mirror 1 (i.e. the mirror at $z = -L/2$) is given in terms of the construction described in Eq. (4.5). [For simplicity, in Eq. (4.5) we omit the overall phase factor $\exp(i\pi/4)$.]

B.4 Converting the fluctuation-dissipation model problem to a static model problem at low frequencies

In Sec. 4.2.2, we claim that at low frequencies, the elastic response of a cylinder to an oscillating pressure profile can be reconstructed to a good approximation using the static response of a cylinder to the same pressure profile, in an appropriate accelerating frame.

Briefly speaking, a relationship between the dynamic and static problems exists because, when the oscillations are sufficiently slow, the effect of dynamical terms in the elastic equations of motions can be neglected. These physical considerations have been discussed at greater length elsewhere, cf. Sec. II of LT [4].

In this appendix, we demonstrate in more technical detail precisely how to establish the desired (approximate) relationship. We assume only basic familiarity with elasticity, on the level of Blandford and Thorne [5].

B.4.1 General quasistatic approach for the elastic response to an oscillating surface stress

We wish to solve the elasticity equations, which can be expressed as

$$\rho \partial_t^2 y_a + \nabla^b T_{ab} = 0, \quad (\text{B.18})$$

for a test-mass mirror subject to an oscillating surface stress on its inner face (i.e. on the surface $z = 0$, where the optic axis of the advanced LIGO arm cavity is the z axis),

$$T_{az}(t, r, z = 0) \equiv \bar{T}_{az}(r, z = 0) \cos(\omega t), \quad (\text{B.19})$$

and otherwise subject to no other stresses. [Here, T_{ab} is given in terms of y_a by Eq. (4.10b).]

B.4.1.1 Step 1: Express the problem in the accelerated frame

We can better understand the response of the mirror substrate if we go into a frame comoving with the center of mass of the test mass mirror. The test mass experiences a net force $F_a \cos(\omega t)$, determined by the surface stress:

$$F_a \equiv -\cos(\omega t) \int d^2 r \bar{T}_{az}(\vec{r}, z = 0). \quad (\text{B.20})$$

In response to this net force, the mirror center of mass R_{cm} accelerates:

$$\vec{R}_{\text{cm}} = -\frac{F_a}{M\omega^2} \cos(\omega t). \quad (\text{B.21})$$

Therefore, to go to the comoving frame, we perform the following transformation: (i) define the comoving displacement field $y'_a = y_a - R_{\text{cm}}$:

$$y'_a \equiv y_a + \frac{F_a}{M\omega^2} \cos(\omega t); \quad (\text{B.22})$$

(ii) define the comoving stress-energy tensor as $T'_{ab} = T_{ab}(y')$. In terms of these two new quantities, the elasticity equations in the accelerated (primed) frame are

$$\rho \partial_t^2 y'_a + \nabla^b T'_{ab} = -\rho \frac{\vec{F}_a}{M} \cos(\omega t). \quad (\text{B.23})$$

The right hand side is simply the inertial force associated with working in an accelerated frame.

The boundary conditions remain unchanged.

B.4.1.2 Step 2: Factor out all sinusoidal dependence

In the accelerated frame, we now assume all quantities oscillate sinusoidally in response to the sinusoidally-oscillating pressure profile and inertial force:

$$\begin{aligned} y'_a &= \bar{y}_a \cos(\omega t) , \\ T'_{ab} &= \bar{T}_{ab} \cos(\omega t) . \end{aligned}$$

Substituting these expressions into Eq. (B.23), we find we can reconstruct a solution to the dynamic solution by solving the static partial differential equation

$$-\rho\omega^2\bar{y}_a + \nabla^b\bar{T}_{ab} = -\rho\frac{\vec{F}_a}{M} \quad (\text{B.24})$$

subject to force-free boundary conditions on all surfaces except the top surface, which is subjected to a *constant* pressure profile $\bar{T}_{az}(\vec{r}, z = 0)$.

B.4.1.3 Step 3: Approximate the problem as static in the accelerated frame

Finally, at sufficiently low frequencies (i.e. frequencies so low that sound crosses the cylinder many times within one period, as is the case for advanced LIGO; cf. notes 2 and 3), the first term in the accelerated-frame elasticity equations [i.e. the term $\propto y'_a\omega^2$ in Eq. (B.23)] can be neglected (cf., e.g., LT Sec. II). The remaining problem [i.e. Eq. (4.10)] can be interpreted precisely as equations for the static elastic response of a solid to an imposed pressure profile.

B.4.1.4 Summary

To summarize, then, we find an approximate solution for the elastic response of a solid to an imposed surface stress T_{az} at $z = 0$ by

$$y_a = \cos(\omega t) \left[-\rho\frac{\vec{F}_a}{M\omega^2} + \bar{y}_a \right] \quad (\text{B.25a})$$

where \bar{y}_a is obtained as a solution to the static elastic equations in an accelerated frame

$$\nabla^b\bar{T}_{ab} = -\rho\frac{\vec{F}_a}{M} \quad (\text{B.25b})$$

subject to the effective static surface stress \bar{T}_{az} and bulk acceleration \vec{F}_a/M given by

$$T_{az} = \bar{T}_{az} \cos(\omega t) \quad (\text{B.25c})$$

$$\vec{F}_a \equiv -\int d^2r \bar{T}_{az}(\vec{r}, z = 0) \quad (\text{B.25d})$$

To summarize our conclusions on physical grounds, because the elastic response occurs much more rapidly than the surface pressure profile changes (i.e. because the sound crossing time is much shorter than the oscillation period of the imposed force), we can effectively treat the elastic response as instantaneous. The elastic solid moves slowly through a sequence of static configurations.

This quasistatic approximation, however, must be performed with care. If one neglects the dynamical terms entirely, as other authors working on this and related subjects have done ², then one finds elasticity equations without the bulk acceleration term. That equation is inconsistent with the static boundary conditions we impose.

B.4.2 Using the quasistatic elastic solution to simplify our thermoelastic noise integral

We can apply the quasistatic elastic solution we just developed [Eq. (B.25)] to find the thermoelastic integral I_A [Eq. (4.9)] associated with our specific elastic model problem, where the surface stress imposed has (i) $T_{az} = 0$ unless $a = z$ and (ii) $T_{zz} = -\cos(\omega t)F_o P(r)$ [Eq. (4.10c)]. Since the accelerated-frame transformation does not change Θ (i.e. $\Theta = \Theta'$) and since $\langle \cos^2(\omega t) \rangle = 1/2$, we conclude that

$$I_A = \frac{2}{F_o^2} \int d^3r \langle |\nabla\Theta|^2 \rangle \approx \frac{1}{F_o^2} \int d^3r |\nabla\bar{\Theta}|^2, \quad (\text{B.26})$$

where the approximation neglects terms which are small in the quasistatic limit (cf. notes 3 and 2).

B.5 On the completeness of basis states for an arm cavity with two identical, infinite mirrors

In this appendix, we prove that the eigenproblem for two identical infinite mirrors [Eq. (4.13)] admits a complete set of orthogonal eigensolutions. In the text, we use the resulting eigensolutions as a basis for building perturbative expansions.

Our proof relies on the observation that the free paraxial propagator [Eq. (4.11a)] happens to be in one-to-one correspondence with the propagator for a quantum-mechanical free particle in two dimensions, cf. Eq. (3.3) of [6]. In particular, we can represent G_L [Eq. (4.11a)] as

$$G_L = e^{iHL} \quad (\text{B.27})$$

²Because of the fluctuation-dissipation theorem, many other authors working on thermal and thermoelastic noise (e.g., Cerdonio and Conti [19]; Liu and Thorne [4]; ...) must solve a similar or identical elastic problem to deduce the effects of noise. And most choose to approach it using a similar quasistatic approximation, ignoring any dynamical effects. However, most make the quasistatic approximation before they go to the accelerated frame, rather than after. Only Liu and Thorne have correctly accounted for the effects of acceleration, though they add those effects in by hand later.

for H a hermitian operator.

Now, when the two mirrors are identical (i.e. $G_1 = G_2$) and infinite (i.e. $T_1 = T_2 = \mathbf{1}$), we can represent the eigenequation (4.13) in terms of the alternative variable $u' = G_1^{-1/2}u$:

$$\begin{aligned} \eta e^{2ikL} u' &= \left[G_1^{-1/2} e^{iHL} G_1^{1/2} \right]^2 u' \\ &= \left[e^{i(G_1^{-1/2} H G_1^{1/2})L} \right]^2 u' . \end{aligned}$$

[Since G_1 is a particularly trivial operator — it merely multiplies the value of $\psi(\vec{r})$ by a position-dependent phase — we have no difficulty in constructing its square root $G_1^{1/2}$.] The operator in parentheses ($G_1^{-1/2} H G_1^{1/2}$) is simply a unitary transformation of a hermitian operator, and is therefore also hermitian. Therefore, it admits a complete set of orthogonal states. Using those complete states, we diagonalize the above equation for a complete set of solutions u_k . These solutions then map, via the unitary transformation $G_1^{1/2}$, to a complete set of orthogonal eigensolutions to Eq. (4.13).

B.6 Second-order perturbation theory for eigenstates of an individual arm cavity

In this appendix, we develop second-order perturbation expansions which relate changes in shape of one specific mirror of a symmetric cavity (the ETM) to changes in the resonant states in general (and to changes of the the ground state in particular) of that cavity.

Since the precise form of the eigenequation depends on how one chooses to represent the state (i.e. at what plane, in what direction), many alternative and equivalent perturbation expansions can be derived³. The form presented in this appendix is that developed by O’Shaughnessy; Sergey Strigin and Sergey Vyatchanin performed their calculations using an independently-derived (but provably equivalent) approach.

As in the text, in this section we make heavy use of standard quantum-mechanics operator notation for states and inner products, a notation described briefly in Section 4.1.3 (cf., e.g., [7, 8, 9]).

B.6.1 Preliminaries: Setting up notation for the expansion

We study the effect of changes in height of mirror 2 on the solutions to the eigenequation (4.13) in the case of a *symmetric* cavity with *infinite* mirrors. For clarity and simplicity of notation, we

³The relationship between these expressions need not be transparent. For example, one can represent the eigenequation using states represented “halfway” through a reflection of one of the mirrors (a choice which conveniently renders the resulting equation always perfectly symmetric). The transformation between our representation and this one depends on the mirror height. Thus, when the mirror heights are perturbed, the relationship between these two representations involves a unitary transformation that depends on the perturbation.

redefine the eigenequation problem we perturb still further, into the following expression:

$$\eta |\psi\rangle = GTG |\psi\rangle . \quad (\text{B.28a})$$

$$G \equiv e^{-ikL} G_1 G_+ , \quad (\text{B.28b})$$

$$T \equiv e^{-i\delta_2} \equiv 1 + \delta T , \quad (\text{B.28c})$$

$$\delta_2 \equiv 2k\delta h_2 . \quad (\text{B.28d})$$

Here, δh_2 is the change in the (inward-pointing) height of the far mirror (mirror 2) and $\delta T \equiv T - 1$. [The operator T , used only in this section, should not be confused with the truncation operators defined in Eq. (4.11).]

B.6.2 Perturbation theory expansion, expressed using operators

We construct a perturbation theory expansion by expanding both sides of the eigenequation Eq. (B.28) in series, giving

$$\begin{aligned} (\eta_o + \varepsilon\eta_1 + \varepsilon^2\eta_2 + \dots) (\psi_o + \varepsilon\psi_1 + \dots) \\ = G(1 + \delta T)G (\psi_o + \varepsilon\psi_1 + \varepsilon^2\psi_2) , \end{aligned} \quad (\text{B.29})$$

and then matching orders on both sides; in this expression, ε is a formal perturbation parameter, added to rescale the change in the reflection operator: $\delta T(\varepsilon) = \varepsilon\delta T$. When we work out order matching, we find (for P a projection orthogonal to the ground state ψ_o , also denoted $|0\rangle$):

$$\eta = \eta_o + \langle 0 | \delta\tilde{T} | 0 \rangle \quad (\text{B.30a})$$

$$+ \left\langle 0 \left| \delta\tilde{T} P \frac{1}{\eta_o - G^2} P \delta\tilde{T} \right| 0 \right\rangle ,$$

$$\psi = |0\rangle + \frac{1}{\eta_o - G^2} P \delta\tilde{T} |0\rangle \quad (\text{B.30b})$$

$$+ \frac{1}{\eta_o - G^2} P \left[\delta\tilde{T} - \langle 0 | \delta\tilde{T} | 0 \rangle \right] \frac{1}{\eta_o - G^2} P \delta\tilde{T} |0\rangle ,$$

where the operation \tilde{O} on an operator O is defined by

$$\tilde{O} \equiv GOG . \quad (\text{B.31})$$

We can furthermore substitute into the above expression the expansion

$$\delta\tilde{T} \approx G \left(-i\delta_2 - \frac{1}{2}\delta_2^2 \right) G = -i\tilde{\delta}_2 - \frac{1}{2}G\delta_2^2G$$

to give us the final form of the second-order expansion of the state

$$\begin{aligned}\eta &= \eta_o - i \langle 0 | \tilde{\delta}_2 | 0 \rangle - \frac{1}{2} \langle 0 | G \delta_2^2 G | 0 \rangle + \left\langle 0 \left| \tilde{\delta}_2 P \frac{1}{\eta_o - G^2} P \tilde{\delta}_2 \right| 0 \right\rangle + O(\delta_2^3) . \\ \psi &= |0\rangle - i \frac{1}{\eta_o - G^2} P \tilde{\delta}_2 |0\rangle - \frac{1}{2} \frac{1}{\eta_o - G^2} P G \delta_2^2 G |0\rangle - \frac{1}{\eta_o - G^2} P \left[\tilde{\delta}_2 - \langle 0 | \tilde{\delta}_2 | 0 \rangle \right] \frac{1}{\eta_o - G^2} P \tilde{\delta}_2 |0\rangle + O(\delta_2^3) .\end{aligned}$$

B.6.3 Perturbation theory expansion, expressed using basis states

As a practical matter, we compute the perturbation series expansion using not the operators themselves, but rather through a finite collection of matrix elements of the relevant operators relative to basis states. Therefore, we insert the identity operator, represented as a sum over all basis states (i.e. $\mathbf{1} = \sum_k |k\rangle \langle k|$), at several points in the above expression. Since the basis states are eigenvectors of the propagation operator G and since the operator G is unitary, we know

$$G |k\rangle = \lambda_k |k\rangle , \quad (\text{B.33a})$$

$$\langle k| G = \lambda_k^* \langle k| . \quad (\text{B.33b})$$

Therefore, we conclude that the eigenvalue changes as

$$\begin{aligned}\eta &= \eta_o \left(1 - i \langle 0 | \delta_2 | 0 \rangle \right. \\ &\quad \left. - \frac{1}{2} \langle 0 | \delta_2^2 | 0 \rangle - \sum_{k \neq 0} \frac{\eta_k |\langle 0 | \delta_2 | k \rangle|^2}{\eta_o - \eta_k} \right) + O(\delta_2^3)\end{aligned} \quad (\text{B.34a})$$

while the state changes according to the expansion

$$\begin{aligned}|\psi\rangle &= |0\rangle - i \sum_{k \neq 0} \frac{\lambda_o \lambda_k}{\eta_o - \eta_k} |k\rangle \langle k | \delta_2 | 0 \rangle \\ &+ \sum_{k \neq 0} |k\rangle \frac{\lambda_o \lambda_k}{\eta_o - \eta_k} \left[-\frac{1}{2} \langle k | \delta_2^2 | 0 \rangle + \frac{\eta_o}{\eta_o - \eta_k} \langle k | \delta_2 | 0 \rangle \langle 0 | \delta_2 | 0 \rangle - \sum_{p \neq 0} \frac{\eta_p}{\eta_o - \eta_p} \langle k | \delta_2 | p \rangle \langle p | \delta_2 | 0 \rangle \right] + O(\delta_2^3) .\end{aligned} \quad (\text{B.34b})$$

B.7 Perturbation theory for cavities bounded by two identical spherical mirrors

For cavities bounded by spherical mirrors – that is, mirrors with height function $h_{1,2} = r^2/2\mathcal{R}_{1,2}$ – the eigenfunctions are known and of simple, tractable Hermite-Gauss form. We can therefore perform perturbation-theory calculations analytically. In this appendix, we describe these basis states and their application to perturbation theory in greater detail.

B.7.1 Background: Notation and definitions

A spherical mirror of height $h(r)$ is uniquely characterized by its radius of curvature \mathcal{R} : $h = r^2/2\mathcal{R}$. We consider two identical such mirrors, placed symmetrically at $\pm L/2$. We drive this cavity so it has $kL/2\pi \gg 1$ wavelengths between the mirrors (so kL is to an excellent approximation independent of mirror shape).

B.7.1.1 Coordinate representation of eigenfunctions

The eigenfunctions for a cavity bounded by two identical spherical mirrors are known [10, 11, 12, 13, 14]:

$$\begin{aligned} \psi_{m,n} &= \sqrt{\frac{1}{2^{m+n}m!n!}} H_m\left(\frac{x}{A(z)}\right) H_n\left(\frac{y}{A(z)}\right) \\ &\times \frac{1}{\sqrt{\pi}A(z)} \exp\left[-\frac{r^2}{2A(z)^2}\left(1 - i\frac{z}{\bar{z}}\right)\right] \\ &\times \exp\left[-i(N+1)\tan^{-1}\frac{z}{\bar{z}} + ikz\right]. \end{aligned} \quad (\text{B.35})$$

Here H_m are Hermite polynomials, $A_o \equiv \sqrt{\bar{z}/k}$, $N \equiv m + n$, $g \equiv 1 - L/\mathcal{R}$, $\theta = \cos^{-1}g$, $k = (n\pi + \theta)/L$, and

$$A(z)^2 \equiv A_o^2\left(1 + (z/\bar{z})^2\right), \quad (\text{B.36a})$$

$$\bar{z} \equiv \frac{L}{2}\sqrt{\frac{1+g}{1-g}}. \quad (\text{B.36b})$$

In particular, at the mirror faces at $z_{\pm} = \pm L/2$, we have

$$z_{\pm}/\bar{z} = \pm\sqrt{\frac{1-g}{1+g}}, \quad (\text{B.37a})$$

$$\tan^{-1}(z_{\pm}/\bar{z}) = \pm\frac{1}{2}\cos^{-1}g = \pm\theta/2, \quad (\text{B.37b})$$

$$A(z_{\pm}) = \sqrt{\frac{L}{k}}\frac{1}{(1-g^2)^{1/4}}. \quad (\text{B.37c})$$

B.7.1.2 Basis as eigensolutions of cavity

After some algebra, one can verify these states $\psi_{m,n}$ are solutions to the round-trip eigenequation Eq. (4.13) [or equivalently Eq. (4.29)] with $h_1(r) = h_2(r) = r^2/2\mathcal{R}$, and

$$\lambda_{mn} = \sqrt{\eta_{mn}} = e^{i\theta(m+n)/2}. \quad (\text{B.38})$$

B.7.1.3 Establishing a quantum-mechanical correspondence

For technical reasons not discussed further here, the eigenfunctions of a cavity bounded by spherical mirrors correspond directly to the states of a 2-dimensional quantum-mechanical scalar nonrelativistic particle in a quadratic potential (i.e. a 2-d quantum simple harmonic oscillator, or SHO). Like the SHO, the states are highly degenerate. If we index the states using their cartesian symmetry properties [i.e. express them as a Hermite-Gauss basis, as in Eq. (B.35)], these quantum states of a SHO have the form $|m, n\rangle = \Psi_m(x)\Psi_n(y)$ for

$$\Psi_n(x) \equiv \frac{1}{[\sqrt{\pi}A]^{1/2}} \frac{1}{\sqrt{2^n n!}} H_n\left(\frac{x}{A}\right) e^{-x^2/2A^2} \quad (\text{B.39})$$

for A some length scale. Therefore, if we use $A = A(L/2)$, we find

$$\psi_{m,n}(x, y) = e^{i\Phi_N} \Psi_m(x)\Psi_n(y) \quad (\text{B.40a})$$

$$\Phi_N = (\text{independent of } N, x, y) + (N + 1)\theta/2 \quad (\text{B.40b})$$

B.7.1.4 Creation and annihilation operators for optical states

In quantum mechanics, many computations involving simple harmonic oscillator states can be rewritten in a more readable form by the use of creation and annihilation operators [7, 8, 9]. These simplifications occur because we can relate any excited state $|0\rangle$ to the ground state via the action of these operators:

$$|m\rangle_{\text{sho}} = \frac{(a_x^\dagger)^m}{\sqrt{m!}} |0\rangle_{\text{sho}} \quad (\text{B.41})$$

$$a_x^\dagger \equiv \frac{1}{\sqrt{2}} \left(\frac{x}{A} - A \frac{d}{dx} \right) \quad (\text{B.42})$$

In fact, because the overall phase factor $\exp i\Phi_N$ is independent of x and y , the same creation and annihilation techniques can be applied to manage optical states and their inner products in the Hermite-Gauss basis.

B.7.2 Tabulating useful matrix elements

When performing perturbation-theory calculations, we need the matrix elements of various operators (i.e. x and x^2) relative to the Hermite-Gauss basis. These elements can be expressed as follows:

$$\begin{aligned}
\langle m, n | \quad x \quad | \bar{m}, \bar{n} \rangle_{opt} & \quad (B.43a) \\
= e^{i(\bar{N}-N)\theta/2} \left\langle m, n \left| \frac{A(a_x + a_x^\dagger)}{\sqrt{2}} \right| \bar{m}, \bar{n} \right\rangle_{sho} \\
= \frac{A}{\sqrt{2}} \delta_{n\bar{n}} \left[e^{i\theta/2} \sqrt{\bar{m}} \delta_{m, \bar{m}-1} + e^{-i\theta/2} \sqrt{\bar{m}} \delta_{\bar{m}, m-1} \right]
\end{aligned}$$

$$\begin{aligned}
\langle m, n | \quad x^2 \quad | \bar{m}, \bar{n} \rangle_{opt} & \quad (B.43b) \\
= e^{i(\bar{N}-N)\theta/2} \left\langle m, n \left| \frac{A^2(a_x + a_x^\dagger)^2}{2} \right| \bar{m}, \bar{n} \right\rangle_{sho} \\
= \frac{A^2}{2} \delta_{n\bar{n}} \left[e^{i\theta} \sqrt{m(m+1)} \delta_{m, \bar{m}-2} \right. \\
\quad \left. + (2m+1) \delta_{m, \bar{m}} \right. \\
\quad \left. + e^{-i\theta} \sqrt{\bar{m}(\bar{m}+1)} \delta_{\bar{m}, m-2} \right]
\end{aligned}$$

In the above, we use the subscript *opt* to denote an optical state and *sho* to denote a SHO state.

B.7.3 Effects of tilt on beam state

If we tilt mirror 2 about the y axis, the mirror's height changes according to $\delta_2 = 2kx\phi$ [cf. Eqs. (4.62) and (4.31)]. Using that expression in the two perturbative expansion expressions, Eqs. (4.33b) and (4.33c), and taking account of the above tools, we can find explicit expressions for the first and second order changes in the beam state evaluated at mirror 1.

$$\left| \psi^{(1)} \right\rangle = -|1, 0\rangle (\phi\sqrt{Lk}) \frac{e^{-i\theta/2}}{\sqrt{2}} \frac{1}{(1-g^2)^{3/4}} \quad (B.44)$$

$$\left| \psi^{(2)} \right\rangle = -|2, 0\rangle (\phi\sqrt{kL})^2 \frac{e^{i\theta}}{(1-g)\sqrt{2(1-g^2)}} \quad (B.45)$$

In performing the above calculation, we have explicitly substituted in $A = A(z_\pm)$ from Eq. (B.37c) and $e^{i\theta} = g + i\sqrt{1-g^2}$ [derived from Eq. (B.37b)] when appropriate.

B.7.4 Effect of displacement on beam state

When the mirror is displaced through a distance d in the x direction, the mirror height is perturbed from $h(r) = r^2/2\mathcal{R}$ into

$$h'(r) = h(\vec{r} + s\hat{x}) = h(r) + \frac{xs}{\mathcal{R}} + \frac{s^2}{2\mathcal{R}}. \quad (\text{B.46})$$

We can in a straightforward manner insert $\delta h_{\text{disp}} = h' - h$ into the perturbative expansion for the state [Eq. (B.34b)]; note that δh has terms both first and second order in ψ ; and thereby deduce the appropriate first and second order changes in the optical state of the cavity.

We will not provide a comprehensive calculation here. Instead, we note that the first-order changes in height δh due to tilt and displacement are identical, involving merely a change in factor:

$$\phi\sqrt{Lk} \rightarrow (s/b)(1-g). \quad (\text{B.47})$$

Therefore, we deduce the first-order effect of a displacement of mirror 2 is

$$|\psi^{(1)}\rangle = -|1,0\rangle (d/b) \frac{e^{-i\theta/2}}{\sqrt{2}} \frac{(1-g)}{(1-g^2)^{3/4}} \quad (\text{B.48})$$

B.8 Postprocessing performed on finite-element computations used in thermoelastic noise integrals

We used a commercial finite-element code, FEMLAB [15], to solve the fluctuation-dissipation-motivated axisymmetric elastic model problem presented in Eq. (4.10). This code expresses its solution for y^a in terms of a piecewise linear function. Unfortunately, the thermoelastic noise integral depends on second derivatives of y_a . Therefore, even though the FEMLAB program includes many useful tools to evaluate physically interesting quantities, we had to evaluate the thermoelastic noise integral ourselves. We therefore sent the FEMLAB output for y^a to a two-dimensional MATLAB array, then used our own MATLAB routines to generate expressions for $|\nabla\theta|^2$ and thus I_A .

The two groups performing this computation (Caltech and Moscow) each developed relatively independent postprocessing codes. In this appendix, we describe the Caltech postprocessing code.

B.8.1 Caltech postprocessing code

The Caltech group cobbled together a rough set of MATLAB routines to convert the finite-element representation of the response y^a to useful values for I_A . While this code contains more than a few kludges used to get it to work, we found the combination of the finite-element solution and this postprocessing routine gave answers for I_A that agreed well with and converged to the values of I_A

for known, analytic solutions (i.e. the LT cylindrical solutions).

B.8.1.1 Postprocessing technique

1. *Preliminaries:* We used the FEMLAB finite-element code, with some number N elements, to find the elastic response of a frustum of front radius R_1 , back radius R_2 , and thickness H . We increased N until we felt comfortable with the resulting response curve. The FEMLAB code worked in cylindrical coordinates (i.e. r, z, ϕ), and assumed axisymmetry.
2. *Output to rectangular grid:* The code allowed us to extract the values of the displacement vector (components $U = y_r$ and $V = y_z$) at any point. We used this ability to obtain values of U and V on a rectangular grid. The code naturally provided smoothly extrapolated values of U and V when we asked for a point outside the volume simulated.

The total number of points on the rectangular grid was chosen to be comparable to N . Specifically, the numbers of points in the r and z directions were chosen as $\sqrt{NR_1/H}$ and $\sqrt{NH/R_1}$, respectively.

3. *Select cutoff region for grid:* Because we used a grid containing too large a region, we must provide a cutoff filter to select only those gridpoints which contain physical values. Furthermore, as a practical matter, the FEMLAB code gave odd results⁴ when we evaluated the response at the edges of the computational domain. We therefore chose a cutoff filter which eliminated all exterior and “near-to-the-boundary” points.

This filter was applied repeatedly (i.e. during each derivative process). Because the filter was such a universal ingredient to each subsequent action, we will not mention each occurrence on which it is used in the following postprocessing.

4. *Compute expansion:* Using the relationship between Θ and y^a , namely

$$\Theta = \partial_z V + \partial_r U + U/r \tag{B.49}$$

we compute θ . Derivatives are formed as centered differences, which are then interpolated (or, for the endpoints, extrapolated) back to the gridpoints. The values of U/r at $r = 0$ are found by extrapolating the values of U/r for $r \neq 0$.

[To circumvent problems that arose due to dividing by small numbers near $r = 0$, we typically erased the values for θ we obtained on four gridlines near $r = 0$ and replaced them by extrapolated values of the region immediately outside.]

⁴In other words, we found NaN (i.e. “not a number”) answers from our code.

5. *Compute $|\nabla\Theta|^2$* : Next, we computed the two derivatives $\partial_r U$ and $\partial_z V$ and used them to form

$$|\nabla\Theta|^2 = (\partial_r U)^2 + (\partial_z V)^2. \quad (\text{B.50})$$

As before, derivatives were evaluated with centered differences which were interpolated to gridpoints.

6. *Compute I_A* : Finally, we used our own two-dimensional (Simpson's rule) integrator to evaluate the two-dimensional integral

$$I_A = \int r dr dz |\nabla\Theta|^2, \quad (\text{B.51})$$

which provides an explicit form for Eq. (4.9).

B.8.1.2 Testing the result

Our postprocessing code is far from polished. To insure the cutoff filter is operating properly and to otherwise guarantee that the results of the postprocessing code seem physically plausible, we usually plotted $|\nabla\theta|^2$ to verify that the integrand is indeed a well-behaved (i.e. smooth-looking) function.

Because we had a serious limitation on the number of points we could practically employ in a reasonable amount of desktop computing time ($\sim \text{few} \times 10^4$), we did not perform systematic convergence testing. However, what testing we did, corroborated by comparisons between our numerical method and exact analytic solutions, suggests our results are relatively accurate. More critically, our computations agreed well with independent computations performed by Sergey Strigin and Sergey Vyatchanin.

B.9 Thermoelastic noise of half-infinite mirrors

To evaluate the thermoelastic noise associated with a given beam shape $P(r)$, we must evaluate the integral I_A [Eq. (4.9)] given the solution y^a to a model elasticity problem [Eq. (4.10)]. As discussed in Sec. 4.5.1.3, if the mirror is sufficiently large compared to the beam shape $P(r)$, we can effectively treat the mirror as half-infinite (i.e. filling the whole volume $z < 0$) in the elasticity problem. In this case, the bulk acceleration term in the elasticity problem drops out [i.e. $V_A \rightarrow \infty$ in Eq. (4.10a)] and we seek only the elastic response of a half-infinite medium to an imposed surface stress. This last problem has been discussed extensively in the literature — cf., e.g., [16, 17] — and there exist simple fourier-based computational techniques to generate and manage solutions. We apply these known solutions from the literature to evaluate the thermoelastic integral I_A .

B.9.1 Elastic solutions for the expansion (θ)

In the case of half-infinite mirrors, the response y^a to the imposed pressure profile $P(r)$ can be found in the literature [cf. Eqs. (8.18) and (8.19) of Landau and Lifshitz's book on elasticity [16], where, however, the half-infinite volume is chosen *above* the $z = 0$ plane rather than below; see also Nakagawa et al. [17], especially their Appendix A]. These expressions allow us to explicitly relate the expansion θ to the imposed pressure profile $P(r)$:

$$\Theta(\vec{r}, z) = \int G^{(\Theta)}(\vec{r}, z; r') P(r') d^2 r' \quad (\text{B.52a})$$

$$G^{(\Theta)}(\vec{r}, z; \vec{r}_o) = -\frac{(1+\sigma)(1-2\sigma)zH(-z)}{2\pi E |(\vec{r}-\vec{r}_o)^2 + z^2|^{3/2}} \quad (\text{B.52b})$$

where $H(x)$ is a step function which is 1 when $x > 0$ and 0 otherwise.

Because we have complete transverse translation symmetry, we can make our results more tractable by fourier-transforming in the transverse dimensions:

$$\tilde{\Theta}(K, z) \equiv \int e^{-i\vec{K}\cdot\vec{R}} \Theta(R, z) d^2 R, \quad (\text{B.53})$$

$$\tilde{P}(\vec{K}) \equiv \int e^{-i\vec{K}\cdot\vec{R}} P(R) d^2 R. \quad (\text{B.54})$$

For example, the convolution relating light intensity profile to the associated elastic response, Eq. (B.52a), can be re-expressed as

$$\tilde{\Theta}(K, z) = G^{(\Theta)}(z, \vec{K}) \tilde{P}(\vec{K}) \quad (\text{B.55})$$

$$\tilde{G}^{(\Theta)}(z, \vec{K}) = -\frac{(1+\sigma)(1-2\sigma)}{2\pi E} e^{-|Kz|}. \quad (\text{B.56})$$

B.9.2 Thermoelastic integral I_A

Inserting the solution discussed above into Eq. (4.9) and using fourier-transform techniques to simplify the resulting integral, we find

$$I_A = \left(\frac{(1+\sigma)(1-2\sigma)}{2\pi E} \right)^2 \int d^2 \vec{K} |K| |\tilde{P}(K)|^2. \quad (\text{B.57})$$

B.10 Configurations with nonidentical frustum mirrors and mesa-like beams

Physically the ETM and ITM need not be identical: while the ITM back face size is constrained to be above a certain radius (for a given mesa-beam D) by diffraction losses, the ETM back face

has no size restriction. Therefore, if we allow different mirrors, we introduce more parameters and therefore more possibilities for finding a mirror and beam configuration with very low noise.

Unfortunately, strictly speaking the mesa beams presented in this paper are designed specifically for identical cavities. Therefore, in the main text, when we explored how thermoelastic noise varied with mirror dimensions and shape, we restricted attention to cavities bounded by *identical* mirrors.

Nonetheless, we expect that *generic* flat-topped beams will always closely resemble the mesa-beam profile, with roughly the mesa-beam diffraction losses. Therefore, this appendix generalizes the computations of Sec. 4.6 to include distinct ETM and ITM mirrors.

More specifically, we (i) assume that, for any D_1 and D_2 , there is some flat-topped beam whose intensity very closely mimics the mesa-beam intensity profile $|u_{\text{mesa}}(r, D_1)|^2$ at mirror 1 and $|u_{\text{mesa}}(r, D_2)|^2$ at mirror 2; (ii) select D_1 , D_2 , R_1 , and R_2 so the clipping approximation to the diffraction losses is 10 ppm at each mirror face (Sec. 4.2.4); and (iii) use the resulting beam intensity profiles to compute the thermoelastic noise integrals I_1 and I_2 (Sec. 4.5.1). The power spectrum of thermoelastic noise for an interferometer with two identical arm cavities, each of which has an ITM like mirror 1 and an ETM mirror 2 can be expressed relative to the baseline [cf. Eq. (4.7)] by

$$S_h/S_h^{\text{BL}} = \frac{I_{\text{ITM}} + I_{\text{ETM}}}{2I_{\text{BL}}}. \quad (\text{B.58})$$

B.10.1 Thermoelastic noise integral for ETM designs

The ETM mirror has only two design constraints: (i) its mass must be 40 kg and (ii) the diffraction losses off its front face must be less than or equal to 10 ppm. Table B.1 shows the value of the thermoelastic noise integral for various values of R_{p1} , R_{p2} , and D , chosen to satisfy the various constraints (mirror mass 40 kg; diffraction losses off the front face of 10 ppm).

While we have complete freedom to adjust the back face size of the ETM, from the table, we see the back face size matters little so long as it is less than $R_{p2} \lesssim 9$ cm. By fitting a quadratic to $(D/b, I/I_{\text{BL}})$ [for only those ETM entries with $R_{p2} \lesssim 9$ cm], we estimate the optimal ETM dimensions; our results are tabulated in Table B.2.

In Table B.2 we also combine the optimal ETM with the optimal ITM and use Eq. (B.58) to express the thermoelastic noise of a nonidentical-mirror configuration relative to the baseline. When the ETM can be different from the ITM, the optimal thermoelastic noise is significantly reduced (because the ETM inner radius can be larger).

Table B.1: The thermoelastic integral I for a frustum end test mass (ETM) and a Mexican-Hat Beam, in units of $I_{\text{BL}} = 2.57 \times 10^{-28} \text{s}^4 \text{g}^{-2} \text{cm}^{-1}$. The values of I/I_{BL} are estimated to be accurate to within one per cent.

R_1	$R_{p1}[\text{cm}]$	$R_{p2}[\text{cm}]$	$H[\text{cm}]$	D/b	I/I_{BL}	$\mathcal{L}_0[\text{ppm}]^1$
$R_{p1} - 8\text{mm}$	17.11	2.00	28.85	4.00	0.175	10
$R_{p1} - 8\text{mm}$	17.11	8.00	19.34	4.00	0.175	10
$R_{p1} - 8\text{mm}$	17.11	12.00	14.87	4.00	0.192	10
$R_{p1} - 8\text{mm}$	19.58	3.00	21.17	5.00	0.134	10
$R_{p1} - 8\text{mm}$	19.58	9.00	14.91	5.00	0.133	10
$R_{p1} - 8\text{mm}$	19.58	13.00	11.83	5.00	0.180	10
$R_{p1} - 8\text{mm}$	19.98	5.00	18.22	5.15	0.133	10
$R_{p1} - 8\text{mm}$	22.08	3.00	16.97	6.00	0.156	10
$R_{p1} - 8\text{mm}$	22.08	9.00	12.45	6.00	0.157	10
$R_{p1} - 8\text{mm}$	22.08	15.00	9.15	6.00	0.310	10
R_{p1}	13.94	5.00	32.93	3.00	0.333	10
R_{p1}	13.94	7.50	26.80	3.00	0.340	10
R_{p1}	13.94	10.00	21.96	3.00	0.345	10
R_{p1}	16.37	2.00	31.34	4.00	0.161	10
R_{p1}	16.37	4.00	27.33	4.00	0.160	10
R_{p1}	16.37	6.00	23.74	4.00	0.160	10
R_{p1}	16.37	8.00	20.63	4.00	0.160	10
R_{p1}	16.37	10.00	17.96	4.00	0.162	10
R_{p1}	16.37	12.00	15.70	4.00	0.173	10
R_{p1}	16.37	14.00	13.78	4.00	0.207	10
R_{p1}	18.85	1.00	25.45	5.00	0.112	10
R_{p1}	18.85	3.00	22.69	5.00	0.112	10
R_{p1}	18.85	5.00	20.12	5.00	0.112	10
R_{p1}	18.85	7.00	17.81	5.00	0.111	10
R_{p1}	18.85	9.00	15.56	5.00	0.113	10
R_{p1}	18.85	11.00	13.97	5.00	0.119	10
R_{p1}	18.85	13.00	12.41	5.00	0.144	10
R_{p1}	19.86	5.00	18.42	5.42	0.107	10
R_{p1}	21.36	3.00	18.04	6.00	0.116	10
R_{p1}	21.36	5.00	16.24	6.00	0.115	10
R_{p1}	21.36	7.00	15.39	6.00	0.115	10
R_{p1}	21.36	9.00	14.58	6.00	0.116	10
R_{p1}	21.36	12.00	11.15	6.00	0.133	10
R_{p1}	21.36	15.00	9.55	6.00	0.224	10

Table B.2: Optimized test-mass and light beam configurations, their thermoelastic noise compared to the baseline. [A subset of this table appears as Table I in MBI [18], and as Table 4.4 in this paper.]

Test Masses $\{R_{p1}, R_{p2}; H\}$	Beam Shape	$\left(\frac{S_h}{S_{BL}^h}\right)_{TE}$
BL: cylinders, $R = R_p - 8\text{mm}$ {15.7, 15.7; 13.0}	BL: Gaussian $r_o = 4.23\text{cm}$	1.000
BL: cylinders, $R = R_p - 8\text{mm}$ {15.7, 15.7; 13.0}	mesa $D/b = 3.73$	0.364
identical frustums, $R = R_p - 8\text{mm}$ {17.11, 12.88, 14.06}	mesa $D/b = 4.00$	0.207
different cones, $R = R_p - 8\text{mm}$ ITM{17.42, 13.18, 13.51}	MH $D/b = 4.00$	
ETM{19.96, $\lesssim 9.$, $\gtrsim 14.$ }	$D/b = 5.15$	0.170
BL: cylinders, $R = R_p$ {15.7, 15.7; 13.0}	Gaussian $r_o = 4.49\text{cm}$	0.856
BL: cylinders, $R = R_p$ {15.7, 15.7; 13.0}	mesa $D/b = 3.73$	0.290
identical frustums, $R = R_p$ {17.29, 13.04, 13.75}	mesa $D/b = 4.39$	0.162
different cones, $R = R_p$ ITM{17.29, 13.04, 13.75}	MH $D/b = 4.39$	
ETM{19.91, $\lesssim 9.$, $\gtrsim 14.$ }	$D/b = 5.42$	0.135

B.10.2 Aside: Constructing mesa-like beams appropriate to nonidentical mirrors

In this appendix, we assume that for any D_1 and D_2 , there is some flat-topped beam whose intensity closely mimics the mesa-beam intensity profile $|u_{\text{mesa}}(r, D_1)|^2$ at mirror 1 and $|u_{\text{mesa}}(r, D_2)|^2$ at mirror 2. In fact, we may explicitly construct such a beam through a generalization of the mesa-beam proposal (though the authors have not numerically explored this possibility).

The mesa beams constructed in this paper were built by averaging minimal gaussian beams which all travelled *parallel* to the optic axis [cf. Eq. (4.5)]. More generally, we can construct mesa-like beams by averaging gaussians which travel at a position-dependent angle to the optic axis; for example, at mirror 1, we could choose the beams to have angle distribution

$$\phi(r) = \frac{R_2 - R_1}{L} \frac{r}{R_1}. \quad (\text{B.59})$$

The resulting mesa-type beams will have power distributions at the two mirrors which closely resemble the power distributions of mesa beams: (i) they will be flat over a large central region, of size D_1 at mirror 1 and size $D_2 \approx D_1(R_2/R_1)$ at mirror 2; and (ii) they will fall off rapidly outside this region, with a falloff rate nearly the same as that of the minimal gaussian.

Bibliography

- [1] LIGO II Conceptual Project Book, available at URL: <http://www.ligo.caltech.edu/docs/M/M990288-A1.pdf>
- [2] The LIGO Scientific Collaboration (LSC) maintains a website describing plans for the first-generation upgrade: <http://www.ligo.caltech.edu/~ligo2/>
- [3] Peter Fritschel maintained a website which provided a succinct summary of an older advanced LIGO design, including material parameters: <http://www.ligo.caltech.edu/~ligo2/scripts/l2refdes.htm> (quoted 28-June-2003)
- [4] Y. T. Liu and K. Thorne, *Phys. Rev. D* **62**, 1222002.
- [5] R. D. Blandford and K. S. Thorne, *Applications of classical physics*, unpublished. Draft versions are available at URL <http://www.pma.caltech.edu/Courses/ph136/yr2002/index.html>
- [6] R. P. Feynman and A. H. Hibbs, *Quantum mechanics and path integrals*, (McGraw-Hill, 1965).
- [7] P. A. M. Dirac, *The principles of quantum mechanics*, 4th edition (Oxford University Press, 1958)
- [8] L. D. Landau and E. M. Lifshitz, *Quantum mechanics*, 3rd edition (Pergamon Press) 1977
- [9] J. S. Townsend, *A modern approach to quantum mechanics*, (McGraw-Hill, New York) 1992
- [10] A. Siegman, *An Introduction to Lasers and Masers* (McGraw-Hill, 1971).
- [11] A. Siegman, *Lasers* (University Science, Mill Valley, 1986); errata URL: http://www-ee.stanford.edu/~siegman/lasers_book_errata.txt
- [12] A. Yariv, *Quantum electronics*, 3rd edition (John Wiley & Sons, 1989).
- [13] L. Ronchi, *Optical resonators*, in *Laser physics handbook*, vol. 1, F. T. Arecchi and E. O. Schulz-Dubois, eds. (North-Holland Publishing, 1972)
- [14] F. M. Barnes, ed., *Laser Theory*, (IEEE Press, 1972)
- [15] FEMLAB, described at <http://www.comsol.com>

- [16] L. Landau and E. M. Lifshitz, *Theory of elasticity*, 2nd edition (Pergammon Press, 1970).
- [17] N. Nakagawa et al., Phys Rev. D **65**, 102001.
- [18] E. Ambrosio, R. O'Shaughnessy, S. Strigin, K. S. Thorne, and S. Vyatchanin, to be submitted to Phys. Rev. D [i.e. Chapter 3 of this paper]
- [19] M. Cerdonio et al, Phys Rev. D **63**, 082003 (2001).

Appendix C

Appendix for numerical relativity

C.1 Useful identities used in the text

C.1.1 An alternative approach to the group velocity

The eigenequation which defines the natural polarization spaces B_j and their associated eigenvalues ω_j ,

$$A^a p_a v_{j,\alpha} = \omega_j v_{j,\alpha}$$

for each $v_{j,\alpha} \in B_j$ (cf. Sec. 5.2.1), may be alternatively expressed as

$$A^a(x) p_a P_j(x, p) = \omega_j(x, p) P_j(x, p) . \quad (\text{C.1})$$

for P_j the projection operator to the j th eigenspace of $A^a p_a$. Differentiate this expression relative to p_b , then apply P_j from the left, to find

$$P_j A^a P_j = \frac{\partial \omega_j}{\partial p^a} P_j = V_j^a P_j . \quad (\text{C.2})$$

Equivalently, if $v_{j,\alpha}$ are a collection of basis vectors for the j th eigenspace which are orthonormal relative to the inner product generated by S ,

$$(v_{j,\beta}, S A^a v_{j,\alpha}) = \delta_{\alpha\beta} V_j^a . \quad (\text{C.3})$$

C.1.2 The no-rotation condition

In Section 5.2.3.2, we claim we can always find a basis for B_j at each point in the neighborhood of a given ray which satisfies the *no-rotation condition* [Eq. (5.10)]:

$$(v_{j,[\alpha}, S(\partial_t + A^a \partial_a) v_{j,\beta]}) = 0 \quad (\text{C.4})$$

for all α, β . In this section, we demonstrate explicitly how to construct a basis which satisfies both the no-rotation condition and which remains orthonormal.

If the right-hand side term in Eq. (C.4) is not zero in the basis $v_{j,\alpha}$, we attempt to choose a new basis

$$v_{j,\bar{\alpha}} \equiv R_{\bar{\alpha}\alpha} v_{j,\alpha}$$

such that Eq. (5.10) is satisfied by the new basis and moreover such that the new basis is orthonormal. In particular, if the no-rotation condition is satisfied in the new basis, then (because $P_j A^a P_j = V_j^a$) we know

$$\begin{aligned} 0 &= \sum_{\alpha\beta} R_{\bar{\alpha}\alpha} R_{\bar{\beta}\beta} (v_{j,[\alpha}, S(\partial_t + A^a \partial_a) v_{j,\beta]}) \\ &+ R_{[\bar{\alpha}\alpha} \sum_{\alpha} (\partial_t + V_j^a \partial_a) R_{\bar{\beta}]\alpha} . \end{aligned} \quad (\text{C.5})$$

[In the second term above, the antisymmetrization is over only the barred indices $\bar{\alpha}$ and $\bar{\beta}$.] On the other hand, if the basis is orthonormal, then $\sum_{\alpha} R_{\bar{\alpha}\alpha} R_{\bar{\beta}\alpha} = \delta_{\bar{\alpha}\bar{\beta}}$, implying

$$\sum_{\alpha} R_{(\bar{\alpha}\alpha} (\partial_t + V_j^a \partial_a) R_{\bar{\beta})\alpha} = 0 .$$

[In the above, the operator is symmetrized over the indices $\bar{\alpha}$ and $\bar{\beta}$.] Therefore, combining the two, we conclude that if the new basis is both orthonormal and satisfies the no-rotation, the matrix R must satisfy the ordinary differential equation

$$(\partial_t + V_j^a \partial_a) R_{\bar{\alpha}\alpha} = - (v_{j,[\alpha}, S(\partial_t + A^a \partial_a) v_{j,\beta]}) R_{\bar{\alpha}\beta}$$

subject to initial data $R_{\bar{\alpha}\alpha} = \delta_{\bar{\alpha}\alpha}$. Solutions for R and thus $v_{j,\bar{\alpha}}$ exist in the neighborhood of a ray.

C.1.3 Reorganizing inner products for the polarization equation

In this section, we describe how to rearrange matters so the last term in Eq. (5.7c) – namely,

$$\left(v_{j,\alpha}, S(\partial_t + A^a \partial_a) v_{j,\beta} \right) \quad (\text{C.6})$$

– has a simpler form. If we choose a basis that satisfies the no-rotation condition [Eq. (5.10)], the antisymmetric part of this matrix is zero. Further, we may express the symmetric part of this

expression by using the relations

$$\begin{aligned} \partial_t(v_{j,\alpha}, Sv_{j,\beta}) &= (v_{j,\alpha}, S\partial_t v_{j,\beta}) \\ &+ (v_{j,\beta}, S\partial_t v_{j,\alpha}) \\ &+ (v_{j,\alpha}, (\partial_t S)v_{j,\beta}) \end{aligned} \tag{C.7}$$

$$\begin{aligned} \partial_a(v_{j,\alpha}, SA^a v_{j,\beta}) &= (v_{j,\alpha}, SA^a \partial_a v_{j,\beta}) \\ &+ (v_{j,\beta}, SA^a \partial_t v_{j,\alpha}) \\ &+ (v_{j,\alpha}, (\partial_a SA^a) v_{j,\beta}) \end{aligned} \tag{C.8}$$

[where we have observed that SA^a and A^a are symmetric relative to the canonical inner product] and the expressions

$$(v_{j,\alpha}, Sv_{j,\beta}) = \delta_{\alpha\beta} \tag{C.9}$$

$$(v_{j,\alpha}, SA^a v_{j,\beta}) = \delta_{\alpha\beta} V_j^a \tag{C.10}$$

[i.e. orthogonality and Eq. (C.3)]. These relations tell us that, if the no-rotation condition is satisfied,

$$\begin{aligned} (v_{j,\alpha}, S(\partial_t + A^a \partial_a)v_{j,\beta}) &= \\ \frac{1}{2}\delta_{\alpha,\beta}\partial_a V^a(\vec{x}, \vec{k}(x)) &- \frac{1}{2}(v_{j,\alpha}, (\partial_t S + \partial_a SA^a) v_{j,\beta}) \end{aligned} \tag{C.11}$$

Our notation for the first term on the right side (i.e. the divergence of the group velocity) is chosen to emphasize that the derivative ∂_a acts on *all* the dependence on \vec{x} – in particular, on any variation of $k_a = \partial_a \phi$ with \vec{x} .

C.2 Demonstrating that the ray-optics approach provides high-quality approximate solutions to the FOSHLS

For any fixed initial data, the ray-optics solution obtained in Sec. 5.2 will break down at some point along each ray. In this section, we estimate how long a solution obtained by solving Eq. (5.7) can be trusted.

Specifically, in Sec. C.2.1 we express the FOSHLS [Eq. (5.2)] using alternative variables better-suited to describing the geometric-optics solution. Next, in Sec. C.2.2 we survey the various orders of magnitude that arise in the problem. Using those orders of magnitude, in Sec. C.2.3 we estimate the error in Eq. (5.2) that occurs when a geometric optics solution is substituted for u (e.g., we estimate how close the norm of the left side is to zero). Finally, knowing how much error we make when using

a geometric-optics solution, in Sec. C.2.4 we estimate how errors involved in a geometric-optics approximation grow; we therefore discover how long we can trust a purely geometric-optics-based evolution.

C.2.1 Review: Writing equations in terms of ϕ and $d_{l,\alpha}$

In Equations (5.5) and (5.6) we describe how to parameterize the N -dimensional state vector u by using N functions $d_{l,\alpha}$ and one additional function ϕ . If we insert this substitution into the original FOSHLS [Eq. (5.2)], then dot the results against each of the orthonormal basis vectors $v_{l,\alpha}$, we obtain the equations

$$\begin{aligned}
0 &= i(v_{l,\alpha}, S\bar{u}) [\partial_t + V_l^a \partial_a] \phi \\
&\quad + \left(v_{l,\alpha}, S(\partial_t + A^a \partial_a - F)\bar{u} \right) \\
&= id_{l,\alpha} [\partial_t + V_l^a \partial_a] \phi \\
&\quad + \partial_t d_{l,\alpha} + \sum_m \sum_\beta (v_{l,\alpha}, SA^a v_{m,\beta}) \partial_a d_{m,\beta} \\
&\quad + \sum_m \sum_\beta d_{m,\beta} \left(v_{l,\alpha}, S(\partial_t + A^a \partial_a - F)v_{m,\beta} \right).
\end{aligned} \tag{C.12}$$

In the above, we have observed that A^a is symmetric relative to the inner product generated by S and that $v_{j,\alpha}$ is an eigenvector of $A^a \partial_a \phi$ with eigenvalue $\omega_j = V_j^a \partial_a \phi$. We can further reorganize this equation by pulling out all terms that involve d_l explicitly, and also by using Eq. (C.2) to simplify $(v_{l,\alpha}, SA^a v_{l,\beta}) = V_l^a \delta_{\alpha\beta}$:

$$\begin{aligned}
0 &= id_{l,\alpha} [\partial_t + V_l^a \partial_a] \phi + [\partial_t + V_l^a \partial_a] d_{l,\alpha} \\
&\quad + \sum_\beta d_{l,\beta} \left(v_{l,\alpha}, S(\partial_t + A^a \partial_a - F)v_{l,\beta} \right) \\
&\quad + \sum_{m \neq l} \sum_\beta (v_{l,\alpha}, SA^a v_{m,\beta}) \partial_a d_{m,\beta} \\
&\quad + \sum_{m \neq l} \sum_\beta d_{m,\beta} \left(v_{l,\alpha}, S(\partial_t + A^a \partial_a - F)v_{m,\beta} \right).
\end{aligned} \tag{C.13}$$

C.2.2 Natural scales used in order-of-magnitude estimates

To make order-of-magnitude arguments regarding the solution, we need to understand how the natural length scales of the problem enter into it.

Rather than complicate the order-of-magnitude calculation unnecessarily, we shall for simplicity proceed as if there existed only one characteristic speed. In other words, we shall freely convert between space and time units by using the norm of A^a ; for example, we can interpret $\tau_{F,n}|A|$ as a

natural length scale. Finally, for brevity, we shall assume space and time units are chosen so $|A| \sim 1$.

Even with the above simplification, many natural scales arise in the problem, including the magnitude of F ; the natural length and time scales on which F and A vary; and the length scale on which the initial data varies. Again, for simplicity we shall summarize all these scales by only two numbers:

- “*Length*” scale (L) We define the natural “length” scale to be the natural time scale that enters on the right side of Eq. (5.7c). To be explicit, L is the smaller of $|d|/|A||\partial_a d|$ and $1/|F|$.
- “*Variation*” scale (τ_{vary}) The remaining scales do not arise directly in the equation. They affect the propagation of the wave packet only because they determine the rate at which terms in the equation are modulated as the wave packet propagates in space and time. We therefore call the smallest of the remaining scales the *variation scale* (τ_{vary}); its value is the smallest of the length and time scales on which A and F vary.

C.2.3 Degree to which ray-optics solution satisfies the FOSHLS

Using the above rough estimates (L and τ_{vary}) to characterize the magnitude of terms both used and neglected, we find that geometric-optics solutions [Eq. (5.7)] very nearly satisfy the full FOSHLS [Eq. (5.2), alternatively expressed as Eq. (C.13)]. To be explicit, when we insert a geometric-optics solution which propagates in the j th polarization [i.e. a solution to Eq. (5.7)] into Eq. (C.13), we find the following:

$$\begin{aligned}
0 &= id_{l,\alpha}(\omega_l - \omega_j) + [\partial_t + V_l^a \partial_a] d_{l,\alpha} \\
&+ \sum_{\beta} d_{l,\beta} \left(v_{l,\alpha}, S(\partial_t + A^a \partial_a - F)v_{l,\beta} \right) \\
&+ \sum_{m \neq l} \sum_{\beta} (v_{l,\alpha}, SA^a v_{m,\beta}) \partial_a d_{m,\beta} \\
&+ \sum_{m \neq l} \sum_{\beta} d_{m,\beta} \left(v_{l,\alpha}, S(\partial_t + A^a \partial_a - F)v_{m,\beta} \right)
\end{aligned} \tag{C.14}$$

[Here, we have used Eq. (5.7a) and the definition of ω_j to simplify the first term.]

We have two circumstances:

- When $l = j$, the first three terms precisely cancel [see Eqs. (5.7a) and (5.7c)]. The only terms remaining are of order $d_{m,\beta}$ $m \neq j$.
- On the other hand, when $l \neq j$, the first term does not cancel. Rather, it is large, because $\partial_a \phi$ is large (i.e. we are in the short-wavelength limit), and the ω_l are proportional to $\partial_a \phi$.

For brevity, assume the eigenvalues of A^a are of comparable magnitude, so to order of magnitude $\omega_j \sim \omega_l \sim \omega \sim \omega_j - \omega_l$. We may then express the equation when $l \neq j$ as

$$0 = \mathcal{O}(\omega d_{l\beta}) + \mathcal{O}(d_{j,\alpha}/L)$$

The second terms will force the first terms, generally, to be nonzero.

From the second case, we know that when $l \neq j$, $|d_l| \sim \mathcal{O}(|d_j|/L\omega)$. Combining this result with the first equation, we conclude that when we use our trial solution, we are ignoring terms of order $|d_j|/L^2\omega$ when $l = j$ and terms of order $|d_j|/L$ when $l \neq j$.

C.2.4 Length of time ray-optics solution can be trusted

To estimate the integrated effects the neglected terms have on the diagonal and off-diagonal polarization amplitudes $d_{j,\alpha}$ and $d_{l,\alpha}$, respectively, we integrate the previous equations.

When $l \neq j$, we have a DE of form

$$\frac{d}{ds}d_{l,\beta} + id_{l,\beta}\Delta\omega + \mathcal{O}(|d_j|/L) = 0$$

where we neglect smaller terms, where $d/ds = \partial_t + V_j^a \partial_a$ represents the derivative along a characteristic, and where $\Delta\omega = \omega_l - \omega_j \sim \omega$. Since we limit attention to ω very large (i.e. $\omega\tau_{\text{vary}} \gg 1$), we may ignore the weak effects of any time variation of L and treat it as constant. Since $|d_j|$ varies along the characteristic much more slowly than does $\exp(-i\omega s)$, we find that after an affine length T , $|d_l|$ will be of order

$$|d_l| \sim T|\partial_s d_j|/\omega L \sim |d_j|T/L^2\omega. \quad (\text{C.15})$$

(Here, I assume $L^2\omega$ is suitably averaged, as L will vary as the path evolves.) Similarly, when $l = j$, we ignore terms of order $|d_j|/L^2\omega$. We have a DE of form

$$\frac{d}{ds}d_{j,\alpha} + (\text{known, real}) + \mathcal{O}(|d_j|/L^2\omega) = 0$$

Therefore, integrating along an affine length T of the ray, we expect errors in the $d_{j\alpha}$'s of relative magnitude

$$\epsilon_{\text{amp}} = T/L^2\omega \quad (\text{C.16})$$

when ϵ_{amp} is small. In both cases, we see the neglected terms will be smaller than $|d_j|$ by magnitude ϵ_{amp} .

If we are simulating a *given* system, with fixed initial data, we can only trust a solution out to time $T \sim L^2\omega$. However, for any compact region of any characteristic (i.e. for any fixed T), we can

always choose ω sufficiently large so the relative errors ϵ_{amp} is arbitrarily small.

C.3 When do PCWPs exist?

Rather than evolve general wave packets using the full geometric-optics equations, for simplicity in this paper we often restrict attention to prototypical coherent wave packets (PCWPs), which – if they exist – vastly simplify the problem of evolution (cf. Sec. 5.3.3). In this appendix, we try to clarify the conditions under which prototypical coherent wave packet solutions exist as exact or approximate solutions to the geometric-optics equations.

We can better understand under what conditions prototypical coherent wave packets exist if we rewrite the polarization equation [Eq. (5.11)] using the basis $f_j^{(\mu)}$ [Eq. (5.14)]. When we do so, we find PCWP are exact solutions only for special circumstances. However, when some eigenvalue of O_j is large, we find that PCWPs arise naturally as limits of *arbitrary* coherent wave packets.

C.3.1 Rewriting polarization equation in the basis of eigenvectors of O_j

Basis vectors and their components: The basis vectors $f_j^{(\mu)}$ are defined above. Since we express the polarization equation in component form relative to some no-rotation basis, we also need notation for the components $f_{j\alpha}^{(\mu)}$ of these basis vectors relative to the no-rotation basis:

$$f_j^{(\mu)} = \sum_{\alpha} f_{j\alpha}^{(\mu)} v_{j,\alpha} .$$

Dual vectors and their components: The basis vectors $f_j^{(\mu)}$ are not necessarily orthogonal. To facilitate computations, we define a dual basis $\tilde{f}_j^{(\mu)}$ such that

$$\delta^{\mu\nu} = \left(\tilde{f}_j^{(\nu)}, S f_j^{(\mu)} \right) .$$

The dual basis vectors can be expressed in terms of components, denoted $\tilde{f}_{j,\alpha}^{(\mu)}$, relative to the no-rotation basis.

Explicitly rewriting polarization equation: Substituting in the expansion

$$\bar{u} = \sum_{\mu} D_{j\mu} f_j^{(\mu)} \quad \leftrightarrow \quad d_{j\alpha} = \sum_{\mu} D_{j\mu} f_{j\alpha}^{(\mu)}$$

into the polarization equation [Eq. (5.11)], then using the dual vector basis to select specific com-

ponents, we find

$$0 = \left(\partial_t + V_j^a \partial_a + \frac{1}{2} \partial_a V_j^a - o_{j\mu} \right) D_{j\mu} - \sum_{\nu} D_{j\nu} \left[\sum_{\alpha} \left(\tilde{f}_{j\alpha}^{(\mu)}, S (\partial_t + V_j^a \partial_a) f_{j\alpha}^{(\nu)} \right) \right] \quad (\text{C.17})$$

C.3.2 Sufficient conditions for PCWP to be exact solution

By definition, a prototypical coherent wave packet solution associated with the polarization direction $f_j^{(\mu)}$ exists only if there is a solution – exact or approximate – to Eq. (C.17) with all $D_{j\nu} = 0$ except for $\nu = \mu$ [i.e. $D_{j,\mu} \neq 0$]. In other words, for a complete collection of PCWP solutions to exist, one for each μ , the third term in Eq. (C.17) must be diagonal, or zero.

[In fact, for the examples addressed in this paper (i.e. in Sections 5.5 and 5.6, for propagation on the light cone), the third term is in fact exactly zero.]

C.3.3 PCWP as limit of arbitrary rapidly growing coherent wave packet

If the largest eigenvalue $o_{j\nu}$ of O_j is particularly large compared to the third term, then generic solutions to the polarization equation [Eq. (C.17)] will converge to a state with $D_{j\nu} \gg D_{j\mu}$ for $\mu \neq \nu$ (i.e. $w = f_j^{(\nu)}$). In other words, if the largest eigenvalue of O_j is large, then generic wave packets will converge to the PCWP with $w = f_j^{(\nu)}$.

C.4 Bounding the energy growth rate

In Sections 5.2 and 5.3 we introduced wave-packet solutions as solutions which are localized in the neighborhood of a given ray. However, while we obtained expressions for the growth rate of certain specialized wave packet solutions [Sec. 5.4], in the main text of this paper we never provided a strict bound on the growth rate of a wave packet.

In fact, we can bound the instantaneous growth rate of a coherent wave packet [Eq. (5.19)] by a quantity independent of the precise polarization state w of that packet:

$$\frac{1}{E} \frac{dE}{dt} \leq \max_{w \in B_j} \frac{(w, SQw)}{(w, Sw)} \quad (\text{C.18})$$

As Q is symmetric relative to S , it has a spectrum of real eigenvalues, each associated with eigenspaces that are orthogonal relative to S . It follows that if κ_s is the largest eigenvalue of Q ,

$$\frac{1}{E} \frac{dE}{dt} \leq \kappa. \quad (\text{C.19})$$

This procedure follows precisely the same outline as the energy-norm upper bound discussed in LS Eq. (2.17) and (2.18).

This upper bound on the growth rate for all polarizations propagating along a given ray can be used as a line-by-line replacement for the maximum growth rate of PCWPs [Eq. (5.20)] in practical calculations to determine the largest amplification possible by a wave packet propagating in the future domain of dependence [e.g., Sections 5.7 and 5.8]

C.5 Rays optics and KST 2-parameter formulation

While KST introduce a very large family of symmetric hyperbolic systems, they emphasize (and limit their calculations to) a simple 2-parameter subset. This two parameter system has both physical characteristic speeds and a simple principal part (i.e. simple form for A^a). As a result, the algebra required for its ray-optics limit (i.e. computations of ω_j , etc) proves particularly simple.

C.5.1 Generally

The KST system has as variables the tensors g_{ab} , P_{ab} , M_{kab} defined over 3-space, for a total of $6+6+18=30$ fields.

C.5.1.1 Principal part and symmetrizer

The principal part has the form [KST Eq. (2.59), along with the definition of $\hat{\partial}_o$ in KST Eq. (2.10)]:

$$(\partial_t - \beta^a \partial_a)g_{ij} \simeq 0 \tag{C.20a}$$

$$(\partial_t - \beta^a \partial_a)P_{ij} + Ng^{ab}\partial_a M_{bij} \simeq 0 \tag{C.20b}$$

$$(\partial_t - \beta^a \partial_a)M_{kij} + N\partial_k P_{ij} \simeq 0 \tag{C.20c}$$

After linearizing about a background solution, this principal part and a choice of representation for the fields (i.e. u) gives us the explicit form for A^a . We may represent the result as

$$A^a = -\beta^a \mathbf{1} + NA_o^a \tag{C.21}$$

for $\mathbf{1}$ the identity operator and A_o^a a construction which depends only on the background metric g and the choice of field ordering used in going to a matrix representation (i.e. the order of the fields in u).

This principal part is symmetric hyperbolic, using as symmetrizer (for example) LS Eq. (3.67):

$$(u, Su) = g^{a\bar{a}}g^{b\bar{b}}dg_{ab}dg_{\bar{a}\bar{b}} + g^{a\bar{a}}g^{b\bar{b}}dP_{ab}dP_{\bar{a}\bar{b}} \\ + g^{a\bar{a}}g^{b\bar{b}}g^{k\bar{k}}dM_{kab}dM_{\bar{k}\bar{a}\bar{b}}. \quad (\text{C.22})$$

This symmetrizer (represented here in LS notation) amounts to nothing more than the naturally-constructed sum of squares of components of g , P , and M .

C.5.1.2 Eigenvalues and group velocity

From the principal part, we can deduce the three possible eigenvalues:

$$\omega_s(x, p) = -\beta^a p_a + sN\sqrt{g^{ab}p_a p_b} \quad (\text{C.23})$$

where $s = 0, \pm 1$. From this expression we obtain the group velocities

$$v_s^a(x, p) = -\beta^a + sN\hat{p}^a \quad (\text{C.24})$$

where $\hat{p}^a \equiv g^{ab}p_b/\sqrt{g^{rs}p_r p_s}$.

C.5.1.3 Eigenfields and projection operators

KST tabulate the eigenfields of the principal-part operator [Eq. (C.20)] in KST Eq. (2.61) and the surrounding text. These expressions yield the following orthonormal basis vectors for the three eigenspaces of $A^a\hat{p}_a$:

$$v_{o,g,ab} = g_{ab} \quad (\text{C.25a})$$

$$v_{o,x,ab} = \frac{[M_{qab}\hat{x}^q - \hat{p}_u\hat{x}^u\hat{p}^q M_{qab}]}{\sqrt{1 - (p^a\hat{x}_a)^2}} \quad (\text{C.25b})$$

$$v_{o,y,ab} = \frac{[M_{qab}\hat{y}^q - \hat{p}_u\hat{y}^u\hat{p}^q M_{qab}]}{\sqrt{1 - (p^a\hat{y}_a)^2}} \quad (\text{C.25c})$$

$$v_{\pm,ab} = \frac{1}{\sqrt{2}}[P_{ab} \pm \hat{p}^q M_{qab}] \quad (\text{C.25d})$$

where we treat symbols for the fields as basis vectors in the space of fields, and where \hat{x} and \hat{y} are two 3-vector fields not parallel to \hat{p} and which are orthonormal relative to the metric g_{ab} at each point.

C.5.2 Special case: Flat spatial metric

When the KST system is applied to a time-independent solution with a flat spatial metric, the algebra simplifies substantially. For example, the symmetrizer S [Eq. (C.22)] is the identity operator on the set of fields. The inner product generated by S is therefore constant in space and time.

C.5.2.1 Simplifying general polarization equation

Since we fix the basis vector convention by Eq. (C.25), we must use the polarization equation in the form of Eq. (5.7c) (i.e. we do not generically expect the no-rotation condition to hold). We therefore must evaluate the term

$$\begin{aligned} (v_{j,\beta}, S(\partial_t + A^a \partial_a)v_{j,\beta}) &= (v_{j,\beta}, S(\partial_t - \beta^a \partial_a)v_{j,\beta}) \\ &\quad + N (v_{j,\beta}, SA_o^a \partial_a v_{j,\beta}) \end{aligned}$$

Since we know how the basis vectors change as a function of the congruence direction $\hat{k} = \hat{p}$ [Eqs. (C.25a–C.25d)], we can rewrite this expression in terms of our knowledge of the congruence and β^a .

For example, for the fields propagating forward along the congruence at unit speed ($j = s = \pm 1$), we have

$$(v_{s,ab}, SA_o^a \partial_a v_{s,cd}) = \left(\frac{1}{2} s \partial_q \hat{p}^q \right) \delta_{ac} \delta_{bd}, \quad (\text{C.26})$$

$$(v_{s,ab}, S(\partial_t - \beta^a \partial_a)v_{s,cd}) = 0. \quad (\text{C.27})$$

C.5.2.2 Simplifying the energy equations

The only new quantity needed to evaluate the energy equation is $\partial_a SA^a$. To evaluate $\partial_a SA^a$, we note S is the identity, so we just differentiate the result we obtain from Eq. (C.21):

$$\partial_a SA^a = -\partial_a \beta^a \mathbf{1} + (\partial_a N) A_o^a.$$

Now, if we take the inner product over the fields, we arrive at the expression

$$(v_{j,\alpha}, (\partial_a SA^a)v_{j,\beta}) = \delta_{\alpha\beta} \left[-\partial_a \beta^a + (\partial_a N) s_j \frac{p^a}{|p|} \right]. \quad (\text{C.28})$$