



ROBUST OPTICAL CONFIGURATIONS FOR
REDUCING THE IMPACT OF SPATIAL LASER BEAM
EFFECTS IN GRAVITATIONAL WAVE DETECTORS

By

SAMUEL JOSHUA ROWLINSON

A thesis submitted to
the University of Birmingham
for the degree of
DOCTOR OF PHILOSOPHY

Institute for Gravitational Wave Astronomy
School of Physics and Astronomy
College of Engineering and Physical Sciences
University of Birmingham
May 2021

University of Birmingham Research Archive
e-theses repository



This unpublished thesis/dissertation is under a Creative Commons Attribution 4.0 International (CC BY 4.0) licence.

You are free to:

Share — copy and redistribute the material in any medium or format

Adapt — remix, transform, and build upon the material for any purpose, even commercially.

The licensor cannot revoke these freedoms as long as you follow the license terms.

Under the following terms:



Attribution — You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.

No additional restrictions — You may not apply legal terms or technological measures that legally restrict others from doing anything the license permits.

Notices:

You do not have to comply with the license for elements of the material in the public domain or where your use is permitted by an applicable exception or limitation.

No warranties are given. The license may not give you all of the permissions necessary for your intended use. For example, other rights such as publicity, privacy, or moral rights may limit how you use the material.

Unless otherwise stated, any material in this thesis/dissertation that is cited to a third-party source is not included in the terms of this licence. Please refer to the original source(s) for licencing conditions of any quotes, images or other material cited to a third party.

ABSTRACT

The gravitational wave detections of the past five years were a culmination of decades worth of research. To increase the detection rates and horizon of detectability going forwards, much more work is yet required for identifying the sources of noise in ground-based, laser interferometer observatories, understanding how these can couple into the differential arm motion and, hence, developing strategies by which such noise can be reduced; for both current and future detectors. One of the principle subjects in this regard is the impact of the transverse, spatial properties of laser beams on the sensitivity of gravitational wave detectors.

In this thesis I will highlight my contributions to the detector instrumentation field on this subject. A central theme underlying this work is the impact that larger beams, planned for future detectors, will have on the optical design of such facilities and, thus, the overall sensitivity of these new observatories. With regards to this, a novel design for mode matching telescopes in the arms of each Einstein Telescope interferometer is analysed for feasibility and adaptive mode matching potential. The results of this work indicate that such a design could be viable, and initial design values are provided as a good starting point for future trade-off studies. The effect that larger beams may have on the alignment to longitudinal coupling noise is also studied, in the context of third-generation gravitational wave detectors. Conclusions from these analyses indicate that the negative impact this has can be counteracted through careful tuning of the optical design.

Due to the complicated nature of the physics involved with these, and other closely related, topics, it is vital to be able to computationally model such phenomena in a fast and reliable way. As a result of this, running parallel to the aforementioned studies, in this thesis I will also address the task of building essential modelling tools in this regard. My contributions to a new open-source package, FINESSE 3, designed to tackle complex interferometer modelling, are detailed; with particular attention paid to my work on implementing the tools for simulating spatial beam effects in interferometers. The uses of these newly developed tools are given throughout the research tasks within this document. The final chapter of this thesis is dedicated to the overall design of FINESSE 3, highlighting how these features address the need, in the instrumentation community, for such a tool going forwards.

STATEMENT OF ORIGINALITY

The contents of this thesis consist of my own research work conducted, during the course of my PhD, between September 2017 and May 2021.

Chapter 1 contains a brief summary of the current state of the gravitational wave instrumentation field, including an introduction to the design of ground-based detectors and an overview of their current sensitivity limitations. This chapter also introduces the framework for describing light field interactions at optics for interferometry and methods by which these can be modelled numerically. These topics provide a background for the physics covered in the subsequent chapters.

Chapter 2 details my work on implementing the current higher-order spatial mode features of FINESSE 3 — a new software package for simulating interferometers in the frequency domain. A major aspect of this chapter is dedicated to the new beam tracing library I implemented, which was used to produce the results found in both Chapters 3 and 4.

Chapter 3 describes my work on investigating the feasibility of a new recycling cavity design for the Einstein Telescope (ET) in which mode matching telescopes are placed in the interferometer arms. This work was published in: *Samuel Rowlinson et al. “Feasibility Study of Beam Expanding Telescopes in the Interferometer Arms for the Einstein Telescope,” Phys. Rev. D 103 (2 Jan. 2021) [1]*. In this chapter I demonstrate that the use of beam expansion telescopes in the detector arms is viable, and I provide a potential initial design for such a mode matching system for both the low-frequency (ET-LF) and high-frequency (ET-HF) detectors.

Chapter 4 provides a follow-up analysis consisting of my work towards analysing the adaptive mode matching potential of the previous arm telescope design in ET-HF. This chapter consists of an investigation into possible actuation strategies via utilisation of the arm telescope mirrors. In this study I demonstrate that actuation on the telescope mirror ZM2 could provide a means for correcting small ITM thermal lens induced mode mismatches. Additionally, I provide an estimate of a suitable initial phase space region for the SRC mode — predicated upon actuators at ZM2 and the recycling mirrors, using ranges based on existing technology at aLIGO.

Chapter 5 describes my work on projecting how alignment to longitudinal noise coupling will scale with larger beam sizes for future detector design. In this chapter I provide an analytical result yielding a nominal scaling relation for this coupling as the beam size to the power of six. This result is backed up with my modelling work on a full aLIGO detector

file to validate this derivation. A key outcome of this chapter is to show that this nominal scaling with the beam size only provides part of the full picture, as the increased coupling to first order modes will also increase alignment signals allowing for a reduction in the residual misalignment.

Chapter 6 gives a summary of the overarching design of FINESSE 3, a package which is a complete re-work (from the ground up) of the original software. Development of this package started at the beginning of my PhD and it has been a focal point of my work throughout. I detail my key role in the development of FINESSE 3 in this chapter whilst also highlighting the overall structure of the software and the motivating reasons behind the design choices made therein.

Appendix A delves into the technical details of my implementation of the beam tracing library in FINESSE 3, which was first introduced in Chapter 2.

Appendix B provides my analytic derivations for the WS phase space overlap quantity used in Chapter 4. These are given as the overlap on an infinite plane and also the overlap for a finite circular aperture, for completion.

Appendix C gives the mathematics behind the scaling of the alignment to longitudinal noise coupling with the beam size. This work consists of my derivation of the field couplings for a DC and AC misaligned Fabry-Perot cavity in terms of the fundamental and first order modes at carrier and signal frequencies. These couplings are expressed in terms of the beam size of the cavity eigenmode to provide useful formulae which are then used directly in Chapter 5.

ACKNOWLEDGMENTS

Firstly I'd like to thank my supervisor, Prof. Andreas Freise, for offering me the opportunity to study for this PhD, and for all the helpful advice and supervision along the way. I am also grateful to all those who helped me in any way with my research, of whom there were many including: Dr. C. M. Lowry, Dr. H. Miao, Dr. A. Dmitriev, Dr. T. Zhang, Dr. A. Jones, Dr. D. Brown, Dr S. Leavey, Dr. A. Allocca, Dr. D. Töyrä, Dr. A. Green and many more. In particular, from this list, I want to especially thank Conor, Artemiy, Aaron and Teng for many useful and interesting discussions which helped me to understand some tricky topics and opened up new directions for my various research tasks. Furthermore, a big thank you to: Mischa Sallé, Philip Jones, Sean Leavey, Aaron Jones, Daniel Brown and Antonella Bianchi for proof-reading, and providing useful feedback on, this thesis. Thank you also to my viva examiners, Keiko Kokeyama and Jon Goldwin.

A major aspect of my PhD was the development of FINESSE 3, which is something that I thoroughly enjoyed contributing to and generally being a part of. Working closely with the core team of developers is what made this such a fun task and so I'd very much like to acknowledge: Daniel Brown, Sean Leavey, Philip Jones, Aaron Jones, Andreas Freise and Mischa Sallé, in this regard. It has been an honour to work with you all over the past 3.5 years. Whichever path I end up traversing next, I wish this project the very best for the future.

It has been a privilege to work with all the members of the Institute of Gravitational Wave Astronomy in Birmingham, and especially those of you within Ifolab. For all of your advice, suggestions and friendship, thank you to: Philip, Aaron, Chiara, Anna, Sam, Artemiy, Joe, Amit, Jiri, Kazuhiro, Daniel, Miguel, Andreas, Conor, Haixing and Denis. I'm sure I missed a few names (and apologies if I did!) but each and every member of the group helped to make my PhD experience more enjoyable. A special thanks goes also to David Stops for your incredibly helpful IT support and close friendship with the group as a whole; many an interesting conversation was had during the (plentiful) coffee breaks in Physics West!

Most importantly, a big thanks goes out to my friends and family. In particular, I am eternally grateful for the unwavering support of my Mom and Dad — thank you for always being there, especially over the past year! Lastly, as strange as it may sound, I'd like to acknowledge the wildlife (especially the crows and ravens), and nature in general, around me. Going for long walks in the countryside has been a lifeline for me during the past couple of months whilst writing up this thesis.

As a final remark: life, and work in particular, throughout the ongoing pandemic has been

both strange and challenging — for all of us. There is no exception in my case, as completing the final year of my PhD through extended periods of lockdown was a stressful and difficult experience; having curtailed several plans along the way. I do not begrudge this necessitous current way of life, however, as it has undoubtedly prevented the already-staggering toll of the pandemic from growing needlessly further. My only hopes, as naive as they may be, are that the vital importance of science, and being thoughtful and considerate of others (for you never know what they may be going through), are two things that society remembers in the aftermath of this global tragedy.

This thesis makes extensive use of Matplotlib [2] for generating the plots shown, and ComponentLibrary [3] for producing the optical layout diagrams.

“We are lost, aren’t we, most perspicacious of pathfinders?”

Lurk Snichtongue, in *Gotrek and Felix, Dragonslayer; by William King* [4]

Contents

	Page
Abstract	i
Statement of Originality	iii
Acknowledgments	v
List of Figures	xii
List of Tables	xvii
List of Source Codes	xvii
Acronyms	xviii
FINESSE Nomenclature	xxi
1 Introduction	1
1.1 Interferometers as gravitational wave detectors	2
1.1.1 Arm cavities	3
1.1.2 Recycling cavities	4
1.1.3 Mode-cleaner cavities	4
1.2 Overview of current detector sensitivity limitations	5
1.2.1 Quantum vacuum noise	6
1.2.2 Coating thermal noise	7
1.2.3 Seismic and suspension noise	8
1.3 Mathematical description of light fields for laser interferometry	8
1.3.1 The plane-wave picture	9
1.3.2 Gaussian beams and Higher-Order spatial Modes	9
1.3.2.1 The wave equation and paraxial approximation	10
1.3.2.2 Properties of Gaussian beams	11
1.3.2.3 Spatial distortions as Higher-Order Hermite-Gauss modes	14
1.3.2.4 The Gaussian beam parameter	16
1.4 Modelling interferometers in the frequency-domain	17
1.4.1 Time-independent light field amplitudes	18
1.4.2 Component field interactions as coupling matrices	18

1.4.2.1	Propagation over a space	19
1.4.2.2	Coupling matrices at a mirror	20
1.4.3	Solving optical interferometer couplings via sparse matrix inversion	21
1.5	Thesis overview	22
2	Modelling Higher-Order Spatial Modes in Gravitational Wave Detectors with FINESSE 3	25
2.1	Motivation and current software review	26
2.2	Beam tracing	27
2.2.1	ABCD matrices and beam parameter transformation	28
2.2.2	Optical cavities	29
2.2.2.1	Eigenmodes	29
2.2.2.2	Stability of optical cavities	31
2.2.2.3	Round-trip Gouy phase and mode spacing	32
2.2.2.4	Example: The Fabry-Perot cavity	33
2.2.3	Propagating beams through an interferometer	35
2.2.3.1	Tree structure	36
2.2.3.2	Mode mismatches	39
2.2.3.3	Symmetric vs asymmetric tracing	41
2.2.3.4	Importance of defining cavity modes	44
2.2.3.5	Order of beam tracing	45
2.2.3.6	Overlapping cavities	47
2.2.4	Beam tracing as a library	48
2.2.4.1	Symbolic beam propagation	48
2.2.4.2	Matrix-less simulations	50
2.3	Mode scattering matrices	50
2.3.1	Coupling coefficients and the overlap integral	51
2.3.2	Computing coupling coefficients in FINESSE 3	52
2.3.2.1	Gouy phase shifts and coupling coefficients	52
2.3.2.2	Example: Scattering matrices at a mirror	53
2.3.2.3	Performance	55
2.3.3	Selecting the mode indices to model	57
2.3.3.1	Misalignments and mode mismatches as additions of HOMs	57
2.3.3.2	Example: Scanning yaw angle of a beam splitter	59
2.3.4	Future work	60
2.4	Capturing beam images	61
2.4.1	CCDs	61
2.4.2	Complex cameras	61
2.4.3	Dimensionality of camera-type detectors	62
2.4.4	Example: Mode scan of misaligned aLIGO arm cavity	62
2.4.5	Performance	63
2.5	Summary	64

3	Feasibility Study of Beam Expanding Telescopes in the Interferometer Arms for the Einstein Telescope	67
3.1	Motivation	68
3.2	Arm cavity eigenmode and telescope constraints	70
3.2.1	Telescope parameter constraints	71
3.3	ET-LF arm telescope design	73
3.3.1	Flat Z mirrors	73
3.3.2	Curved telescope mirrors, and no lens at ITM	73
3.3.3	Curved telescope mirrors with a lens at the ITM	75
3.4	ET-HF arm telescope design	79
3.4.1	Further flexibility of the telescope design	81
3.5	Parameter sensitivity and mode matching	82
3.5.1	Mode matching in the presence of thermal lensing	83
3.6	Summary	84
4	Adaptive Mode Matching with Arm Telescopes in the Einstein Telescope	87
4.1	WS phase space	88
4.2	Core ET mode mismatches in WS space	90
4.2.1	Quantifying the thermal lens focal lengths	92
4.3	Types of actuation for adaptive mode matching	92
4.4	Determining actuation effectiveness in WS space	93
4.5	Suitability of telescope mirror actuation	94
4.5.1	RoC actuators at ZM1 and ZM2	94
4.5.2	Shifting the positions of ZM1 and ZM2	96
4.5.3	Estimating the available phase space regions	97
4.6	Summary	98
5	Scaling of Alignment Noise Coupling with Larger Beams	101
5.1	Static and oscillatory misalignments	102
5.2	Approximating the scaling of alignment noise coupling for larger beams . . .	104
5.2.1	Coupling coefficient from HG_{00} to HG_{01}	105
5.2.2	Resonance enhancement factor	108
5.2.3	Scaling of the alignment noise coupling	110
5.3	Numerical experiment for scaling validation with FINESSE	111
5.3.1	Selecting a range of arm cavity beam sizes	111
5.3.2	Magnitudes of DC and AC misalignments	113
5.3.3	Mode matching the arm cavities to the recycling cavities	114
5.3.4	Detecting the alignment coupling signal	115
5.3.4.1	Higher order couplings	116
5.4	Implications and potential for suppression	119
5.4.1	Stronger alignment signals	119
5.4.2	Scaling with the other dependent variables	121
5.5	Summary	122

6	FINESSE 3 Software Design	123
6.1	Motivation	124
6.1.1	Modular structure and API	125
6.2	Cython extensions	127
6.3	Simulation workspace pattern	128
6.4	Program flow	129
6.5	Documentation	131
6.5.1	API documentation	133
6.6	Current status and future work	134
6.6.1	Longer-term plans and ideas	134
6.7	Summary	135
7	Summary and Conclusions	137
A	Beam Tracing with Trees	141
A.1	Planting a forest	143
A.2	Constructing the changing trace forest	144
A.3	Finding potentially varying mode mismatch couplings	146
A.4	Dealing with unstable tracing regions	147
A.4.1	What happens when there are no stable trace dependencies?	148
B	Analytic solution to WS phase space overlap integrals	151
B.1	Overlap on an infinite plane	152
B.2	Computing the overlap for a finite circular aperture	153
C	Field couplings for a DC and AC misaligned Fabry-Perot cavity	157
C.1	HG ₀₀ field couplings at DC	159
C.2	HG ₀₁ field couplings at DC	159
C.3	HG ₀₁ field couplings at AC	161
C.4	Computing the signal beat at the photodiode	163
C.4.1	Error estimation	164
	Bibliography	167

List of Figures

1.1	A simplified schematic showing the optical design of the Advanced LIGO detectors with each of the main optical cavities highlighted.	3
1.2	Current design noise budget for advanced LIGO, with corresponding estimates of BNS and BBH detection ranges.	5
1.3	The transverse and longitudinal profiles of a Gaussian beam, illustrating the beam waist, with the near-field region around it, and the far-field regime beyond the waist.	13
1.4	A grid of Hermite-Gauss modes showing their spatial distributions up to an index of two, for both the tangential and sagittal planes.	16
1.5	A representation of the couplings for the bi-directional propagation of a field over a free space.	19
1.6	Local field couplings at a mirror, via the plane-wave picture, in terms of how FINESSE handles such interactions of the light with the surface.	21
2.1	A conceptual view of the transformation of a beam from one modal basis to another, via a simplified optical diffraction process.	28
2.2	Cavity round-trip Gouy phase versus g -factor.	31
2.3	A representation of the round-trip ABCD matrix of a Fabry-Perot cavity, complemented by the profile of the wavefront RoC for the eigenmode.	33
2.4	The different regions of Fabry-Perot cavity stability as a function of the mirror curvatures.	35
2.5	Key cavities constituting the core of an aLIGO-type interferometer.	36
2.6	Fabry-Perot cavity <code>TraceTree</code>	37
2.7	Representation of a “trace forest” for the core of a dual-recycled Michelson interferometer.	38
2.8	Intersection of two separate trace trees at a beam splitter, showing how mode mismatches can be represented in a modal simulation.	39
2.9	Contrast defect for a simple Fabry-Perot Michelson showing scattering into HOMs.	40
2.10	Schematic rationalising symmetric beam tracing operations, using the example of a mirror with matched modal bases across both surfaces.	42
2.11	A Michelson interferometer with flat mirrors exemplifying the symmetric vs asymmetric tracing behaviour.	43
2.12	Difference in circulating power with cavity tracing versus no cavity tracing when scanning the radii of curvature of the cavity mirrors. Figures of merit for the mode mismatch in both cases are also given for illustration.	44

2.13	Difference in output power of detuned Michelson for swapped beam tracing priorities between the input laser mode and cavity modes.	46
2.14	Overlapping recycling cavities in the central interferometer of aLIGO.	47
2.15	An example of using symbolic beam propagation to compute tuning of the aLIGO PRC telescope.	49
2.16	Transforming coordinate systems across two separate beam segments, with an angle of misalignment between the two frames.	51
2.17	Scattering matrices at a mirror.	55
2.18	Performance comparisons of scattering matrix calculations implemented in FINESSE 3, versus Pykat and FINESSE 2.	56
2.19	Comparing computation times for a misaligned beam splitter simulation using the mode selection pattern feature of FINESSE 3.	60
2.20	CCD intensity profiles for a misaligned circulating beam in an aLIGO arm cavity.	63
2.21	A comparison of simulation run-times for calculating a CCD image of a beam, based on the FINESSE 3 implementation versus the previous version.	64
3.1	A general overview of the ET observatory layout, showing the xylophone configuration with a separation of the low and high frequency detectors.	68
3.2	ET tunnel layout displaying a schematic of a corner station of the facility. An example configuration of the arm telescopes is given for illustration.	69
3.3	Schematic of an ET signal recycling cavity with an arm telescope and a lens at the ITM.	72
3.4	Waist size and position of beam matched to an ET-LF arm cavity as function of ITM lens focal length.	74
3.6	ET arm telescope mirror RoC combinations yielding, relatively long, stable SRC with 6 mm waist.	75
3.5	Round-trip Gouy phase in the SRC as a function of the curvatures of the telescope mirrors. Plotted for two separate intra-telescope lengths to highlight the improved stability for relatively large distances.	76
3.7	Round-trip Gouy phase in the SRC, as a function of the telescope mirror curvatures, with an ITM lens of focal length $f = 75$ m.	77
3.8	ET arm telescope mirror RoC combinations yielding, relatively short, stable SRC with 6 mm waist.	78
3.9	Propagation of the ET-LF SRC beam from the ITM to the SRM, showing both the beam sizes at each optic and the accumulated Gouy phase over the cavity.	79
3.10	Propagation of the ET-HF SRC beam from the ITM to the SRM, showing both the beam sizes at each optic and the accumulated Gouy phase over the cavity.	80
3.11	Decreasing the nominal solution distance from ZM1 to BS to show consistent SRC stability.	80

LIST OF FIGURES

3.12	Varying the RoC of the ET-HF ZM1 telescope mirror around the nominal design value and plotting the resulting changes in the spot size on the central beam splitter, and the round-trip Gouy phase of the SRC.	82
3.13	SRC mode changes against small changes in telescope distances and optic curvatures.	83
3.14	Simultaneous changes in ET-HF telescope mirror RoCs required for correcting thermal lensing induced mode mismatches.	84
4.1	WS phase space contours of overlap, against a primary mode given as the SRC mode, of the nominal ET-HF configuration.	90
4.2	Mismatches propagated to the BS in ET-HF as a result of different core optic distortions.	91
4.3	Actuations on the ZM RoCs in WS phase space as compared to other optic distortions.	95
4.4	The effect on the SRC mode given shifts in the positions of the telescope mirrors. Shows that, for this design, such shifts do not have any significant impact and are therefore likely not suitable for adaptive mode matching purposes.	96
4.5	A Gaussian kernel density estimate of the initial region, in WS phase space, which the ET-HF SRC mode can occupy in order for actuators on ZM2 and SRM RoCs to recover mode matching with the arm cavities to 99.9%.	98
5.1	A schematic of a type of alignment coupling in a Fabry-Perot cavity, in which the input mirror is given a static tilt and the end mirror tilts at some signal frequency.	102
5.2	The modal picture, in both DC and signal sideband regimes, for a statically and dynamically misaligned Fabry-Perot cavity. This provides a view of the couplings at the mirrors and the corresponding field conversions.	103
5.3	A comparison of the magnitude of the coupling coefficient k_{01} for an analytic approximation against FINESSE 3.	108
5.4	The g -factors and beam sizes for an aLIGO arm cavity as a function of mirror RoCs. Together these form a representation of the arm cavity geometry regimes used for the alignment coupling analysis.	112
5.5	Fractional misalignment angle contours versus beam sizes on aLIGO arm cavity mirrors and angles of misalignment of a cavity mirror.	113
5.6	The mode mismatches from the aLIGO arm cavities to recycling cavities when sweeping over a range of arm cavity beam sizes. Correspondingly, the round-trip Gouy phases of the recycling cavities are given when shifting the recycling cavity telescope mirror RoCs to mode match back to the arm cavities.	114
5.7	A flow chart giving a broad overview of the steps taken in a FINESSE simulation for computing the alignment coupling as a function of the beam size on the arm cavity mirrors of an aLIGO file.	115
5.8	The alignment to longitudinal coupling signal, for an aLIGO type GW detector, as a function of the beam sizes impinging on the symmetric arm cavity mirrors.	116

5.9	Additional HOM couplings present in the beam size to alignment coupling simulation, against the beam size on the arm cavity mirrors. Demonstrates the small mode mismatch and higher-order misalignment terms present. . . .	117
5.10	Mode mismatches between the core aLIGO cavities as a function of the spot size on the arm cavity mirrors, after minimising the mismatch between the arm and recycling cavities.	118
6.1	The modular structure of FINESSE 3, shown as a “network” of the sub-modules and packages. Each module is separated into a “domain” predicated on the type of operation the code performs.	126
6.2	A representation of the workspace pattern for a typical component of a FINESSE 3 model and simulation. This provides an insight into how optimisations are performed when entering into a simulation state.	129
6.3	The idea to implementation process for creating a custom component in FINESSE 3, highlighting the steps (with relevant options) necessary to write a custom connector type object which can then be used in conjunction with simulations.	130
6.4	An overview of the procedure followed by FINESSE 3 from creating the model, either via the API or KatScript, to running an arbitrary parameter sweep on the model.	132
6.5	An example from the API documentation, displaying the documentation page for the ABCD method of the mirror class. This demonstrates the rich features of the modern, Sphinx-based manual; allowing for embedded vector graphics and rendered equations.	133
A.1	A simplified representation of the <code>TraceTree</code> structure implemented in FINESSE 3. This data structure is used as the core object through which a beam trace from a specific dependency is performed.	142
A.2	A conceptual overview showing how <code>TraceForest</code> containers are planted from a list of ordered trace dependencies in FINESSE 3.	144
A.3	Examples of changing <code>TraceForest</code> structures for simulations where different geometric parameters, at varied locations in an interferometer, are swept. This also highlights how customising the beam tracing order can affect the determined changing modal basis paths of a simulation.	145
A.4	Zooming in on intersecting trace trees from different dependencies at a beam splitter, highlighting how FINESSE 3 can pre-determine which node couplings have changing mode mismatch values during a simulation.	146
A.5	An hypothetical geometric parameter scan with arbitrary unstable cavity regions, showing how FINESSE 3 handles beam tracing, at these data points, via the contingent <code>TraceForest</code> pattern.	147
A.6	Masking of simulation detector outputs in unstable cavity regimes where no other stable trace dependency exists in the model. Uses an aLIGO arm cavity, entering into two different unstable zones, as an example.	149

B.1	Dimensionless overlap, in terms of the beam size and Gaussian defocus, of a primary mode with itself for different ratios of a circular aperture radius to the beam size.	155
C.1	Reference configuration for a DC plus AC misaligned Fabry-Perot cavity. The circulating carrier fields at each location inside the cavity are shown, along with the “injected” fields as pertaining to the tilt couplings.	158
C.2	A comparison of the analytical approximation of the amplitude of the circulating HG_1 field at DC, versus the results of a FINESSE 3 simulation.	162
C.3	An estimation on the error for the derived analytical approximation of the alignment signal, as a function of the tilt angles of the cavity mirrors. Demonstrates that this approach is valid only for small misalignment angles.	165

List of Tables

2.1	Differences in HOM couplings, on reflection from a mirror, using symmetric vs asymmetric tracing.	43
3.1	A summary of the key parameters of the ET-LF and ET-HF arm cavities. . .	71
3.2	Baseline solution data obtained for both the ET-LF and ET-HF arm telescope parameters. The beam sizes and Gouy phase accumulations in each segment are also shown.	78
4.1	Values for the Rayleigh range in each segment of the signal recycling cavity of ET-HF, based on our initial arm telescope design.	97
5.1	Approximations of the scaling of the alignment to longitudinal noise coupling for third generation GW detectors, given as ratios to the estimated aLIGO value.	121

List of source codes

1	Source code for a symbolic beam propagation example required to calculate beam size and accumulated Gouy phase changes, given a change in optic positions in a beam-expansion telescope.	49
2	Algorithm implemented in FINESSE 3 for propagating a beam recursively through a <code>TraceTree</code>	143

Acronyms

BS Beam splitter.

DARM Differential arm length.

ET Einstein Telescope.

ET-HF High Frequency detector(s) of the Einstein Telescope.

ET-LF Low Frequency detector(s) of the Einstein Telescope.

ETM End Test Mass.

FSR Free spectral range.

GW Gravitational Wave.

HG Hermite-Gauss.

HOM Higher-Order Mode.

IFO Interferometer.

IMC Input Mode Cleaner.

ITM Input Test Mass.

LG Laguerre-Gauss.

LIGO Laser Interferometer Gravitational wave Observatory.

OMC Output Mode Cleaner.

PR(2/3) Power recycling cavity telescope mirrors (2/3) in aLIGO.

PRC Power Recycling Cavity.

Acronyms

PRM Power Recycling Mirror.

RoC Radius of Curvature.

SR(2/3) Signal recycling cavity telescope mirrors (2/3) in aLIGO.

SRC Signal Recycling Cavity.

SRM Signal Recycling Mirror.

TEM Transverse Electro-magnetic Mode.

ZM(1/2) ET arm telescope mirror (1/2).

FINESSE Nomenclature

component.port.direction Syntax for an Optical Node. For example, `ITM.p2.o` corresponds to the *output* node of the *second* port of a mirror component called *ITM*.

Gauss A manually defined beam parameter at a specific node of the optical network..

maxtem Maximum TEM order up to which mode indices are included.

Node Represents a specific connection at a component. Mathematically, a node corresponds to a single equation in the full interferometer matrix..

Port A collection of all the nodes at a specific point on, or surface of, a component.

Trace dependency Beam tracing dependency object that a specific optical node relies upon. This is either an optical Cavity or a Gauss instance..

TraceForest The container of **TraceTree** structures, used for performing beam traces through a full optical network and essentially all other tracing related tasks performed in FINESSE 3..

TraceTree A data structure representing an optical node with links to the physically coupled nodes, views to the associated ABCD matrices and references to the specific trace dependency that the node relies upon..

Chapter One

Introduction

The recent first detections of gravitational waves (GWs) from both binary black hole (BBH) [5] and binary neutron star (BNS) [6, 7] mergers herald a new era of modern astronomy. These detections have not only solidified our understanding of General Relativity, they also have the potential to open up new, unexplored avenues of astronomical observations — such as being able to probe the inner structure of neutron stars via a deeper predictability of their equations of state [8, 9]. Tens of GW candidates have now been detected [10, 11] from BBH mergers, with the majority found during the third (and latest, as of writing) observing run (O3) [12]. In addition to the Advanced LIGO (aLIGO) [13] detectors, Advanced VIRGO [14] became a permanent addition to the GW detector network during O3 whilst KAGRA [15] is expected to join before the fourth observing run (O4), improving the sky localisation prospects of GW signal sources [16]. Despite these successes, enhancements to the sensitivity of current and future GW detectors are still required to increase the rates of detection [17] going forward, particularly for BNS mergers, and the distance to which such detections can be reliably made.

Detecting GWs is not a simple task. The initial observations seen in the last few years were a culmination of decades worth of effort, with research and construction expertise spanning most continents. Understanding why gravitational waves are so difficult to detect requires quantification of the amplitude of these waves. This amplitude, as a scalar value, can be defined by [18]

$$h(r) = \frac{2G}{c^4} \frac{1}{r} \frac{d^2 I}{dt^2}, \quad (1.1)$$

which is a dimensionless quantity most commonly referred to as the *strain* of the gravitational wave — a name chosen due to its description of how space-time is stretched and squashed by the passing of a GW. The terms of this equation are as follows: G is the gravitational constant, c is the speed of light, r gives the distance to the source and I represents the

quadrupole moment of the source (which depends upon the density of the system). The factor G/c^4 is incredibly small and, coupled with a typical distance of $r \sim 10$ Mpc, gives

$$h \sim 10^{-68} \frac{d^2 I}{dt^2}, \quad (1.2)$$

in SI units. Correspondingly, for a BNS inspiral with neutron star masses of $1.4 M_\odot$, one obtains a strain of $h \sim 10^{-21}$ (see Section 2.4.2 of [18]).

Due to the strain effect a GW has on space-time, one can characterise the deviation, δL , in a length measurement L by [18]

$$\delta L = hL. \quad (1.3)$$

If we apply the length of the aLIGO detector arms ($L \sim 4$ km) and the strain amplitude from above to this equation, we obtain a deviation $\delta L \sim 10^{-18}$ m. It is this magnitude of distance change that GW detectors must (and indeed can) measure in order to detect gravitational waves — a significant challenge given such a tiny value.

1.1 Interferometers as gravitational wave detectors

Laser interferometry provides an incredibly precise method for measuring small distance differentials, making it an ideal system through which GWs can be detected. The design, and implementation, of current GW detectors employs a modified Michelson interferometer [13], where the underlying principle of operation is that a passing GW will effect a differential change in the optical path length between the interferometer arms — realised as a phase shift in the light. This relative phase shift then results in a change to the optical power at the output, commonly referred to as the anti-symmetric (AS) port, of the interferometer; due to the interference of the beams from the arms.

The optical design of the Advanced LIGO detectors is shown in Figure 1.1. This design forms the basis of the current, state-of-the-art LIGO Livingston Observatory (LLO) and LIGO Hanford Observatory (LHO) facilities. In this figure we show the key optics of the configuration — consisting of the input path, which includes the phase-modulator (EOM) and the input mode-cleaner (IMC) cavity, the core interferometer, including the recycling cavities (PRC, SRC) and Fabry-Perot arm cavities (XARM, YARM), and finally also the output beam path containing the output mode-cleaner (OMC) cavity and detection photodiodes (DCPDs). Not included in this figure are the squeezer and related optics, which were installed for O3 to suppress quantum noise at high frequencies (see §1.2.1 for more details). The radio-frequency (RF) sidebands produced by the EOM, at 9 MHz and 45 MHz, propagate through the central interferometer. These sideband fields are picked off at special

ports on reflection from the PRC (giving the REFL signal) and transmission through PR2 (yielding the POP signal), and are then used as error signals for sensing and control of the length and alignment degrees of freedom in the interferometer [19]. The length degrees of freedom refer to the distances between the core optics in terms of their positions along the optical axis, whilst the alignment variables are the angular pitch (rotation around the x -axis) and yaw (rotation around the y -axis) motions of these optics.




Figure 1.1: A schematic of the aLIGO detectors showing the main input optics, the core interferometer and the key output path components. Highlighted prominently on this figure are each of the optical cavities present in this type of GW detector — i.e. the arm cavities, recycling cavities and mode-cleaner cavities. Also shown is the electro-optic modulator (EOM), which generates radio-frequency (RF) sidebands used for controlling the interferometer, and the output DC photodiodes (DCPDs) which detect a signal which depends upon the differential arm length (DARM); i.e. the degree of freedom that contains the GW information.

We will now discuss the role played by each of the main optical cavities, shown in Figure 1.1, over the course of the following sections. The details are intended to be brief; one can find much more in-depth analyses of these sub-systems in the literature [13, 20, 21].

1.1.1 Arm cavities

The Fabry-Perot cavities in the arms of the interferometer serve several purposes. First and foremost, the presence of these optical resonators allow the light field to build up in the

arms hence increasing the storage time of photons in the interferometer. In addition the arm cavities, by their very nature, increase the power in the interferometer arms which acts to increase the signal-to-noise ratio (SNR) via relative suppression of the shot noise. This is particularly important for the noise profile of the detectors at high frequencies, as we will see in §1.2.

A secondary benefit of using arm cavities is the ability they provide for shifting the highest concentration of power away from the central beam splitter. Without resonators in the arms the power on the central beam splitter is relatively high (when attempting to achieve similar levels of shot noise reduction), which then induces thermal lensing in the beam splitter substrate (a phenomenon identified and characterised at GEO600 [22]).

1.1.2 Recycling cavities

Present at the core of the interferometer are a set of overlapping optical cavities, broadly termed the “recycling” cavities; this arrangement is often called dual-recycling. These cavities can be separated into two distinct types: the power recycling cavities (PRC) and signal recycling cavities (SRC). Note that we use plurality here as the PRC really consists of the power recycling cavities from the PRM to ITMX (often denoted as PRX) and also from PRM to ITMY (denoted as PRY); and similarly for the signal recycling paths, from the SRM to each input test mass.

The power recycling cavity, as its name implies, is responsible for increasing the power in the carrier field which then circulates throughout the core (through the symmetric port of the central beam splitter) and arms of the interferometer [23]. Again, this increase in power is useful for increasing the SNR due to a smaller amount of shot noise relative to the field (which is “generated” from the main carrier light via sidebands at the GW frequency) containing the GW information.

Signal recycling is a newer feature of the current GW detectors, having been implemented in aLIGO [13] whilst power recycling was a feature of initial LIGO. This type of field recycling is currently used to enable the detector bandwidth to be broadened over a larger range of frequencies [23, 24]. It should be noted that this type of recycling also has the potential to be used in a different way in the future, whereby the peak sensitivity at a particular frequency is increased rather than operating in a broadband resonant sideband extraction scheme.

1.1.3 Mode-cleaner cavities

The final two cavities highlighted in Figure 1.1 are the mode-cleaner cavities, IMC and OMC. The former sits in the input path, prior to the power recycling cavity, whilst the latter lies just before the output photodiodes. Both cavities fulfill similar roles but with subtle variations in their tasks. In both cases, these cavities clean the spatial profile of the

beam such that transmission of light in the fundamental carrier mode (HG_{00} in the case of aLIGO) is maximised, inherently minimising any power in higher-order modes. See §1.3.2.3 for details on the spatial properties of laser beams. The OMC is also tuned such that the RF sidebands are filtered out from the carrier light, ensuring that only the fields containing the GW signal reach the output photodiodes. Meanwhile the IMC is tuned in such a way as to allow transmission of these RF fields, as these are used for length and alignment sensing required for controlling the core optics of the interferometer. In addition, the IMC serves as a frequency stabiliser for the carrier field and also provides passive polarisation filtering [25].

1.2 Overview of current detector sensitivity limitations

In this section we will give a very brief overview of the sensitivity limitations of the current detector facilities, in terms of the primary noise sources in different frequency regimes of the detection band. For a more thorough discussion of the current status of all noise sources, see, for example, [26, 27].

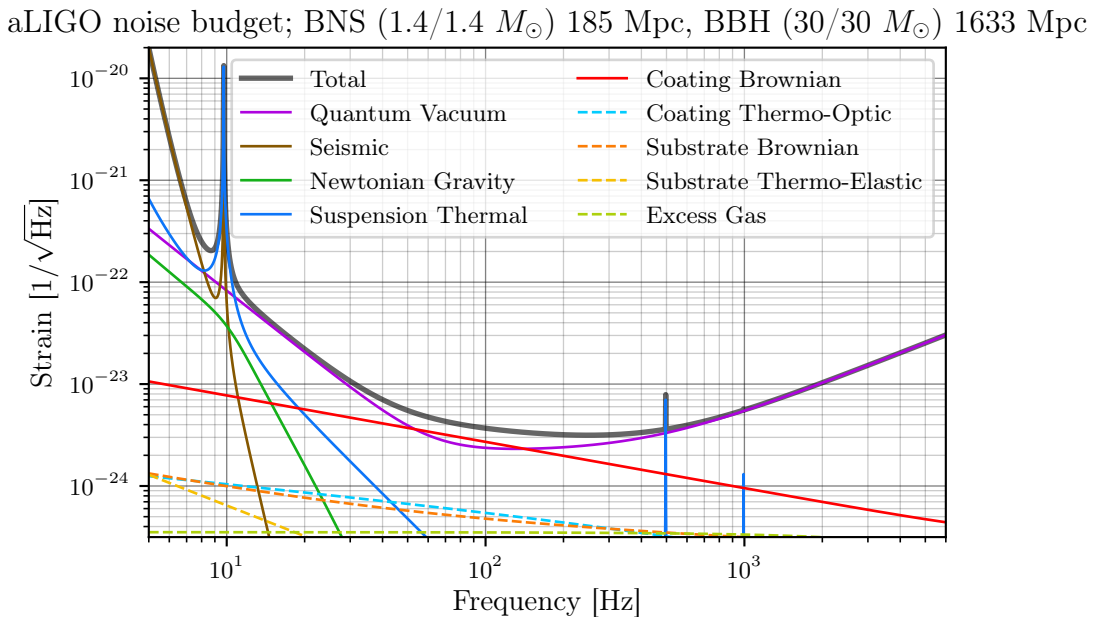


Figure 1.2: This figure shows the current noise budget for aLIGO. The primary noise sources are given as solid lines whilst other, less important, sources are shown as dashed lines. The solid grey line gives the total spectral density of the noise; i.e. the sum of all the individual traces. This noise budget was produced using `pygwinc` [28]. The BNS and BBH ranges were calculated from this design curve using the `inspiral_range` package [29] (see also [30, 31]).

Figure 1.2 shows the current (design) noise budget for aLIGO, as produced by the software package `pygwinc` [28]. Displayed on this figure are the key noise sources, plotted in terms of

their amplitude spectral density, over a range of frequencies describing the detection band of these facilities. The first trace, given by the solid grey line, represents the sum of all the individual noise sources — giving the total noise budget, often referred to as the “design curve” of the detector(s). The data at the top of the figure are the inspiral ranges, computed from this design curve, for both BNS and BBH mergers, assuming some fixed standard masses for each as shown [32, 30]. We will now briefly discuss the key noise sources, i.e. those which currently limit the sensitivity, from Figure 1.2, along with ongoing and future plans for reducing the impact of these sources.

1.2.1 Quantum vacuum noise

The predominant noise source across most of the detection band for current GW detectors is the quantum vacuum noise [33]; often simply referred to as the “quantum noise”. From Figure 1.2 one can observe that this dominates the design noise curve in the regions $f \in [12, 50]$ Hz and $f \gtrsim 150$ Hz. The quantum noise itself consists of two principle effects — quantum radiation pressure noise (QRPN), which is important at low frequencies ($f \lesssim 100$ Hz), and shot noise, dominating at high frequencies ($f \gtrsim 100$ Hz). The QRPN is caused by small amplitude fluctuations of the light field in the arms of the interferometer, thereby inducing displacement noise at the optics via the net force of these fluctuations [34]. Meanwhile, the shot noise is an effect of Poisson statistics [35] arising from the uncertainty in the arrival time of photons at the output photodiode. This latter noise source can be suppressed via utilising a higher power laser beam, as the increased power in the arm cavities commensurately increases the ratio of the signal to the shot noise; this can also be interpreted as a simple statistical effect whereby increasing the photon count of the signal then lessens the effect of the “photon-counting” uncertainty. Increased power, however, will also amplify the amplitude fluctuations of the field, thus increasing the QRPN. This is not an issue given the current practical limits of the aLIGO detectors, as technical noise sources instead dominate at the frequencies at which the QRPN would be expected to be prevalent; as outlined broadly in Section III of [26].

Reduction of the quantum noise, at high frequencies, was achieved in O3 via the application of “squeezing” [36] - a novel technique which is an area of highly active research due to its potential to significantly improve both current and future detector sensitivities. A simple description of squeezing is as follows: a vacuum state may be prepared in such a way that the quantum uncertainty is not equal in the amplitude and phase quadratures of the field [37]. By injecting such a state, squeezed in the phase-quadrature, into the interferometer, the shot noise at the output photodiodes can be pushed below the standard quantum limit (SQL); thus increasing the detector sensitivity at high frequencies¹. A side-effect of such a scheme however, comes from the uncertainty principle. Squeezing the phase-quadrature

¹The astrophysical motivation behind increasing particularly the high-frequency sensitivity of GW detectors, lies in the information behind the BNS merger stage being encoded in GWs at such frequencies. Obtaining these data will offer the ability to probe the equation of state of neutron stars. See, for example, [9, 38].

of the vacuum field necessitates that the amplitude quadrature is then in a state of “anti-squeezing”. The implication of this is that the amplitude fluctuations of the field in the arms of the interferometer increase when injecting such a squeezed field — leading to larger QRPN effects and, therefore, a reduction in the sensitivity at lower frequencies. Again, this is not an issue currently due to technical noises dominating at these frequencies; however it may become a problem in the future as these control-induced noises are reduced. A proposed scheme for solving this issue is known as frequency-dependent squeezing [39, 40, 41]. In short, this is a technique which can allow for amplitude-quadrature squeezing at low frequencies and phase-quadrature squeezing at high frequencies; thereby providing broad-band improvements to the detector sensitivity. Frequency-dependent squeezing can be achieved via reflection of the injected vacuum state from a filter cavity which has a pole at the cross-over frequency of QRPN and shot noise; effectively achieving phase-rotation of the squeezed field. This type of quantum noise reduction scheme is planned for the upcoming aLIGO upgrade, A+ [42, 43], as well as third generation GW detectors, e.g. the Einstein Telescope (ET) [44] and Cosmic Explorer (CE) [45].

1.2.2 Coating thermal noise

Displayed prominently via the red trace on Figure 1.2 is the coating Brownian noise. This is a type of thermal noise, associated with the coatings of the optics, which can produce fluctuations of the mirror surfaces [46] — thereby inducing displacement noise which may contribute to the differential arm (DARM) motion. As shown on the aLIGO noise budget, this noise source dominates the design curve at frequencies $f \in [50, 150]$ Hz. The coating thermal noise is predominantly an effect for the arm cavity mirrors and, in particular, the ITMs. This is due to the high optical power passing through these mirrors, which causes the mirror surface (and substrate) to heat up. The Brownian noise, in terms of the amplitude spectral density, scales with mirror temperature as \sqrt{T} ; see Equation (1.4) below. Reducing the Brownian thermal noise is an area of active research involving, for example, investigations of different types of optic coatings [47, 48].

Suppression of the coating thermal noise via larger beams, for future ground-based GW detectors, is a topic which surfaces several times throughout this thesis — in particular in Chapters 3 and 5. Therefore, it makes sense to quantify the thermal noise contributions explicitly here for reference. The *power* spectral density of the thermal noise, $S_x(f)$, is given by [49]

$$S_x(f) = \frac{2k_B T}{\pi^2 f Y_{\text{sub}}} \frac{d}{w^2} \left(\frac{Y_{\text{coat}}}{Y_{\text{sub}}} \phi_{\parallel} + \frac{Y_{\text{sub}}}{Y_{\text{coat}}} \phi_{\perp} \right), \quad (1.4)$$

where k_B is Boltzmann’s constant, T is the temperature of the mirror, Y_{sub} and Y_{coat} are the Young’s moduli of the substrate and coating materials, respectively, and d is the thickness of the mirror coating. The values ϕ_{\parallel} and ϕ_{\perp} are the mechanical loss-angles of the coating

parallel and perpendicular to the layers of the coating [47], respectively. The key variable for us here is w , which is the radius of the beam (see §1.3.2.2) impinging on the mirror coating. Given the $1/w$ dependence of the *amplitude* spectral density ($\sqrt{S_x(f)}$) on the thermal noise, we can note that increasing the beam size on the arm cavity mirrors will improve the sensitivity of detectors; particularly in the mid-frequency band where this noise source dominates. See the introduction of Chapter 5 for an overview of the detector proposals which include larger beams impinging on (larger) arm cavity mirrors.

1.2.3 Seismic and suspension noise

The other key traces in Figure 1.2, at the lower end of the frequency spectrum, are the seismic (brown) and suspension thermal (blue) noises. Seismic noise refers to the coupling of ground motion to motion at the test masses. This is minimised, in aLIGO, via the deployment of a four-stage (quadruple pendulum) suspension system where the test masses form the final stage of this system [50, 51]. Even with this advanced scheme, however, seismic noise remains a dominant noise source at $f \lesssim 8$ Hz. Research is currently ongoing into the design of new seismic isolation systems to use for third-generation detectors; see [52].

Similarly to the brownian motion of the coating, thermally induced displacement of the suspension system can also occur — translating into motion at the test mass. This is known as the suspension thermal noise, which can be computed via application of the fluctuation-dissipation theorem [53, 54]. In addition to new isolation schemes, research is also ongoing for reducing specifically this type of noise in both current and future detectors. The KAGRA detector [15] was the first modern GW detector to (recently, as of writing) demonstrate cryogenic cooling of test masses for this purpose [55]. Furthermore, considering future detector design, the low-frequency (LF) interferometers of ET (denoted as ET-LF) are currently planned to also use cryogenics (see Sections 6.10.2 and 6.11 of [44]) for cooling the test-masses and suspensions in order to suppress the suspension thermal noise, as well as the thermal noise of the mirror coatings and substrates.

1.3 Mathematical description of light fields for laser interferometry

As a theoretical background to the concepts discussed throughout this thesis, we will now describe the laser beams used in GW detectors in terms of their mathematical properties. In §1.3.1 the description of light fields using a plane-wave basis is given. This is mostly just to provide a foundation, including common notation, for the more complete description using higher-order modes (HOMs) found in §1.3.2. The vast majority of topics in this thesis are based on this latter field picture, due to the capability for modelling transverse beam shape effects that such a framework provides.

1.3.1 The plane-wave picture

The time-varying electric field component of an infinite, plane-parallel and monochromatic (with angular frequency ω) electromagnetic wave can be defined by

$$E(z, t) = E_0 \cos(\omega t - kz + \varphi), \quad (1.5)$$

with $E_0 \equiv E_0(t)$ as the amplitude of the field and $k = \omega/c$ as the wavenumber, whilst z describes the propagation of the field along the optical axis (i.e. perpendicular to the wavefront in this case). The quantity $\varphi \equiv \varphi(t)$ is an absolute phase offset. It is almost always more convenient to express the field as a complex quantity, and so we can recast it as follows (noting that each variable in Equation (1.5) is a real number):

$$\begin{aligned} E(z, t) &= \Re\{E_0 \exp(i(\omega t - kz + \varphi))\} \\ &= \Re\{E'_0 \exp(i(\omega t - kz))\}, \end{aligned} \quad (1.6)$$

where $E'_0 \equiv E'_0(t)$ is then given by

$$E'_0 = E_0 \exp(i\varphi). \quad (1.7)$$

We can use Equation (1.6) to then express our field as

$$a(z, t) = E'_0 \exp(i(\omega t - kz)), \quad (1.8)$$

whilst noting that the only term with any physical meaning here is $\Re\{a(z, t)\}$ — i.e. the field amplitude. The full terms, $a(z, t)$ and E'_0 , are “complex amplitudes” which simply provide a clean mathematical notation that is easier to work with. Throughout this document, such complex amplitudes will typically be denoted with the symbols a and b when referring to single-field amplitudes, and E when referring to sums over these fields.

1.3.2 Gaussian beams and Higher-Order spatial Modes

Whilst the description of laser beams via infinite plane-parallel wavefronts is sufficient for simpler analyses of interferometric systems, it fails to pick up some of the more complicated couplings which exist in practice. In reality the spatial profile of a laser beam is not flat but can be approximated as a Gaussian distribution. The mathematical description of beams as Gaussian modes stems from the use of spherical mirror optical cavities in interferometry;

more details on this can be found later in §2.2.2. In §1.3.2.2 of this section, we will give an overview of the properties of beams described in a Gaussian basis.

Simply using a pure Gaussian mode to describe beam shapes is, however, still not enough to capture the physics involved with complex interferometric systems. In a real interferometer we have, for example, misalignments of optical axes, mode mismatches between resonators, clipping due to finite optic sizes (often referred to as apertures) and imperfections of optical surfaces. All of these effects lead to distortions of the beam away from the ideal Gaussian model. To mathematically describe such distortions we use a sum of spatial components called higher order modes (HOMs). A brief introduction to this concept, in the specific context of the Hermite-Gauss (HG) modal basis, is given in §1.3.2.3.

The physics detailed in this section are intended to cover only those aspects of HOMs which will be depended upon in the following chapters. There are several sources which contain more thorough reviews of HOMs and Gaussian beams, these include [56, 57, 58, 59]. Note, also, that we will focus only on the HG mode basis; another set of solutions to the paraxial wave equation are the Laguerre-Gauss (LG) modes, details of which can be found in, for example, [60, 61, 62].

1.3.2.1 The wave equation and paraxial approximation

The general wave equation for electromagnetic waves, as derived from the Maxwell equations in a vacuum, can be written as [63]

$$\ddot{\mathbf{E}} = c^2 \nabla^2 \mathbf{E}, \quad (1.9)$$

where $\mathbf{E} \equiv \mathbf{E}(t; x, y, z)$ is the temporal and spatial varying electric field vector. We can assume that the transverse profile of the electric field (i.e. the beam shape) is described by a time-independent spatial distribution function $u(x, y, z)$, such that the normalised electric field, as a scalar quantity, is then

$$E(t; x, y, z) = u(x, y, z) \exp(-i(kz - \omega t)), \quad (1.10)$$

with $k = \omega/c$ as the wavenumber and ω as the angular frequency of the field. Substituting this field into Equation (1.9), and computing the time differential, yields

$$-\frac{\omega^2}{c^2} u(x, y, z) \exp(-i(kz - \omega t)) = \nabla^2 (u(x, y, z) \exp(-i(kz - \omega t))). \quad (1.11)$$

Re-arranging this, expanding the Laplacian and evaluating the derivatives of the z dependent

terms then gives,

$$\left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}\right) u(x, y, z) + \frac{\partial^2 u(x, y, z)}{\partial z^2} - 2ik \frac{\partial u(x, y, z)}{\partial z} = 0, \quad (1.12)$$

where we have already noted that $-k^2 + \omega^2/c^2 = 0$ (by definition) and cancelled out the common $\exp(-i(kz - \omega t))$ terms of the electric field vector.

We can now make an assumption, known as the paraxial approximation [64, 65], that the spatial profile $u(x, y, z)$ is a function which varies slowly over the beam axis z . This allows us to state that

$$\left|\frac{\partial^2 u(x, y, z)}{\partial z^2}\right| \ll \left|2k \frac{\partial u(x, y, z)}{\partial z}\right|, \quad (1.13)$$

so that we can neglect the left-hand side term of this in Equation (1.12). It then trivially follows that we can write the partial differential equation of the spatial distribution as

$$\left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}\right) u(x, y, z) - 2ik \frac{\partial u(x, y, z)}{\partial z} = 0, \quad (1.14)$$

which is the paraxial wave equation. A general solution to this equation in the form of a linear superposition of HG modes is shown in §1.3.2.3. First, however, we will discuss the properties of the lowest order of this modal solution set — namely the pure Gaussian beam.

1.3.2.2 Properties of Gaussian beams

Gaussian beams exhibit a classic radial profile in the transverse (x - y) plane, characterised by an intensity distribution $I(r)$, at a radius r from the centre of the beam, of [56]

$$I(r) = I(0) \exp\left(-\frac{2r^2}{w^2(z)}\right); \quad \text{with } I(0) = \frac{2P}{\pi w^2(z)}, \quad (1.15)$$

where w is the radius of the beam². This is defined as the radius at which the intensity drops off by $1/e^2$ from the maximum intensity $I(0)$. The quantity P is simply the total power of the beam. The beam size evolves with the distance the beam propagates over the z -axis and is at a minimum (termed the “waist size” and denoted with w_0) when the Gaussian

²Note that throughout this thesis we will typically refer to this quantity as simply the “beam size”, as it is often written in the literature (e.g. [56, 58]).

beam is at its waist position (a coordinate denoted as z_0). A pure Gaussian beam, of a known wavelength, is completely defined in terms of the w_0 and z_0 parameters — with these quantities a number of other useful properties can be derived.

The concept of a Gaussian beam, with a fixed waist size and position, is shown in an illustrative sense in Figure 1.3. In this figure there is the added detail of a potentially astigmatic beam, which is a beam consisting of different radii in the x - z (tangential) and y - z (sagittal) planes³. For the equations which follow in this section we will assume a non-astigmatic beam, however such quantities can also be applied equally to the two separate planes; as implied by the symbols given on Figure 1.3. An important aspect to note from this figure is that the profile of a Gaussian beam can be separated into two “longitudinal” regimes, namely the *near-field* and *far-field* regimes. The former is the zone near to the beam waist and is quantified by the Rayleigh range z_R , which is defined as

$$z_R = \frac{\pi w_0^2}{\lambda}, \quad (1.16)$$

where w_0 is the waist size and λ is the wavelength of the beam. The latter regime is then defined by a position $z \gg z_R$ — i.e. far beyond the waist position of the beam. We can use the Rayleigh range, in combination with the waist position, to define an expression yielding the size of the beam at any point along the z -axis [58]:

$$w(z) = w_0 \sqrt{1 + \left(\frac{z - z_0}{z_R} \right)^2}. \quad (1.17)$$

Again, this is shown conceptually in Figure 1.3.

A key quantity also displayed in Figure 1.3 is the radius of curvature of the beam wavefront. This is defined as the “surface” across the transverse plane of the beam which has a uniform phase [56]. The radius of curvature can be computed as

$$R_C(z) = z - z_0 + \frac{z_R^2}{z - z_0}. \quad (1.18)$$

This property is particularly important as it is necessary to quantify when designing optical resonators, and optics for interferometry in general. For standing wave cavities, the radius of curvature of optical surfaces (e.g. of mirrors and beam splitters) must be equal to $R_C(z)$, at a given position z along the beam axis, for the surface to be mode matched to the beam at that point. The concept of mode matching will be covered in more detail in §2.2.3.2. Another

³In some cases in the literature this is referred to as an “elliptical” beam, whilst an astigmatic beam is explicitly the case where the waist positions differ between the two transverse planes. Throughout this thesis, however, we will use the term “astigmatism” to refer to both of these scenarios.




Figure 1.3: An illustration of the transverse and longitudinal profiles of a Gaussian beam. The cross-sections (in the x - y plane) show the radial intensity distribution of such a beam, whilst the propagation over the z -axis displays the near-field and far-field regimes; as described in the text. At the longitudinal location $z = z_0$ the beam is at its waist-position; with the radius of the beam given by $w_{0,x}$ and $w_{0,y}$ in the x - z (tangential) plane and y - z (sagittal) plane, respectively. After propagating some distance $z_1 - z_0$, these radii values are then $w_x(z_1)$ and $w_y(z_1)$ as shown. The angles $\Theta_{x|y}$ are the divergence (or diffraction) angles of the beam in each plane, representing the far-field angle between the z -axis and the beam radius. Note that the left-handed coordinate system (traversing with the beam, in the positive z direction) used here corresponds to the convention we use in FINESSE 3.

useful quantity, inexorably linked to $R_C(z)$, is the curvature of the wavefront, defined by

$$S(z) = \frac{1}{R_C(z)} = \frac{(z - z_0)}{(z - z_0)^2 + z_R^2}. \quad (1.19)$$

The curvature of the beam is characterised by the following values:

$$S(z) = \begin{cases} 0, & \text{if } z = z_0 \\ 1/2z_R, & \text{at } z - z_0 = z_R \\ 1/z, & \text{when } z \gg z_R, \end{cases} \quad (1.20)$$

which corresponds to an anti-symmetric (around z_0) function which is linear around zero in the near-field regime and falls-off as the reciprocal of the propagated distance from the waist in the far-field. The curvature value of zero at the beam waist indicates that the wavefront is flat at this point. As we will see in Chapter 4 the curvature is a useful property, in combination with the beam size, for crafting mode mismatch contours and visualising actuator responses in this variable space.

Finally, the last symbol shown on Figure 1.3 is the divergence (also called the diffraction or far-field) angle of the beam — denoted by Θ . This quantity represents the angle between the propagation axis and the axis constructed via a projection from the waist to the beam size in the far-field (see the solid black lines extending from the waist position on Figure 1.3). The divergence angle can be computed by

$$\Theta = \arctan\left(\frac{w_0}{z_R}\right) \approx \frac{\lambda}{\pi w_0}, \quad (1.21)$$

where the approximation here assumes that this angle is small⁴. This angle is important especially for modelling HOMs, as it forms the metric by which one can estimate the order of scattering into higher order modes given the misalignment angle between two optical axes. We will touch upon this further in §2.3.3.1, whilst in Chapter 5 a practical example of this, in terms of the misalignments of an aLIGO arm cavity, is given.

1.3.2.3 Spatial distortions as Higher-Order Hermite-Gauss modes

Solutions to Equation (1.14) can take many forms. In this section (and throughout this thesis) we will refer to only one such family of solutions known as HG modes. We can write the HG-based solution to the paraxial wave equation as a linear combination of these modes [58]:

$$u(x, y, z) = \sum_{n,m} a_{nm} u_{nm}(x, y, z) \quad (1.22)$$

where n , m are the tangential and sagittal mode indices, respectively, and a_{nm} are the complex amplitudes of each mode. The quantities $u_{nm}(x, y, z)$ are then the spatial profiles of each HG _{nm} mode. The order of any given mode is defined as the sum of its mode indices, i.e. $n + m$. The orthonormality of the HG modes means that we can separate the spatial distribution by plane as follows [57]:

$$u_{nm}(x, y, z) = u_n(x, z) u_m(y, z) \quad (1.23)$$

⁴In fact, a re-cast form of the paraxial approximation given in §1.3.2.1 *requires* that the divergence angle is sufficiently small (i.e. $\Theta \ll 1$).

with $u_n(x, z)$ then defined as [56]

$$u_n(x, z) = \left(\frac{2}{\pi}\right)^{\frac{1}{4}} \left(\frac{\exp(i(2n+1)\psi_x(z))}{2^n n! w_x(z)}\right)^{\frac{1}{2}} H_n\left(\frac{\sqrt{2}x}{w_x(z)}\right) \exp\left(-i\frac{kx^2}{2R_{C,x}(z)} - \frac{x^2}{w_x^2(z)}\right), \quad (1.24)$$

where $H_n(x')$ are the physicist's Hermite polynomials and $w_x(z)$, $R_{C,x}(z)$ are the explicitly tangential plane selections of the beam size and radius of curvature of the wavefront (as defined in §1.3.2.2). The expression for $u_m(y, z)$ is equivalent in form, swapping $n \rightarrow m$, $x \rightarrow y$ and $w_x, R_{C,x}, \psi_x \rightarrow w_y, R_{C,y}, \psi_y$. Again, for a non-astigmatic beam one can simply define $w(z) = w_x(z) = w_y(z)$; and similarly for the other properties.

The quantity $\psi_{x|y}(z)$ is the Gouy phase [66] which is defined by

$$\psi_{x|y}(z) = \arctan\left(\frac{z - z_{0,[x|y]}}{z_{R,[x|y]}}\right). \quad (1.25)$$

This describes the additional phase picked up by a Gaussian beam as compared to a plane-wave field. For higher-order modes in general, the extra phase accumulated is

$$\Psi_{nm}(z) = \left(\frac{1}{2} + n\right) \psi_x(z) + \left(\frac{1}{2} + m\right) \psi_y(z), \quad (1.26)$$

or for a non-astigmatic beam with $\psi(z) = \psi_x(z) = \psi_y(z)$,

$$\Psi_{nm}(z) = (1 + n + m) \psi(z), \quad (1.27)$$

i.e. proportional to the *order* of the HG mode. This higher order phase accumulation term shows up in the equations, in the first exponential term, inherently when taking the product of $u_n(x, z)$ and $u_m(y, z)$ as Equation (1.23) requires. The mode order dependency of the HOM phase accumulation is an important concept as it indicates differing phases for resonance conditions of disparate modes. This is a phenomenon which is exploited by mode cleaner cavities, for allowing transmission of certain modes whilst filtering out others; as implied in §1.1.3.

The transverse shapes of $u_{nm}(x, y, z)$, for a single point z on the beam axis, are displayed in Figure 1.4 — using HG modes up to the second order in either plane. These profiles show the *intensity* distributions of each mode, i.e. the quantity $u_{nm}(x, y, z) u_{nm}^*(x, y, z)$; see §2.4.1 for a more general definition including the complex amplitudes of the fields.

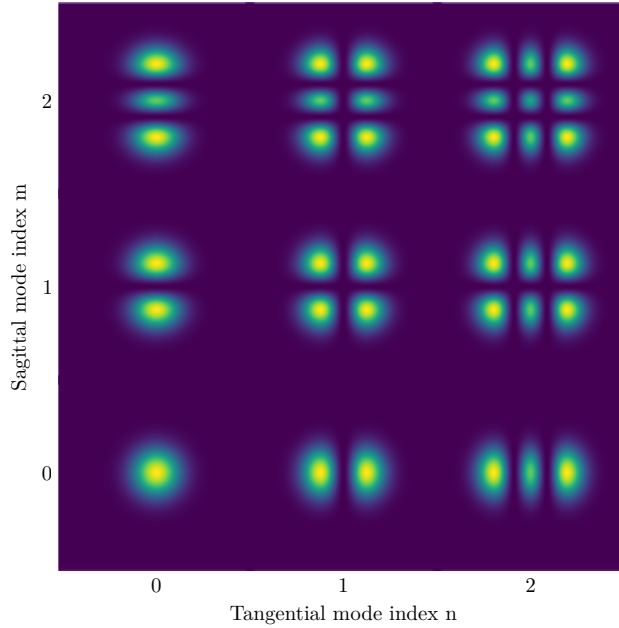


Figure 1.4: This figure shows a grid of the intensity patterns of HG modes up to a mode index of two for both the tangential mode index (on the x -axis), n , and sagittal mode index (on the y -axis), m . The mode index, in either plane, corresponds to the number of “dark-fringes” in that plane of the profiles of each spatial mode. Note that the HG_{00} mode is just a simple Gaussian distribution. This set of profiles was produced using the beam imaging capabilities of FINESSE 3, see §2.4 for further details.

1.3.2.4 The Gaussian beam parameter

A very powerful encapsulation of the properties of a Gaussian beam is the “beam parameter” $q_{x|y}(z)$, defined as [59]

$$\frac{1}{q_{x|y}(z)} = \frac{1}{R_{C,[x|y]}(z)} - \frac{i\lambda}{\pi w_{x|y}^2(z)}, \quad (1.28)$$

or often more conveniently:

$$q_{x|y}(z) = (z - z_{0,[x|y]}) + iz_{R,[x|y]}. \quad (1.29)$$

This is a complex number consisting of the distance to the waist position as the real part and the Rayleigh range as the imaginary part. Recall from §1.3.2.2 that one can completely define a Gaussian beam with the waist size and position, thus with Equation (1.29) we have a quantity which encapsulates both of these properties in one complex variable. Each of these properties defined in §1.3.2.2 can be re-written in terms of the Gaussian beam parameter

and, equally, the spatial profile $u_{nm}(x, y, z)$ can also be re-defined in terms of this. Section 9.8 of [56] gives a good summary of these re-formatted terms. Note that each z term in $u_{nm}(x, y, z)$ will then be implicitly included in the beam parameter $q(z)$, often leading to an alternative form of this function appearing as $u_{nm}(x, y; q)$ — or, more explicitly for an astigmatic beam, $u_{nm}(x, y; \vec{q}) \equiv u_{nm}(x, y; q_x, q_y)$. Such nomenclature will be used later on in §2.3, when dealing with coupling coefficients between HG modes.

In addition to providing a concise and complete description of a pure Gaussian beam, the beam parameter is also useful for defining a mathematical treatment of the interactions of Gaussian beams with optics. See §2.2.1 for details on this framework.

1.4 Modelling interferometers in the frequency-domain

GW detectors are designed to operate in a steady-state, where control loops act to keep the interferometer locked at operating points. In addition, the interaction of light fields with the optics used in these detectors can be well-described by linear couplings. Furthermore, the laser sources themselves are incredibly stable (at a wavelength of $\lambda = 1064$ nm in the case of aLIGO) [67, 68] which enables us to model the main, carrier light field as being monochromatic. We can assume that the frequency of the light field produced by these lasers does not vary with time. It is for these reasons that modelling interferometers for GW detectors is typically best suited to frequency-domain analyses. As we will see throughout the course of this thesis, the size of the system of linear equations (as well as the length and complexity of the equations themselves) grows quickly with the number of components, frequencies and HOMs included in a model. Thus, it has become a necessity to build interferometer simulation tools for numerically modelling such configurations — especially when considering complicated spatial effects involving full detector models. Fortunately, computational modelling lends itself well to solving the types of linear couplings associated with frequency-domain interferometry. Calculating solutions to sets of linear equations is a well-established field of scientific computing [69, 70, 71], with many fast matrix-based algorithms having been developed for performing such tasks [72].

One of the most widely-used interferometer simulation tools in the GW detector community is FINESSE [73] along with the Python package Pykat [74] which acts as a wrapper and set of utilities around this software. FINESSE is a frequency-domain modelling program which is used for solving optical couplings in interferometric systems. At its heart, the code boils down to the inversion of an (often large) sparse matrix of couplings between fields, at different locations in the modelled interferometer, in order to resolve these field amplitudes; this will be detailed briefly in §1.4.3.

A significant portion of this thesis is dedicated to the topic of the new version of this software, FINESSE 3 [75], which has been under development since the start of my PhD. This new version will address challenges faced in modelling future GW detectors. See Chapter 2 for

details on my implementation of HOM features in FINESSE 3, and Chapter 6 for an overview (including motivations) of the design of this new FINESSE version. The underlying principles of solving optical couplings in interferometers, in the frequency-domain, given in this section apply equally to both FINESSE 3 and older versions of the software.

1.4.1 Time-independent light field amplitudes

In §1.3.1 we gave a general formula for the electric field component of a (monochromatic, planar) electromagnetic wave via Equation (1.8). By entering into the frequency-domain, we can (by definition) state a fixed time for any analyses. Naturally, it makes sense to choose $t = 0$ such that this equation simply becomes

$$E(z) = E'_0 \exp(-ikz). \quad (1.30)$$

Similarly in the modal picture, using HG modes (see §1.3.2.3), our electric field equation is then

$$E(x, y, z) = \exp(-ikz) \sum_{n,m} a_{nm} u_{nm}(x, y, z) \quad (1.31)$$

where each symbol here corresponds to those given previously in Equation (1.22). These equations form the framework of optical field descriptions for steady-state interferometer simulations in plane wave and modal bases, respectively.

1.4.2 Component field interactions as coupling matrices

We will now summarise a convenient representation of optical field couplings, in the form of matrices where each element describes how a single field, as measured at a specific point, couples into another field. This section will treat only plane wave, monochromatic couplings to keep the equations compact. In addition, we will look at only propagation over free spaces and couplings at a surface which is at a normal to the optical axis (simply termed a mirror in FINESSE nomenclature); these concepts can be equally applied to other local systems, such as couplings at lenses or non-normal incidence surfaces (i.e. beam splitters). An expanded treatment, including HOMs, for the coupling matrices of a mirror is given later on in §2.3.2.2.

1.4.2.1 Propagation over a space

Figure 1.5 gives a picture of the fields present in a free space type system as used by FINESSE. The input fields, i.e. the fields propagating into the space, are denoted as a_1 and b_1 , whilst the output field, traversing out of the space, are given as a_2 and b_2 . If the length of the space is L , and it can be described by a medium with refractive index n_r , then, as Equation (1.30) implies, the couplings can be defined by

$$a_2 = \exp(-ikn_r L) a_1, \quad \text{and} \quad b_2 = \exp(-ikn_r L) b_1. \quad (1.32)$$

Note that the equations for the propagation in both directions are equal. From now on, in this section, the symbol b_{in} represents both a_1 and b_1 whilst b_{out} represents a_2 and b_2 . As we are working in the frequency-domain, it is often useful to rewrite this explicitly in terms of the frequency of the light field:

$$b_{\text{out}} = \exp\left(-i\frac{2\pi}{\lambda}n_r L\right) b_{\text{in}} = \exp\left(-i\frac{\omega n_r L}{c}\right) b_{\text{in}}, \quad (1.33)$$

where ω is then the angular frequency of the field and c is the speed of light.




Figure 1.5: A representation of the couplings for the propagation of the “input” fields, a_1 and b_1 , to the “output” fields, a_2 and b_2 , over a free space of length L described by a medium with refractive index n_r . In terms of FINESSE, a space can be some distance between two optics (where we typically set $n_r = 1$) or it can represent the substrate of, e.g., a mirror which then takes on a refractive index value corresponding to the substrate material.

In FINESSE all field frequencies are given as offsets, $\Delta\omega = \omega - \omega_0$, from the default frequency ω_0 (which is determined from the default wavelength of the model, i.e. $\omega_0 = 2\pi c/\lambda_0$). This is necessary to avoid numerical errors which would arise from using (large) absolute frequency values and also provides a convenient way of enabling default resonance conditions. As a

result of this, we can recast Equation (1.33) as⁵

$$b_{\text{out}} = \exp\left(-i\frac{\Delta\omega n_r L}{c}\right) b_{\text{in}}. \quad (1.34)$$

We can write these simple, linear equations in a compact form using a notation referred to as the “coupling matrix”. In this case:

$$\begin{pmatrix} a_2 \\ b_2 \end{pmatrix} = \begin{pmatrix} 0 & s_{12} \\ s_{21} & 0 \end{pmatrix} \begin{pmatrix} a_1 \\ b_1 \end{pmatrix}, \quad (1.35)$$

where $s_{12} = s_{21} = \exp(-i\Delta\omega n_r L/c)$. Moving to the modal picture simply requires including the necessary Gouy phase terms in the exponent of these coefficients; this is predicated on the convention, used in FINESSE, that mode mismatches occur only at component boundaries rather than across spaces. More details on this can be found in §2.3.2.1.

1.4.2.2 Coupling matrices at a mirror

The other local field couplings we will cover now are those at a mirror — defined in FINESSE as a thin surface, with associated reflectivity, transmissivity and loss values, which is normal to the optical axis. In addition to the transmission couplings a mirror also sees reflection interactions and, hence, each output field can then be represented as a linear sum of the two corresponding input fields. This is shown conceptually in Figure 1.6 using a flat surface with the input fields as a_1 and b_1 , and outputs as a_2 and b_2 .

Similarly to §1.4.2.1, we can describe these interactions using a matrix of individual couplings:

$$\begin{pmatrix} a_2 \\ b_2 \end{pmatrix} = \begin{pmatrix} m_{11} & m_{21} \\ m_{12} & m_{22} \end{pmatrix} \begin{pmatrix} a_1 \\ b_1 \end{pmatrix}. \quad (1.36)$$

The elements m_{12} , m_{21} are the transmission couplings from the front to back surface and vice-versa, and are equally given by $m_{12} = m_{21} = it$; where t is the amplitude transmissivity of the mirror. This employs the FINESSE convention of a 90 degree phase shift on transmission through a surface (see Section 3.3.1 of [76]).

⁵There is a subtlety here in the way in which FINESSE handles resonance conditions. Spaces are always, by default, resonant for the carrier light. Shifting these resonances involves tuning the microscopic positions of surfaces; as we will see in §1.4.2.2. Good summaries of these conventions are documented in Section 3.3.2 of [76] and Section 2.5 of [56].




Figure 1.6: This figure shows a conceptual overview of the local field couplings at a mirror (which is treated as being infinitesimally thin), via the plane-wave picture, in terms of how FINESSE handles such interactions of the light with the surface. The quantity ϕ is the “tuning” of the mirror, which describes the microscopic displacement as a unitless value. This value is inexorably linked to the wavelength of the beam, whereby $\phi = 360^\circ$ means that the mirror is displaced by one wavelength. In FINESSE, the convention used for the direction of ϕ is that positive values are those tunings which traverse with the normal angle to the front surface of the mirror; as shown. The variables r and t shown on this figure are the *amplitude* reflectivity and transmissivity of the mirror, respectively.

The couplings on reflection are slightly more complicated as they depend upon the microscopic displacement of the surface — a quantity referred to as the “tuning” (ϕ), described in the caption of Figure 1.6. In terms of the angular frequency, ω , of the light field, these reflection couplings can be written as [76]

$$m_{11} = r \exp\left(i 2\phi \frac{\omega}{\omega_0}\right), \quad \text{and} \quad m_{22} = r \exp\left(-i 2\phi \frac{\omega}{\omega_0}\right), \quad (1.37)$$

where r is the amplitude reflectivity of the mirror and ω_0 is the default carrier field frequency. In §2.3.2.2 we will investigate how the inclusion of HOMs expands upon this coupling matrix picture for light interactions at a mirror.

1.4.3 Solving optical interferometer couplings via sparse matrix inversion

From the coupling matrix framework, it follows logically that we can express the steady-state optical field couplings in an interferometer via a single system of linear equations. This allows us to define one equation to be solved in order to obtain the complex field amplitudes

at each node of the configuration:

$$\mathbf{M}_{\text{IFO}} \vec{x} = \vec{y}. \quad (1.38)$$

Here \mathbf{M}_{IFO} is the full interferometer coupling matrix which describes all the couplings between fields (at each frequency and of each HOM) at all the optical nodes of the model. The vector \vec{x} represents the solution vector containing each of these complex field amplitudes, meanwhile \vec{y} is the vector of right-hand sides (RHS) which consists of the injected field quantities; i.e. from input laser beams in the model. Solving Equation (1.38) requires inverting the full interferometer matrix, i.e:

$$\vec{x} = \mathbf{M}_{\text{IFO}}^{-1} \vec{y}. \quad (1.39)$$

This is a computationally expensive process, especially for simulations involving many nodes, field frequencies and HOMs as the order of \mathbf{M}_{IFO} (i.e. number of rows, or columns as this matrix is always square) scales linearly with each of these. Fortunately, this interferometer matrix is inherently sparse⁶ due to, in general, no cross couplings existing between fields of nodes which are separated by more than one space (as implied by the coupling matrices of the previous section). This allows us to store the matrix via a sparse data structure format for memory efficiency. In FINESSE 3 we use the compressed sparse column (CSC) format for improved performance, compared to other sparse formats, in terms of column slicing and computing the products of the matrix with a vector [77, 78]. In addition, this format is compatible with KLU [79] which is a well-established and efficient library for solving sparse systems of linear equations. This library is employed by FINESSE as the core tool for calculating Equation (1.39) as required.

A good overview of how the interferometer matrix is structured in FINESSE, in terms of the individual coupling matrices of the components included in a model, can be found in Section 1.3.3 of [80].

1.5 Thesis overview

The central theme of this thesis lies in the impact of the transverse properties of laser beams for future gravitational wave detectors — with particular focus placed on consequences for the Einstein Telescope. More specifically, the topic of larger beams in such detectors is an important subject throughout this document. The impact this will have is a key motivation

⁶In fact \mathbf{M}_{IFO} tends to be very sparse, especially for large configurations such as representations of full GW detector files. For a typical aLIGO file, with 5 frequency bins and no HOMs, a rough estimate of the sparsity is 99.9% (or a density of just 0.1%). This sparsity value would typically be even larger when including HOMs.

of Chapter 3, in which a novel design for mode matching telescopes in the arms of the ET detector interferometers is introduced and analysed. The results of this chapter, also published in my paper “*Feasibility Study of Beam Expanding Telescopes in the Interferometer Arms for the Einstein Telescope*” [1], demonstrate that it is (to first order) viable to employ such a recycling cavity design. These results also provide a good starting point from which future studies on, e.g., the astigmatism requirements for such a design can be carried out. As a follow up to this initial telescope design, Chapter 4 then includes an analysis of the feasibility of using the arm telescope mirrors for adaptive mode matching. A key conclusion, from this chapter, is that actuating on the radius of curvature of the mirror ZM2 provides a potential method for correcting for small thermal lens distortions (up to around $f_{\text{th}} \sim 65$ km) at the ITM in ET-HF. Another outcome from this analysis is a characterisation of the initial area of phase space that the ET-HF SRC mode could occupy, whilst retaining the ability to recover mode matching to the arm cavities, via actuations on the ZM2 optics and recycling mirrors.

Another important aspect related to the topic of larger beams is the potential impact on the alignment to longitudinal noise coupling. This is the subject of Chapter 5. In this chapter an analytical treatment of this noise coupling is given in terms of its scaling with the size of the beams in the arm cavities of an interferometer; where a nominal w^6 relation is derived. A numerical validation of this scaling relation is also provided, using a full aLIGO detector file to model the beam size to alignment coupling transfer function via FINESSE. Despite the problematic-looking beam size to the power of six result, this chapter ends by clarifying that this increased noise coupling can be overcome due to the increased strength of alignment signals that comes with this, and the potential to tune the design via, e.g., the arm cavity finesse to significantly reduce this scaling. Order of magnitude estimates for the “true” alignment to longitudinal coupling, as ratios to aLIGO, are provided to further show that the real impact of increased beam sizes is minimal for most of the currently planned third-generation GW detectors.

Due to the complexity of accurately modelling the spatial mode effects discussed both in this thesis and otherwise, it is vital (for any research where HOMs are involved) that a reliable, fast and well-documented set of free and open-source simulation tools is available. A major aspect of this thesis is therefore dedicated to my contributions towards the new version of FINESSE 3 — a new interferometer modelling software package. The overarching design of this suite is discussed in Chapter 6; including the key motivating reasons for developing this set of tools and how the design chosen should future-proof the software for many years to come. As this thesis focuses principally on HOMs, a whole chapter (see Chapter 2) is allocated for delving into my implementation of modal simulation features in FINESSE 3. In particular, the new beam tracing library is discussed in detail — the utility that this already provides is highlighted further again in Chapters 3 and 4 where these new tools were used for my aforementioned ET design work. Additional outcomes from Chapter 2 demonstrate that the HOM modelling capabilities of FINESSE 3 provide significant performance improvements over previous versions of the software — an important development given the need for fast turn-around of simulation results, particularly for detector commissioning related tasks.

Chapter Two

Modelling Higher-Order Spatial Modes in Gravitational Wave Detectors with FINESSE 3

Parts of this chapter contain modified text and figures from my work documenting the Higher-Order Mode features [81] of FINESSE 3.

Higher-order modes can arise from several sources in gravitational wave detectors. Misalignments of the optical axes of each cavity (see §§ 1.1.1 to 1.1.3), as well as mode mismatches (see §2.2.3.2) between these cavities, will lead to scattering into HOMs. In particular, the latter effect can specifically arise from thermal aberrations in the mirror substrates (particularly the ITMs) caused by non-uniform temperature distributions as a result of localised heating by the beam [82]. This thermal lensing is especially strong for high power laser beams [83]; an important consideration due to the necessity for future GW detectors to go to higher powers for suppressing shot noise (see §1.2 for a very brief overview of current sensitivity limits).

As HOMs typically do not contain any of the GW signal, any light in these modes can be considered as “junk-light” as it will be detected at the output photodiode and lead to increased shot noise; causing a degradation of the detector sensitivity in the high-frequency regime. In addition, this light can couple back into the fundamental carrier light via, e.g., alignment noise coupling; as we will investigate later in Chapter 5. The suppression of this “alignment to longitudinal” coupling via alignment control schemes [84] can also contribute to the low-mid frequency regime of the detection band [26]. Other detector phenomena which involve HOMs include (but are not limited to) parametric instabilities [85, 86, 87], point absorbers [88] and the degradation of squeezing due to mode mismatches [89, 90]. All of these effects are complicated and can affect, and / or be affected by, very high mode orders.

Due to these, and other, HOM related issues in GW detectors it is important to understand

the overall effect that they can have on detector sensitivity and how these repercussions can be mitigated. The complicated, and sometimes untractable, analytics involved with HOMs often means that modelling them via simulations is the best solution for understanding their effects in the interferometer. In this chapter, we will cover how HOMs can be modelled via Hermite-Gauss (HG) modes, using coupling coefficients, in the context of our software package FINESSE 3. In §2.2 we will explore how the modal bases are set at each node of an interferometer network, and why this choice of beam parameter can be important for computing HOM couplings. Then we will show how coupling coefficients are calculated and used to model scattering of light into HOMs via modal scattering matrices in §2.3. Finally, the detection of HOM content in light fields is discussed, in the context of capturing images of the beam shape, in §2.4. First, however, it is important to understand the capabilities of current interferometer modelling software with regards to modal simulations — and why specific changes and improvements have been made in FINESSE 3 to build upon this existing framework.

2.1 Motivation and current software review

There are several existing interferometer modelling tools in the GW instrumentation community, each of which is designed with specific tasks and goals in mind. As an example, Optickle [91, 92] is an opto-mechanical modelling tool which can be used to simulate interferometers with a single light wavelength; in addition to RF sidebands. This tool is designed for quick interferometer studies where HOM effects are not an issue — thus, Optickle only models light fields up to the first HG order. As a result of this Optickle is not the best tool to use for modelling the spatial properties of beams, especially in the case of mode mismatches (see §2.3 for more information on this topic). A different class of simulation type which can handle additional HOM effects is FFT propagation. The main packages employing this method of simulation are SIS [93], FOGP [94] and OSCAR [95] — all of which are Matlab based FFT suites. This type of simulation is suitable for modelling simple cavities which include HOM considerations. In particular, it is useful for modelling scenarios which are not easily covered by modal-type simulations (more on this class of simulation in a moment) such as: diffraction losses in both stable and unstable optical cavities [96] and intensity profiles of beams in unstable resonators. The disadvantage of such codes is that they are often not general (e.g. OSCAR can only deal with linear, Fabry-Perot cavities) and the execution times of such simulations tend to be longer than modal simulation types.

A prominent tool, of this modal type, used by the community is FINESSE 2 [73], in combination with the “wrapper” and utilities package Pykat [74]. This is the predecessor software to the new version (FINESSE 3) which is discussed in the context of modal simulations in this chapter and then more extensively in Chapter 6. FINESSE 2 allows for modelling of HG modes up to an “arbitrary” order, taking into account effects such as mode mismatches and misalignments of beams so long as these distortions do not violate the paraxial approximation (see §1.3.2.1). It uses beam tracing and coupling coefficient calculations to perform

modal simulations — the framework of which was used for FINESSE 3 to build upon (as will be outlined later in this chapter). This method of simulating the transverse properties of beams is applicable to many optical interactions and is also fast (as compared to other code types), making it a preferable option for modelling HOMs over, e.g., FFT propagation codes.

One key motivation behind improving the HOM-related capabilities of FINESSE lies in the beam tracing library. There are several reasons for this which will be outlined briefly here. Firstly, it is important that the modal basis used to compute HOM couplings is defined in a clear and customisable way. This is because the value of beam parameters at nodes, and the order in which they get set, can determine the location and magnitude of coupling coefficients at each component in the model. More details, including examples, on this will be given throughout §2.2 (in particular §2.2.3.5), however it is important to note here that FINESSE 3 will, in general, allow for greater customisation and flexibility in how beam traces are performed on a model — both initially and during a modal simulation. Another point of motivation behind an improved beam tracing library is the increasing necessity to be able to easily access model beam parameters, ABCD matrices, cavity stabilities and so on. One such example of the usefulness of these tools can be found in Chapter 3; in which a science task which uses the beam tracing capabilities of FINESSE 3, for the design of arm telescopes in ET, is covered.

Modal interferometer simulations have the potential to be very computationally expensive, especially when the level of mode mismatch and / or misalignment between optical axes is large — requiring high mode orders to converge to the correct solutions for field amplitudes (see §2.3.3.1 for more details). Whilst extensive work has been done [97] to reduce the simulation time of modal simulations in FINESSE 2, more improvements can yet be made to increase performance even further. Several avenues have been explored in this regard and details on the speed-ups obtained with FINESSE 3 thus far can be found in the latter parts of §2.3 and §2.4. In addition, a physically motivated new feature for reducing the number of modes which need to be modelled, for certain types of simulations, is introduced in §2.3.3.

2.2 Beam tracing

Beam tracing, in the context of modal-based interferometer modelling, is the process by which beam parameters (see §1.3.2.4) are propagated through an arbitrary configuration. In order for HOM coupling coefficients to be computed at components (such as mirrors), a modal basis in the form of beam parameters is required at every node of the interferometer network. We will cover coupling coefficient calculations in more detail in §2.3, however the important consideration to note for now is that FINESSE 3 uses these coefficients, in combination with the local coupling matrices of components, to calculate the light fields at each node.

As well as being vital for running any modal simulation, the beam tracing functionality of

FINESSE 3 is also provided as a standalone library — providing tools for executing propagation analyses on any arbitrary interferometer configuration. An exemplary science case of these features is covered by the contents of Chapter 3, the results of which were all obtained via the beam tracing capabilities of FINESSE 3.

In this section we will cover the mathematical framework in which beam tracing is performed (§2.2.1), how optical cavities define their own modal basis (§2.2.2) and the exact mechanism by which FINESSE 3 structures and orders the tracing of beams (§2.2.3). Included in this latter section are the nuances of how overlapping optical cavities are handled and the options available for modifying the priority of specific dependency traces in simulations.

2.2.1 ABCD matrices and beam parameter transformation

A beam parameter is, in general, transformed when a beam passes from one medium into another by interacting with some optical surface or propagating over any distance. We can use the ABCD matrix formalism [56, 58] to treat these transformations — giving us the mechanics to describe how HOMs couple between these interfaces. ABCD matrices represent optical diffraction processes, at optical interfaces or across propagation distances, which have been simplified down to a 2×2 matrix. Mathematically, one can interpret the ABCD matrix as a “ray transfer matrix” [58]:

$$\begin{pmatrix} r_2 \\ \theta_2 \end{pmatrix} = \begin{pmatrix} A & B \\ C & D \end{pmatrix} \begin{pmatrix} r_1 \\ \theta_1 \end{pmatrix} \quad (2.1)$$

acting on a vector $\vec{r} = (r, \theta)$ which describes the position (r) of the wavefront and the angle (θ) relative to the optical axis. The ABCD matrix then defines how a beam is transformed when transitioning to a new frame. This method of representing the transformation of the beam geometry is useful when combined with the Gaussian beam parameter (see §1.3.2.4).




Figure 2.1: Conceptual view of the transformation of a beam defined by the basis q_1 , by a component S , into a new basis q_2 . The refractive indices of the media on either side of the component are shown as n_1 and n_2 , respectively. Note that these indices are often equal, with $n_1 = n_2 = 1$, in interferometer modelling due to an assumed perfect vacuum; except for in special cases such as modelling of substrates.

Consider Figure 2.1, here a beam described by the modal basis q_1 propagates through a medium with index of refraction n_1 . This beam is then incident upon some interaction component S (which can itself be a length of free space), transforming the modal basis into a new beam parameter q_2 in the subsequent medium with a refractive index of n_2 . The equation describing this transformation is known as the *ABCD* law [59] and is defined by

$$q_2 = n_2 \frac{A \frac{q_1}{n_1} + B}{C \frac{q_1}{n_1} + D}, \quad (2.2)$$

where A , B , C and D are the corresponding elements of the *ABCD* matrix associated with the interaction S . This equation forms the fundamental basis of the beam tracing performed by FINESSE 3. It is via Equation (2.2) that beam parameters are transformed when propagating any arbitrary modal basis through a configuration. Each component present in FINESSE 3 has an associated *ABCD* matrix [81] (stored both symbolically and numerically) which is used when computing the beam parameter transformation. For any component which does not transform the modal basis, i.e. thin components with no focusing elements, this *ABCD* matrix is simply an identity matrix. Definitions for the *ABCD* matrices of common components, such as mirrors, beam splitters and thin lenses, can be found in [56, 58, 81].

2.2.2 Optical cavities

Optical cavities, sometimes called resonators or just cavities, store light in a transverse sense as well as longitudinally. In the modal picture, the resonance condition for cavities includes a consideration of the spatial properties of the beam in addition to the wavelength as seen in the plane-wave case. This concept is known as the eigenmode of the cavity and can be defined as the optical field with a spatial profile matching that of the injected field after a round-trip of the cavity [56]. In all cases throughout this thesis we consider only optical cavities with spherical mirrors, such that our eigenmodes can be described by the Gaussian beam parameter (see §1.3.2.4).

2.2.2.1 Eigenmodes

To define the eigenmode mathematically, we can invoke the *ABCD* law given in Equation (2.2). By defining the cavity eigenmode as q_c , and assuming a vacuum inside the cavity such that $n_1 = n_2 = 1$, this becomes

$$\frac{Aq_1 + B}{Cq_1 + D} = q_1 = q_c, \quad (2.3)$$

where A , B , C and D are now the elements of the combined *ABCD* matrices of all the

components in the cavity; known as the round-trip ABCD matrix, M_{RT} . This is computed by simply multiplying each of these matrices together in reverse order (as measured from the end of the chain of interactions), i.e:

$$M_{\text{RT}} = \prod_i^N M_i, \quad (2.4)$$

where N is the number of interactions, across free-spaces or components, in the cavity and M_i are then each of the ABCD matrices associated with these interactions over the cavity round-trip. More specifically, N can be defined in the following way:

$$N = \begin{cases} 2N_c, & \text{if } N_c = N_s \\ 4N_s, & \text{otherwise,} \end{cases} \quad (2.5)$$

where N_c is the number of components inside the cavity and N_s is the number of free-spaces. The former case typically applies to ring [98] (such as the aLIGO input mode-cleaner, IMC) and bow-tie (such as the aLIGO output mode-cleaner, OMC [21]) type cavities, whilst the latter generally applies to linear style cavities.

The equation identifying the eigenmode is then simply, after re-arranging Equation (2.3) for $q_1 = q_c$, equivalent to a form of the quadratic equation:

$$Cq_c^2 + (D - A)q_c - B = 0. \quad (2.6)$$

Due to the nature of this equation, one can note that, mathematically, there are two solutions in general for q_c . In practice there can be only one solution, and this is defined as the q_c which has a positive imaginary part (or Rayleigh range z_R as seen in §1.3.2.4) — i.e. the condition

$$z_R \equiv \pi w_0^2 / \lambda > 0, \quad (2.7)$$

must be satisfied in order for the cavity to be considered “stable”. A more compact description of this condition, in terms of the round-trip ABCD matrix of the cavity, will be given in the next sub-section (§2.2.2.2).

Cavity eigenmodes are calculated in FINESSE 3 via Equation (2.6). These eigenmode values are automatically updated when M_{RT} changes, as discussed briefly in §2.2.2.4 and §2.2.3.1.

2.2.2.2 Stability of optical cavities

As mentioned previously, a cavity is only stable when there exists a complex eigenmode with a real waist-size. This is ensured when the determinant of Equation (2.6) is negative [56]. We can re-cast this requirement into a more useful, and commonly used, form via the g -factor [58], defined as:

$$g = \frac{A + D + 2}{4}, \quad (2.8)$$

where A and D are the usual elements of the round-trip ABCD matrix of the cavity. A cavity is stable when the following condition is satisfied:

$$0 \leq g \leq 1. \quad (2.9)$$

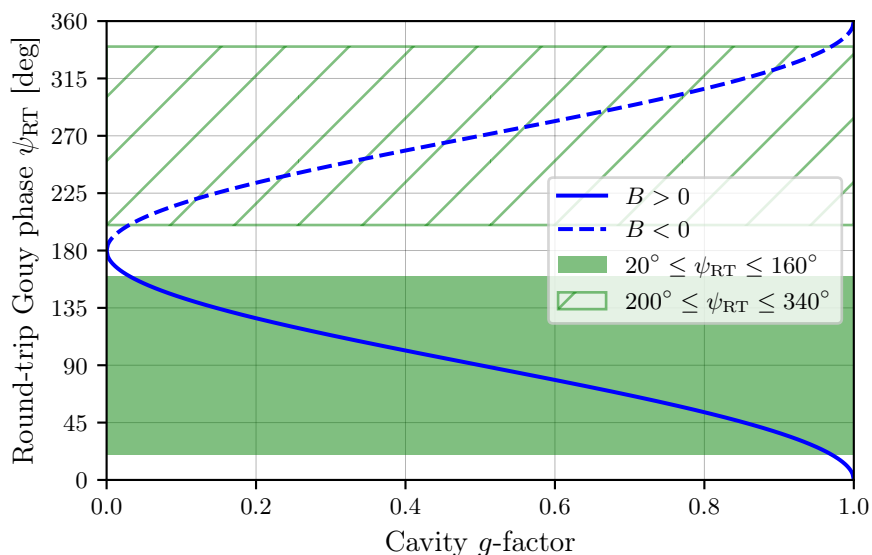


Figure 2.2: The round-trip Gouy phase, ψ_{RT} , plotted against the g -factor of a cavity. Highlighted in the shaded regions are sensible approximations for a stable cavity which is not in a near-unstable regime. The requirement $20^\circ \leq \psi_{\text{RT}} \leq 160^\circ$ will be used later on in Chapter 3 for constraining solution spaces to suitable signal recycling cavity (SRC) designs in ET.

Another useful metric for determining the stability of a cavity is a quantity called the round-trip Gouy phase, ψ_{RT} . This value can be more useful than the g -factor alone as it also provides information on the modal resonance structure of a cavity. The round-trip Gouy phase determines which order of optical modes are resonant within a cavity [56] due to the Gouy phase accumulated by an higher-order mode being proportional to the order, $n + m$, of the mode (see Equation (1.26)). This quantity can be computed using the round-trip ABCD

matrix, and is inexorably linked to the cavity g -factor [99]:

$$\psi_{\text{RT}} = 2 \arccos \left(\text{sign}(B) \sqrt{\frac{A + D + 2}{4}} \right) = 2 \arccos (\text{sign}(B) \sqrt{g}). \quad (2.10)$$

The closer g is to 0 or 1, i.e. $\psi_{\text{RT}} \rightarrow 180^\circ$ or $\psi_{\text{RT}} \rightarrow 0^\circ$, respectively, the smaller the tolerances are on any of the geometric properties of the cavity before it becomes unstable [100, 101]. Due to this, it can be useful to approximate a range of round-trip Gouy phases that give a stable cavity which is not in the near-unstable regime. Such an approximation is given in Figure 2.2. From this figure, one can see that the selected ψ_{RT} regions cover most of the g -space, excluding only the extrema of this range such that near-unstable cavities are avoided.

2.2.2.3 Round-trip Gouy phase and mode spacing

One can equally define the round-trip Gouy phase in terms of the individual Gouy phases measured from the start, position z_1 , of the cavity to the end, position z_2 , (i.e. half a round-trip), via:

$$\psi_{\text{RT}} = 2 (\psi(z_2) - \psi(z_1)) = 2 \left(\arctan \left(\frac{z_2}{z_R} \right) - \arctan \left(\frac{z_1}{z_R} \right) \right), \quad (2.11)$$

where z_R is the Rayleigh range of the corresponding cavity eigenmode q_c . The modal resonance separation, $\delta\phi$, defined as the phase difference in the resonance peaks for each mode order, can be given in terms of ψ_{RT} :

$$\delta\phi = \frac{\psi_{\text{RT}}}{2}, \quad (2.12)$$

This is a useful quantity for verifying the modal resonance structure when performing a cavity scan in FINESSE 3, as shown briefly in §2.2.3.4. Converting this phase shift into a frequency yields a quantity known as the mode-separation frequency, defined as

$$\delta f = \frac{\psi_{\text{RT}}}{2\pi} \Delta\nu, \quad (2.13)$$

where $\Delta\nu$ is the free spectral range (FSR) of the cavity. This value then describes the mode-spacing of contiguous higher-order modes in terms of the frequency separation of their resonance peaks.

2.2.2.4 Example: The Fabry-Perot cavity




Figure 2.3: A representation of the ABCD matrices of a Fabry-Perot cavity. The separate ABCD matrices of the components which make up the cavity are highlighted in different colours; note that the ABCD matrix for the propagation over the length of the cavity is the same in both directions [58]. The round-trip matrix M_{RT} is computed using Equation (2.4) as applied to this configuration; recalling that the order of multiplication is important, such that the product shown here represents the reverse chain of operations from the location of q_1 . Also shown here is the progression of the radius of curvature of the wavefront over the cavity, for illustration. The radius of the curvature of wavefront of the beam defined by the cavity eigenmode must be equal to the radius of curvature of the mirrors of the resonator.

Here we will consider the simple case of a two mirror resonator, i.e. a Fabry-Perot cavity [102], as an example. Figure 2.3 shows a pictorial representation of the ABCD matrices involved in the path of such a cavity, with the corresponding round-trip matrix formula given. Using this equation, and the definitions of the mirror reflection and free-space propagation ABCD matrices also given in the figure, we can derive this round-trip matrix to be

$$M_{\text{RT}} = \left(\begin{array}{cc} 1 - \frac{2L}{R_{C,2}} & 2L - \frac{2L^2}{R_{C,2}} \\ \frac{4L}{R_{C,1}R_{C,2}} - \frac{2}{R_{C,1}} - \frac{2}{R_{C,2}} & 1 + \frac{4L^2}{R_{C,1}R_{C,2}} - \frac{2L}{R_{C,2}} - \frac{4L}{R_{C,1}} \end{array} \right), \quad (2.14)$$

where L is the cavity length and $R_{C,1|2}$ are the RoCs of the cavity mirrors. These round-trip matrices are computed automatically in FINESSE 3 and are updated, both inside the context of a simulation and outside, whenever a dependent parameter (such as a cavity mirror RoC, in this case) changes. The updating of these matrices uses the `TraceTree` data structure of FINESSE 3 for efficiency, more details on this structure will be given in §2.2.3.1.

Applying Equation (2.8) to Equation (2.14), we obtain the following expression for the Fabry-Perot cavity g -factor:

$$g_{\text{FP}} = \left(1 - \frac{L}{R_{C,1}}\right) \left(1 - \frac{L}{R_{C,2}}\right) = g_1 g_2, \quad (2.15)$$

where g_1 and g_2 are then the individual g -factors for the two cavity mirrors. Similarly, we can apply Equation (2.10) to Equation (2.14) to obtain a specialised expression for the round-trip Gouy phase of a Fabry-Perot cavity:

$$\psi_{\text{RT,FP}} = 2 \arccos(\text{sign}(g_2) \sqrt{g_1 g_2}), \quad (2.16)$$

where we have noted that, in this case

$$\text{sign}(B) = \text{sign}\left(2L \left(1 - \frac{L}{R_{C,2}}\right)\right) = \text{sign}(g_2). \quad (2.17)$$

Figure 2.4 shows the different regions in which a Fabry-Perot cavity is stable, as a function of the curvatures of the cavity mirrors. The left plot shows the values of Equation (2.15) whilst the right shows the corresponding round-trip Gouy phases via Equation (2.16). Both cases assume a fixed cavity length. On the left plot we see two special cases of critically stable cavities - confocal, where the RoCs of both cavity mirrors equal the cavity length, and concentric, where both RoCs are equal to *half* the cavity length. The former corresponds to the smallest ratio of the beam size on the mirrors to the waist size whilst the latter represents a cavity with a “point-waist” (i.e. $w_0 = 0$ m). On the right plot the region where the sign of the B element of the round-trip ABCD matrix M_{RT} flips is framed, illustrating where the cavity round-trip Gouy phase ψ_{RT} switches quadrant.

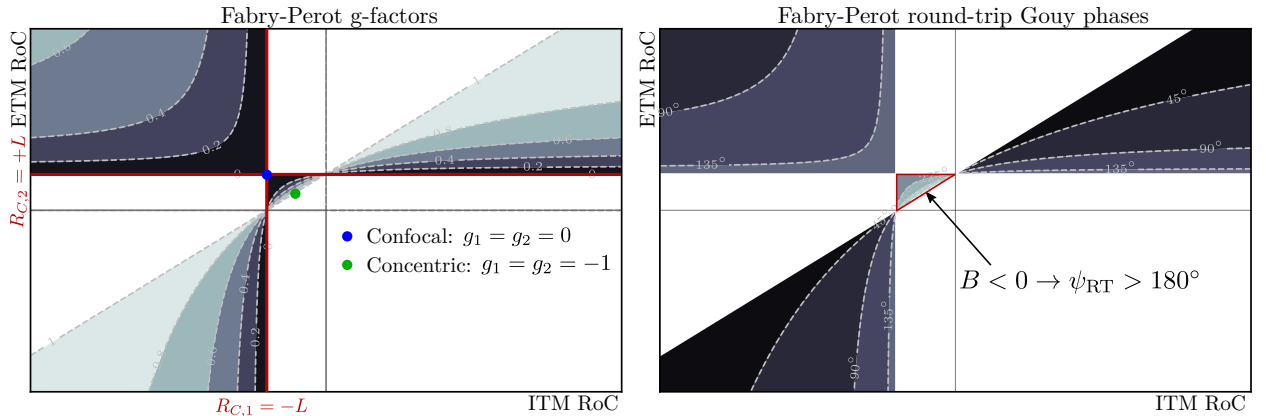


Figure 2.4: Regions of stability for a Fabry-Perot cavity of a fixed length, L , as a function of the RoCs of the mirrors, $R_{C,1}$ and $R_{C,2}$. The left plot gives the g -factors of the cavity and the right gives the values of ψ_{RT} . Note that when either of the mirror RoCs equals the length of the cavity, the corresponding g -factor is zero — leading to a critically stable cavity. Highlighted on the left plot are two special cases of critically stable cavities. All white regions, on both plots, represent unstable cavities with $g > 1$ or $g < 0$.

2.2.3 Propagating beams through an interferometer

Now that we have covered the mathematical framework for transforming beam parameters (§2.2.1) and the basics of optical cavity modes (§2.2.2.1), in this section the mechanics behind the beam tracing algorithms of FINESSE 3 can be explored. An understanding of how FINESSE 3 propagates modal bases throughout a configuration is essential due to the impact this can have on HOM couplings — both in terms of the *magnitude* and *locations* of these couplings. The combination of these affects the mode mismatches (see §2.2.3.2) and hence the quantities of fields in HOMs at various points in the modelled interferometer. In §2.2.3.5 we will discuss the utility that FINESSE 3 provides for defining the locations of such mode mismatches. The contents of this section consist entirely of my work on implementing beam tracing within FINESSE 3. For a more in-depth coverage of the technical implementation of beam tracing in FINESSE 3, refer to Appendix A.

As a practical example, the beam propagation will be discussed in the context of the core aLIGO optics — as shown by Figure 2.5. This figure displays all of the key optical cavities present in the central interferometer. Additional optics, modelling the mirror and beam splitter substrates, as well as the recycling cavity telescope mirrors, will be ignored to reduce the verbosity of these beam propagation examples.

Modelling the modal bases of the recycling cavities is complicated due to the split paths at the central beam splitter. This is an important consideration as it implies overlapping, coupled cavities. The choice of beam parameters in such a configuration is not well-defined and has the potential to lead to unexpected mode mismatches. A brief discussion of how

FINESSE 3 can be used to explicitly declare, and clarify, which cavity to use or is being used is given in §2.2.3.6 on the topic of overlapping cavities.

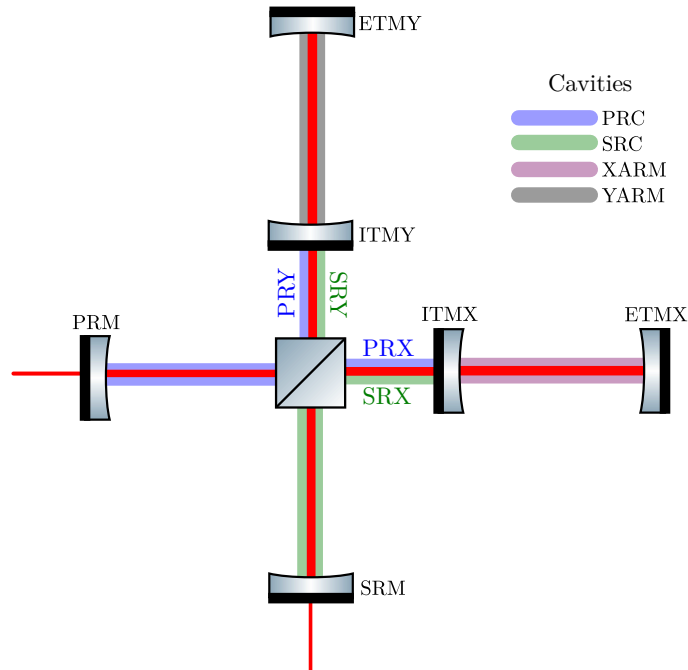


Figure 2.5: A schematic of the core advanced LIGO interferometer optics, with the arm cavities (XARM, YARM) and dual-recycling cavities (PRC, SRC) highlighted. Due to the nature of these recycling cavities, each one must be separated at the beam splitter such that, e.g., the PRC really consists of both the PRX and PRY cavities, as displayed (and similarly for the SRC).

2.2.3.1 Tree structure

FINESSE 3 uses a nodal, directed-graph [103] based system for representing interferometer configurations. In such a system it is assumed that any component can only split the beam across, at most, two paths; e.g. light hitting a beam splitter will be transmitted parallel (and continuously) to the incident beam in one path and reflected orthogonal to the incident beam in the other path. Any geometric offsets in the beams are instead modelled via HOMs using coupling coefficients. Given this, a natural, and convenient, way to represent a propagated beam is via a binary tree [104] type structure — denoted as a `TraceTree` in FINESSE 3 syntax. The `TraceTree` structure stores optical nodes at each tree node and contains references to its “left” and “right” sub-trees (if they exist), with the ABCD matrices for these couplings and the refractive index associated with the optical node also stored for efficiency and convenience.

The `TraceTree` for an optical cavity simply decays to a linked list as there can only ever be one direction in which the beam traverses on a round-trip of any given cavity. This allows for simple and efficient computation of the round-trip ABCD matrix (and, commensurately, the

eigenmode), via Equation (2.4), through recursing the uni-directional “tree”. Performing a beam trace on the internal path of the cavity (defined as an *internal cavity trace*) is equally as simple, where Equation (2.2) is applied at each node of the cavity tree. This is demonstrated pictorially in Figure 2.6 for the simple case of a Fabry-Perot cavity; equivalent to the picture already seen in Figure 2.3. An important aspect of the FINESSE 3 beam tracing is introduced in Figure 2.6 — namely, the storage of beam parameters at the *node* rather than *port* level; a useful distinction which improves on the initial framework which FINESSE 2 provided [76]. This separation of beam parameter directions simplifies the mode mismatch (§2.2.3.2) and coupling coefficient (§2.3) calculations, as well as allowing a new type of “asymmetric” beam tracing (see §2.2.3.3).




Figure 2.6: A representation of the **TraceTree** for a Fabry-Perot cavity. The beam parameters computed and stored for each node are shown on the right, using (2.2) and the ABCD matrices given in Figure 2.3. Note that the default behaviour of beam tracing in FINESSE 3 is so-called “symmetric” tracing — where the beam parameter for nodes opposite to those that have been traversed are reversed, flipping the sign of the distance to the waist. More details on this, and the alternative “asymmetric” tracing, are given in §2.2.3.3.

As indicated via §2.2.2.1, the light fields in most optical cavities can be safely described in terms of their eigenmodes [76]. Furthermore, in GW detectors, the fields are typically dominated by the behaviour of the key cavities (see §1.1). For these reasons, cavities are a sensible choice for using as beam tracing dependencies; i.e. points from which the other beam parameters in a configuration can be obtained. An example highlighting the importance of using cavity modes for modal simulations is given in §2.2.3.4.

Most interferometer configurations will contain several trace trees due to the existence of multiple dependencies (e.g. cavities) — see Figure 2.5 for example. The intersection of trees from different trace dependencies indicate the locations (in terms of the nodal graph structure) at which possible mode mismatches can occur if the dependencies are not mode matched to each other; more details of this will follow in §2.2.3.2. The positions of these mode mismatches determine the point(s) at which light is scattered into HOMs, thereby affecting the results of modal simulations of this type. It is important, consequently, that the structure of these multiple trees is well-defined. In FINESSE 3 these trace trees are stored in a **TraceForest**, where the order in which each tree gets traced is given implicitly by the index of the tree in the forest.




Figure 2.7: A conceptual overview of the “trace forest” structure for a dual-recycled Michelson interferometer. In this schematic, only the arm cavities (XARM and YARM) are enabled as tracing dependencies — this is to illustrate the branching taking place and also to highlight how the modal bases are set at each node when a given cavity, in this case XARM, takes priority. The numbers next to each tree define the beam tracing order.

An example of a `TraceForest` for the configuration shown previously (in Figure 2.5) is given in Figure 2.7; where only the arm cavities are enabled now as documented. In this figure the XARM cavity is given the highest beam tracing priority (see §2.2.3.5 for a brief discussion on the ordering definition), resulting in the beam parameters at the input and output of the Michelson interferometer being set based on the propagation of the eigenmode of XARM. If this tracing priority was switched around, then the beam parameters at these nodes would, consequently, be set via the propagation of the YARM eigenmode. This structuring of beam tracing priorities is highly flexible in FINESSE 3; there will be a brief discussion on this in §2.2.3.5. From the trace trees shown on the figure, one can see that the traces from the two arm cavity trees intersect at the couplings: $\text{BS.p2.i} \rightarrow \text{BS.p4.o}$ and $\text{BS.p2.i} \rightarrow$

BS.p1.o; i.e. the output and input of the Michelson interferometer, respectively. There is then a potential for mismatches, and hence scattering into HOMs, at these locations. Further exploration of this continues in the next section.

For more technical details on the `TraceTree` and `TraceForest` structures, as well as discussions of tree intersections and the implications of these for modal simulations, see Appendix A.

2.2.3.2 Mode mismatches

Consider Figure 2.8 which shows a zoomed-in version of Figure 2.7 at the beam splitter. As mentioned previously, the two trace trees dependent on each of the arm cavities intersect at this component. This figure illustrates the concept of mode mismatches in the context of this example, where the beam parameters at the input ports (two and three) are propagated to port four via Equation (2.2) as usual.

A useful figure-of-merit for the overlap, in terms of the field amplitudes, between two beam parameters, q_1 and q_2 , is (Page 143 of [105])

$$\mathcal{O}(q_1, q_2) = \frac{4|\Im\{q_1\} \Im\{q_2\}|}{|q_1^* - q_2|^2}. \quad (2.18)$$

This quantity returns values $\mathcal{O} \in [0, 1]$, where unity indicates a full mode match between the two beam parameters and zero gives complete mode mismatch. The mode mismatch is then simply $\mathcal{M}(q_1, q_2) = 1 - \mathcal{O}(q_1, q_2)$. Both of these quantities will be used frequently throughout this thesis to quantify the amount of mode mismatch present in various systems.

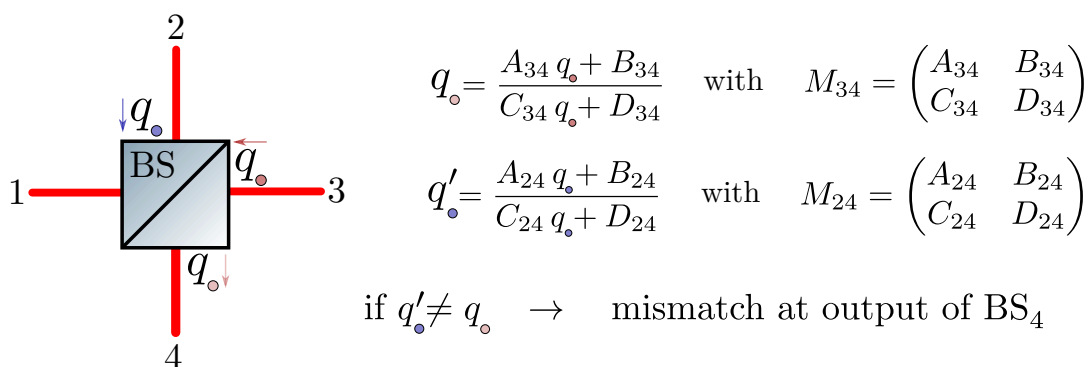


Figure 2.8: Beam tracing of XARM (red) and YARM (blue) dependent trace trees to the beam splitter corresponding to Figure 2.7. As the XARM trace has the highest priority, the beam parameter at BS.p4.o is set from the beam parameter at BS.p3.i as that figure shows. If the beam parameter propagated from the YARM dependent tree to BS.p4.o is not equal to this former beam parameter, then there is a mode mismatch present at this node.

When mode mismatches are present, light will be scattered into HOMs — specifically, into modes spaced by even orders (see §2.3.3.1). Assuming that the carrier light consists of a pure HG_{00} mode, this means that the modes $\text{HG}_{02/20}$, $\text{HG}_{04/40}$ and so on will be created. Smaller values of \mathcal{O} will result in larger magnitudes for scattering into higher-order fields. A simple example of this is shown in Figure 2.9 which is based on the mismatch scenario given in Figure 2.8. In this figure the XARM and YARM cavities are mismatched, thus leading to scattering into HOMs due to the resultant mismatch at the beam splitter. The mismatch values on the x-axis are computed via $\mathcal{M}(q_1, q_2)$, scaled to a percentage. Note that we define the “fractional HOM power” (left-hand y-axis) as the ratio of the power in all HOMs to the total detected power. Expressed mathematically this is

$$F_{\text{HOM}} = \frac{\sum_{n,m} a_{nm} a_{nm}^* [n \neq 0, m \neq 0]}{\sum_{n,m} a_{nm} a_{nm}^*}, \quad (2.19)$$

where a_{nm} are the complex amplitudes for the fields with modes HG_{nm} . As a brief aside, this quantity is simple to compute in FINESSE 3 via the use of mode index masks on photodiodes; whereby only the field amplitudes of non-masked modes contribute to such detector outputs.

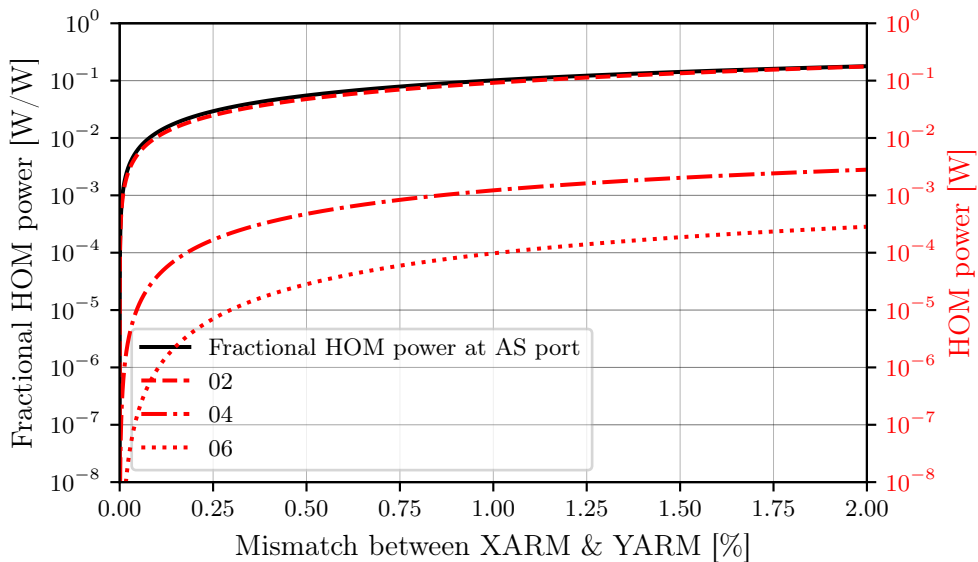


Figure 2.9: Scattering of carrier light into HOMs for a given percentage mismatch between the XARM and YARM cavity modes (as propagated to the beam splitter). This is based on the configurations shown in Figure 2.7 and Figure 2.8. As expected, larger mismatches between these arm modes result in a higher fractional power of HOMs; as detected at the anti-symmetric (AS) port of the beam splitter (i.e. the output of the Michelson interferometer).

2.2.3.3 Symmetric vs asymmetric tracing

As noted previously in Figure 2.6, the default behaviour of the FINESSE 3 beam tracing algorithm is to set the beam parameters at nodes opposite to those traversed by reversing the sign of the distance to waist. In mathematical terms, this operation can simply be expressed as

$$q_{n'} = -q_n^*, \quad (2.20)$$

where n is a node that has been traversed and n' is the corresponding opposite optical node at that traversed port. In this section we will rationalise the logic behind this decision and briefly discuss the limitations; whilst offering an alternative option provided by the FINESSE 3 beam tracing algorithm.

Consider Figure 2.10, which shows the wavefronts of a beam incident on and reflected from a surface orthogonal to the beam axis. The incident beam is described by the beam parameter $q_1 = z' + iz_R$, where z' is the distance from the beam waist to the mirror surface and z_R is the Rayleigh range of this beam. Upon reflection from the surface, the wavefront of the beam is necessarily reversed (as the beam changes direction of travel). This can be treated as the case where the reflected beam has a virtual waist positioned behind the surface (relative to the direction of propagation), at a distance of z' from the surface (i.e. $2z'$ from the real waist position of the incident beam). If we project this coordinate system transformation into that of the incident beams', then we obtain a beam parameter q_2 of the reflected beam:

$$\begin{aligned} q_2 &= z' - z_0 - 2z' + iz_R \\ &= -z' + iz_R \\ &= -q_1^*. \end{aligned} \quad (2.21)$$

It is important to note that this method of setting the beam parameters will only result in complete mode matching, on reflection from a surface, when the beams on both sides of the surface are fully matched to each other. Referring again to Figure 2.10, this means that there will be no mode mismatches in this system only when $q_3 = q_1$; assuming that the refractive indices of the media on both sides of the surface are equal, such that the surface transmission ABCD matrix is the identity matrix. If $q_3 \neq q_1$ (i.e. mismatched beams across the surface) then, along with the expected mode mismatch on transmission, one would also see a mismatch on reflection at the front surface when setting beam parameters in this way, as in general (using Equation (2.2) and the reflection ABCD matrix for mirrors [56])

$$q_2 = \frac{q_1}{1 - \frac{2}{R_C}q_1} \neq -q_1^*. \quad (2.22)$$




Figure 2.10: A schematic showing the wavefront of a beam incident (red) upon a mirror described by beam parameter q_1 . The reflected beam (blue) has a reversed wavefront, with a beam parameter q_2 , which can be projected behind the mirror to a virtual waist as displayed. The beam parameters q_3 and q_4 represent the beam on the other side of the surface, if it exists — see text for details on why these are important when performing symmetric traces.


There is an additional subtlety here in that, as implied earlier, this symmetric tracing will only result in no mode mismatch (referring again to Figure 2.10) if there is indeed a beam parameter propagated from the right hand side too; e.g. if this mirror represents the input mirror of a cavity. Considering a case where we have a laser shining onto this mirror, this provides us with a simple example of how the option of “asymmetric” tracing, in FINESSE 3, can be used to remove spurious mismatches in some configurations. Asymmetric tracing simply means that opposite direction nodes are not given the reverse beam parameter (unless there is no other option), instead the beam parameters at these opposite nodes are computed using the ABCD matrices of their interacting components. In the context of Figure 2.10 again, this means that q_2 will be set using the left-hand side of Equation (2.22) rather than the right. Table 2.1 provides an example for the effect on the HOM couplings for such a configuration when performing the beam tracing with these different algorithms.

Another, more practical, example where asymmetric tracing can help to remove spurious mode mismatches is a simple Michelson interferometer with flat mirrors. Figure 2.11a shows the specific configuration considered here, where the input beam has a waist before the interferometer as indicated. This location of the beam waist leads to a higher degree of mismatch at the mirrors in the arms when performing symmetric beam tracing. In fact, for the specific configuration chosen, this mode mismatch on reflection from the arm mirrors is $\mathcal{M} \sim 95\%$. This is an extremely large mismatch value and, as such, leads to significant coupling into very high order HOMs in the simulation. If, however, we choose to perform the beam tracing asymmetrically instead, then our mode mismatch at these mirrors is $\mathcal{M} = 0\%$. This is due to the reasons discussed above — namely, that the beam parameter on reflection from a flat mirror should be identical to the incident beam parameter according

Table 2.1: Power in even-order (see §2.3.3.1) HG modes, on reflection from the mirror (with $R = 1$) in Figure 2.10, using a 1 mm waist size position 0.5 metres from the mirror surface. The symmetric tracing simulation gives non-zero HOM couplings due to Equation (2.22) whilst the asymmetric tracing version yields no scattering into these HG modes as one would expect. Using the beam parameter stated above, the symmetric tracing gives an overlap between q_1 and q_2 of $\mathcal{O} \approx 0.9$ using Equation (2.18) — i.e. a mismatch of approximately 10%.

Modal power on reflection	Symmetric Tracing	Asymmetric Tracing
HG ₀₀	0.8971 W	1 W
HG ₀₂	0.0462 W	0 W
HG ₀₄	0.0036 W	0 W
HG ₀₆	0.0003 W	0 W

to Equation (2.2). These ideas, and their resulting effect on the power recorded at the Michelson interferometer output, are captured in Figure 2.11b. It should be noted that the maximum TEM orders (`maxtem`) shown on this figure only include modes of even-order due to only mode mismatches being present here; more details on this mode selection feature can be found in §2.3.3.



(a) A “flat Michelson” interferometer where the waist of the input beam is positioned before the beam splitter. (b) Power detected at the Michelson interferometer output photodiode as a function of the tuning of the mirror in the y-arm, mY .

Figure 2.11: A configuration consisting of a Michelson interferometer with flat mirrors is shown in (a). This setup exemplifies the behaviour of symmetric vs asymmetric beam tracing. In (b) we perform simulations in which the tuning of the mirror mY is scanned. The blue traces give the *total* power, as a ratio of the input power, at the output whilst the red traces shows the *modal* power — i.e. excluding the power in the HG₀₀ mode. One can see that for the asymmetric tracing curves we get the expected “light” and “dark” fringe behaviour, where all of the power is always within the fundamental mode. For the symmetric tracing simulations, however, we see that the modal power dominates and a very large maximum mode order would be required to reach the expected power (due to significant power in very high order HG modes).

2.2.3.4 Importance of defining cavity modes

It is important in any analysis involving optical cavities that HOM couplings, for cavity nodes, are computed using their respective cavity eigenmode as the bases. If the beam parameters at nodes inside a cavity get set to values which do not correspond to the associated cavity eigenmode then this can result in mismatches being present where there shouldn't be any and / or mode matching where there should be mismatches in practice. Such a scenario will often lead to unexpected HOM couplings being present when computing power build-ups or transfer functions.




Figure 2.12: This plot highlights the importance of using the cavity eigenmode to define its own basis (blue traces) rather than using a different manual basis defined elsewhere in the system (red traces). The cavity power is plotted against the mirror radii of curvature, showing significant drops for the non-cavity mode simulations due to large mode mismatches inside the cavity. Meanwhile the simulations where a cavity basis is used result in a small mode mismatch which comes from the non-unity overlap of the updated cavity eigenmode with the input beam parameter.

An example of this is shown in Figure 2.12. This plot uses an ET-LF style arm cavity (see Chapter 3, Table 3.1 for details) as the base configuration. The input beam is mode matched to the nominal cavity — i.e. matched when the radii of curvature of both cavity mirrors are $R_c = 5580$ m. In the following simulations we scan around this matched RoC value, setting both cavity mirrors simultaneously such that the cavity remains symmetric. The blue traces show the circulating power, for different maximum mode orders M_T , when tracing via a

cavity object. The red traces give the cavity power when using the input beam parameter as the only beam tracing dependency. In the former case the cavity eigenmode is re-computed at each data point, giving a consistent and optimal modal basis when calculating coupling coefficients at the mirrors. The small dip in the cavity power when shifting away from $R_c = 5580$ m comes from the mode mismatch that results between the cavity eigenmode and the input beam parameter. This mismatch is encapsulated in the corresponding blue trace on the lower sub-plot of Figure 2.12, the values of which were computed using $\mathcal{M}(q_1, q_2)$ (see Equation (2.18)).

When tracing with only the input beam parameter, we get a sharp drop in the cavity power after shifting away from the matched point. This is a result of using a non-optimal mode base such that any geometric changes to the cavity yield mode mismatches on reflection from the cavity mirrors — due to the symmetric beam tracing procedure outlined in §2.2.3.3. The mode mismatch values in this case, shown as the red curve on the lower sub-plot of Figure 2.12, have been computed with q_1 as the input beam parameter, to a cavity mirror, propagated via reflection and q_2 as the output beam parameter stored on reflection; this value was then doubled due to this mismatch taking place on reflection from *both* mirrors.

2.2.3.5 Order of beam tracing

The order in which trace dependencies (i.e. cavities and manually set beam parameters) in a model are propagated can affect the results of modal simulations. We can see immediately from Figure 2.7 that if the priorities of the two cavities in this interferometer were switched, then the beam parameters in the input and output paths would be set based on propagation from the YARM cavity rather than the XARM. Another typical case where this could apply is an identical configuration to this figure, where the only difference now is that a beam parameter is set at the input laser (a common case, as this is often known and well-defined). It is this scenario which we consider and model in Figure 2.13. In this simulation we sweep over the waist size of the laser beam as a deviation to the waist size which mode matches the laser to both the XARM and YARM cavities (i.e. $\Delta w_0 = 0$ is the fully mode matched case). The output *modal* power, at the anti-symmetric port of the beam splitter, is plotted as a fraction of the total power at this point — see Equation (2.19) for how this quantity is calculated. The central beam splitter is detuned slightly such that the interferometer is not operating exactly on the dark fringe.

The blue curve on Figure 2.13 shows the output power when prioritising the input laser beam during the beam tracing — resulting in the beam parameters at all the nodes in the central interferometer being set using the propagation of this beam. Conversely, the red curve uses the XARM cavity as the first tracing dependency whilst the input laser is traced last — resulting in beam parameters being set according to the trace forest structure given in Figure 2.7. Increased power is detected, in both cases, at the interferometer output when the input laser is not mode matched to the arm cavities. This is simply due to scattering into HOMs which accumulate extra phases dependent upon their mode order (i.e. Gouy phase).

Larger mismatches, and thus a greater degree of scattering into HOMs, occurs for the blue curve due to the extra mode mismatches upon reflection from the arm cavity input mirrors (see §2.2.3.3 for details on this). One can correct for mismatches of this type by instead tracing with the cavity modes first. This essentially shifts the mode mismatch locations from the ITMs to the central beam splitter, the latter case being covered by the view shown in Figure 2.8.

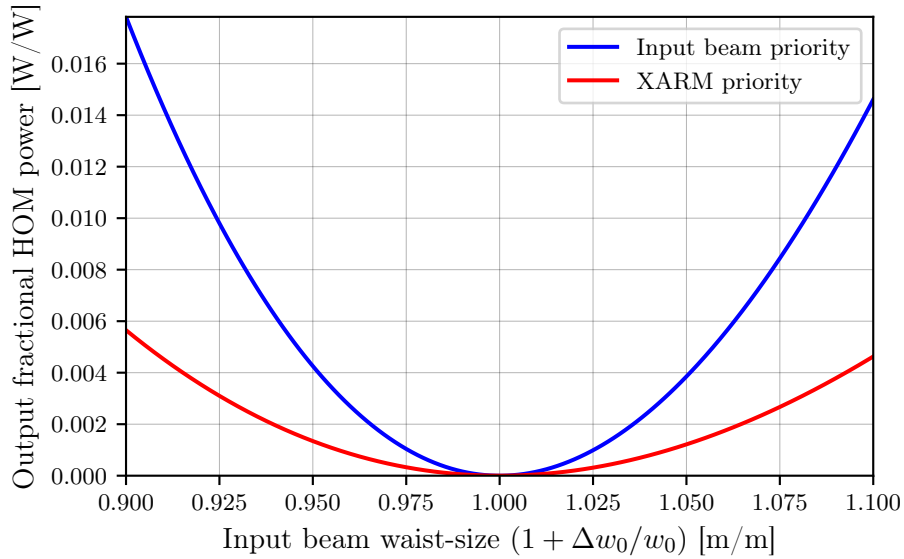


Figure 2.13: This figure shows the difference in the output power of a detuned Michelson interferometer where the beam tracing priorities of the input laser mode and cavity modes are swapped between simulations. These power values are plotted as a function of the input laser waist size such that mismatches will occur between this laser and the cavity eigenmodes when $\Delta w_0 \neq 0$. Prioritising the input laser mode in the beam tracing results in larger mismatches, and hence greater modal scattering, than when tracing one of the cavity modes first. This is due to the symmetric tracing behaviour, which causes mode mismatches on reflections from the input cavity mirrors, described previously in §2.2.3.3.

FINESSE 3 provides a great deal of flexibility in ordering beam traces, where each tracing dependency object has a priority value associated with it. The underlying `TraceForest` built before beam tracing is then constructed via the descending order of dependency priorities. Each priority value can be changed between simulations, allowing for different beam tracing orders to be executed depending upon the simulation being executed. This flexibility provides the user with ways to limit spurious mode mismatches, as we have seen in this section. It also gives fine-tunable control over, and easy access to, the modal bases of each node in a configuration. More details on some further implications of trace dependency ordering can be found in Appendix A.2.

2.2.3.6 Overlapping cavities

In more complicated optical systems cavities can overlap with each other, causing the eigenmode of such a system to not be well-defined. A pertinent example of this is the recycling cavities in an aLIGO-style GW detector. This is something implied previously in Figure 2.5. To resolve the ambiguity as to which modal basis is used when overlapping cavities are present, FINESSE 3 simply uses the aforementioned beam trace ordering logic outlined in §2.2.3.5.

An example, based on the central interferometer of aLIGO from Figure 2.5, is shown in Figure 2.14. In this figure the trace order of the cavities is listed from first to last and the corresponding trace-trees for these cavities are given, colour coded by the specific cavity they depend upon. With this specific chosen trace order, the cavity PRY has no effect on the modal basis used by the nodes inside the recycling cavities. This is because the PRX cavity sets the beam parameters from the PRM to BS whilst, beforehand as it is a higher priority dependency, the beam parameters from BS to ITMY are set from the SRY trace.

One can easily change the order in which these types of overlapping cavities are traced by simply changing the priority values of each cavity in the model. These priority values apply to both the internal cavity traces, as seen in this section, and traces from cavities (external traces) as noted in §2.2.3.5.




Figure 2.14: This figure focuses on the overlapping recycling cavities in the central interferometer of an aLIGO-style detector. Each cavity is listed in the top right by the order in which they will be traced in this particular example. Using this cavity ordering, the trace trees which are constructed for the beam tracing are then given in the lower right. One may note that the PRY cavity has no effect on the beam parameters set in the recycling cavities here, as the other cavity instances cover all the nodes in these cavities already.

2.2.4 Beam tracing as a library

In addition to being the foundation of modal simulations, beam tracing in FINESSE 3 is also provided as a standalone library. The core algorithm (described throughout this section) used for setting beam parameters at each node of a configuration is exposed to the user, allowing convenient access to the modal bases across a model — thereby making it easier to check for mode mismatches, and perform “chained” simulations.

One common problem encountered during the design of GW detectors is the task of building beam expansion telescopes. This was a task previously required for aLIGO in the context of the recycling cavity design [20] and it is also an ongoing area of research for the Einstein Telescope; forming the basis of Chapters 3 and 4. Given the planned larger arm cavity beam sizes for future GW detectors, it is a task which will be just as important in detector design in the coming decades. The design of such telescopes requires knowledge of the beam parameters and accumulated Gouy phase over the full path of the system when testing different configurations.

As a solution to this, and indeed other related tasks, a generic beam propagation function was implemented in FINESSE 3. This tool allows one to propagate an arbitrary beam (i.e. with any initial beam parameter) through any optical path of a configuration — thus yielding a set of the beam parameters, accumulated Gouy phases and ABCD matrices corresponding to this propagation. The optical path is bi-directional, allowing for back-propagations of beams in addition to forward propagation. This function was used extensively throughout Chapter 3 to obtain the results noted there; note in particular Figures 3.9 and 3.10 in that chapter, which correspond to direct plots of the return of this function for the given telescope designs.

Furthermore, the implementation of this tracing functionality is also motivated by the increasing need for such automated tools in many aspects of detector design. For example, the layout of optics for frequency-dependent squeezing relies heavily on mode matching telescopes; see, e.g., [106]. Providing the ability to design such systems with a fast and robust algorithm, alongside a useful interface for observing beam changes over optical paths, is something that the FINESSE 3 beam tracing library offers.

2.2.4.1 Symbolic beam propagation

The beam propagation tool is also capable of storing the symbolic form of beam parameters, ABCD matrices and properties dependent on these at each node of an interferometer model. This is a powerful feature, which allows leveraging of the speed of NumPy [107] to compute properties of the beam as a function of any of its dependent parameters. More details on this can be found in [108].

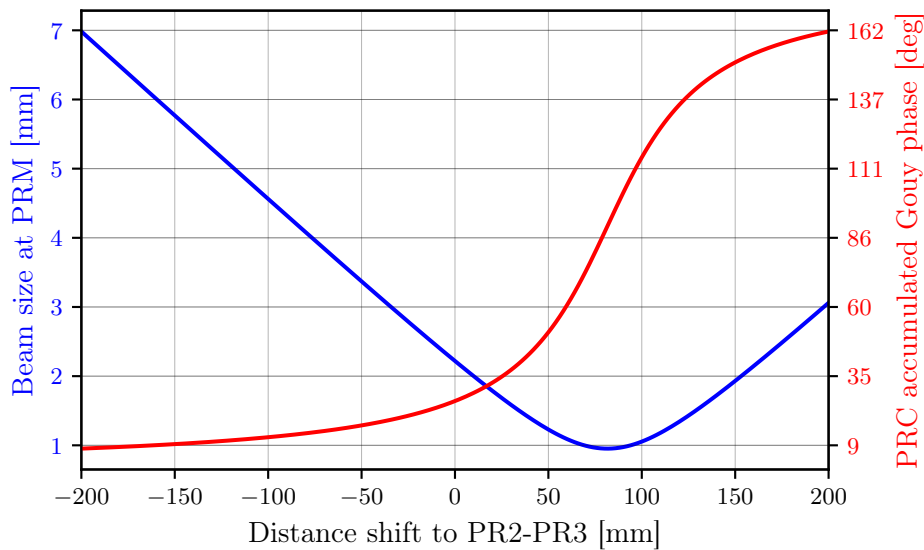


Figure 2.15: This shows an example where symbolic beam propagation has been utilised to demonstrate tuning of the aLIGO PRC telescope, giving both the beam size on the PRM and the total accumulated Gouy phase in the PRC.

```

1  import finesse
2  import numpy as np
3
4  IFO = finesse.Model()
5  # Parse an aLIGO file
6  IFO.parse_file("design.kat")
7
8  # Propagate a beam from ITMY to PRM with symbolic expressions
9  beam = IFO.propagate_beam(IFO.ITMY.p2.i, IFO.PRM.p2.i, symbolic=True)
10 # Symbolic expressions for accumulated Gouy phase and spot-size at PRM
11 ACC_GOUY_SYM = beam.total_acc_gouy
12 W_PRM_SYM = beam.w(beam.end_node)
13
14 # Define a range of distance shifts in mm
15 dchanges = np.linspace(-200, 200, 201)
16 # Make a corresponding array of distance changes
17 # to length between PR2 and PR3 mirrors
18 distances = IFO.spaces.lp2.L.value + dchanges * 1e-3
19
20 # Evaluate the symbolic expressions from above
21 subs = {IFO.spaces.lp2.L: distances}
22 acc_gouys = ACC_GOUY_SYM.eval(subs=subs)
23 ws = W_PRM_SYM.eval(subs=subs)

```

Listing 1: Python source code for calculations which produce the traces shown in Figure 2.15. Note the concise nature of this script and the interoperability with NumPy.

A typical use case of this symbolic beam propagation, in the context of the more generic aforementioned telescope design problem, is given in Figure 2.15. Here we perform a common task of sweeping over the distance between telescope mirrors, in this case the PR2 and PR3 optics in the aLIGO PRC, and calculate the beam size at the PRM along with the Gouy phase accumulated in the cavity. These are important values as the former is limited in the design by the mirror size whilst the latter gives us a metric for the PRC stability (as we saw in §2.2.2.2) and is also constrained to certain values to allow the alignment sensing system to operate. With the FINESSE 3 symbolic beam propagation capabilities, computations such as this can be performed with only a few lines of code — see Listing 1.

There is scope for more work on this feature of FINESSE 3 in the future — including simplification of these symbolic expressions, providing the opportunity to obtain concise, analytic representations of beam parameters, and associated properties, in an optical path.

2.2.4.2 Matrix-less simulations

It is worth noting too that the beam tracing also forms the core part of another type of simulation which can be performed in FINESSE 3 — namely “*matrix-less*” simulations. These are simulations in which the interferometer matrix itself doesn’t need to be solved, simulations of this type are created automatically when the only detectors present in a model are those which do not need the field amplitudes to compute their outputs. Such detectors include, but are not limited to, beam property detectors (which probe a specific property of a beam parameter at a node), accumulated Gouy phase detectors (for summing the Gouy phases accumulated over a path) and mode mismatch detectors (which calculate one minus the overlap, see Equation (2.18), for a coupling at a component). Given that no scattering matrix calculations, component matrix filling nor solving of the full interferometer matrix occurs in such simulations, these are typically very fast to execute.

Matrix-less simulations are covered in terms of the “flow” of the program for such model types in §6.4. That section will also compare the logic of how these simulations are executed with more typical matrix simulations.

2.3 Mode scattering matrices

As mentioned previously in §2.2, FINESSE 3 uses modal coupling coefficients, computed in a Hermite-Gauss basis (see §1.3.2.3), to calculate field amplitudes for higher-order modes. These coupling coefficients describe how the modal components of an electric field defined by a modal basis q_1 couple into the modal components of a field defined by another basis q_2 . This, therefore, has mathematical consequences for when mode mismatches (see §2.2.3.2) are present in a system due to the non-zero coupling to HOMs that then occurs between these differing bases. In addition to providing a framework for computing couplings due

to mismatches, coupling coefficient calculations can also be used for describing *misaligned* optical systems; e.g. where the optical axes of two cavities do not overlap.

In this section we will cover the necessary mathematical details for coupling coefficient calculations (§2.3.1) using the Bayer-Helms [109] formalism, explore how matrices of coupling coefficients (termed “scattering matrices”) are computed in FINESSE 3 (§2.3.2) and briefly discuss how HOM simulations may be optimised by avoiding the computation of unnecessary coupling coefficients (§2.3.3).

2.3.1 Coupling coefficients and the overlap integral




Figure 2.16: The transformation of coordinate systems across two different beam segments. The input beam is described by beam parameter q_1 and is in the coordinate system (x, y, z) whilst the output beam can be described with q_2 and is in a new coordinate system (x', y', z') . This latter coordinate system is obtained by rotating the former system around the y -axis by a misalignment angle γ . Note that if $q_1 = q_2$ then we have a case of pure misalignment with no mismatch. Simplified coupling coefficients for such a scenario can be found in [109].

Consider Figure 2.16, here a beam described by the basis q_1 in a coordinate system (x, y, z) encounters a “boundary” (i.e. typically an optical surface) where the subsequent beam is given by the basis q_2 in a new coordinate system (x', y', z') . This new coordinate system is found by rotating the x - z plane by an angle γ . We can refer to this as optical axis *misalignment*. If the spatial distribution of the input beam is $u_{n'm'}(x, y; q_1)$ and of the output beam is $u_{nm}(x, y; q_2)$, then we can define the amplitude of a given mode, $\text{HG}_{n'm'}$, for

the output beam as [56]

$$u_{nm}(x, y; q_2) \exp(i(\omega t - kz)) = \sum_{nm} k_{nmn'm'} u_{n'm'}(x, y; q_1) \exp(i(\omega t - kz')), \quad (2.23)$$

where the complex values $k_{nmn'm'}$ are then the coupling coefficients from the modes $\text{HG}_{n'm'}$ into HG_{nm} . Here we have assumed that the spatial distributions are described in the Hermite-Gauss basis. Given that HG modes are orthonormal, as described in §1.3.2.3, the computation of these coupling coefficients can be performed via the overlap integral [109]:

$$k_{nmn'm'} = \exp\left(i2kz' \sin^2\left(\frac{\gamma}{2}\right)\right) \iint dx' dy' u_{n'm'} \exp(ikx' \sin \gamma) u_{nm}^*. \quad (2.24)$$

Whilst Equation (2.24) can be calculated via numerical integration methods, this operation can be computationally expensive — especially for cases with both mode mismatches and misalignments. Instead, FINESSE 3 uses an analytic solution to Equation (2.24), provided by the framework from Bayer-Helms [109], to compute coupling coefficients efficiently. An excellent summary of this analytic solution can be found in Section 9.16 of [56].

2.3.2 Computing coupling coefficients in FINESSE 3

Coupling coefficients are calculated in modal FINESSE 3 simulations both initially (based on the starting state of the model) and whenever a dependent parameter, such as a beam parameter or tilt angle, changes during an axis scan. The exact form that these calculations take, in the context of applying them to component coupling matrices, is covered by the example of applying such computations to a mirror component in §2.3.2.2. First, however, we will clarify a subtlety between the Bayer-Helms solution [109] and the FINESSE 3 implementation.

2.3.2.1 Gouy phase shifts and coupling coefficients

A subtle difference between the conventions of FINESSE 3 and the mathematical derivation of the Bayer-Helms coefficients lies in how Gouy phase shifts are applied. In FINESSE 3, this phase shift is included as part of the propagation of a field through a space component. Consequently, the equation describing the evolution of the amplitude of a field through a medium with refractive index n_r and total distance L is [76]

$$b_{\text{out}} = b_{\text{in}} \exp\left(-i\left(\frac{n_r L \Delta\omega}{c} + \left(\frac{1}{2} + n\right) \Psi_x + \left(\frac{1}{2} + m\right) \Psi_y\right)\right), \quad (2.25)$$

where $\Delta\omega$ is the offset to the default frequency (i.e. $\omega - \omega_0$) and Ψ_x and Ψ_y are the Gouy phases accumulated over the space in the tangential and sagittal planes, respectively. These Gouy phase accumulations are defined as

$$\Psi_{[x|y]} = \left| \arctan \left(\frac{z_{\text{out},[x|y]}}{z_{R,\text{out},[x|y]}} \right) - \arctan \left(\frac{z_{\text{in},[x|y]}}{z_{R,\text{in},[x|y]}} \right) \right|. \quad (2.26)$$

The effect of this convention is such that Gouy phase shifts are explicitly included in the amplitude coefficients, a_{nm} , of the field (single frequency, ω):

$$E(t, x, y, z) = \sum_{n,m} a_{nm} u_{nm}(x, y; \vec{q}) \exp(i(\omega t - kz)). \quad (2.27)$$

This is in contrast to the Bayer-Helms coefficients, which have been derived assuming that the Gouy phase is included in the spatial profile, $u_{nm}(x, y; \vec{q})$, of this formula. To rectify this difference in Gouy phase application, each coupling coefficient, $k_{nmn'm'}$, in FINESSE 3 is rotated (in complex space) by an angle

$$\phi_r = \left[\left(\frac{1}{2} + n' \right) \psi_{x,2} + \left(\frac{1}{2} + m' \right) \psi_{y,2} \right] - \left[\left(\frac{1}{2} + n \right) \psi_{x,1} + \left(\frac{1}{2} + m \right) \psi_{y,1} \right], \quad (2.28)$$

where $\psi_{[x|y],1}$ and $\psi_{[x|y],2}$ are the Gouy phases of the input and output beams, respectively. Qualitatively, this is equivalent to removing the Gouy phase of the input beam parameter and adding the Gouy phase of the output beam parameter; an operation referred to as “reversing” the Gouy phase in FINESSE 3 parlance.

A result, and indeed an intended goal, of this convention is that mode mismatches never occur across spaces in FINESSE 3. Instead, they will always occur at component boundaries (as we will see by the example of a mirror in the next section). An advantage to this representation is that computationally expensive scattering matrix calculations can be avoided entirely for space components (which are inevitably numerous, especially in large optical layouts); without any loss of generality, as scattering into HOMs will still occur as necessary.

2.3.2.2 Example: Scattering matrices at a mirror

Here we will explore how coupling coefficients can be used alongside the coupling matrix framework already covered in §1.4.2. To do this, we will use the example of a mirror — observing how the overall coupling matrix for this type of component is modified for modal simulations, and why matrices of coupling coefficients (scattering matrices) are useful constructs. See Figure 2.17 for a pictorial representation of the terms used in this section.

Considering Figure 2.17, and recalling the mirror input-output relations from §1.4.2.2, we can write the full, modal coupling matrix (for fields of a single frequency) as

$$\begin{pmatrix} a_{1o,00} \\ a_{1o,01} \\ \vdots \\ a_{2o,00} \\ a_{2o,01} \\ \vdots \end{pmatrix} = \begin{pmatrix} m_{11} k_{0000} & m_{11} k_{0001} & \dots & m_{21} \rho_{0000} & m_{21} \rho_{0001} & \dots \\ m_{11} k_{0100} & m_{11} k_{0101} & \dots & m_{21} \rho_{0100} & m_{21} \rho_{0101} & \dots \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots \\ m_{12} p_{0000} & m_{12} p_{0001} & \dots & m_{22} \kappa_{0000} & m_{22} \kappa_{0001} & \dots \\ m_{12} p_{0100} & m_{12} p_{0101} & \dots & m_{22} \kappa_{0100} & m_{22} \kappa_{0101} & \dots \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots \end{pmatrix} \begin{pmatrix} a_{1i,00} \\ a_{1i,01} \\ \vdots \\ a_{2i,00} \\ a_{2i,01} \\ \vdots \end{pmatrix}, \quad (2.29)$$

where the coupling coefficients from modes nm into $n'm'$ are noted as k for the front surface reflection, κ for the back surface reflection, p for the front to back surface transmission and ρ for the back to front surface transmission. The terms m_{11} , m_{22} , m_{12} and m_{21} are identical to those given previously in §1.4.2.2. Finally, the quantities $a_{(1|2)[i|o],nm}$ give the field amplitude at the specified node for the corresponding HG_{nm} mode.

From Equation (2.29) it is clear to see that, by using this matrix structure, we can write

$$\begin{aligned} \mathbf{K}_{11} &= \begin{pmatrix} k_{0000} & k_{0001} & \dots \\ k_{0100} & k_{0101} & \dots \\ \vdots & \vdots & \ddots \end{pmatrix} & \mathbf{K}_{22} &= \begin{pmatrix} \kappa_{0000} & \kappa_{0001} & \dots \\ \kappa_{0100} & \kappa_{0101} & \dots \\ \vdots & \vdots & \ddots \end{pmatrix} \\ \mathbf{K}_{12} &= \begin{pmatrix} p_{0000} & p_{0001} & \dots \\ p_{0100} & p_{0101} & \dots \\ \vdots & \vdots & \ddots \end{pmatrix} & \mathbf{K}_{21} &= \begin{pmatrix} \rho_{0000} & \rho_{0001} & \dots \\ \rho_{0100} & \rho_{0101} & \dots \\ \vdots & \vdots & \ddots \end{pmatrix}. \end{aligned} \quad (2.30)$$

The matrices in Equation (2.30) are the scattering matrices for a mirror, and correspond to the symbols of the same names in Figure 2.17. Using these matrices we can write a more compact form of the mirror coupling matrix in Equation (2.29) as

$$M = \begin{pmatrix} m_{11} \mathbf{K}_{11} & m_{21} \mathbf{K}_{21} \\ m_{12} \mathbf{K}_{12} & m_{22} \mathbf{K}_{22} \end{pmatrix}. \quad (2.31)$$

This same procedure can be applied to other components — e.g. beam splitters, lenses — such that the computation of scattering matrices is a generic operation applicable across modal FINESSE 3 simulations.

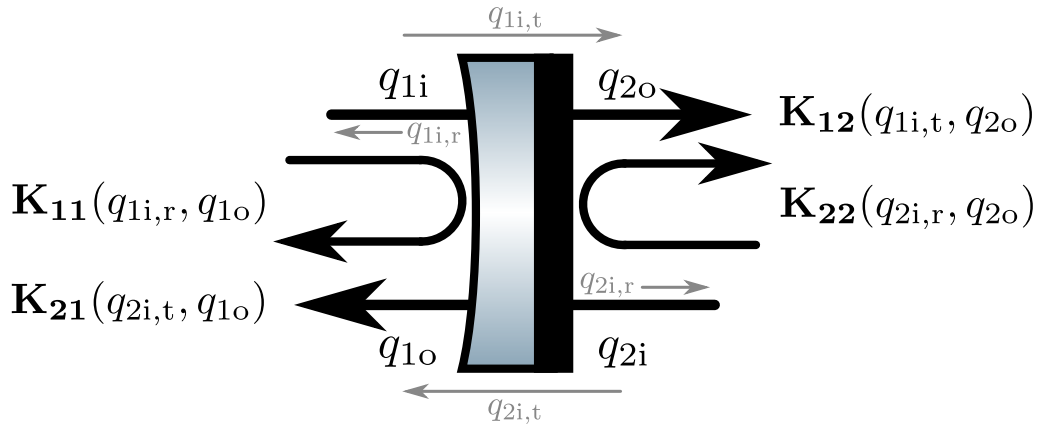


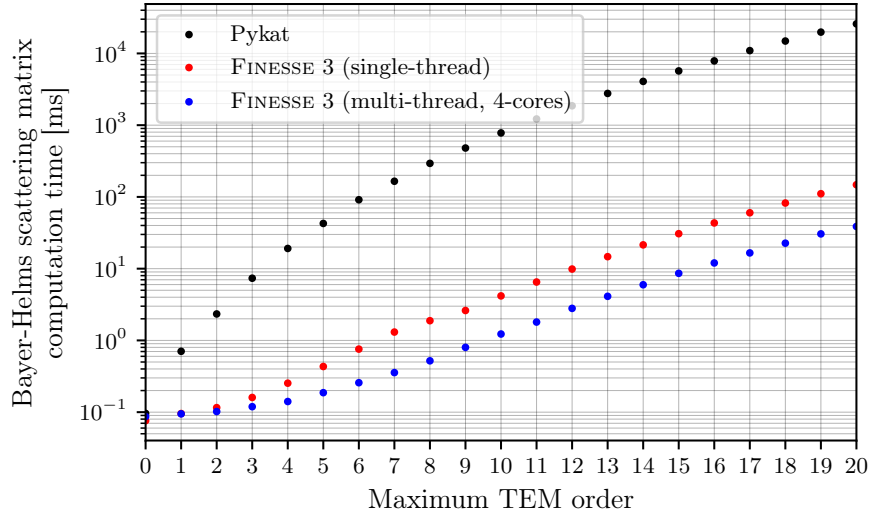
Figure 2.17: Scattering matrices, as a function of their dependent beam parameters, associated with each coupling at a mirror. Each of $q_{1i,t}$, $q_{1i,r}$, $q_{2i,r}$, $q_{2i,t}$ represent the respective input beam parameters transformed according to the reflection (r) and transmission (t) mirror ABCD matrices using Equation (2.2). Note that the \mathbf{K}_{nm} matrices here will also be dependent upon the tilt angle(s) of the mirror relative to the optical axis in the full picture.

2.3.2.3 Performance

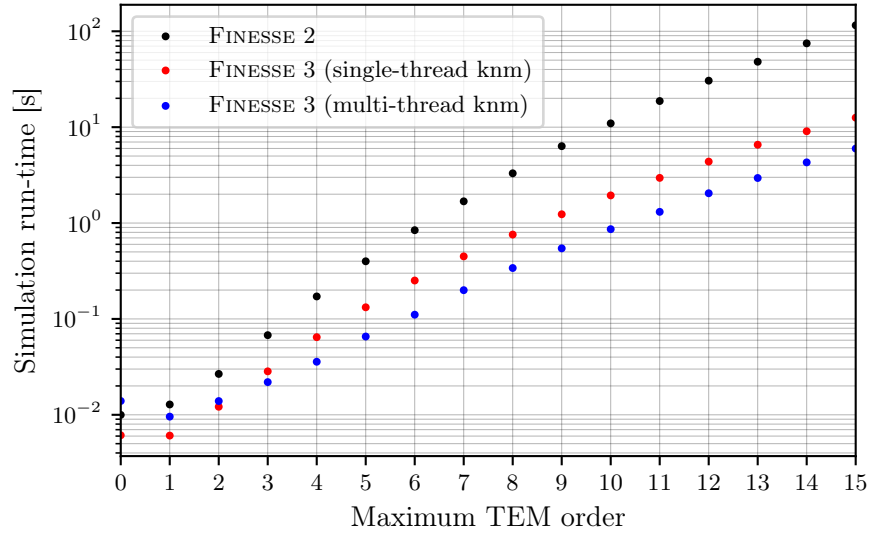
The implementation of Bayer-Helms based coupling coefficient calculations programmed into FINESSE 3 is heavily optimised. Extensive use of variable caching, including contiguous caching of complex powers, is performed to significantly improve the performance of scattering matrix calculations over other implementations of this type of code. A direct example of this is shown in Figure 2.18a, where scattering matrices for a pair of mismatched and misaligned beams are computed for different maximum TEM orders. Note that this maximum order value is a triangular number coefficient, meaning that the total number of HOMs scales as

$$N_{\text{HOMs}} = \frac{(M_T + 1)(M_T + 2)}{2}, \quad (2.32)$$

where M_T is the maximum TEM order (as given on the x -axis of Figure 2.18a). On this plot, we show the same scattering matrix computations using both Pykat [74] and FINESSE 3. The FINESSE 3 benchmarking results are split between single-threaded and multi-threaded modes. The other plot shown in Figure 2.18 gives an indirect comparison for the scattering matrix performance between FINESSE 3 and FINESSE 2. This is done because directly comparing the performance of this code for FINESSE 2 is non-trivial, due to the nature of the code-base. The indirect comparison uses a typical modal simulation of an input beam mode mismatched to a Fabry-Perot cavity whilst sweeping over the tilt misalignment of the input mirror.



(a) Direct comparison of calculation times for scattering matrices, computed via the Bayer-Helms analytic method, in FINESSE 3 versus Pykat [74].



(b) Indirect comparison of scattering matrix computation performance in FINESSE 3 versus FINESSE 2. These data were obtained via a simulation dominated by coupling coefficient calculations.

Figure 2.18: Benchmarking scattering matrix computation times both directly (a) and indirectly (b). The indirect comparison uses a simulation dominated by calculations of coupling coefficients. More specifically, the model it uses constitutes an input beam mismatched to a cavity where we sweep over the tilt of the input mirror of the cavity with an axis scan. This means that coupling coefficients with both mismatch and misalignment constituents must be computed (between all modes up to the given order) at this input mirror at each data point. On both sub-figures the red and blue dots represent single and multi threaded (using a 4-core, 8-thread processor) computations of scattering matrices in FINESSE 3, respectively.

2.3.3 Selecting the mode indices to model

One new feature of FINESSE 3 (not implemented by previous versions of the code) is the ability to explicitly select which mode indices to include before running a modal simulation. This mode-selection feature is physically motivated by the “mode index invariance” of, for example, Equation (2.29). That is to say that it doesn’t matter which modes are included — so long as coupling coefficients are computed between all of them such that the scattering matrices are square in all cases. Non-contiguous jumps between mode indices will not affect the results of a modal simulation, so long as the couplings between the selected modes and filtered-out modes would be zero anyway. This invariance provides the option of significant performance enhancements to specific modal simulations. An example of this will be given in §2.3.3.2. Before we cover this example, however, we will first rationalise the motivation behind this feature and indicate the types of mode selection patterns which can be useful.

2.3.3.1 Misalignments and mode mismatches as additions of HOMs

From §1.3.2.3, we have seen that we can write the spatial profile of Hermite-Gauss modes in the tangential plane as

$$u_n(x, z) = \left(\frac{2}{\pi}\right)^{\frac{1}{4}} \left(\frac{\exp(i(2n+1)\psi(z))}{2^n n! w(z)}\right)^{\frac{1}{2}} H_n\left(\frac{\sqrt{2}x}{w(z)}\right) \exp\left(-i\frac{kx^2}{2R_C(z)} - \frac{x^2}{w^2(z)}\right). \quad (2.33)$$

In this section we will consider “distortions” to this distribution in terms of mode mismatches and misalignments. To simplify the maths here, we can set $z = 0$ such that we are at the waist position of the beam. Doing so reduces Equation (2.33) to

$$u_n(x, z = 0) = \left(\frac{2}{\pi}\right)^{\frac{1}{4}} \left(\frac{1}{2^n n! w_0}\right)^{\frac{1}{2}} H_n\left(\frac{\sqrt{2}x}{w_0}\right) \exp\left(-\frac{x^2}{w_0^2}\right), \quad (2.34)$$

where we have noted that the Gouy phase $\psi(z = 0) = 0$ (using the definition via Equation (1.25)) and the wavefront radius of curvature at the waist is $R_c(z = 0) = \infty$. The fundamental mode ($n = 0$), i.e. a pure Gaussian profile, at the waist can then be written as

$$u_0(x) = \left(\frac{2}{\pi w_0^2}\right)^{\frac{1}{4}} \exp\left(-\frac{x^2}{w_0^2}\right). \quad (2.35)$$

It is from Equation (2.35) that we can now deduce couplings to HOMs when applying defects.

Firstly we will consider the case of misalignment in the tangential plane. We can describe this as [110, 56]:

$$u_{\text{tilt}}(x) = u_0(x) \exp(i k x \sin \alpha) \approx u_0(x) \exp(i k \alpha x), \quad (2.36)$$

where α is the angle of rotation of the beam from the initial optical axis and is assumed to be small. We can expand the exponential in Equation (2.36) to give us

$$u_{\text{tilt}}(x) = u_0(x) \sum_{j=0}^{\infty} \frac{(i k \alpha x)^j}{j!}. \quad (2.37)$$

From this equation we can see that small tilt angles yield HOM terms in $\sim \alpha^j x^j$ with $j \in \mathbb{Z}_{>0}$. This means that a tilt in the tangential plane can be described as additions of modes HG_{j0} ; and, similarly for a tilt in the sagittal plane, modes HG_{0j} . One may then select only these modes, up to a given order (determined by the magnitude of the misalignment), in FINESSE 3 if modelling a system with either of these tilts — assuming that the incident beam of such a system is an aligned, fundamental Gaussian (i.e. HG_{00}). An example of the former case (tangential plane tilts) is given in §2.3.3.2.

We can perform a similar treatment in the case of a mode mismatch. Note that beams can be mismatched in both waist size and / or waist position, however we will only cover the former case; the same principles apply to the latter case, one can find a treatment of this in [110]. We can describe a small fractional offset, ϵ , to the original beam waist size (w_0 in the equations above) via:

$$w'_0 = (1 + \epsilon)w_0. \quad (2.38)$$

This allows us to express the spatial profile of the waist size mismatched beam as

$$\begin{aligned} u_{\text{size}}(x) &= \left(\frac{2}{\pi w_0^2}\right)^{\frac{1}{4}} \sqrt{1 + \epsilon} \exp\left(-\frac{x^2}{w_0^2}(1 + \epsilon)^2\right) \\ &= \left(\frac{2}{\pi w_0^2}\right)^{\frac{1}{4}} \exp\left(-\frac{x^2}{w_0^2}\right) \sqrt{1 + \epsilon} \exp\left(-\frac{\epsilon x^2}{w_0^2}(2 + \epsilon)\right) \\ &= u_0(x) \sqrt{1 + \epsilon} \exp\left(-\frac{\epsilon x^2}{w_0^2}(2 + \epsilon)\right) \\ &\approx u_0(x) \sqrt{1 + \epsilon} \exp\left(-\frac{2\epsilon}{w_0^2}x^2\right), \end{aligned} \quad (2.39)$$

where we have neglected the ϵ^2 term in the exponential by assuming this value is small. Again, we can expand this exponential yielding:

$$u_{\text{size}}(x) = u_0(x) \sqrt{1 + \epsilon} \sum_{j=0}^{\infty} (-1)^j \frac{\left(\frac{2\epsilon}{w_0^2} x^2\right)^j}{j!}. \quad (2.40)$$

From this equation then, we see that terms in $\sim \epsilon^j x^{2j}$, with $j \in \mathbb{Z}_{>0}$ as before, are obtained. This translates to couplings only to *evenly-spaced* modes. Consequently, if the incident beam is described purely by HG_{00} , then a mismatch (in waist size) with a different beam will lead to scattering into the $\text{HG}_{(2j)0}$ modes. Note that this waist-size mismatch applies equally to the sagittal plane, so that $\text{HG}_{0(2j)}$ modes (as well as the ‘‘cross-modes’’ $\text{HG}_{(2nj)(2mj)}$ where $n, m \in \mathbb{Z}_{>0}$) are also produced.

In both cases of misalignment and mismatch, one can note that the coefficients of higher order x terms scale as c^j (where $c = \alpha$ for misalignment and $c = \epsilon$ for waist size mismatch). Thus, for larger misalignments and mismatches, greater HOM orders are required to describe the resulting imperfection(s) between the beams. Note that as this analysis (and FINESSE 3) uses HG modes, any misalignment and mismatch must obey the paraxial approximation (see §1.3.2.1) in order to be able to be modelled accurately. Furthermore, the analytical framework provided by Bayer-Helms [109] also assumes that the order of mode mismatch and / or beam axis misalignment between two frames is relatively small. See Section 4.9 of [76] for a good summary of the limits to the paraxial approximation as pertaining to the coupling of HOMs.

2.3.3.2 Example: Scanning yaw angle of a beam splitter

An exemplary case of how mode-selection can significantly improve the performance of certain modal simulations is highlighted in Figure 2.19. This is based on a model where the light power on reflection from a yawing (i.e. rotating around the y -axis) beam splitter is detected. Included in the plot are computation times for this simulation using all modes up to a given maximum order M_T (see Equation (2.32)); for both FINESSE 2 (red dots) and FINESSE 3 (black dots). Compare these data to the blue markers, which show simulation run-times in FINESSE 3 where only tangential plane modes (i.e. HG_{n0} with $n \in [0, M_T]$) are included now; using the information provided in §2.3.3.1. The results of the simulation are the same in all cases, as tilting the beam splitter in this way will only scatter the incident light (assuming this is purely HG_{00}) into tangential modes; if the plane of rotation was swapped to rotating around the y -axis then only sagittal modes would be produced.

Even with this small example, given by Figure 2.19, the performance benefits are very significant. At a maximum mode order $M_T = 15$, one can see that the mode selection simulation is ~ 200 times faster than FINESSE 2 and ~ 20 times faster than FINESSE 3 with

all modes up to M_T modelled.

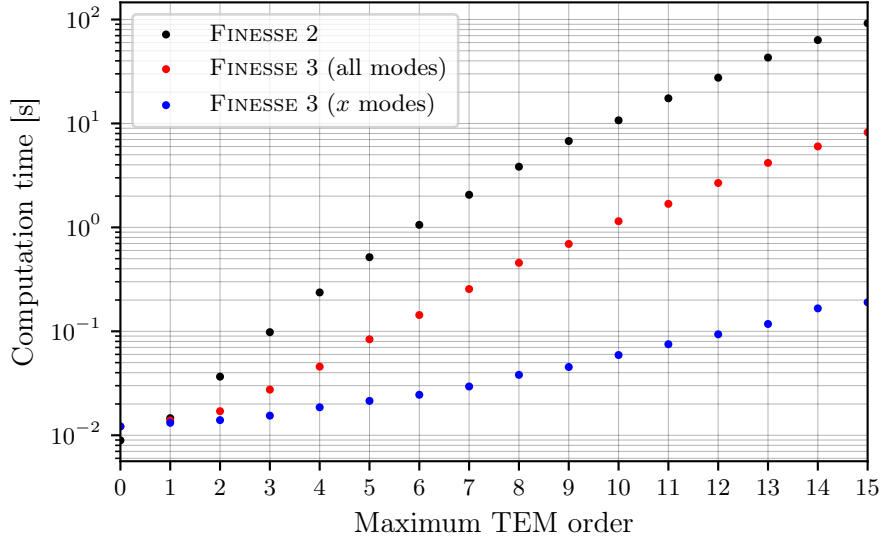


Figure 2.19: Benchmarking the performance of a modal simulation detecting power on reflection from a beam splitter where the yaw angle is scanned. Both the FINESSE 2 and red-dot FINESSE 3 data are from simulations in which no explicit mode selection is performed, whereas the blue-dot FINESSE 3 data selects only tangential modes due to the nature of this simulation.

2.3.4 Future work

The planned work for the near-future, as pertaining to coupling coefficients in FINESSE 3, is to implement scattering matrix calculations for surface maps and clipping due to finite apertures. Both of these features were previously implemented for FINESSE 2, however they use purely numerical integration routines for computing these types of coupling coefficients. This is an area for potential improvement, both in terms of performance and documenting the calculations, in which the FINESSE 3 implementation ideally would include analytic solutions for modal scattering due to surface distortions and apertures. For the latter case, there are possible recipes in the literature [60] which can be followed.

Another prospective improvement is the ability to define small lateral displacements of optics. This is a feature which was not implemented in FINESSE 2, however solutions exist in both the Bayer-Helms (see, in particular, Section III B of [109]) and other (e.g. Section 3.3.3 in [60]) formalisms for describing translations of beams in terms of coupling coefficients. This will be a particularly useful feature to implement, as it would avoid the current necessity to mimic beam shifts through the use of an off-axis telescope with tilt angles applied to both optics (see, for example, [111, 112]).

2.4 Capturing beam images

One way of probing the modal content of light fields in FINESSE 3 is to detect the beam image itself. This detection is separated into two camera types — CCDs and “complex cameras”. The former computes the intensity per unit area (acting like a CCD camera in an experimental setup), see §2.4.1, whilst the latter calculates the amplitude and phase, per pixel, of the beam at a specific frequency, see §2.4.2. Each method relies on using the complex field amplitudes, a_{nm} , at the probed location(s) in the configuration. As these quantities are computed via the modal scattering framework (see §2.3), it follows that the beam images themselves are formed via the computed coupling coefficients.

2.4.1 CCDs

The detector output function for a single pixel at coordinates (x, y) of a CCD type camera in FINESSE 3 is [76]

$$s(x, y) = \sum_{ij} \sum_{nm} u_{nm}(x, y; \vec{\mathbf{q}}) u_{nm}^*(x, y; \vec{\mathbf{q}}) a_{inm} a_{jnm}^* \quad \text{with} \quad \{i, j | i, j \in \mathbb{N}_0 \wedge \omega_i = \omega_j\}, \quad (2.41)$$

where ω is the (angular) frequency of a field and a_{knm} are the amplitude coefficients for the HG_{nm} field at frequency ω_k . The quantity $u_{nm}(x, y; \vec{\mathbf{q}}) \equiv u_{nm}(x, y; q_x, q_y)$ is the full spatial profile function of a HG_{nm} mode (see §1.3.2.3 for details). The value of $s(x, y)$ is then a real number representing the intensity per unit area of the beam.

2.4.2 Complex cameras

Complex type cameras detect the amplitude and phase of a field at a single frequency, as a function of the coordinates (x, y) of each pixel. Commensurately, the equation describing this output is [76]:

$$z(x, y) = \sum_j \sum_{nm} u_{nm}(x, y; \vec{\mathbf{q}}) a_{jnm} \quad \text{with} \quad \{j | j \in \mathbb{N}_0 \wedge \omega_j = \omega_p\}, \quad (2.42)$$

where ω_p is the “probing” frequency — i.e. the frequency of the field to detect. Note that, as opposed to CCD type detectors, the output value, $z(x, y)$, now is a complex number as it contains both the amplitude and phase information for the probed field at the given coordinates.

2.4.3 Dimensionality of camera-type detectors

A significant improvement to beam imaging in FINESSE 3, as compared to previous versions, is the ability to detect the intensity¹ with different dimensions of detector. Each camera probe is separated into three categories — `Pixel`, `ScanLine` and `Image` classes. With `Pixel` cameras, one can detect the intensity at a single coordinate in the transverse plane; defaulting to the central pixel if not specified. The `ScanLine` cameras provide the ability to detect this quantity over a slice of the beam profile, and, naturally, `Image` type cameras compute the full beam profile over both the x and y axes of the transverse plane. A consequence of this is the relative ease for now executing higher-order simulations over multiple parameters whilst calculating a specific dimension of a beam profile. The simulation example in the next sub-section, §2.4.4, is one such demonstration of this. In that example, the beam profile is generated in “one-pass” with a single parameter sweep. This is in contrast to FINESSE 2 where one would need to explicitly scan over both axes of the camera to produce each beam image. Executing a single parameter sweep, whilst using `Image` type cameras in FINESSE 3, also enables one to automatically produce animated plots of the resulting beam profile over the scanned value. An example of this functionality can be found in [113].

In addition to providing a convenient interface, this separation by dimensionality also yields significant performance enhancements. The detector output implementations of the `Image` and `ScanLine` camera types leverage OpenMP [114] for executing multi-threaded computations of Equations (2.41) and (2.42) over the given coordinate data. A benchmarking example of the FINESSE 3 CCD image camera implementation, versus FINESSE 2, is given in §2.4.5.

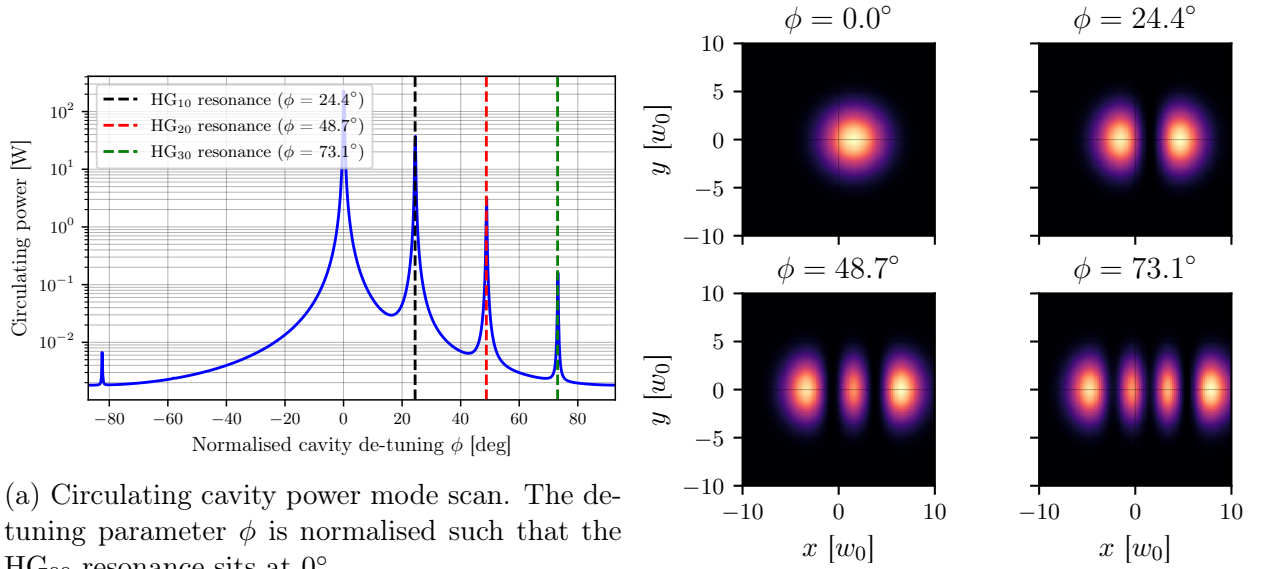
2.4.4 Example: Mode scan of misaligned aLIGO arm cavity

A typical use-case of camera probes in FINESSE is for verification of the beam shape when dealing with a misaligned and / or mismatched system. As an example, in this section we will look at an aLIGO arm cavity which has an optical axis misaligned, in the tangential plane, to the incident beam by $0.8 \mu\text{rad}$ — approximately 3% of the divergence angle ($\Theta \approx 28 \mu\text{rad}$) of the cavity eigenmode. This misalignment will excite HG_{n0} modes, as described by §2.3.3.1. Note that this misalignment angle is fairly extreme, this is just for illustration and is not indicative of the magnitude of residual misalignment observed in the aLIGO arm cavities.

Figure 2.20 demonstrates this example, performed by FINESSE 3 using a photodiode and CCD to detect the circulating power and intensity profile, respectively. Figure 2.20a shows the circulating power when scanning the de-tuning of the cavity. This de-tuning is given as a normalised quantity such that the HG_{00} resonance lies at $\phi = 0^\circ$. From the resulting resonance structure, we obtain the expected mode-spacing of $\delta\phi \sim 24.4^\circ$ — equivalent to the

¹When we say “intensity” in this part, note that one may interchange this freely with “amplitude and phase” as outlined by the detector types earlier.

value calculated via Equation (2.12). Corresponding to this cavity scan, in Figure 2.20b we plot the intensity profiles of the circulating beam at the detunings of each HG_{n0} resonance peak. The beam shapes correspond to the expected profiles seen previously in §1.3.2.3. The small offsets of each intensity profile from $x = 0$ are implicitly due to the misalignment in the tangential plane.



(a) Circulating cavity power mode scan. The detuning parameter ϕ is normalised such that the HG_{00} resonance sits at 0° .

(b) Intensity profiles of the circulating beam, at specified detunings.

Figure 2.20: A cavity scan of an aLIGO cavity with a misalignment of $0.8 \mu\text{rad}$ at the end mirror. On the left we see the expected modal resonance structure, showing scattering into tangential HG modes. In the right sub-plot the intensity profiles of the beam, at each resonance peak, are given — again showing the expected shape corresponding to the resonance of each mode. Note that the axes of the intensity profiles are given in units of the cavity waist size. This modal simulation was performed to maximum order of 8, with only tangential modes selected (see §2.3.3 for mode selection details).

2.4.5 Performance

Again, as with scattering matrix calculations (see §2.3.2.3), my implementation of beam image detection, via Equations (2.41) and (2.42), in FINESSE 3 is well-optimised. As a short benchmarking example, Figure 2.21 shows the run-times of a simulation where the intensity profile of a fundamental Gaussian beam (HG_{00}) is computed for increasing image resolutions. On the resulting plot we compare FINESSE 3 to FINESSE 2 for the same simulations. One can see that FINESSE 3 is consistently and significantly faster than FINESSE 2 for this type of simulation.

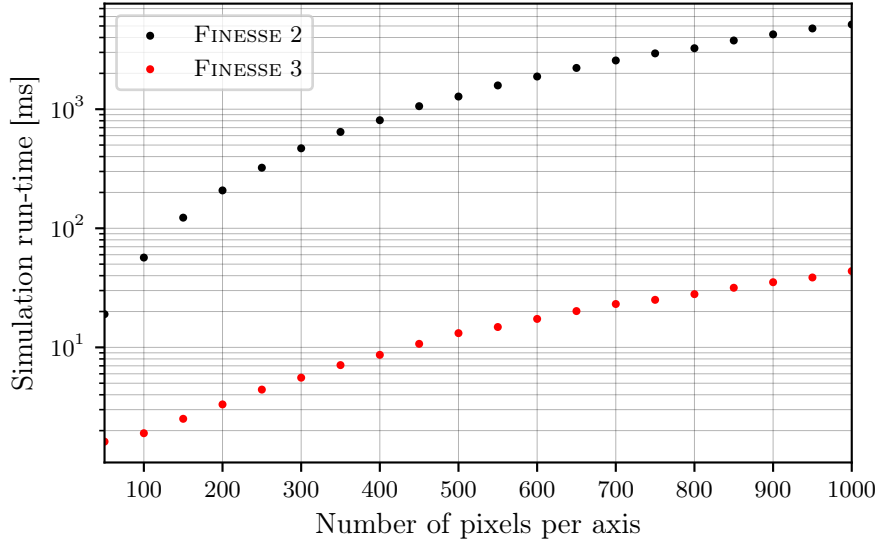


Figure 2.21: A comparison of execution times for simulating the intensity profile of an HG₀₀ beam, using different pixel resolutions. The FINESSE 3 simulations are consistently ~ 100 times faster than FINESSE 2 in this case, as a result of the optimised camera output implementation.

2.5 Summary

The presence of HOMs in GW detectors is an issue which requires careful analysis for gaining an understanding of the underlying physics. Due to the inherently messy nature of the analytics involved with HOMs, modelling these extra fields via interferometer simulation tools is required in most cases. The current software packages in the detector community used for these types of modal modelling tasks provide an excellent toolset in this regard. However, improvements, optimisations and new features can always be built upon these existing frameworks. In this chapter I have introduced some of the HOM simulation capabilities of one such new package — FINESSE 3 — which builds upon the work of previous FINESSE, and related software, versions [73, 74].

A short summary of the aspects of HOM modelling tools I have developed are as follows. For improved modelling capabilities my implementation of asymmetric tracing (§2.2.3.3) and flexible dependency ordering (§2.2.3.5) provide the means for greater control of mode mismatches. On the topic of speeding up HOM simulations, my work on implementing efficient coupling coefficient calculations (§2.3.2.3), beam imaging computations (§2.4.5) and the mode selection feature (§2.3.3) provide significant performance improvements over FINESSE 2. Finally, as utilities outside of typical simulations, the beam tracing tools (§2.2.4) I implemented allow for easy propagation analysis of arbitrary paths in an interferometer-type setup; we will cover such a use case in the next chapter. The symbolic beam tracing framework also provides the potential for further development in the future — for example,

for limiting the impact of numerical errors on mode mismatches by way of simplification of analytically expressed terms.

One of the key developments in FINESSE 3 is the implementation of a flexible and customisable beam tracing library. This library uses the well-established framework of beam parameter propagation via ABCD matrices as well as the analytics behind optical cavity eigenmodes, as covered in §2.2.1 and §2.2.2, respectively. My implementation of the beam tracing is covered in more detail in §2.2.3. Specifically, in this section we showed how the `TraceForest` and `TraceTree` structures are useful representations of the beam propagation paths in an interferometer. A key aspect of FINESSE 3 beam tracing is the ability to clearly define the priorities of propagation from different dependencies. In §2.2.3.5 we showed that this is important due to the effect that these trace orders can have on the results of certain modal simulations. The addition of the asymmetric tracing feature delivers a way to limit, and in some cases remove as we saw in Table 2.1 and Figure 2.11, simulation-specific mode mismatches which would not exist in a real experimental set-up.

Calculating coupling coefficients, from one modal basis into another, underpins the core of how FINESSE 3 models HOM effects in interferometers. Thus, it is important that these computations are well-defined and efficiently implemented. My work in this area focuses towards the latter consideration here. As we saw earlier in §2.3.2.3, my implementation of Bayer-Helms [109] based scattering matrix calculations provide significant performance improvements over previous iterations. On the topic of clearly detailing the nuances of this implementation, §2.3.2.1 provides in-depth definitions of the “reverse Gouy” operation carried out by FINESSE 3. We will recall the analytics of coupling coefficients in Chapter 5, for the specific case of a beam misaligned to the optical axis of an aLIGO arm cavity.

Capturing images of beams in an interferometer can be a good method for checking the modal content of the beam. We have shown, in §2.4, how this can be achieved in FINESSE 3 — both in terms of the intensity profile (as an analog to a CCD in an experiment) and the amplitude, phase of the field. Due to the complicated nature of these calculations, requiring the computation of the spatial distribution $u_{nm}(x, y; q)$ for each coordinate, it is important for these to be well-optimised. This is especially true of simulations requiring lots of modes (e.g. for large mismatches and/or misalignments). It was shown in §2.4.5 that my implementation of these calculations provides a significant speed-up over previous iterations of the FINESSE software.

In summary, in this section we have outlined the importance of modelling HOMs, how FINESSE 3 performs modal simulations and, importantly, improvements that I have made to modal modelling capabilities. Over the next few chapters we will explore specific use cases of FINESSE 3 for simulating modal effects in GW detectors, in terms of design related tasks for the Einstein Telescope.

Chapter Three

Feasibility Study of Beam Expanding Telescopes in the Interferometer Arms for the Einstein Telescope

This chapter is a reformatted and edited version of my paper: Samuel Rowlinson, Artemiy Dmitriev, Aaron Jones, Teng Zhang, and Andreas Freise, “Feasibility Study of Beam Expanding Telescopes in the Interferometer Arms for the Einstein Telescope” [1]. I am the principal author of this paper and was the lead of all of the modelling work. The beam tracing tools used to carry out this modelling work were also developed by me, the details of which can be found in §2.2. Figures 3.5 and 3.10 were not included in this paper and have been added here for extra information. In addition, §3.4.1 was added as a further exploration of the ET-HF telescope design flexibility.

The Einstein Telescope (ET) is a proposed *third-generation* gravitational wave detector [44]. Once constructed, ET will provide an unprecedented level of sensitivity enabling: precise tests of General Relativity, studies of compact binary coalescences involving both intermediate black holes and neutron stars, and will be sensitive enough to test several dark matter candidates [115]. ET combines a unique layout and design combining well-proven concepts from current gravitational-wave detectors with new technology. Figure 3.1 shows a sketch with the basic features of the layout: the ET observatory is composed of three detectors that together form an equilateral triangle. Each detector consists of two interferometers, one low-frequency detector (ET-LF) with its sensitivity optimised for low frequencies from 3 Hz to 30 Hz and another high-frequency detector (ET-HF) with its sensitivity optimised for high frequencies from 30 Hz to 10 kHz.

Similarly to current generation gravitational wave detectors, i.e. Advanced LIGO [13] and Advanced Virgo [14], each interferometer in ET is a Fabry-Perot Michelson interferometer with power and signal recycling cavities (PRC & SRC) for arm cavity power enhancement and shaping the signal response [116], respectively. This interferometer configuration pre-

sented in [44] represents the initial detector anticipated to be installed, with upgrades and refinements to be implemented over several decades. The details of the design of the initial detector will be prepared, in sync with research and development of the required technology, over the next years.




Figure 3.1: The diagram on the top right shows a general overview of the ET observatory layout, with 3 detectors forming a equilateral triangle of 10 km length. Bottom left is a sketch that shows that each detector consists of two interferometers, one optimised for high frequencies (HF) and one for low frequencies (LF). The core interferometer layout is based on a Michelson interferometer with Fabry-Perot cavities in the arms and recycling.

3.1 Motivation

The Fabry-Perot cavities in the arms of the detectors are designed to have large beam sizes on the test mass mirrors, i.e. the input test mass (ITM) and end test mass (ETM), for reducing the impact of thermal noise of the optics on the detector sensitivity. Beam-expander telescopes are used to match input light with smaller beam diameters to the larger beams in the arm cavities. In Advanced LIGO such telescopes are located between the beam splitter and the recycling mirrors [20], and in Advanced Virgo similar telescopes are part of the input-output optics outside the main interferometer [117].

However, in the ET the beam sizes on the main mirrors are significantly larger, requiring a very large substrate for the central beam splitter, larger than the main optics, due to the angle of incidence of 60 deg. In this chapter we investigate the feasibility of an alternative layout with beam-expander telescopes located between the main beam splitter (BS) and the

arm cavities. Such telescopes would provide smaller beam sizes in the central interferometer formed by the beam splitter, power recycling mirror (PRM) and signal recycling mirror (SRM), allowing for the use of much smaller optical components. This not only reduces the cost and complexity of these optics and their suspension systems, but also simplifies the mitigation of secondary reflections and scattered light, and reduces the effect of beam jitter [118]. Figure 3.2 shows a sketch of the optical layout in the lower left corner of the ET triangle, including possible locations for the beam expansion telescopes. An additional advantage of positioning the telescopes between the beam splitter and arm cavity is that Z-shaped telescopes provide flexibility in beam steering, for example they provide the possibility to steer the ET-HF beam around the suspension system of the ET-LF ITM, and they decouple the angle of incidence on the main beam splitter from the angle between the long interferometer arms [119, 120].




Figure 3.2: This is a sketch of the lower left corner of the triangle, showing an example implementation of the optical layout, the vacuum system for the main optics and the corresponding cavern layout. In particular this shows the possible location of Z-shaped telescope systems for ET-LF and ET-HF detectors. In this example the telescopes have been placed to achieve an angle of incidence of 45° on the ET-HF beam splitter.

In this chapter we analyse different arm telescope designs with the target of achieving a 6 mm large beam waist at the main beam splitter. The constraint on the size of this waist (and thus the spot size on the beam splitter) is rationalised in §3.2, where other constraints stemming from the design of the central interferometer and the arm cavities [44] are discussed as well. Regions of interest in our parameter searches, covered in §3.3 for ET-LF and in §3.4 for ET-HF, are defined as regimes where the telescope configuration gives this 6 mm waist positioned at the beam splitter, whilst the SRC is stable.

In §3.5, we present an analysis on the sensitivity of the ET-LF telescope solution to each free parameter. Following on from this, a preliminary study quantifying the necessary active

changes (for mode matching) to the telescope mirror curvatures in the presence of thermal lensing in the ITM, is given. The results in this chapter were obtained using beam parameter propagation [59] via the symbolic ABCD matrix capabilities of FINESSE 3 [75]. The methodology involved grid-based searches of the parameter space to find regions satisfying the aforementioned requirements. In this context, grid-based searches are defined as analyses of high dimensional scans over any given free parameters. An independent verification of the results was performed using a new analytical framework for optimisation of beam expansion telescopes in coupled cavities [121].

3.2 Arm cavity eigenmode and telescope constraints

The maximal beam sizes on test masses is set via the tolerance on power lost through clipping at the mirror edge. This clipping loss (in terms of power) due to a finite aperture for LG modes¹ is given by [56]

$$l_{\text{clip}} = 1 - p!(p + |l|)! \sum_{m=0}^p \sum_{n=0}^p \frac{(-1)^{n+m}}{(p-n)!(p-m)! (|l|+n)! (|l|+m)! n! m!} \frac{1}{\gamma(|l| + n + m + 1, X)}, \quad (3.1)$$

where $X = 2R^2/w^2$, with R as the mirror radius and w as the radius of the beam impinging on the mirror, and γ is the lower incomplete gamma function.

The HG₀₀ mode is equivalent to the LG₀₀ mode, allowing us to use Equation (3.1) for the ET-D optics (setting $p = l = 0$). Given this, we can express the clipping loss for the mirrors in both the ET-LF and ET-HF arm cavities as simply

$$l_{\text{clip}} = \exp\left(-\frac{2R^2}{w^2}\right). \quad (3.2)$$

The relation then, for the maximum beam size w_{max} permissible to achieve a minimum clipping loss l_{clip} is given by:

$$w_{\text{max}} = \sqrt{\frac{2}{\log\left(\frac{1}{l_{\text{clip}}}\right)}} R. \quad (3.3)$$

The resulting beam sizes, and corresponding mirror radii of curvature, for ET-HF and ET-LF

¹We use a formula for LG modes here as clipping losses due to circular apertures are more simply expressed in terms of LG symmetry, as opposed to using an equation based on HG modes.

that satisfy the requirement² of a clipping loss of ~ 1 ppm are listed in Table 3.1. Note that the diameters of the mirrors shown in this table were taken directly from [122, 44] but do not necessarily represent final design values. The arm cavity parameters shown in Table 3.1 are used to define the arm cavity model serving as the starting point of the beam propagation analyses in §3.3.

Table 3.1: A summary of the key parameters of the ET-LF and ET-HF arm cavities. Note that we have assumed symmetric arm cavities for simplicity here.

	ET-HF	ET-LF
Wavelength (λ)	1064 nm	1550 nm
Cavity length (L)	10 km	10 km
Free spectral range ($\Delta\nu$)	15 kHz	15 kHz
ITM/ETM diameter (M_d)	62 cm	45 cm
ITM/ETM curvature (R_C)	5070 m	5580 m
Beam radius on ITM/ETM (w)	12.0 cm	9.0 cm
Beam radius at cavity waist (w_0)	1.42 cm	2.90 cm
Rayleigh range (z_R)	591 m	1702 m
Distance to waist from ITM (z_0)	5 km	5 km
Cavity stability factor (g)	0.95	0.63
Round-trip Gouy phase (ψ_{RT})	333°	285°
Mode separation frequency (δf)	1.1 kHz	3.1 kHz

3.2.1 Telescope parameter constraints

A simplified schematic of one beam-expander telescope is shown in Fig. 3.3. In order to reduce the impact of thermal noise on the sensitivity, ET-LF will make use of cryogenics to cool down the test masses to ~ 10 to 20 K. Cryoshields of around 40 m length will be used along the beam before and after the cryogenic mirrors [122], placing a lower limit on the distance between the telescope mirrors and the ITM. Thus, commensurately setting a lower limit for the SRC length for ET-LF. In §3.3.2 and §3.3.3 this distance is kept fixed at the current design value of 52.5 m.

The picture is different for ET-HF, where a signal recycling cavity length of 100 m or less is preferred in order to improve the quantum-noise limited higher-frequency sensitivity [123].

²The total losses are typically dominated by scattering and absorption, however this standard clipping loss value (which is small by comparison) is used to allow for small beam mis-centering.

ET-HF does not use cryogenics, so that the lower limit on the SRC length is given only by the minimum distance allowed between the vacuum tanks.

As stated in §3.1, we targeted a waist size of 6 mm at the beam splitter. The final design for the size of the beam on the central beam splitter will be based on a trade-off study including the following considerations. In Advanced LIGO the main beam splitter sits between the telescopes and the arm cavities [20] — resulting in a spot size on the central beam splitter comparable to the beam size on the ITMs; and thus requiring a large beam splitter to avoid clipping losses. In contrast, in this work, we propose placing the telescopes between the beam splitter and the arms, to allow use of a smaller beam splitter, an idea briefly discussed in [124]. Assuming a beam splitter radius of 15 cm and 60 degree intersection angle of two incident beams, sub part per million clipping losses are achieved with beam sizes smaller than ~ 10 mm, setting an upper bound on the acceptable beam size at the beam splitter.




Figure 3.3: Schematic of an ET arm telescope with a lens at the ITM. This type of configuration is used for the analyses in this section, where the Z mirrors are flat in §3.3.1 whilst the lens has an infinite focal length in §3.3.2.

Similarly to GEO600 [125], ET-HF will operate with \sim kW levels of power on the central beam splitter [122]. As such, this circulating power will induce thermal lensing in the beam splitter substrate, as is the case for GEO600 [22]. This thermal lensing causes an undesirable excitation of higher-order modes and reduces the interferometric visibility. The strength of this lensing is related to the beam intensity and thus reduced with larger beams at the central beam splitter. For consistency between both solutions, a waist size of 6 mm at the central beam splitter was targeted for both ET-LF and ET-HF. It is important to note, however, that, in the case of ET-HF, this may contribute undesirably to scattering into higher-order modes. The strength of this effect will require further studies.

Note that although the targeted beam size on the beam splitter is 15 to 20 times smaller than that on the test masses of ET-LF and ET-HF, respectively, the thermal noise contribution from the main beam splitter is still much smaller than that from arm cavity mirrors, taking into account the arm cavity finesse, at around 900, and the fact that the beam splitter will have fewer coating layers and a smaller substrate volume.

3.3 ET-LF arm telescope design

In this section we analyse potential arm telescope configurations focusing on achieving a stable SRC whilst keeping a waist size of $w_0 \sim 6$ mm near to the beam splitter. The analyses in this section are performed for ET-LF. In this and the following sections we will describe beam-expander telescopes with curved mirrors with a non-normal angle of incidence. With the commonly used spherical mirrors, such telescopes would suffer from astigmatism which would reduce the mode matching in the interferometer. Note that the following computations assume spherical mirrors with small angles of incidence and negligible astigmatism. The schematic layout shown in Figure 3.2 however implies relatively large angles of incidence; requiring non-spherical mirrors. If the final optical design will include significant angles of incidence, our results provide a good starting point for designing the required non-spherical surfaces, based on the desired beam parameters along the optical path.

3.3.1 Flat Z mirrors

The simplest beam-expander configuration possible is one based on a lens at the ITM whilst keeping the Z mirrors flat (i.e. they are just steering mirrors). Using our target waist size, we can deduce an ITM lens focal length value for this setup, which is shown in Figure 3.4. Also displayed in that figure are the distances to the beam waist from the ITM. This figure demonstrates that a distance of ~ 1 km is necessary for maintaining a waist, of the appropriate size, at (or near to) the beam splitter. Given the impracticality of this distance (implying a comparable SRC length), we can reject this type of configuration.

3.3.2 Curved telescope mirrors, and no lens at ITM

Allowing the mirrors in the Z-configuration to have some curvature gives another type of beam-expander configuration which can be explored for feasibility. Based on the results in the single lens case we expect ZM2 to have a positive RoC for converging the beam rapidly and ZM1 to have a negative RoC in order to achieve a small beam size on the beam splitter over a short distance.

Another criterion comes from the stability of the SRC. The setup has significant number of free degrees of freedom, *i.e.* RoCs of ZM1 and ZM2, the telescope distance (the distance between ZM1 and ZM2) and other free spaces in the SRC, which determine the round trip Gouy phase. But the basic behaviour of this system can be understood intuitively when combining a basic understanding of beam propagation and our simulation results. Firstly, the accumulated Gouy phase contribution from ITM to ZM2 can be ignored, since this distance is much smaller than the Rayleigh range of the beam from the arm cavities which




Figure 3.4: The waist-size (blue) and distance to waist from ITM (red) of a beam matched to the arm cavity for different ITM lens focal lengths. Highlighted on this plot is the focal length value which corresponds to our targeted waist size. Note that even with some flexibility on the waist size of a few millimeter, the required distance from the ITM to the waist is still on the order of at least 0.5 km.

is ~ 1.7 km as shown in Table. 3.1 and the Gouy phase is calculated as,

$$\psi = \arctan\left(\frac{z}{z_R}\right). \quad (3.4)$$

Secondly, the main Gouy phase contribution comes from the distance from SRM to the telescope, because the Gouy phase changes faster near to the beam waist position, see Eq. 3.4. For our analyses, a very short distance from the waist to the SRM of 10 m is assumed. Finally, a minimal distance ~ 100 m between the beam-expander mirrors helps to reduce the overall SRC length to ~ 250 m whilst retaining a 6 mm beam waist.

Figure 3.5 displays the round-trip Gouy phase in this SRC over the curvatures of the telescope mirrors, plotted at two distinct intra-telescope distances (between ZM1 and ZM2) of 50 m and 100 m, to identify the regions of interest. In Figure 3.5a the interesting parameter space is defined by a ~ 6 mm waist size and round-trip Gouy phase $\psi_{\text{RT}} \in [20^\circ, 160^\circ]$ or $\psi_{\text{RT}} \in [200^\circ, 340^\circ]$; i.e. representing a stable SRC which is not near-unstable (see Figure 2.2 in §2.2.2.2). This leads to a distance to the waist of the order of > 300 m. To reduce this distance the analysis was repeated for a longer distance between the Z mirrors as shown in Figure 3.5b. Here the region of interest, defined as above, corresponds to a distance to the waist from ZM1 of less than 150 m.

The results shown in Figure 3.6 were obtained after the wide parameter space search demon-

strated in Figure 3.5. This figure shows that a stable SRC is possible, however, a relatively long SRC is required – on the order of 250 m on average - whilst the RoC of ZM1 is small relative to the ZM2 RoC. We can trade off the achievable SRC length and the required curvature of ZM1, i.e. a shorter SRC can be achieved by reducing the radius of curvature of this mirror – this ultimately comes down to a design choice, based on other design parameters outside the scope of this work.




Figure 3.6: ZM1 and ZM2 RoC combinations yielding a 6 mm waist size where the position of this waist is less than 150 m from ZM1. The blue trace shows the round-trip Gouy phase, ψ_{RT} , in the SRC. All the values in this trace satisfy the condition $20^\circ \leq \psi_{RT} \leq 160^\circ$ so that every solution shown in this figure represents a stable SRC. The red trace gives the corresponding length of the SRC for each of the RoC combinations.

3.3.3 Curved telescope mirrors with a lens at the ITM

The previous solutions, shown in Figure 3.6, lead to a relatively long SRC, > 100 m. In this section we will investigate how adding a lens at the ITM can be used to reduce the length of the SRC and relax the requirements on the radius of curvature of ZM1.

We produced a set of animated plots of the round-trip Gouy phase to gain an intuition of how the solution regions evolve. Similarly to §3.3.2, the content shown here is based on these wide parameter searches and thus only shows the conclusions. We find short focal lengths ($f \lesssim 100$ m) that result in a real waist beyond the ZM1 optic give solutions satisfying our requirements. A focal length of $f = 75$ m was chosen for a more in-depth analysis. Note that any focal length comparable to this value will result in similar telescope behaviour but with slightly different solutions for the ZM1 and ZM2 curvatures and distances to the waist.




 (a) ψ_{RT} with 50 m telescope length.

 (b) ψ_{RT} with 100 m telescope length.

Figure 3.5: Round-trip Gouy phase in the SRC. The z contours here represent the distance to the waist as measured from the ZM1 optic. All white regions are masked out as they describe configurations in which there is a real waist between ZM1 / ZM2 or a virtual waist behind ZM2. Plot (a) shows that there are two distinct regions in which the $w_0 = 6$ mm contour intersects with a stable round-trip Gouy phase - see the blue regions in the upper-left and lower-right quadrants. The lower sub-figure, (b), then shows the same upper-left quadrant for an increased distance between the Z mirrors, where this larger distance leads to smaller distances to the waist from ZM1. This also corresponds to smaller round-trip Gouy phase values due to the majority of the Gouy phase being accumulated from ZM1 to the waist, allowing for a stable SRC in a wider range of curvatures as can be seen from (b). Note that in (a) the small region where the round-trip Gouy phase is 720° occurs due to two beam waists being present in the, unstable, SRC for configurations of this type.

Using this focal length, and a distance between the Z mirrors of 50 m, Figure 3.7 was produced – giving the SRC round-trip Gouy phase over the ZM RoCs, with contours for the 6 mm waist size and distances to the waist (from ZM1) of 50 and 150 m overlaid on the plot. These contours, along with the color-map, frame the region which provides potential configurations for achieving a stable SRC of a suitable length. It is immediately apparent from this figure that the range of possible ZM1 curvatures which give solutions is much larger than for the configuration with no lens at the ITM. Here the lens takes the role of focusing the beam such that both ZM1 and ZM2 have negative RoCs and act together to collimate the beam from the arm cavity. By inspecting Figure 3.7 we find that this solution region approximately corresponds to ZM1 RoC $\in [-130 \text{ m}, -30 \text{ m}]$, ZM2 RoC $\in [-70 \text{ m}, -90 \text{ m}]$. This region is shown in Figure 3.8.




Figure 3.7: SRC round-trip Gouy phase for an ITM lens of focal length $f = 75 \text{ m}$ and a distance between the Z mirrors of 50 m. Note that the solutions now require the curvatures of ZM1 and ZM2 to both be negative - this is because the ITM lens, of the focal length used here, is responsible for focusing the beam to a waist from the arm cavity. ZM1 collimates the beam going towards the beam splitter, whilst ZM2 acts as a “beam expander” to prevent the beam (as propagated from the arm cavity) from focusing down to a waist too quickly.

We can use Figure 3.8 to pick a reference solution for the telescope parameters. One such solution set is given in Table 3.2, where the curvature combinations were chosen such that a relatively short SRC is obtained (important for ET-HF, see §3.2.1) whilst the edges of the solution range are avoided (i.e. avoiding a near-unstable SRC). The resulting g -factor of the signal recycling cavity for this solution set is $g \sim 0.37$. In addition, this solution yields $> 99.9\%$ mode matching between the SRC and arm cavity. A discussion of the ET-HF results, also stated in Table 3.2, is given in §3.4.




Figure 3.8: The key take-away from this plot is that by introducing a lens, shorter SRC lengths can be obtained along with the option of less demanding (i.e. larger) ZM1 RoCs. Taking ZM1 RoC = -40 m as an example data point, we see from this plot that this corresponds to an SRC length of approximately 170 m and a round-trip Gouy phase of approx. 110° . Compare this to the same ZM1 RoC value on Figure 3.6 which gives around 310 m and 155° for these quantities, respectively.

Table 3.2: Parameters of the telescopes chosen using Figure 3.8, with values for the beam size and accumulated Gouy phase associated with these given at each optic in the configuration. Note that the computed values have been given to 2 significant figures to avoid unnecessary precision at this stage. Where appropriate, values for LF and HF have been given separately. See Figure 3.9 for a visual representation of these data for ET-LF. The focal length of the ITM lens, in both cases, is 75 m.

Optic		SRM	BS	ZM1	ZM2		
RoC [m]	LF	-9410	inf	-50	-82.5		
	HF	-630			-63.2		
Beam radius [mm]	LF	6.1	6.2	8.9	30		
	HF	6.3	6.4	8.3	38		
Space		SRM-BS	BS-ZM1	ZM1-ZM2	ZM2-ITM		
Length [m]	LF	10	70	50	52.5		
	HF			80			
Gouy phase [deg]	LF	7.5	39	5.3	0.6	Total accumulated	52
	HF	4.8	26	4.9	0.2	Gouy phase [deg]	36




Figure 3.9: The ET-LF telescope design from the arm cavity to the SRM is shown; where the beam size and accumulated Gouy phase are plotted over the distance propagated through the SRC. Note that the beam is roughly collimated from ZM1 to the SRM; this is an important consideration for Figure 3.11.

3.4 ET-HF arm telescope design

The ET-LF solution given in Table 3.2 was used as a starting point for a similar analysis on the ET-HF arm telescope. The results are then given in the same table, denoted with HF to distinguish the values from LF where appropriate. The larger beam size (impinging on the arm cavity mirrors, see Table 3.1), and shorter wavelength of ET-HF, result in the requirement for a longer telescope length when considering the same waist size target of 6 mm. This increased length requirement can potentially be relaxed by decreasing this waist size target, however the required trade-off study is beyond the scope of this work. The solution given for ET-HF in Table 3.2 is optimised given this waist (and stable SRC) requirement. Note, also, that the ET-HF solution uses the same ITM lens focal length ($f = 75$ m) as the ET-LF solution found in §3.3.3. The resulting g -factor of the signal recycling cavity for this solution set is $g \sim 0.65$. The solution stated also yields $> 99.9\%$ mode matching between the SRC and arm cavity. A trace of the ET-HF SRC beam is shown in Figure 3.10.




Figure 3.10: The ET-HF telescope design from the arm cavity to the SRM is shown; where the beam size and accumulated Gouy phase are plotted over the distance propagated through the SRC. Again, similarly to Figure 3.9, the beam is approximately collimated from the ZM1 optic to the SRM.




Figure 3.11: The accumulated round-trip Gouy phase in the SRC and radius of the beam impinging on BS for both ET-LF and ET-HF where the distance between the ZM1 and BS is decreased from its nominal value given in Table 3.2. All other telescope parameters are kept constant.

Given that both the solutions for ET-LF and ET-HF result in a beam that is roughly collimated between ZM1 and the SRM (see Figures 3.9 and 3.10), we can alter the distance between BS and ZM1 without affecting the beam size on the beam splitter by much more than a few hundred microns. This is demonstrated in Figure 3.11. Even by reducing this distance significantly, e.g. to 10 m, a stable SRC can still be obtained for both ET-LF and ET-HF (with $\psi_{\text{RT}} \approx 60$ and $\psi_{\text{RT}} \approx 40$ degrees, respectively); where the beam size on the beam splitter would then be around 5.7 mm and 5.9 mm for ET-LF and ET-HF, respectively. Of particular importance to ET-HF, this could allow for a nominal reduction in the SRC length from 210 m to around 150 m whilst keeping the other telescope parameters, shown in Table 3.2, constant.

3.4.1 Further flexibility of the telescope design

As implied in §3.2.1, the small beam size on the central BS is an area of potential concern for ET-HF due to the increased thermal lensing this could induce in the substrate. This thermal distortion could then increase the contrast defect thereby reducing the detector sensitivity. For this reason, it can be interesting to investigate the flexibility of this initial ET-HF telescope design with regards to the spot size on the BS versus the stability of the recycling cavities.

There are a number of ways in which an increased beam size on the BS can be achieved, in terms of varying the previously obtained telescope design. One of the simplest investigations in this regard, however, is performed by scanning the RoC of the telescope mirror ZM1 and observing the effect this has on the beam size in the central interferometer as well as the round-trip Gouy phase of the SRC. This is shown in Figure 3.12. This plot gives the beam sizes on the BS (solid blue trace) and SRM (dashed blue trace), and the round-trip Gouy phase of the SRC (red curve), as a function of a scan of the ZM1 RoC away from its nominal design value. It is clear from the resulting data that one may reduce the (absolute) ZM1 RoC to increase the beam size on the BS. This is to be expected due to the stronger divergence effect that such a RoC would bring. A side effect of this is that the SRC is pushed further towards instability, as the values of the round-trip Gouy phase suggest. However, even with a change in the ZM1 RoC from -50 m to, e.g., -40 m the spot size at the beam splitter can be pushed to around 12 mm whilst the round-trip Gouy phase would then be approximately 40° ; i.e. still comfortably within a region of geometric stability.

A trade-off analysis which takes into account more concrete requirements for the beam size at the BS will be needed for ET-HF in the future. Such an investigation can take place once a study into the potential thermal lensing impact of small beams in the central interferometer is carried out. The results shown here, in the mean time, provide a picture of this telescope design flexibility in terms of tuning the geometric parameters to fit new requirements.




Figure 3.12: This figure shows the effect that changing the ZM1 RoC design value has on the beam sizes on the BS and SRM, along with the resultant changes to the SRC round-trip Gouy phase. The nominal design value, $R_C = 50$ m, is shown via the dashed green line. We can see from this plot that by decreasing the (absolute) RoC value of ZM1, a larger beam can be attained at the BS whilst still maintaining a stable SRC.


3.5 Parameter sensitivity and mode matching

Taking the results found in the end of §3.3.3 as our baseline configuration, we can determine the critical parameters of, for example, our ET-HF telescope design. These can be determined by deviating the key free parameters of the system to observe the effect on the SRC waist size, round-trip Gouy phase and mode matching to the arm cavity. Figure 3.13 displays the results of such an analysis, where the left plots give the aforementioned target parameters as a function of the distances between the optics whilst the right plots are based on deviations in the radii of curvature of the Z mirrors and the focal length of the ITM lens.

The mode matching quantity shown in Figure 3.13 is defined by the “overlap” (\mathcal{O}) figure-of-merit given previously via Equation (2.18) in §2.2.3.2. In this case q_1 represents the eigenmode of the SRC propagated to the arm cavity and q_2 is the arm cavity mode itself. Recall that this quantity returns values $\mathcal{O} \in [0, 1]$, where unity indicates a full mode match between the two beam parameters and zero gives complete mode mismatch.

From Figure 3.13 we can deduce, in terms of the optic geometries, that this telescope configuration is most sensitive to the focal length of the lens at the ITM. Thus the next step of this analysis will focus solely on the ITM lens focal length changes as a result of thermal aberrations in the beam. Thermal lensing is investigated, in particular, due to it, potentially, being responsible for the largest effective changes to the ITM lens focal length — as shown

at the end of the next section.



(a) SRC mode properties as a function of small changes in the key SRC free-space lengths.

(b) SRC mode properties as a function of small changes in the telescope optic curvatures.

Figure 3.13: SRC waist-size, round-trip Gouy phase and mode overlap with the arm cavity as functions of the key distances between optics, (a), and the curvatures of the telescope optics, (b). Each deviation is given in relative terms where a value of zero corresponds to the (ET-HF) baseline value given in Table 3.2. From (a) we can see that the distance between the ITM and ZM2 is the most critical length. Whilst, in (b), we find that the focal length of the ITM lens is the critical parameter in terms of the optic geometries. Note that the mode matching values in the lower sub-plot were computed via Equation (2.18).

3.5.1 Mode matching in the presence of thermal lensing

Surface deformation and refractive index differentials, caused by temperature distributions in the mirror substrates, lead to thermal aberrations (lensing) in the beam in ground-based GW detectors [82]. This thermal lensing results in mode mismatches between the arm and recycling cavities. In terms of the optics present in our configuration, the thermal lensing acts to modify the effective focal length of the ITM lens, thereby altering the geometry of the beam in the signal recycling cavity (see the black traces in Figure 3.13 for this effect in a broad sense). To minimise these distortions, adaptive optics are required [126, 127]. This is one of the tasks which could be performed by the arm telescopes of the ET detectors, thus in this section we will quantify the required deviations to the telescope mirror curvatures, for recovering mode matching to the arm cavities, in the presence of varied thermal lens focal lengths. In this case the ET-HF telescope (plus arm cavity) configuration is used, as this detector is designed to operate at high power [44, 122] where thermal lensing will be more prevalent.

Figure 3.14 quantifies the necessary modifications to the radii of curvature of ZM1 and ZM2 in order to recover $> 99.9\%$ mode matching of the SRC to the arm cavity; for an assumed

range of thermal lens focal lengths of $f_{\text{th}} \in [100 \text{ km}, 15 \text{ km}]$. Note that for a strong focal length of 15 km from the thermal lens, the effective focal length of the ITM lens reduces to $f \approx 74.63 \text{ m}$; i.e. a deviation of about 0.5% from the target value of 75 m noted in §3.5. This focal length distortion results in a mode mismatch, between the arm cavity and signal recycling cavity, of approximately 20%. At this extreme thermal lens, the modifications to the Z mirror radii of curvature indicated by Figure 3.14 reduce this mode mismatch to effectively 0%.




Figure 3.14: Simultaneous changes in Z mirror RoCs required for recovering “complete” mode matching from the SRC to the arm cavity. At the extreme of $f_{\text{th}} = 15 \text{ km}$ on this plot, the required deviation in the RoC of ZM1 is approximately 1.4% whilst for ZM2 it is 1.9%.

A more in-depth analysis of the suitability of these ET arm telescope mirrors as actuation points, for more general mode mismatch corrections, continues in Chapter 4.

3.6 Summary

In this chapter we investigated arm telescope configurations, for both ET-LF and ET-HF, which are suitable for the optical layout in the ET 2020 design update [44]. These telescope configurations are motivated by smaller beams on the optics in the central part of the interferometer. The beam expanders also provide the ability to steer the ET-HF beam around the ET-LF ITM suspension systems; with the added benefits of decoupling the angle of incidence on the beam splitter from the beam axes in the interferometer arms. Our requirements for this telescope can be summarised as targeting a 6 mm waist size positioned at the main beam splitter whilst maintaining a stable SRC (quantified approximately as having a length

of the same order as the Rayleigh range of the beam). Further details on these requirements were given in §3.2.1.

We demonstrated that it is possible to achieve a stable SRC, of a sensible length, with telescopes in the arms of both the ET-LF and ET-HF interferometers. Reducing the length of the SRC, in accordance with [123], can be attained via the introduction of a lens at the ITM for pre-focusing the beam from the arm cavity. Our baseline solutions for such a configuration are given in Table 3.2. Further reductions to the length of the SRC, whilst changing the spot size on the beam splitter by only a few hundred microns, are possible via decreasing the distance from ZM1 to BS - see Figure 3.11 for details.

Our baseline configuration for ET-HF was analysed in §3.5 where we found that the focal length of the ITM lens is the critical parameter in terms of the sensitivity for mode matching and SRC stability. However, in §3.5.1, we found that effective changes in this focal length due to thermal lensing can be compensated with actuation on the curvatures of the telescope mirrors. In particular, we saw that the mode mismatch (of approximately 20%) due to a strong thermal lens, with $f_{\text{th}} \sim 15$ km, can be fully corrected with changes of approximately 1.4% and 1.9% in the RoCs of ZM1 and ZM2, respectively.

The results presented here provide evidence for the feasibility of beam-expander telescopes in the interferometer arms of ET and provide essential input for trade-off studies of the optical layout. Further studies are required to study other aspects of this setup, in particular the effects of astigmatism in specific telescope implementations, and the possible negative impact on the contrast defect due to having the telescopes in a configuration that allows differential beam tuning.

Chapter Four

Adaptive Mode Matching with Arm Telescopes in the Einstein Telescope

Following on from the previous chapter, we can now use our initial ET arm telescope (and, by extension, recycling cavity) design to investigate the potential for active mode matching with the telescope mirrors; referred to as ZM1 and ZM2 as before. In §3.5.1 we provided an initial analysis of how these mirrors could correct for thermal lensing in the input test masses (ITMs) of ET-HF. This is, however, only one of a number of distortions which can lead to mode mismatches in the interferometer; and, additionally, it does not necessarily tell us that actuation of the Z mirrors is practical. In this chapter we will address this question of actuation suitability in terms of the types of mode mismatch that may occur in the central interferometer. Note that we will focus solely on ET-HF for these analyses, due to both the planned higher operating power for this detector (inducing stronger thermal lensing effects, as briefly discussed in §3.5.1, see also §4.2.1) and a better understanding of actuation limits and surface deformations of fused silica optics.

In addition to adaptive optics [126], it is also interesting to consider actuation strategies for dealing with static mode mismatches. An example of this could be the deviation of procured optics from design parameters during the manufacturing process — e.g. a difference in the measured focal length of an ITM lens versus the design specification. In this such example one can predict, based on the results shown in Figure 3.14, that the Z mirrors could be used to correct for this type of static distortion. We will, however, explore this idea more carefully (along with adaptive matching) in §4.5.

To assist in the visualisation of the mode mismatches and mirror actuations, we will make extensive use of the “WS phase space” technique first introduced in [128] and later expanded upon by [129]. This concept will be briefly described in §4.1, including an analytic expression for the dimensionless modal power overlap (derived in Appendix B) which was not given in these previously cited works.

All of the results shown in this section have been obtained using FINESSE 3, including the implementation of the analytic WS phase space overlap function as noted in §4.1. The plots showing traces in phase space for optic distortions and telescope mirror actuations rely on the beam tracing functionality of FINESSE 3, which was covered extensively in §2.2.

4.1 WS phase space

As we saw previously in §1.3.2.4, the Gaussian beam parameter can be defined via Equation (1.28). We can “re-cast” the real-part of this equation as a quantity known as the Gaussian defocus S :

$$S = \frac{1}{R_c} = \frac{z}{z^2 + z_R^2}, \quad (4.1)$$

where R_c is the wavefront radius of curvature, z is the distance to the waist and z_R is the Rayleigh range. Note that, as we can see from the latter part of Equation (4.1), this quantity is simply equivalent to the *curvature* of the wavefront — a property of Gaussian beams which was also introduced in §1.3.2.2. Using the defocus, we can then simply write the beam parameter, q , as

$$\frac{1}{q} = S - \frac{i\lambda}{\pi W^2}, \quad (4.2)$$

where W is the beam radius¹. In this way, a pure Gaussian beam is completely defined in terms of the W and S parameters. Similarly to §2.2.3.2 then, where we defined an overlap quantity via Equation (2.18) in terms of two beam parameters, we can define a metric from which we can obtain the loss due to mismatch between two modes in WS phase space; termed “WS space” for short notation. From [129], we can define this quantity, $O(W, S)$, as a dimensionless “power” equivalent expression via a product of overlap integrals:

$$O(W, S) = \iint E(W, S) E(W_P, S_P)^* dx dy \times \iint E(W, S)^* E(W_P, S_P) dx dy, \quad (4.3)$$

where (W_P, S_P) describe the point in the phase space for some primary mode (to which we

¹Note that we typically denote this beam size quantity as w , however in this chapter we will use W to be consistent with the literature [128, 129].

compare all other modes given by each of (W, S)). These overlap integrals act on the fields,

$$E(W, S) = \sqrt{\frac{2}{\pi}} \frac{1}{W} \exp\left(- (x^2 + y^2) \left(\frac{1}{W^2} + ik\frac{S}{2}\right)\right). \quad (4.4)$$

Solving Equation (4.3) analytically (see Appendix B), we obtain:

$$O(W, S) = \frac{4}{W^2 W_P^2} \frac{1}{\left(\frac{1}{W^2} + \frac{1}{W_P^2}\right)^2 + \frac{k^2}{4} (S - S_P)^2}. \quad (4.5)$$

This form of the WS space overlap function has been implemented in FINESSE 3 as a utility tool [81]. Similarly to Equation (2.18), this equation yields values $O(W, S) \in [0, 1]$ where a value of unity indicates full mode matching between the mode (W, S) and the primary mode (W_P, S_P) .

An example of how we will use Equation (4.5) is shown in Figure 4.1; where contours of overlap computed via this equation are plotted as a function of both W and S . This figure uses the ET-HF arm telescope and SRC design, from Table 3.2 in §3.4 of the previous chapter, as a baseline configuration. Here the signal recycling cavity (SRC) and arm cavity (ARM) eigenmodes are propagated to the same longitudinal plane of the interferometer, which is, in this case, the position at which the beam splitter (BS) sits. We take the SRC mode as the primary mode and use Equation (4.5) to plot the contours of overlap with this mode. One can see, as expected from the results of the previous chapter, that the ARM mode is mode matched to the SRC mode with $> 99.9\%$ overlap when considering the system with no distortions.

In addition to providing a concrete example of WS space in the context of our ET recycling cavity design, Figure 4.1 also represents the foundation from which the mode matching and actuation analyses in this chapter are built. Any optic distortion, or actuation, or longitudinal position shift will be propagated, via the beam tracing formalism described in §2.2, to the plane of the beam splitter for comparison to the original SRC mode. The use of the beam splitter as the reference plane for these WS space calculations is motivated by §3.2, in which we asserted that the beam size at, and waist position near to, the beam splitter is a key constraint for ensuring a stable SRC with a sensible spot size on the BS. It is worth noting that, as highlighted in Section IV C of [129], one can easily propagate this phase space to any other plane in the interferometer; which can be useful if performing comparisons at that new plane assists with some other mode matching analysis.




Figure 4.1: Contours of overlap, computed via Equation (4.5), against the SRC mode for the ET-HF configuration described by Table 3.2. The primary mode, by definition, always occupies the point of 100% overlap on these types of plots. For comparison, the ARM mode is also propagated to the same plane as the SRC mode (the beam splitter in this case). As we explicitly designed the arm telescopes to mode match the SRC to the arm cavities, the overlap between these two modes is almost perfect ($\gg 99.9\%$) as expected.

4.2 Core ET mode mismatches in WS space

The mode mismatches we will explore can be summarised as distortions to the recycling mirror curvatures, distortions to the arm cavity mirror curvatures and thermal lensing in the ITMs of the arm cavities. This latter phenomenon can be described as modifying the effective focal length of the lens in the ITM (by applying the thin lens formula) as noted in §3.5.1. The effect of shifting the curvatures, as well as positions, of the telescope mirrors are to be referred to broadly as “actuators” on the telescope mirrors; these will be investigated, by comparison to the mode mismatches, in §4.5.

Other distortions (static or dynamic) can, of course, occur in the interferometer — such as deviations to the mode cleaner cavity mirrors. We will only focus on the “core” interferometer optics of ET-HF for these analyses, however, as mode mismatches between the arm and recycling cavities will have the biggest effect in terms of scattering into HOMs in the central interferometer.




Figure 4.2: Mode mismatches due to RoC distortions of the core optics in ET-HF, propagated to the plane of the beam splitter. The primary mode to which the comparisons are made is the SRC mode as before in Figure 4.1. A key take-away of the plot is the angle between each distortion displacement vector. We can see that the area, in phase space, covered by these mismatches is potentially very large as a result of the near-orthogonality of the thermal lens and recycling mirror lines. Note that the lengths of each displacement vector are currently just for illustrative purposes.

Examples of each of the aforementioned distortions are shown on Figure 4.2. Note that the lengths of each trace are purely illustrative in this case and do not represent typical or expected values for the mode mismatches due to each curvature deformation. See §4.5 for more realistic values as compared to the possible actuation spaces. The dashed traces on this figure represent distortions to the mirror curvatures, grouped by colors according to the locations of the mirrors. The dash-dotted line gives the mode mismatch as a result of thermal lensing in the ITM. The modal basis changes according to each of these variations are all propagated to the longitudinal plane of the beam splitter as before. The purple and blue traces give distortions to the high-reflective (HR) surfaces of the ITM and ETM, respectively, resulting in changes to the arm cavity mode. Due to the Gouy phase separation of $\Delta\psi \sim 36^\circ$ between the ITM lens and recycling mirrors (see Table 3.2), the displacement vectors for the thermal lens and recycling mirror distortions are almost orthogonal (details on why this is can be found in §4.4 and [129]). The implication of this is that the area covered by these distortions, in WS space, is potentially quite large — this is dependent also upon the magnitude of these distortions. As a result, to correct for these mismatches one would require a set of actuators which also, potentially, cover a large (and coincident) area of phase space. We will explore some of the theory behind this actuation effectiveness in §4.4, then investigate the suitability of the telescope mirrors as actuation points for compensating for the different distortions. First, however, a projection of the extreme values of the thermal

lensing, for both ET-LF and ET-HF, is provided.

4.2.1 Quantifying the thermal lens focal lengths

We can predict the most extreme values of the thermal lens focal lengths f_{th} , for both ET-LF and ET-HF, by using a simple model (see page 40 of [130]):

$$f_{th} = \frac{\pi k_{th} w^2}{P_{abs} \beta}, \quad (4.6)$$

where k_{th} is the thermal conductivity of the substrate, w is the beam size, P_{abs} is the total optical power absorbed in the substrate and $\beta \equiv dn/dT$ is the change in refractive index n with the temperature T .

For ET-LF, which will use silicon for the test masses, we have $k_{th} \sim 5 \times 10^3 \text{ W m}^{-1} \text{ K}^{-1}$, $\beta \sim 2 \times 10^{-5} \text{ K}^{-1}$ (see page 160 of [130]), $w \sim 9 \text{ cm}$ (from Table 3.1) and $P_{abs} = 18 \text{ kW} \times 1 \text{ ppm} \times 2$ where 18 kW is the projected ET-LF arm cavity power (see Table 6.1 of [44]), a coating absorption of 1 ppm is assumed and the factor of two assumes equivalent power absorbed in the substrate and coating [131]. Using these values we obtain $f_{th} \sim 180 \times 10^3 \text{ km}$. This is a very small thermal lens effect, corresponding to an effective dioptr² shift of approximately 5 nD. For this reason, it is not so interesting to investigate possible adaptive mode matching compensation of such thermally induced distortions in ET-LF.

Conversely for ET-HF, in which the substrate material is projected to be fused silica, the values for the dependencies of Equation (4.6) are as follows: $k_{th} \sim 1.4 \text{ W m}^{-1} \text{ K}^{-1}$, $\beta \sim 1 \times 10^{-5} \text{ K}^{-1}$ [132], $w \sim 12 \text{ cm}$ (see Table 3.1) and $P_{abs} = 3 \text{ MW} \times 1 \text{ ppm} \times 2$ (again, see Table 6.1 of [44]). Given these values the corresponding focal length for ET-HF is then $f_{th} \sim 1 \text{ km}$. A thermal lens of this magnitude will require a compensation system via heating of the edges of the test masses or compensation plates adjacent to the ITMs.

4.3 Types of actuation for adaptive mode matching

In this section we will briefly cover the different types of wavefront actuator used for adaptive mode matching, and consider those which would be practical at the arm telescope mirrors of the ET detectors. A more complete overview of actuators in aLIGO can be found in [126]. The key types of wavefront actuation employed by current GW detectors rely on the following effects: thermo-elastic deformation of surface curvatures and thermo-refractive gradients in

²A dioptr² is a unit of measurement (denoted with the symbol D) for lenses which is equivalent to the reciprocal of the focal length. For surfaces (i.e. mirrors and beam splitters), a dioptr² is simply twice the reciprocal of the radius of curvature.

optic substrates. The geometry of the former actuator is to heat the rear of a mirror causing the front surface to bend (thermo-elastically) and actuate on the incident beam on reflection. Conversely, the latter actuator relies on the thermo-refractive effect whereby a thermal lens is introduced in a mirror substrate via a CO₂ laser beam projected onto the rear surface of the mirror; this then provides actuation on the incident beam on transmission.

Given that the ET arm telescope mirrors are designed to be used on reflection (refer back to Figure 3.3 for the proposed recycling cavity configuration), the actuator geometry is assumed to consist of rear panel heaters for inducing thermo-elastic deformations of the fronts of the mirror surfaces. Actuation on the recycling mirrors assumes a thermo-refractive gradient effect, based on the analysis given in Appendix A of [129]. Finally, wavefront actuation at the ET-HF test masses will assume ring heater actuators of a similar design to aLIGO (see Section 4 of [126]). Projected actuation ranges, based on the existing technology, for each of these cases are given in §4.5.1. As we will see in this section, the use of realistic values for the actuator ranges, and the utility that they provide, further demonstrate the feasibility of the proposed telescope design.

4.4 Determining actuation effectiveness in WS space

In Section IV D of [129], Perecca *et al* demonstrate that, in general, the degree of orthogonality γ between two actuators (acting on the RoC of optics) placed at longitudinal planes separated by a Gouy phase $\Delta\psi$, is

$$\gamma = |\sin(2\Delta\psi)|. \quad (4.7)$$

In this context, the degree of orthogonality defines the ability of a pair of RoC actuators (of fixed ranges) for actuating on orthogonal quadratures. The maxima (γ is unity, indicating actuator orthogonality) and minima (γ is zero, indicating actuator degeneracy) of this function occur at

$$\gamma = \begin{cases} 1, & \text{when } \Delta\psi = \pi n + \pi/4; \quad n \in \mathbb{Z} \\ 0, & \text{when } \Delta\psi = \pi n + \pi m/2; \quad n, m \in \mathbb{Z}. \end{cases} \quad (4.8)$$

This corresponds to a Gouy phase separation of $\Delta\psi = 45^\circ$ giving orthogonality and $\Delta\psi = 0^\circ$ or $\Delta\psi = 90^\circ$ yielding degeneracy. As the Gouy phase separation of longitudinal planes defines the angles between their distortion displacement vectors in WS space, something we saw in §4.2 with the recycling mirror versus thermal lens traces, we can readily infer the degree of orthogonality from such phase space plots. In fact, if the angle of separation of two vectors in WS space is ϕ then the degree of orthogonality between these is simply given by $\gamma = |\sin \phi|$.

The Gouy phase separation between the actuation longitudinal planes is not enough to completely infer the effectiveness of a pair of actuators. One must also consider the actuation strength (or range) — i.e. the maximum possible shift in the defocus of the beam upon actuation on the optic. This value controls the length of the displacement vector associated with the actuation in WS space. Together, the angle between the displacement vectors and the lengths of these vectors map out an area in phase space which then defines the available zone in which the actuators are able to operate. Pairs of actuators with larger actuation ranges, and a greater degree of orthogonality, project to a larger area of phase space; and, thus, a better mode matching capability.

4.5 Suitability of telescope mirror actuation

In this section we will explore the feasibility of correcting for the distortions shown in §4.2 using the telescope mirrors (ZM1, ZM2) of an ET-HF interferometer arm. The information provided by §4.4 will assist in determining how suitable RoC actuators at these optics are (see §4.5.1) and also the effect of shifting the positions of these telescope mirrors (see §4.5.2). The results of these actuation analyses will then be used to guide a recipe for obtaining an estimate of the initial region of phase space that the SRC mode could occupy, whilst maintaining the ability to recover mode matching to the arm cavities. This will be covered in §4.5.3.

4.5.1 RoC actuators at ZM1 and ZM2

Similarly to the other optic distortions, to obtain traces in phase space for the ZM RoC actuations we scan over these degrees of freedom separately and compare the resulting beam parameters at the plane of the beam splitter; using the original SRC mode as the primary mode as before. The results of such a simulation are shown in Figure 4.3. On this plot the actuation traces for ZM1 and ZM2 are shown as solid blue and red lines, respectively. Each of the other optic distortions (which result in mode mismatches), previously shown in Figure 4.2 of §4.2, are also given here too for comparison.

The Gouy phase accumulated between the two telescope mirrors is $\sim 4.9^\circ$ (using Table 3.2) and so one would expect to see a relatively small angle of separation between the traces for ZM1 and ZM2. This is exactly what we see on Figure 4.3. Further to this, we can also see that the ZM2 (solid red) and thermal lens (dash-dotted green) traces lie almost parallel to each other. Again this is to be expected due to almost no Gouy phase being accumulated from the ITM to ZM2. The actuation ranges shown on Figure 4.3 are as follows: $40 \mu\text{D}$ for the test masses (from pp. 13-14 of [126]), 50mD for the recycling mirrors (see Section III B.2 of [129]) and $150 \mu\text{D}$ for the telescope mirrors (based on [133, 134]). These values use the aLIGO optics and compensation systems as a reference, the actual values will depend

upon the design of each of these optics in ET (as well as the design of actuation systems) — however these should provide a good starting point for this type of feasibility analysis.




Figure 4.3: This figure shows the actuation displacement traces for the telescope mirrors, ZM1 (blue) and ZM2 (red), as compared to distortions (or actuations) on the other core optics. The shapes at the end of each trace describe the approximate limit on the relevant actuation ranges; see the text for details, including references, on these values. From the ZM1, ZM2 actuation traces we can note the relative degeneracy between these two potential actuator locations, arising due to the small Gouy phase separation in these longitudinal planes. The range of thermal lens focal lengths is given for comparison to the ZM2 actuation, it is not indicative of projected values for such thermal distortions.

The triangles at the extrema of the ZM1 and ZM2 actuation traces in Figure 4.3 represent the maximum actuation strength from above. We can note, from these, that actuation on ZM1 is largely ineffectual relative to actuating on ZM2. This is to be expected, given the results shown in Figure 3.13b of §3.5. The green dash-dotted curve shows the mismatches as a result of a range of thermal lens focal lengths of $f_{\text{th}} \in [\infty, 65]$ km. In §4.2.1 we found that the predicted thermal lens for ET-HF is $f_{\text{th}} \sim 1$ km; a distortion which requires thermal compensation as noted previously. Actuation on ZM2 could then be used to perform fine tuning of the mode matching for such distortions — allowing the initial actuator to be coarse. The thermal lens focal length range given above highlights the extent of this fine tuning potential by comparison to the ZM2 actuation trace in Figure 4.3.

4.5.2 Shifting the positions of ZM1 and ZM2

Another potential option which we consider here, just for completion, is the possibility of actively shifting the relative positions of the telescope mirrors — again, for the purpose of adaptive mode matching. A corresponding representation of this is shown in Figure 4.4, where the longitudinal locations of ZM1 and ZM2 are shifted. The blue trace shows the effect of increasing the distance between the central BS and ZM1, the red line represents the distance between the two telescope mirrors and, finally, the green curve is as a result of moving ZM2 relative to the ITM (of either arm). The corresponding triangles are the maxima of the distance shifts used; 1 cm for each deviation in this case.




Figure 4.4: The effect on the SRC mode given shifts in the positions of the telescope mirrors. Each distance deviation is given an upper limit of 1 cm. This shows that, for this particular design, such shifts do not have any significant impact and are therefore likely not suitable for adaptive mode matching purposes.

It is clear from this figure that shifting the positions of the telescope optics has no significant impact on the SRC mode — particularly for the blue and red traces. The green trace (ITM to ZM2 shift) has the most meaningful effect, as shown, however this is still negligible in comparison to RoC actuation on ZM2 as indicated by the previous sub-section (see §4.5.1). One can observe why this is the case by inspecting Table 4.1, in which the Rayleigh range values across each segment of the SRC are given. The large z_R values for the spaces BS-ZM1 and ZM1-ZM2 imply that one would need (impractical) distance shifts of several metres (or many tens of metres in the former case) to have any significant impact on the mode matching.

Table 4.1: The values of the Rayleigh range in each segment of the ET-HF SRC, based on our initial arm telescope design from §3.4. Distances between the optics in this recycling cavity are also given for completion. See Table 3.2 for other details on the SRC design, such as accumulated Gouy phases.

	SRM-BS	BS-ZM1	ZM1-ZM2	ZM2-ITM
Distance z [m]	10	70	80	52.5
Rayleigh range z_R [m]	114	114	2.46	0.137

4.5.3 Estimating the available phase space regions

Returning to the surface RoC actuation picture, we can begin to build up an estimate for the available region (as represented in WS space) that the ET-HF SRC mode could occupy whilst maintaining the potential to recover $\geq 99.9\%$ mode matching to the arm cavities. To perform such an analysis, we will use the results obtained from §4.5.1. More specifically, we assume that our available actuators are those acting on the ZM2 telescope mirror and on the recycling mirrors, independently. The estimated ranges of these actuations were given previously; in the legend of Figure 4.3.

The recipe used for obtaining an available region in phase space was adapted from the process outlined in Section V B of [129]. In this analysis we will propagate our phase space to the output of the SRM, rather than using the previous representation at the beam splitter. Given that the beam is approximately collimated in the ZM1 to SRM path (as shown in §3.4), this makes little difference to the parameters and, hence, to the mode matching contours. The logic behind propagating to this plane will become clear from the following outline of the recipe. For each mode that exists within the 99.9% mode matched contour, at the longitudinal plane of the SRM, we propagate this mode backwards to the ZM2 optic (in either arm). The first step of this propagation involves a reflection ABCD matrix transform at the SRM. In this backwards propagation, we apply actuations to the SRM and ZM2 mirrors — sweeping over each combination within the given ranges. The resulting mode at ZM2 is then propagated back to the SRM using the non-actuated RoC values for each mirror. This then gives us a new beam parameter, at the SRM, which represents a mode that can be displaced back into the 99.9% mode matched contour by actuating on it with the corresponding ZM2, SRM actuation combination. The results of this analysis are given in Figure 4.5. Note that the choice of the SRM as the primary plane here should be intuitive. In particular, using this plane as the reference simplifies the above process — as it only requires a single backwards and forwards propagation per mode due to the SRM and ZM2 RoCs affecting both of these in a single pass. The beam propagation function of FINESSE 3, covered in §2.2.4, was employed as a core tool for performing the aforementioned reverse and forward mode traces required for this analysis.

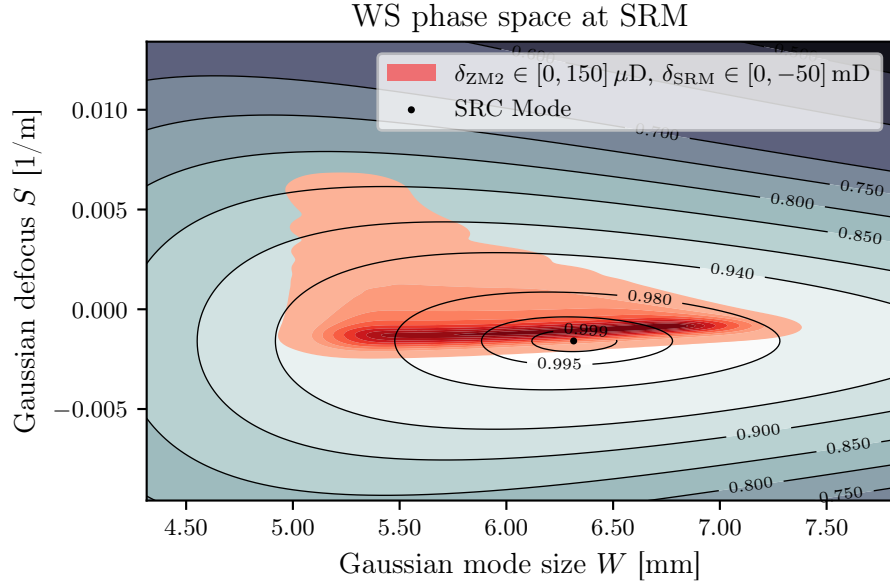


Figure 4.5: This figure shows an estimate of the initial region of WS space that the SRC mode could occupy, with actuation on the RoCs of ZM2 and SRM of the specified strengths. Modes in this highlighted zone can be displaced to a mode matching value of $\geq 99.9\%$ through actuations on these optics; up to the specified limits. See the text for a description of the recipe used to obtain this region. Darker red tones indicate modes which can be more readily shifted back into a mode matched state. More specifically, there exists a higher relative number of actuation value combinations, which result in at least 99.9% mode matching, for these modes when compared to the lighter red regions.

The red distribution shown in Figure 4.5 represents a Gaussian kernel density estimate (KDE) of the propagated mode data. The darker tones of this colour-map indicate SRC mode values which were produced more frequently as a result of the procedure outlined above. Commensurately, for such regions, there exists a larger number of combinations of actuation values for displacing these modes back into the 99.9% mode matched region; as compared to those modes in the distribution with a lighter red tone. Consequently, the darker red regions represent SRC configurations which require less adjustment to the nominal parameters (i.e. RoCs of SRM and telescope mirrors), to achieve mode matching to the arm cavities. This KDE was produced via utilisation of the distributions module of the software package seaborn [135].

4.6 Summary

Adaptive optics are important in GW detectors for dealing with, for example, thermal distortions that can arise due to high laser power slightly deforming optical surfaces and substrates. In the case of the Einstein Telescope, it is ET-HF which will be most affected by such phe-

nomena; where the impact of these distortions may be much greater than the effects observed in current facilities, due to the significantly increased beam power. The novel recycling cavity design in the ET interferometers, involving telescopes in the arms (as discussed in detail in Chapter 3), also presents an acute challenge with regards to reducing real-time losses due to mode mismatches. For these reasons, it is important to develop a robust set of strategies for actuating on the necessary optics in order to keep the interferometer well mode matched.

In this chapter we covered an initial study in this regard. The recycling cavity design, including arm telescopes, of ET-HF from the preceding chapter was analysed in terms of the suitability of actuation on the key optics for recovering from mode mismatches. To perform this work, we made extensive use of the WS phase space technique developed by Perecca *et al* [129]. An extension to this in the form of a new overlap function, see Equation (4.5), and Appendix B for the derivation, was introduced and used for calculating and plotting the contour plots shown throughout this chapter. We saw in §4.5 that, due to the small Gouy phase separation of the ITM and ZM2 longitudinal planes, the actuation trace for ZM2 in phase space lies almost parallel to the curve representing thermal lens distortions at the ITM. As a result of this, actuation on this optic could be suitable for correcting for small thermal distortions — up to a thermal lens focal length of approximately 65 km assuming an actuation range of $150 \mu\text{D}$; see Figure 4.3 for the details. Using this figure we also noted that, in this initial configuration, actuation on the telescope mirror ZM1 would be relatively ineffectual. Furthermore, in §4.5.2, we showed that small displacements of the telescope optics would also not have any significant impact on the real-time mode matching. However, the position of ZM2 could potentially be shifted as a pre-matching step during installation to correct for small static offsets of curvatures from design values.

Finally, using the results obtained above, in §4.5.3 we estimated the initial region in phase space that the SRC mode could occupy; whilst retaining the ability to recover to an acceptable level of mode matching. This was predicated on the existence of actuators at ZM2 and the recycling mirrors, with actuation ranges of $150 \mu\text{D}$ and 50mD , respectively. The resulting distribution highlighted the zones of phase space from which it would then be easier to recover back to a mode matched state.

Further studies on the actuation strategies for the ET interferometers will be required, however these analyses provide a good starting point and concrete recipes for performing such investigations. Any changes to the underlying recycling cavity designs will still be able to rely upon the processes described and used here. One outstanding challenge, which remains to be tackled thoroughly, is the task of quantifying the requirements for the telescope design and actuation on the corresponding mirrors when taking into account differential arm effects. This will be an important issue for ensuring that any contrast defect can be minimised.

Chapter Five

Scaling of Alignment Noise Coupling with Larger Beams

Alignment control noise is one of the dominant noise sources, in current GW detectors, in the low-mid frequency regime [26, 136]. In particular the angular control scheme in aLIGO, used to keep the interferometer optics aligned, contributes a significant amount of noise to the differential arm (DARM) motion at frequencies below 20 Hz [26, 84]. A good visual indication of how alignment control affects DARM can be found in Figure 2 of [26], where measured data for the full noise budgets (including the angular control noise) are given for both the LIGO Hanford Observatory (LHO) and LIGO Livingston Observatory (LLO) for the third observing run (O3). Note that the reduction of angular control noise is a topic of ongoing research within the detector community [52].

The alignment of the optics will be particularly important for future detectors operating at higher powers, e.g. ET-HF, where radiation pressure could be more of a concern. The increased circulating power in the arm cavities of such detectors can potentially drive larger torques on the test masses due to this radiation pressure effect [137, 138]. This is an effect which has been studied in detail [139] and can lead to a decrease in the stability of the arm cavities when they are in a “soft” mode regime due to the mirrors being pushed further out of alignment. It should be noted, however, that future detectors will likely use heavier test masses to somewhat nullify this increased radiation pressure effect; for example, see Sections 6.1 and 6.5 of [44] where the masses of the mirror for ET are briefly discussed.

Due to the effect alignment control has on the overall detector noise, it is therefore important to understand how the optic misalignments themselves couple into longitudinal signals and therefore why they are necessary to control, and, in particular, how this is relevant for future ground-based GW detectors. In this chapter we will specifically focus on the question of how larger beams could affect alignment noise coupling, and hence consider the implications this has for future alignment control schemes. The topic of larger beam sizes in future detectors is a common one throughout this thesis, as we saw in Chapter 3 with regards

to the implications for recycling cavity designs in ET, and therefore it is investigated here as part of this concurrent, and important, research topic. To briefly highlight why larger beams are vital to understand, one can note that proposals for essentially all future ground-based GW detectors include larger mirrors as a key upgrade. These include, for example, LIGO Voyager [140, 27], ET [44], Cosmic Explorer (CE) [45], Advanced Virgo Plus upgrades [141] and the proposed Neutron Star Extreme Matter Observatory (NEMO) [142]. A good summary of the baseline parameters, which includes the projected beam sizes, for a selection of these future detectors can be found in Table 2, Section LT-3.4.2, of [27].

From prior research [143] we expect that, in general, the alignment to longitudinal noise coupling will increase significantly with increased beam sizes. By investigating the mathematical details of this coupling in more detail, however, we can reveal which parameters can be tuned to offer suppression of this noise; this topic will be discussed in §5.4. This is a key goal of this chapter and, thus, a significant portion of it will be dedicated to clarifying the analytics involved with this type of noise coupling (see §5.2). As a supplement to these analytical scalings, Appendix C provides an in-depth derivation of the field couplings for the type of misalignment we are considering (highlighted in §5.1). To validate the derived scaling relationship, we will then apply this type of misalignment to a full aLIGO model in §5.3, using FINESSE to simulate the resulting couplings.

5.1 Static and oscillatory misalignments




Figure 5.1: This is a schematic of the type of alignment coupling being considered, where a static (DC) misalignment, from the optical axis of the incident beam, of β_{DC} is applied at the input mirror (ITM) of a cavity whilst the end mirror is tilted (in the same rotational degree of freedom) by β_{AC} at a signal frequency of ω_m . See Appendix C for a more detailed picture including the fields present in this type of system.

Figure 5.1 shows the type of alignment noise coupling we are considering in this chapter. This involves a static (DC) misalignment of the input mirror of a Fabry-Perot cavity from an incident beam. In addition, the end mirror of the cavity tilts at some “signal” frequency ω_m . The alignment fluctuations in this system are then a sum of the static (DC) and oscillatory (AC) terms [144]. The combination of these two misalignments cause a coupling from the incident mode (assumed to be a pure HG_{00} mode) to the HG_{01} mode at the frequency ω_m . This latter field then couples back into the fundamental mode at this sideband frequency — a phenomenon often called alignment to longitudinal coupling. A more in-depth look at the specific field couplings at each frequency can be found in Appendix C. Note that here we are considering only small mirror tilts in the *sagittal* plane such that, firstly, the generated couplings are only to first order and, secondly, HG_{01} modes are scattered into and not HG_{10} . This type of analysis works equally well when considering only small tilts in the tangential plane too, in which case the latter modes are generated instead of the former.

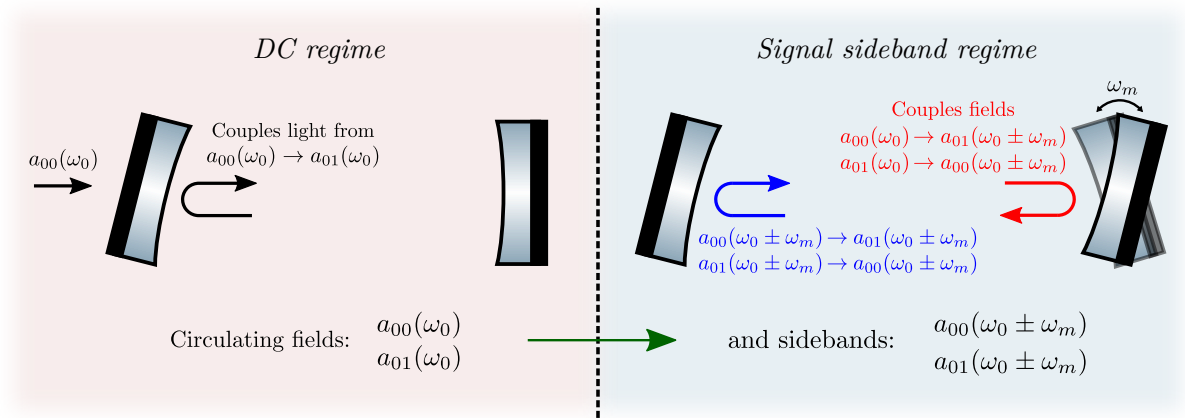


Figure 5.2: A representation of Figure 5.1 in the modal picture. Here the specific field couplings in both the carrier (DC) and signal sideband (AC) regimes are shown. In this picture, as the carrier fields are resolved initially, we can treat the DC model as an initial system where only the static tilt is applied at the input mirror. This results in the circulation of fields in both the HG_{00} and HG_{01} modes in the cavity. These fields are then applied to the signal picture, such that the tilt modulation applied to the end mirror produces signal sidebands (at the modulation frequency ω_m) of these fields as shown.

The view given in Figure 5.2 describes the exact couplings we are considering in more explicit detail. In this figure the fields present are separated into two regimes — the DC (carrier light) regime and AC (signal sideband) regime — corresponding to the two distinct solving steps performed by a modal-based simulation like FINESSE. Note that in this modal picture, the carrier fields (left) are necessarily resolved before calculating the signal sideband fields (right). It should be noted that the principle of *small* tilt modulations (to the end mirror, in the context of this figure) is important here. Through this we are able to neglect second-order frequency shifts, e.g. $\omega_0 + \omega_m \rightarrow \omega_0 + 2\omega_m$, by assuming that the field quantities at these higher modulations are negligible.

Using the field couplings from Figure 5.2, we can determine the form of the signal which will be detected at the photodiode (of the corresponding system in Figure 5.1). One important aspect to note, in relation to this signal, is that HG_{00} and HG_{01} modes do not beat on the photodiode. This means that, in first order, without any modulated light (from the AC couplings) no alignment fluctuations would be visible as noise in the GW channel. It is for this reason that it is important to include AC misalignment terms, rather than just static misalignments, for modelling this effect. We can write the signal detected at the photodiode as

$$P \sim a_{00}^*(\omega_0) a_{00}(\omega_0 + \omega_m) + a_{00}(\omega_0) a_{00}^*(\omega_0 - \omega_m) + a_{01}^*(\omega_0) a_{01}(\omega_0 + \omega_m) + a_{01}(\omega_0) a_{01}^*(\omega_0 - \omega_m), \quad (5.1)$$

where the quantity $a_{nm}(\omega)$ represents the field described by the mode HG_{nm} at a frequency ω ; with ω_0 as the carrier frequency. One can see from Equation (5.1) that it is important to include the static misalignment at the input mirror for modelling this alignment coupling, otherwise the field $a_{01}(\omega_0) = 0$ and, commensurately, both terms $a_{00}(\pm\omega_m) = 0$; leading to a “null” response when only an AC misalignment is modelled.

The scattering of light in the field HG_{01} at the carrier frequency back into HG_{00} light at the signal frequency is the key effect here; i.e. the terms $a_{00}(\omega_0 + \omega_m)$ and $a_{00}^*(\omega_0 - \omega_m)$ of Equation (5.1), which are “generated” according to Figure 5.2. If at any point the oscillation frequency ω_m matches a GW signal frequency, then this will cause coupling of light into the GW channel — mimicking a longitudinal signal.

5.2 Approximating the scaling of alignment noise coupling for larger beams

In this section we will approximate the alignment noise coupling as a function of the spot size of the beam impinging on the cavity mirrors from Figure 5.1. We will find from this that this type of alignment noise, coupled into the longitudinal GW channel, scales as w^6 ; where w is the symmetric beam size in the cavity. This work is based on [143] where this approximation was originally introduced — we will derive this result in more detail in this section. As a follow-up to this, we will then validate this approximation using an aLIGO FINESSE model in §5.3.

A derivation of the coupling relations for each field, up to first order, associated with Figure 5.1 can be found in Appendix C. We will briefly use some results from this appendix to rationalise why the specific quantities in this section are expressed in terms of the beam size. In this appendix, we find that the circulating carrier HG_{01} field, a_{01} , and signal sideband

(either upper or lower) field, \bar{a}_{01} , can be computed as

$$a_{01} \approx \sqrt{\frac{R}{1-R}} \frac{Rk_{01} \exp(-i(\psi_{\text{RT}} - \frac{\pi}{2}))}{1 - R \exp(-i\psi_{\text{RT}})} \quad \text{and} \quad \bar{a}_{01} \approx \sqrt{\frac{R}{1-R}} \frac{\kappa_{01} \exp(-i(\frac{\psi_{\text{RT}}}{2} - \frac{\pi}{2}))}{1 - R \exp(-i\psi_{\text{RT}})}, \quad (5.2)$$

where R is the power reflectivity of both cavity mirrors and ψ_{RT} is the cavity round-trip Gouy phase. The quantities k_{01} , κ_{01} are the coupling coefficients from the HG_{00} mode to the HG_{01} mode on reflection from the input mirror (rear surface) and end mirror (front surface), respectively. It is important, therefore, to know how these coupling coefficients scale with the beam size — this will be investigated in the next sub-section, §5.2.1. One can observe that the denominators for these fields are common factors, known as the resonance enhancement factor d , and in this case are denoted as

$$d_{\text{HG}_{01}} = \frac{1}{1 - R \exp(-i\psi_{\text{RT}})}, \quad (5.3)$$

i.e. the resonance enhancement factor for the HG_{01} mode. Again, to determine the amplitude scaling of the alignment coupling with beam size, it is important to understand how $|d_{\text{HG}_{01}}|$ changes with w too. This is explored in §5.2.2. A summary of the analytic treatment, including all these factors, is given in §5.2.3.

5.2.1 Coupling coefficient from HG_{00} to HG_{01}

Equation (5.2) defines the terms which multiply to give the signal response as detected at the photodiode of Figure 5.1. It is clear from these terms that the couplings from the HG_{00} to HG_{01} mode are important. For this reason, here we will derive an expression for the corresponding coupling coefficient k_{0100} (see the definition of $k_{nmn'm'}$ in §2.3.1) as a function of the misalignment angle, γ , and the beam size w . Note that this coupling coefficient applies for both the static (DC) and oscillatory (AC) pictures, as indicated by Figure 5.2. In the former case this coefficient is dependent upon γ_{DC} whilst in the latter it is dependent upon $\gamma_{\text{AC}} \sim \sin(\omega_m t)$. The derivation of this coupling coefficient is identical otherwise. Also note that, as shown in [109], we can use the orthogonality of Hermite-Gauss modes to separate $k_{nmn'm'}$ into $k_{nmn'm'} = k_{nn'} k_{mm'}$. This means that the coupling coefficient we will derive here can be given as $k_{00} k_{01}$.

As we are considering a case of only angular misalignment with no mode mismatches, we can use a significantly simplified form of the coupling coefficients. In Section III C of [109]

an expression for the partial coupling coefficient $k_{nn'}$ for a mode matched system is given as

$$k_{nn'} = (-1)^{n'} E^{(x)} \sqrt{n!n'} \sum_{\mu=0}^{\min(n,n')} \frac{(-1)^\mu \bar{X}^{n-\mu} X^{n'-\mu}}{\mu!(n-\mu)!(n'-\mu)!}, \quad (5.4)$$

where

$$X = - \left(\frac{\bar{z}_2}{\bar{z}_R} + i \right) \frac{\gamma}{\bar{\gamma}_0}, \quad \bar{X} = X^*, \quad E^{(x)} = \exp \left(-\frac{X\bar{X}}{2} \right) \quad (5.5)$$

assuming that the lateral displacement terms are zero. The terms \bar{z}_2 , \bar{z}_R and $\bar{\gamma}_0$ represent the distance to the waist, Rayleigh range and divergence angle of the beam in the second medium. Given that we are only considering couplings on reflection from the cavity mirrors, this second medium is equivalent to the first — thus meaning that these quantities are equal to the equivalent quantities for the cavity itself. Using this information, and noting that the divergence angle is

$$\bar{\gamma}_0 = \frac{\lambda}{\pi w_0} = \frac{w_0}{z_R}, \quad (5.6)$$

we can re-write the term X from Equation (5.5) as

$$X = - \left(\frac{z}{z_R} + i \right) \frac{\gamma z_R}{w_0} = - (z + iz_R) \frac{\gamma}{w_0}, \quad (5.7)$$

where z is the distance to the waist, z_R is the Rayleigh range and w_0 is the waist size of the cavity mode. Furthermore, we can write the exponential misalignment term as

$$\begin{aligned} E^{(x)} &= \exp \left(-\frac{X\bar{X}}{2} \right) = \exp \left(-\frac{1}{2} \frac{\gamma^2}{w_0^2} (z^2 + z_R^2) \right) \\ &= 1 - \frac{1}{2} \left(\frac{\gamma^2}{w_0^2} (z^2 + z_R^2) \right) + O(\gamma^{2n}). \end{aligned} \quad (5.8)$$

By assuming that γ is small, and noting that only terms which are linear in γ need to be considered for the alignment coupling transfer function, we can see that $E^{(x)} \approx 1 + O(\gamma^2) \approx 1$ as there are no terms linear in γ . Using this, and resolving the sum in Equation (5.4), we find a simple expression for the partial coupling coefficient from the fundamental mode to

n' -th order modes as

$$k_{0n'} = \frac{(-X)^{n'}}{\sqrt{n'}}. \quad (5.9)$$

Applying this to the first order mode $n' = 1$, as is our focus here, we obtain

$$k_{01} = -X = (z + iz_R) \frac{\gamma}{w_0}. \quad (5.10)$$

At this point we can note that $k_{00} = 1$, implying that the coupling coefficient from HG₀₀ to HG₀₁ will be given purely by the expression for the partial coefficient k_{01} . Taking the modulus of Equation (5.10), we obtain

$$\begin{aligned} |k_{01}| &= \frac{\gamma}{w_0} |z + iz_R| \\ &= \frac{\gamma}{w_0} \sqrt{z_R^2 \left(1 + \left(\frac{z}{z_R} \right)^2 \right)} \end{aligned} \quad (5.11)$$

from here we can note, from Equation (1.17), that $\sqrt{1 + (z/z_R)^2} = w(z)/w_0$, which lets us write

$$\begin{aligned} |k_{01}| &= \frac{\gamma}{w_0} \frac{w}{w_0} z_R \\ &= \frac{\pi}{\lambda} \gamma w. \end{aligned} \quad (5.12)$$

Commensurately, we can express k_{01} as

$$k_{01} = \frac{\pi}{\lambda} \gamma w \exp(i\varphi), \quad (5.13)$$

where φ is the angle of the cavity eigenmode — i.e. $\varphi = \arctan(z_R/z)$, or for a symmetric cavity $\varphi = \arctan(\cot(\psi_{\text{RT}}/4))$ with ψ_{RT} as the round-trip Gouy phase of the cavity.

Figure 5.3 shows how the magnitude of Equation (5.13) compares against the values computed using FINESSE 3 (which does not use the approximations noted in this section). These values are plotted as a function of the misalignment angle γ which is given as a fraction of the divergence angle $\bar{\gamma}_0$ of the beam. The model used for this specific example is a sym-

metric ET-HF arm cavity; see Table 3.1 for the data representing this type of configuration. One can see from Figure 5.3 that Equation (5.13) is only valid for very small tilt angles ($\gamma/\bar{\gamma}_0 \lesssim 0.02$).




Figure 5.3: A comparison of $|k_{01}|$ using FINESSE 3, black line, versus the approximation in Equation (5.13), red dashed line. The values are given as a function of the fractional misalignment angle — i.e. the angle of misalignment over the divergence angle of the beam. This figure implies that the analytic approximation derived in this section is only valid for very small tilt angles.

5.2.2 Resonance enhancement factor

The reciprocal of the denominator of a transfer function (on transmission or reflection) for a Fabry-Perot cavity is known as the resonance enhancement (or suppression) factor. For a cavity of length L , this quantity is defined by [143]

$$d = \frac{1}{1 - r_1 r_2 \exp(-i(2kL + (1 + n + m)\psi_{\text{RT}}))}, \quad (5.14)$$

where r_1, r_2 are the amplitude reflectivities of the input and end mirrors, respectively, and ψ_{RT} is the round-trip Gouy phase of the cavity (see Equation (2.11) for the relevant definition in this case). The values n and m represent the indices of any given HG mode as usual. We can assume that the cavity is tuned such that the carrier HG₀₀ mode is resonant — i.e. $2kL + \psi_{\text{RT}} = 2\pi N$ with $N \in \mathbb{Z}$. This then simplifies the exponential term of Equation (5.14)

as follows:

$$\begin{aligned}\exp(-i(2kL + (1 + n + m)\psi_{\text{RT}})) &= \exp(-i(2kL + \psi_{\text{RT}})) \exp(-i((n + m)\psi_{\text{RT}})) \\ &= \exp(-i((n + m)\psi_{\text{RT}})),\end{aligned}\tag{5.15}$$

which, in turn, reduces d to,

$$d = \frac{1}{1 - r_1 r_2 \exp(-i(n + m)\psi_{\text{RT}})}.\tag{5.16}$$

The magnitude of d can then be computed as,

$$|d| = \sqrt{\frac{1}{1 + R_1 R_2 - 2r_1 r_2 \cos((n + m)\psi_{\text{RT}})}},\tag{5.17}$$

where R_1 and R_2 are now the *power* reflectivities of each cavity mirror. The magnitude of the enhancement factor for the HG₀₁ mode follows:

$$|d|_{\text{HG}_{01}} = \sqrt{\frac{1}{1 + R_1 R_2 - 2r_1 r_2 \cos \psi_{\text{RT}}}}.\tag{5.18}$$

Now we recall, using Equation (2.10), that for a symmetric Fabry-Perot cavity the round-trip Gouy phase can be written as

$$\psi_{\text{RT}}(g) = 2 \arccos(g),\tag{5.19}$$

where $g_1 = g_2 = g$ is the stability factor of both cavity mirrors; as introduced in §2.2.2.4. We note that

$$\cos \psi_{\text{RT}} = \cos(2 \arccos g) = 2 \cos^2(\arccos g) - 1 = 2g^2 - 1,\tag{5.20}$$

and use this to write Equation (5.18) as

$$|d|_{\text{HG}_{01}} = \sqrt{\frac{1}{1 + R_1 R_2 - 2r_1 r_2 (2g^2 - 1)}}.\tag{5.21}$$

We can make the assumption now that the cavity is of high finesse such that $R_1 R_2 \approx 1$, this then yields

$$\begin{aligned} |d|_{\text{HG}_{01}} &\approx \sqrt{\frac{1}{2 - 2(2g^2 - 1)}} \\ &= \frac{1}{2} \sqrt{\frac{1}{1 - g^2}}. \end{aligned} \quad (5.22)$$

The g -factor can be expressed as a function of the beam size on the mirrors using

$$g(w) = \sqrt{1 - \left(\frac{L\lambda}{\pi w^2}\right)^2}, \quad (5.23)$$

which then allows us to write

$$|d|_{\text{HG}_{01}} \approx \frac{1}{2} \frac{\pi}{L\lambda} w^2. \quad (5.24)$$

From Equation (5.24) we can deduce that, for a high-finesse cavity of fixed length L , the resonance enhancement factor scales as the square of the radius of the beam impinging on the cavity mirrors for first order HG modes. The consequence of this is that the HG_{01} (and HG_{10}) field amplitude for the cavity is enhanced by a factor of the beam size squared. As a numeric example, the ratio of Equation (5.24) for the ET-HF to aLIGO arm cavities would be approximately two — using data from Table 3.1 and Table 1 of [145], respectively.

5.2.3 Scaling of the alignment noise coupling

A full analytic treatment of the field couplings for this system can be found in Appendix C, here we will simply state the final result obtained. It was determined in this appendix that one can express the beating between the HG_{01} fields at DC and AC, reflected from the cavity, as

$$|a_{01}^*(\omega_0) a_{01}(\omega_0 + \omega_m)| \approx \frac{R^2}{2} \left(\frac{\pi}{\lambda}\right)^4 \frac{w^6}{L^2} |\beta_{\text{DC}} \beta_{\text{AC}}| |i_0|^2, \quad (5.25)$$

where R is the power-reflectivity of the mirrors (assuming an impedance matched cavity), L is the cavity length and i_0 is the amplitude of the HG_{00} mode, at the carrier frequency, incident upon the cavity. A comprehensive list of the assumptions made can be found at the

start of the aforementioned appendix. The values β_{DC} and β_{AC} are the angles as given in Figure 5.1.

From Equation (5.25) we obtain the previously noted (see §5.2) beam size to the power of six scaling. Note that this expression also depends upon the cavity length (which is inexorably linked to the beam size), the mirror reflectivities and the wavelength of the beam. Considerations following from these dependencies will be discussed in §5.4. First however, we will validate and compare this beam size scaling relation to a numerical simulation using a full aLIGO model — a scenario in which it is not simple to prepare an analytic formula.

5.3 Numerical experiment for scaling validation with FINESSE

To validate the approximate scaling relation given in Equation (5.25) of §5.2.3, we will show the results of a simulation using a full advanced LIGO model in this section; see Figure 1.1 in §1.1 for the reference configuration. Misalignments of the type previously described, and illustrated in Figure 5.1, are applied to both of the arm cavities present in the model. The exact magnitudes of these misalignments will be discussed in §5.3.2. First, however, we will explore a suitable range of beam sizes on which to test the alignment coupling relation.

5.3.1 Selecting a range of arm cavity beam sizes


The current aLIGO design values, for both arm cavities, utilise RoCs of 1934 m and 2245 m for the ITM and ETM, respectively [145]. These values, coupled with the arm cavity length of $L \sim 4$ km, yield beam sizes of $w_{\text{ITM}} \sim 5.3$ cm and $w_{\text{ETM}} \sim 6.2$ cm; with a cavity g -factor of $g \sim 0.83$ [146]. In this section we select a beam size space to use, focusing on values which yield beams larger than the current design. For simplicity, and consistency with the analytics shown in §5.2, the following analyses will assume symmetric arm cavities. As a result of this the data points for the cavity geometries, which give our different beam sizes, must be shifted from the design values given above; whilst remaining consistent with the near-concentric regime of the arm cavities.

From Equation (5.23) we know that the beam size w at the cavity mirrors for a symmetric cavity is given by


$$w^2 = \frac{L\lambda}{\pi} \sqrt{\frac{1}{1 - g_c}}, \quad (5.26)$$

where L is the cavity length and $g_c = g^2$ is the stability factor of the cavity. For the

simulations in this section we will keep L fixed, therefore the only parameter which can be varied to alter w is the radius of curvature of the cavity mirrors. As g_c depends on this radius of curvature, and w increases as g_c increases, a large range of beam sizes can be obtained by using a range of cavity stabilities from the current (design) g -factor, given above, to $g_c \rightarrow 1$. An appropriate way to visualise this is by plotting the arm cavity g_c values as a function of the cavity mirror RoCs. This will show us a suitable range of RoCs yielding a symmetric arm cavity which sits in the concave-concave regime. Figure 5.4a demonstrates this with g_c plotted as the colour-map. Highlighted on this plot are two contours — the first corresponds to the current g -factor from the design, whilst the second is the critically stable factor corresponding to the “maximum” beam size on the cavity mirrors. Also given on this plot is the location at which the current design sits, as noted at the start of this section. The red line represents a symmetric arm cavity in the near-concentric regime, with g -factors stretching from the current design to the critically stable contour. Corresponding to this line of symmetric RoCs, in Figure 5.4b, are the beam sizes impinging on the arm cavity mirrors for each of these values. It is this range of curvatures, and associated beam sizes, which will be used throughout the rest of this analysis.



(a) Arm cavity g -factors against the mirror RoCs.



(b) Radii of the beam impinging on the cavity mirrors, corresponding to the symmetric RoC values given on red line of Figure 5.4a.

Figure 5.4: This figure gives a representation of the arm cavity geometry regimes used for the alignment coupling analyses. In (a) the cavity stability factors are shown, with a corresponding line (highlighted in red) giving symmetric arm cavities in the near-concentric regime. Plot (b) then displays the beam sizes on the cavity mirrors corresponding to this selected symmetric cavity space. The point corresponding to the current aLIGO design values for the cavity mirror RoCs is given in (a) for illustrative purposes.

5.3.2 Magnitudes of DC and AC misalignments

One key aspect of the simulation to consider is the size of the misalignment angles on the cavity mirrors. We recognise from §5.1, and Equation (5.1) in particular, that the static misalignment of the ITMs for the arm cavities must be non-zero in order for the detected alignment response to be non-zero. This is not the only consideration, however, as Figure 5.3 previously demonstrated. In that figure we saw that our derived analytic approximation for the alignment scaling depends upon the misalignment angles being very small compared to the divergence angle of the beam. Refer back to Figure 5.1 for the symbols β_{DC} and β_{AC} used in this section.

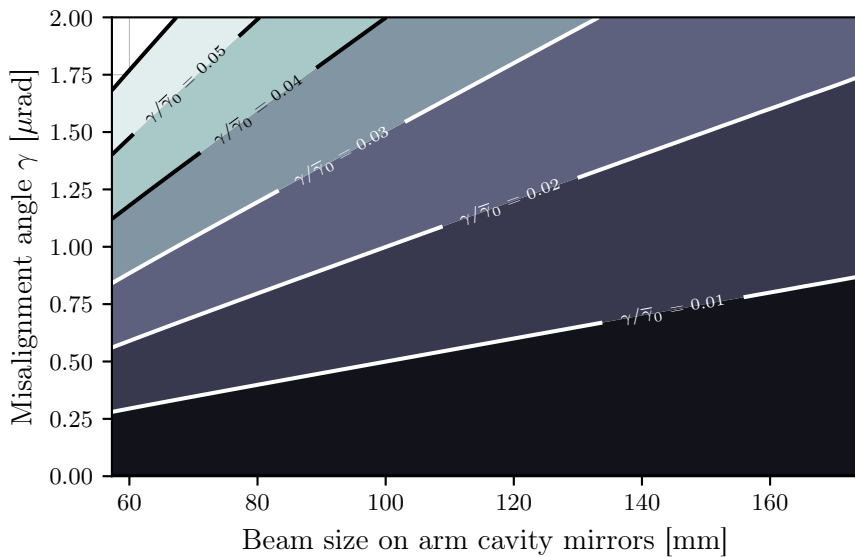



Figure 5.5: This figure shows the fractional angular misalignment (γ divided by the divergence angle of the cavity mode, $\bar{\gamma}_0$) as a function of the beam sizes selected via Figure 5.4. Recall from Figure 5.3 that values $\gamma/\bar{\gamma}_0 \gtrsim 0.02$ result in significant errors for the analytic approximation of $|k_{01}|$. We can avoid this issue by staying in a regime where the misalignment angle is significantly smaller than the divergence angle of the cavity, for all values of the cavity eigenmode.

We will now rationalise the choice of the oscillatory misalignment angle β_{AC} ; the same logic here applies to the static misalignment angle, which was set to $\beta_{\text{DC}} = 1$ nrad using the error estimation analysis of Appendix C.4.1. The divergence angle of the cavity mode scales as the inverse of the waist size, as noted from Equation (5.6), and the waist size decreases as the beam size at the mirrors increases [146]. Thus, the divergence angle increases as the beam size w increases — implying that our limit on the misalignment angle is constrained more at smaller values of w . This is demonstrated in Figure 5.5 where the beam sizes obtained from Figure 5.4b are used to find contours of the quantity $\gamma/\bar{\gamma}_0$, also as a quantity of the misalignment angle γ . We see from this plot that ideally, to avoid issues with a break-down of the approximation given by Equation (5.13), the value of the angle β_{AC} should be chosen

such that it lies in the black region of the filled contours. It is also important to use a value large enough to avoid numerical errors due to very small detected alignment coupling signals in the simulation. For these reasons, a value of $\beta_{AC} = 0.1 \mu\text{rad}$ was chosen; this corresponds to an actual misalignment angle, between the coordinate systems, of $\gamma = 0.2 \mu\text{rad}$.

5.3.3 Mode matching the arm cavities to the recycling cavities

A side effect, which must be taken into consideration when modelling these types of geometric changes in a GW detector model, is that significant mode mismatches will occur in the central interferometer if no compensations are made. In this case, as the eigenmodes of the arm cavities change whilst scanning over the arm cavity mirror RoCs, the mode mismatches to the recycling cavities will increase as we shift further from the operating points. These mode mismatches would cause significant couplings to HOMs at the output of the interferometer, as noted in §2.2.3.2. The magnitude of this issue is highlighted in Figure 5.6a where the mismatch between the arm cavities and each recycling cavity is given as a function of the pre-determined range of beam sizes on the arm cavity mirrors. One can see from this figure that, at the most extreme beam size shift, the mismatch from the arm cavity to both recycling cavities hits $\sim 75\%$.



(a) Arm to recycling cavity mode mismatches. (b) Recycling cavity round-trip Gouy phases.

Figure 5.6: These figures together provide the picture of the mode mismatch issue when shifting the arm cavity eigenmodes (denoted as ARMs) to obtain larger beam sizes and the effect of the matched solutions on the recycling cavity stabilities. In (a) the mode mismatches, computed via one minus Equation (2.18) as a percentage, between the arm cavities and recycling cavities are given as a function of the beam size range previously discussed. To correct for these mode mismatches, the radii of curvature of PR2, PR3 and SR2, SR3 were optimised to minimise these mismatch values. The resulting effect on the recycling cavity round-trip Gouy phase, ψ_{RT} , is shown in (b).

To compensate for these mismatches, prior to detecting the alignment coupling signal we

optimised the radii of curvature of the recycling cavity telescope mirrors — namely PR2, PR3 of the PRC and SR2, SR3 of the SRC; see §1.1.2 again for the locations of these optics. The RoCs of these mirrors were optimised such that the mismatch values between the arm and corresponding recycling cavities were minimised; see [147] for the optimisation routine used for this. The effect of the RoC shifts to these mirrors, for each data point, is shown on Figure 5.6b in terms of the round-trip Gouy phase (and hence, stability) of the recycling cavities. One can observe from this plot that, as might be expected, increasing the beam size at the arm cavity mirrors pushes the recycling cavities (and especially the SRC) further towards instability due to the larger beam then present in the central interferometer.

5.3.4 Detecting the alignment coupling signal

Now that we have covered the prerequisite steps necessary for computing the alignment signal with suitable parameters, and in an optimal modal basis, we can discuss the exact simulation performed and its results. A flow chart detailing the steps undertaken by the FINESSE simulation used in this analysis is displayed in Figure 5.7. Boxes 1 and 2 in this figure were covered by §5.3.2, whilst A and B were discussed in §5.3.1 and §5.3.3, respectively. The next step, after mode matching the central interferometer and the arms, is to ensure that the interferometer is at a sensible operating point — at a basic level this means keeping the arm cavities on resonance for the carrier HG_{00} mode and minimising the SRC power in this mode. This is the step shown in box C, whereby the locking loops which control each of the degrees of freedom of the interferometer (i.e. DARM, CARM, MICH, PRCL and SRCL [148]) are executed and the resulting feedback is applied to the optics. Step D then simply involves detecting the output signal, which is computed via Equation (5.1), by placing amplitude detectors for the HG_{01} mode (at probing frequencies of both ω_0 and ω_m) after the signal recycling mirror.

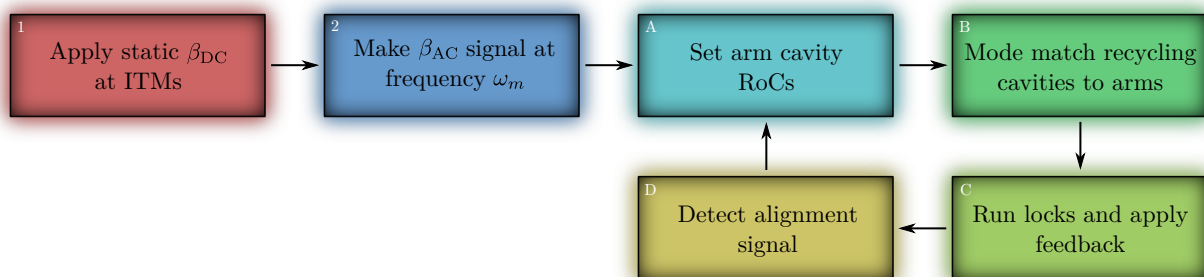


Figure 5.7: An overview of the simulation steps undertaken for modelling the alignment coupling signal, as outlined and detailed throughout this section.

The results of this simulation are captured in Figure 5.8. In this figure the beam sizes impinging on the symmetric arm cavity mirrors are given on the x-axis, whilst the y-axis gives the alignment coupling to photodiode signal detected on transmission through the SRC.

One can see from the resulting trace, and the curve fitted to it, that we obtain an alignment scaling approximately in-line with the expected relation in §5.2.3. The reason behind the relatively poorer fitting function quality at small values of the beam size (as compared to larger beams), lies in how the divergence angle of the cavity eigenmode scales with the spot size. As noted in §5.3.2, the divergence angle increases as the beam size at the cavity mirror increases. A consequence of this is that the ratio of, e.g., the static misalignment angle to the divergence angle increases for smaller beams — leading to a deterioration in the derived scaling relation, as noted in §5.2.1 (and in Appendix C). This argument applies equally to the oscillatory misalignment too as the magnitude of this tilt angle term is also fixed during the simulation.




Figure 5.8: This figure shows the alignment to longitudinal coupling signal as a function of the beam sizes impinging on the arm cavity mirrors. The dashed red line gives the power-law curve fit to these data, indicating a scaling of $P \sim w^6$ as expected. Note that the fit function given is in the units where w is in metres. The poorer quality of fit for smaller beam sizes was hinted at in §5.3.2. As the divergence angle of the cavity mode scales with the size of the beam on the mirrors, the ratio of the misalignment angle (considering either residual or AC terms) to this divergence angle thus increases for smaller beams.

5.3.4.1 Higher order couplings

There can be a few additional reasons as to why the blue data curve shown in this Figure 5.8 does not match the scaling relation exactly. Firstly, in §5.2.1 we assumed that the misalignment angle γ is small such that non-linear terms in this variable were neglected. Whilst it is true that the misalignment angle chosen was small, these higher order coupling terms are still non-zero. Further to this, our scaling relation only considers terms in the first order HG mode. In the simulation, and in practice, the couplings to higher modes (e.g. HG₀₂) will also be non-zero; and thus contribute to the detected signal at the output. Lastly, we also

assumed in §5.2.2 that the arm cavities were of large enough finesse such that the product of the mirror reflectivities is unity. This is also an approximation which is not applied in the simulation.

An additional consideration is the presence of mode mismatches in the simulated interferometer. Whilst the recycling cavities are explicitly mode matched to the arm cavities, as detailed in §5.3.3, the optimisation routine employed for this mode matching is not perfect, and, consequently, small mismatches between the arms and recycling cavities do occur. As a result of this, a very small portion of the light is scattered into the $\text{HG}_{02|20}$ modes, and ends up propagating to the anti-symmetric port of the interferometer. Again, power in these modes complicates the scaling (with beam size) of the field then detected at this point. To optimally alleviate the effect of mode mismatches between the cavities, the IMC and OMC were both removed from the model when performing the simulation which produced Figure 5.8; preventing the need for an extra step in matching these mode cleaner cavities to the recycling cavities.

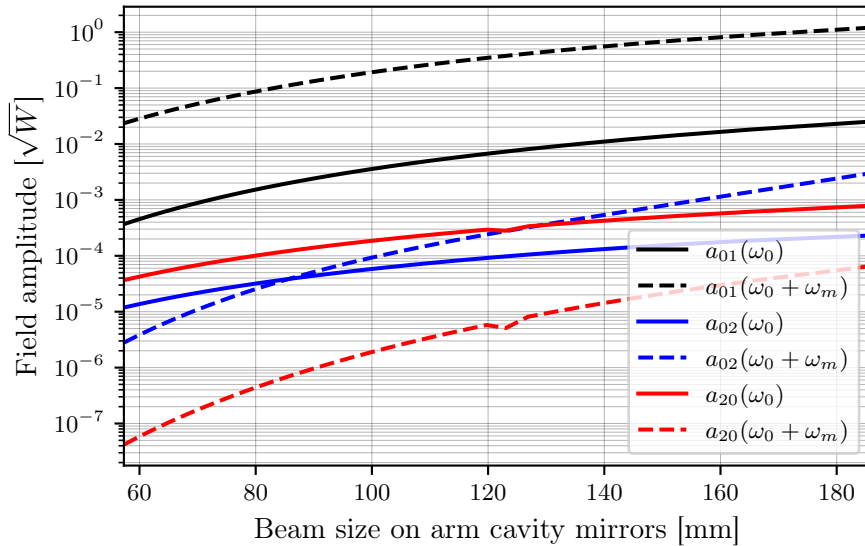


Figure 5.9: Plotted here are the amplitudes of various HOM fields, on transmission through the SRC, as detected by the FINESSE simulation used to obtain the results in Figure 5.8. The product of the black traces gives the coupling shown in that figure. The red traces show the amplitude in the HG_{20} mode, for both carrier and sideband fields, which are a result of mode mismatches in the interferometer; between the arm and recycling cavities (including inter-dependent mismatches) as noted in the text. The blue traces are more complicated as these comprise both mode mismatch and higher order misalignment contributions.

The combinations of the above effects are demonstrated in Figures 5.9 and 5.10. The former provides evidence for the couplings discussed in this section through the contributions seen in the red and blue traces. The red data show field amplitudes for the HG_{20} mode, at both carrier and sideband frequencies. These field amplitudes can only be present due to mode

mismatches in the interferometer; as there are no yaw misalignments applied to any of the optics. Meanwhile, the blue curves give amplitudes in the HG_{02} mode; again for the fields at both carrier and sideband frequencies. The power in these fields is more complicated as they will consist of contributions from both higher order pitch misalignment terms as well as the aforementioned mode mismatches (due to mode mismatches scattering into even order modes, as detailed in §2.3.3.1). Figure 5.10 gives a detailed break-down of the mode mismatches, during the simulation, between each core cavity within the modelled aLIGO file. Combinations of these mismatches lead to scattering into even order modes which result in the red traces of Figure 5.9, as well as contributions to the blue traces on that figure. All of the shown mismatch values are relatively small, reaching a maximum of only $\sim 3 \times 10^{-5} \%$ at the largest spot size, such that the degradative effect that this has on the simulation results is not an overly limiting factor. However, the net result of all the above effects leads to an observed numerical scaling in the alignment to longitudinal coupling which does not quite obey a power law of w^6 .




Figure 5.10: The mode mismatches between all the core cavities in the aLIGO file, as a function of the spot size on the arm cavity mirrors. The y -axis values were computed via unity minus Equation (2.18), averaged over the tangential and sagittal planes, as a percentage. These traces represent the mismatches *after* employing the optimisation routine mentioned in §5.3.3. The largest mismatches here are between the split SRCs in the X and Y arms (magenta), the split PRC and SRC in the Y and X arms, respectively, (cyan), and the XARM cavity and the SRC in this arm (dashed blue). Essentially all of these mismatches increase with the spot size on the arm cavity mirrors, however this is likely just a side-effect of the aforementioned optimisation routine.

5.4 Implications and potential for suppression

Whilst the approximate w^6 scaling of alignment to longitudinal coupling may seem alarming at first, there are some subtleties and nuances in this scaling relation which can serve to lessen its impact. For convenience, we will again show the key equations which constitute this scaling relationship. Firstly, the magnitude of the coupling coefficient from HG₀₀ to HG₀₁ which is given by

$$|k_{01}| \approx \frac{\pi}{\lambda} \gamma w, \quad (5.27)$$

where λ is the beam wavelength, γ is the misalignment angle between the incident beam and the cavity axis and w is the size of the beam as before. Secondly, the resonance enhancement factor for the first order mode is the other contributing factor:

$$|d|_{HG_{01}} \approx \frac{1}{2} \frac{\pi}{L\lambda} w^2, \quad (5.28)$$

where L is the length of the cavity. It is also useful to re-state the full scaling relationship derived from these previous quantities, we will now denote this as $P_{\text{refl},01}$ in accordance with the nomenclature used in Appendix C:

$$P_{\text{refl},01} \approx \frac{R^2}{2} \left(\frac{\pi}{\lambda}\right)^4 \frac{w^6}{L^2} |\beta_{\text{DC}} \beta_{\text{AC}}| |i_0|^2, \quad (5.29)$$

where R is the reflectivity of both cavity mirrors, the β factors are the tilt angles (in the pitch degree of freedom) of the mirrors and i_0 is the amplitude of the input field in the HG₀₀ mode. We will now consider each of these equations individually.

5.4.1 Stronger alignment signals

First and foremost, Equation (5.27) shows us a linear scaling of k_{01} with the beam size. This, in turn, causes stronger coupling into first order HG modes for larger beams, which will also increase the power in the alignment signals [143] used in alignment control schemes [26, 84]. The alignment sensing in GW detectors uses split photodiodes for detecting beats between different mode orders [149]. Consequently, in the context of this coupling, such photodiodes would detect the beat of the HG₀₀ and HG₀₁ modes (at both DC and AC frequencies). From the derivations in Appendix C we can determine that this detected beat is effectively the square root of Equation (5.29), as we only have one contribution of the k_{01} coefficient and the resonance enhancement factor $d_{HG_{01}}$ for such a signal. We will denote this quantity as

$\bar{P}_{\text{refl},01}$ which is thus defined as

$$\bar{P}_{\text{refl},01} \approx \frac{R}{\sqrt{2}} \left(\frac{\pi}{\lambda}\right)^2 \frac{w^3}{L} |\beta| |i_0|, \quad (5.30)$$

where β can be either β_{DC} or β_{AC} depending upon which frequency we are looking at. As this beating quantity is used in the alignment scheme, this effectively allows for a reduction in the residual DC misalignment by this value — essentially suppressing the “static” misalignment term. Note that this assumes a well-behaved control loop whose performance is limited by the sensor noise. Commensurately, we can argue that the “real” scaling of the alignment to longitudinal noise coupling is effectively equal to Equation (5.30); i.e. the corresponding and equal left-over term of Equation (5.29) that is then *not* suppressed by the alignment scheme. This modified scaling relation then only has a w^3 dependence.

As a result of this, it is interesting then to express an order of magnitude version of Equation (5.30) as a function of the parameters which are expected (or known) to change for future GW detectors:

$$\bar{P}_{\text{refl},01} \sim \frac{1}{\lambda^2} \frac{1}{L} w^3. \quad (5.31)$$

Using this equation, and the projected values for each of these variables for third generation detectors (from Table 2, Section LT-3.4.2, of [27]), we can estimate the approximate scaling of the alignment noise coupling in each case. This is shown in Table 5.1, where values of Equation (5.31) are given for the planned third generation detectors, normalised to the value for aLIGO (using approximate values $\lambda \sim 1 \mu\text{m}$, $L \sim 4 \text{ km}$ and $w \sim 5 \text{ cm}$). The values in this table provide a glimpse into how the alignment noise may scale for these future detectors. When inspecting the ET detector values, it is important to note that ET-LF is the more important case as the alignment control noise dominates at lower frequencies; the scaling for ET-HF is less important, but still needs to be considered. The scaling ratios of 1.04 for ET-LF and 1.38 for CE (pess) indicate that the actual alignment noise scaling is far less concerning than the initial w^6 relation. In addition, for CE the absolute scaling is estimated to actually decrease relative to aLIGO — due to the longer wavelength coupled with the much longer arm cavity, whilst the beam size increase is relatively small by comparison. The only detector for which this absolute alignment noise scaling increase is very significant is ET-HF, where this large ratio arises due to the projected beam size being relatively large (compared to aLIGO) in comparison to the arm length.

It is important to remember that the alignment noise coupling ratios given here are just estimates based on simplified analytics. A more precise prediction of how this noise may scale for future detectors will require thorough simulation work using established model files for each detector.

Table 5.1: This table gives approximations of the scaling of the alignment to longitudinal noise coupling for future detectors, determined via Equation (5.31), as ratios to the estimated aLIGO value. Note that these are order of magnitude estimates and do not represent precise projections of the alignment noise scaling for the listed detectors. CE (pess) refers to the “pessimistic” initial design scenario for Cosmic Explorer [27].

	aLIGO	ET-LF	ET-HF	CE (pess)	CE
Beam Wavelength [μm]	1	1.5	1	1	1.5
Arm Length [km]	4	10	10	40	40
Beam Size (on mirrors) [cm]	5	9	12	12	14
Normalised $\bar{P}_{\text{refl},01}$	1	1.04	5.53	1.38	0.98

5.4.2 Scaling with the other dependent variables

Another important point comes from the assumption made in §5.2.2, namely that the cavity is of very high finesse such that (in the case of an impedance matched cavity) $R^2 \sim 1$. By lowering the finesse of the cavity sufficiently, this approximation is no longer valid and thus the resonance enhancement factor for the first order mode does not then scale as w^2 . The scaling is more complicated when removing this assumption, however an idea of the effect can be seen from Equation (5.29) whereby reducing the value of R (hence lowering the cavity finesse) decreases $P_{\text{refl},01}$, commensurately.

We can continue to look at Equation (5.29) for hints at how the various dependent parameters by themselves provide different scaling relationships for the alignment coupling. Firstly, we can note the nominal $1/\lambda^4$ dependence in this equation. By substituting in the beam size from Equation (5.26), we obtain

$$P_{\text{refl},01} \approx \frac{R^2}{2} \frac{1}{L} \left(\frac{\pi}{\lambda}\right)^3 \left(\frac{1}{1-g_c}\right)^{\frac{3}{2}} |\beta_{\text{DC}} \beta_{\text{AC}}| |i_0|^2, \quad (5.32)$$

where $g_c = (1 - L/R_C)^2$ is the g -factor of the cavity. In other words, the scaling we obtain with the wavelength of the beam is $P_{\text{refl},01} \sim 1/\lambda^3$ — and, thus, using a longer wavelength beam would act to suppress this coupling. Furthermore, Equation (5.32) provides us with an approximate scaling of this alignment coupling against the length of the cavity L . If we assume that the stability factor, g_c , of the cavity remains constant for a given design then one can note that this scaling then has a dependence of $P_{\text{refl},01} \sim 1/L$. Thus, by increasing L (and the radius of curvature of the cavity mirrors by the same amount to keep g_c fixed) one could theoretically suppress this alignment scaling. It is important to note however, that both of the choices of L and λ are fundamental to the design of a ground-based GW detector as the values of these parameters have consequences across many aspects of the interferometer. These include, for example, considerations for the optic coatings in the case of the beam wavelength, and Sidles-Sigg instabilities [139] for the cavity length — as

increasing this whilst keeping g_c fixed would lead to a larger waist size of the arm cavity mode.

More work remains to be done on this topic of alignment noise scaling for future detectors. In particular, the impact these types of scaling relations have on alignment control schemes will be important to understand. This is due to the aforementioned issue of the alignment control system itself being a dominant noise source at low frequencies in current detectors; rather than the noise couplings due to misalignments themselves.

5.5 Summary

Alignment noise is currently an important source of noise to understand in GW detectors. This statement will hold equally true for future detectors, such as ET-HF, where higher power could potentially result in stronger coupling to pitch and yaw motions on the arm cavity test masses. Given that future detectors will also typically utilise larger beams to reduce coating Brownian thermal noise [27], it is important to understand how the alignment noise can scale with the beam size. We have investigated such an issue in this chapter, including my work, extended from [143], on characterising the analytically based approximation for alignment to longitudinal coupling (§5.2), followed by a numerical validation of the ensuing w^6 scaling relation in §5.3. This modelling work was carried out by myself, using FINESSE to obtain the results highlighted in this chapter.

We saw in §5.3.4 that the results of the alignment coupling simulation broadly agree with the approximate analytical scaling derived. The slight deviations in the simulation data were discussed in this section. These deviations can be briefly summarised as contributions of higher order misalignment terms, and HG mode couplings, being assumed to be negligible in the analytical approximation — whilst in the simulation (and in practice) this is not quite true.

In §5.4 we discussed the implications of this scaling. An important point here is that the increased coupling into first order modes for larger beams also provides the potential for stronger alignment signals, yielding a benefit to future alignment control systems. This section also highlights that this scaling can be suppressed through careful tuning of the arm cavity design; for example, lowering the cavity finesse would assist in dampening the effect of this increased beam size scaling.

There is scope for future work on this specific aspect of alignment coupling. In particular, a concrete recipe for tuning the cavity design to optimally suppress the increased beam size scaling would be an interesting direction of development. This could be done by carefully selecting values of the scaling dependent parameters, L (cavity length), λ (beam wavelength) and R (reflectivities of the mirrors), to potentially obtain a scaling relation of reduced impact.

Chapter Six

FINESSE 3 Software Design

The use and development of FINESSE 3 has been a common theme throughout this thesis. In Chapter 2 we covered the motivation behind developing this software for modelling HOM effects in interferometers, including examples of the improvements built upon previous versions in this regard. We have also seen use cases for FINESSE 3 involving specific science tasks in Chapter 3, for a feasibility study on a novel telescope design for ET, and Chapter 4, in which the modelling of adaptive mode matching for ET is investigated in the context of this telescope design. The results behind several mode matching related figures in Chapter 5 also utilised FINESSE 3. In this chapter we will focus on the overarching software design of FINESSE 3 as a whole, with specific attention paid to the motivation behind the design choices therein.

The FINESSE software package was originally conceptualised and implemented in an initial form in 1997 [150], and has been used extensively in the GW detector instrumentation community ever since [151, 152]. It remains one of the major software suites in the field to this day. The new version, FINESSE 3, which is the topic of this chapter, is a complete re-development of both the original software, and its eventual wrapper and utility code Pykat [74], from the ground up. Development on FINESSE 3 began around three years ago (in late 2017) and has continued with a core team of around five developers (see the GitLab page which hosts the project [75] and the documentation [81]); a number which is expected to grow in the coming years. Over the course of my PhD, I became lead developer on the following aspects of the software: the extension from simple plane waves to the use of optics and beams with transverse spatial properties, integration with Cython and the documentation. My implementation of the HOM modelling capabilities of FINESSE 3 are given in detail in Chapter 2, however in summary these include, but are not limited to, a new beam tracer (providing, for example, convenient access to beam parameters and mode mismatches), efficient computation of coupling coefficients for modal calculations and improved performance of transverse beam profile imaging. The integration with Cython is vital for performing efficient calculations, improving the performance of the code significantly over a pure Python implementation; and even resulting in faster simulations when compared

to the previous C implementation of FINESSE 2. The use of Python as the main language for FINESSE 3 initially gave cause for concern, due to the potential performance limitations, at the beginning of the development of this package. However, our work on “cythonising” the key back-end code (and, in particular, my contributions in the form of efficient HOM implementations) confirmed this package structure as a sensible choice. See §6.2 for more details on the Cython extensions. Finally, the documentation is implemented using state-of-the-art development tools — providing a useful, and modern, set of pages complete with an automated API reference, details on the physics behind the code and examples on using it. This documentation will allow users to build complicated models more easily and also provides a platform for more people to contribute to the project in the future. My overall contributions to FINESSE 3 in general have been extensive, covering almost all aspects of the code and documentation¹ — where necessary, I will highlight these specific contributions.

We will begin this chapter by discussing one of the key design choices for FINESSE 3, namely the use of Python as the high-level package and interface to the software, in §6.1.1. Tightly coupled to this is the use of Cython [154] for performing the necessary low-level calculations, including the solving of the interferometer matrix. This separation of concerns will be explored briefly in §6.2. As a follow-up to the use of Cython, we will discuss a core concept behind FINESSE 3 which is the use of a “workspace” pattern — for optimising simulations as well as providing a convenient interface for users to extend the code for their own requirements. Combining these topics together, we will then give an overview of the “flow” of the software in terms of running a simulation, from making the model to accessing the results, in §6.4 — highlighting the steps taken, with links back to Chapter 2 where appropriate. Finally, in §6.5, the documentation [81] of FINESSE 3 will be briefly discussed; this a topic in which my contributions were particularly important, being the lead developer of this aspect of the software.

6.1 Motivation

Several motivating reasons lie behind the development of a new and improved version of the popular and widely used interferometer modelling tool FINESSE. The key factors of these will be discussed in this section. The motivation behind improving the modal aspects of FINESSE was covered extensively already in §2.1, and so is not repeated here.

One of the major points behind the development of FINESSE 3 is a purely structural consideration, whereby each element of the code-base is separated into its own sub-module based on the type of physics tasks it is required to perform. This creates a clean package which is easy to work with, as highlighted in §6.1.1, and, just as importantly, is flexible for users and developers to extend to suit their own requirements necessary for completing some science task; this will be discussed briefly in §6.3.

¹As of writing, over 900 commits to the FINESSE 3 repository have been authored by myself [153].

The use of a modern code-base was, and is, a key point in the implementation of FINESSE 3. In addition to the modular structure of the source code itself, we now also have a dedicated test-suite (utilising the `pytest` [155] framework to write and run all of our tests) and a set of documentation pages included in the same project (see the host page at [75]). Both of these aspects of the software are factored into the development cycle via the continuous integration capabilities of GitLab, whereby all the tests are executed and documentation is built on a per-commit basis to the main repository. These development tools are consistent with any major, modern, open-source software package and should help to future-proof FINESSE 3 for the next couple of decades. This is important from a scientific perspective due to the robustness that the continuous testing and documenting will provide to users relying on the software; especially in the case of detector commissioning tasks which often require quick response for modelling calls.

More technical reasons for developing FINESSE 3 include [156] (but are not limited to): the ability to model closed-loop control systems with electrical couplings, the inclusion of mechanical components with explicit motional degrees-of-freedom (e.g. pitch and yaw for suspension systems) and the implementation of beating between distinct carrier light fields; required for modelling proposed schemes for limiting quantum noise (such as the unstable opto-mechanical filter [157], and EPR squeezing [158]). All of these additional features and improvements will assist in modelling efforts for the design and commissioning of third-generation GW detectors going forward. It is worth noting also, as mentioned in Chapter 2, that the improved beam tracing capabilities of FINESSE 3 have already been used for science tasks in the form of arm telescope feasibility studies; as detailed in Chapter 3 and, by extension, Chapter 4.

6.1.1 Modular structure and API

The use of Python as the packaging and API language naturally leads to the modular structure of the FINESSE 3 code-base. As is typical with a well-structured Python package, FINESSE 3 consists of several sub-modules and packages, each of which is targeted towards a specific aspect of the code — both in terms of the physics involved (e.g. the `finesse.knm` sub-package contains code for coupling coefficient calculations) and utility tools (e.g. `finesse.plotting` is made up of classes and functions for conveniently plotting results from simulations). This clean separation of concerns is advantageous in several regards. Firstly, it reduces development time due to being inherently more understandable as a code-base. This applies to both the core developers of FINESSE 3 and users who want to extend the code for their modelling requirements; see §6.3. Secondly, it improves the testing framework by allowing for functional and unit tests on separate parts of the code. Finally, this modular structure lends itself well to the documentation side of the project — again, making it clearer as to where each feature lives in the code-base and providing the means to document the equations therein in a clean way. See §6.5 for details on the documentation.

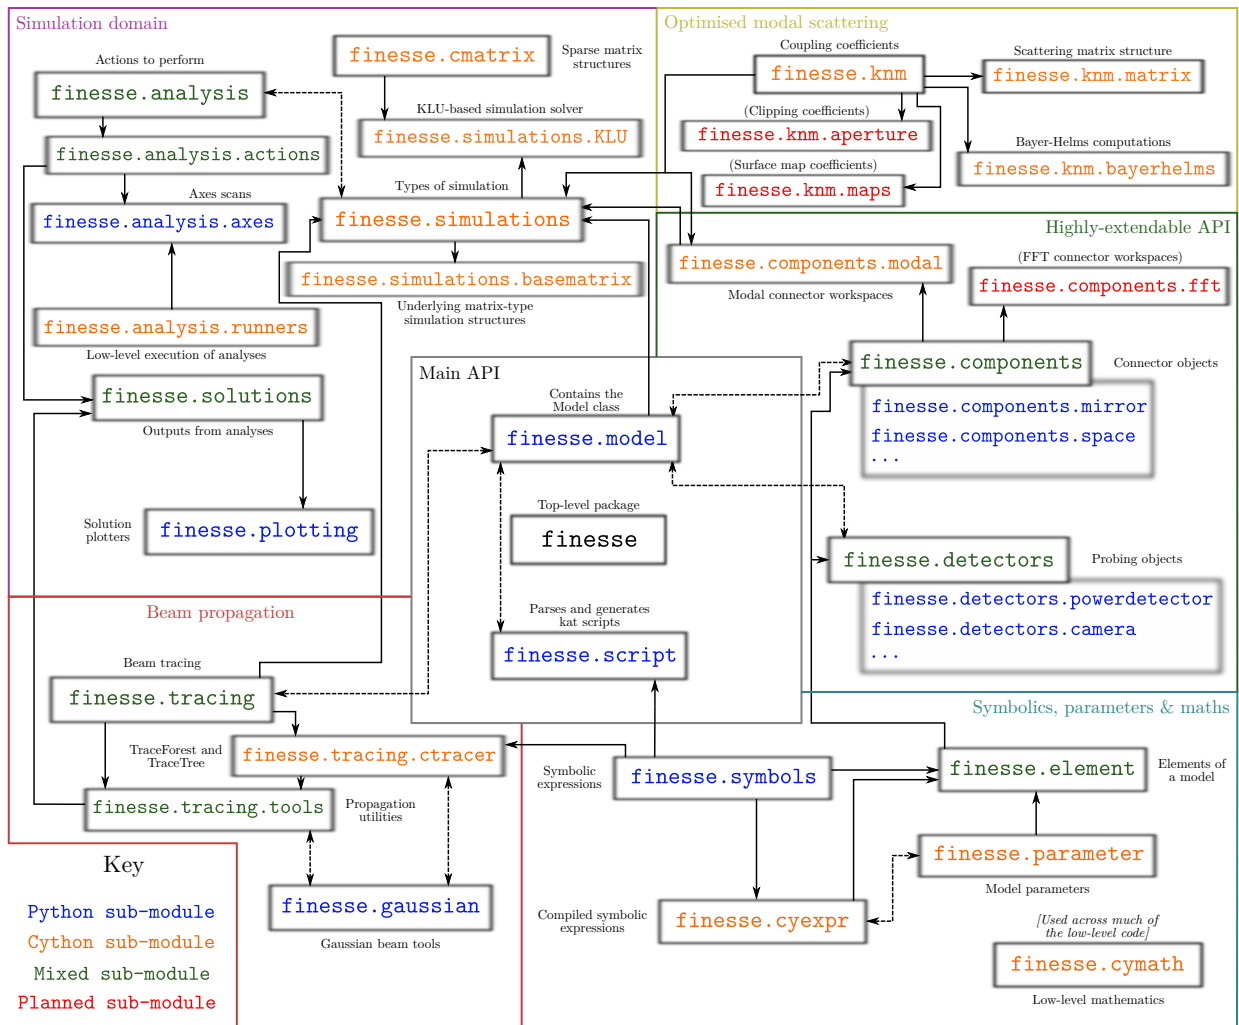


Figure 6.1: A reference chart showing a (partial) network representation of the modular structure of FINESSE 3; based on the current status of the code-base. Common sub-modules and packages have been grouped together in the boxes of different colours, predicated on the “domain” in which these parts of the code operate. For example, the “Simulation domain” shown in the upper-left includes the analyses, matrix solvers, solution objects and plotting modules, as these all pertain to a typical simulation process in FINESSE 3. One may note that the domain containing component and detector code has been marked as “highly-extendable”. We will touch upon this concept in §6.3.

Figure 6.1 gives a network representation of this modular structure employed by FINESSE 3. The information contained in this figure is relatively densely encoded, however the important point is to realise the separation of parts of the code-base according to the “domain” in which they operate. One key point to note is that the “model” and “script” modules are core to the package as a whole. These represent the object used to store the configuration and the process by which `kat`-script files are parsed into these configuration models, respectively. Each sub-module shown on Figure 6.1 is colour-coded in terms of the type of module it is

— i.e. Python, Cython or a mix of the two. The red sub-modules are those which have not yet been implemented, but are planned for the future. The solid arrows represent one-way communication between any two modules (e.g. the simulation code relies on the beam tracing module, but not vice-versa). On the other hand, dashed lines mean two-way communication between the sub-modules. Note that, in the interests of reducing complexity, Figure 6.1 only displays modules associated with the optical to optical couplings in FINESSE 3. Other sub-modules which deal with mechanical and electric couplings also exist but have not been included in this figure.

6.2 Cython extensions

Cython [154] is an “intermediate-level” extension to the Python programming language, providing the ability to write Python-esque code, with extensive optional optimisations, in modules which are then compiled into C code; typically in the form of shared object files. The codes in these files are then exposed to Python, and other Cython modules, in a well-defined way; as detailed in a concise manner in [159]².

This ability to write optimised code with hooks into the Python application programming interface (API) is ideal for FINESSE 3. It enables the code-base to remain consistent with the modular structure detailed in §6.1.1 whilst facilitating the use of highly-streamlined C code. By making extensive use of explicit typing, contiguous memory-views (and pointers to C arrays) and fast C function calls (including in-lined calls), we are able to achieve C level speed for our simulation code. We covered some examples of the impressive performance of FINESSE 3, in the context of modal simulations, when compared to FINESSE 2 (a pure C code-base) in Chapter 2 — see specifically, §2.3.2.3, §2.3.3 and §2.4.5.

The reason *why* FINESSE 3 requires such well-optimised code lies in the complexity of the calculations performed and, inexorably coupled to this, the “tight” loops over quantities such as the number of HOMs. We saw a glimpse of the complexity of HOM calculations specifically in §2.3.1, where the computation of coupling coefficients relies on an analytic solution, derived by Bayer-Helms [109], which contains three-fold sums with sizes proportional to the number of modes. Furthermore, the solving of the sparse interferometer matrix itself requires interaction with a C-based optimised solver, KLU [79], for performance reasons. This is due to the number of equations N_{eq} , i.e. the order of the interferometer matrix, scaling as

$$N_{\text{eq}} = N_{\text{nodes}} \times N_{\text{freqs}} \times N_{\text{HOMs}}, \quad (6.1)$$

which is simply the product of the number of nodes, frequencies and HOMs in the system. An expression for the number of HOMs, as a function of the maximum modelled mode order,

²This source also provides a good overview of the different ways of declaring functions in Cython, including benchmarking results using a typical numerical example.

can be found in Equation (2.32); note, however, that this expression does not take into account explicit mode selection as detailed in §2.3.3. For a typical GW detector file, we have $N_{\text{nodes}} \sim 400$, $N_{\text{freqs}} \sim 5$ (at DC, assuming carrier plus RF sidebands) — giving $N_{\text{eq}} \sim 2000$ even before including any HOMs³. Such a sparse matrix requires heavily optimised code for both filling and solving.

One especially useful feature of Cython extensions is that they essentially operate from the same Python code as any pure Python module — i.e. one can write a naive extension in exactly the same syntax as a Python file. This is, and will continue to be, a particularly important aspect of developing FINESSE 3, as one can easily convert any Python code to a Cython extension if a bottleneck in this piece of code is identified. Indeed, even by simply using the same code (which is then compiled via Cython) one can, potentially, obtain significant speed-ups (see [160] for example) even before using typed variables and other low-level features. This provides the ability to develop in a pattern of “scaled-optimisations” — avoiding any premature optimisations whilst continually drilling down into the Cython extensions with lower and lower level code when necessary (after profiling flags up bottlenecks in such code).

6.3 Simulation workspace pattern

A central part of the FINESSE 3 pipeline is the “workspace” pattern used for optimising simulations and providing a clean separation between the Python API and C back-end. When a simulation is triggered, each component (also referred to as “connectors”) in the model will provide a low-level version of itself known as its workspace. These workspaces then store the component connection information, required for assigning the correct sub-matrix view of the full interferometer matrix, as well as other data — such as the C level versions of each parameter of that connector and, if the simulation is modal, scattering matrices (see §2.3) for each coupling associated with the component connections. A visual representation of this can be found in Figure 6.2.

In this figure a broad overview of the sequence of events for workspace initialisation, as related to a simulation, is given. The process shown here applies equally to workspaces for detectors in FINESSE 3, albeit much simpler as these typically only need low-level output functions to be assigned to them. One can observe from Figure 6.2 that workspaces have a lifetime limited to the lifetime of the associated simulation — control is essentially handed back to the connector object once the simulation is complete, allowing its parameters to be changed via the Python API once again.

An important cornerstone of the workspace pattern is its extensibility through custom connector creation. This means that any user wanting to extend FINESSE 3 with an extra

³Here we have used the aLIGO design kat-script file (see, for example, Appendix E of [80]) where the optical nodes were counted using the network model property [103] in FINESSE 3.

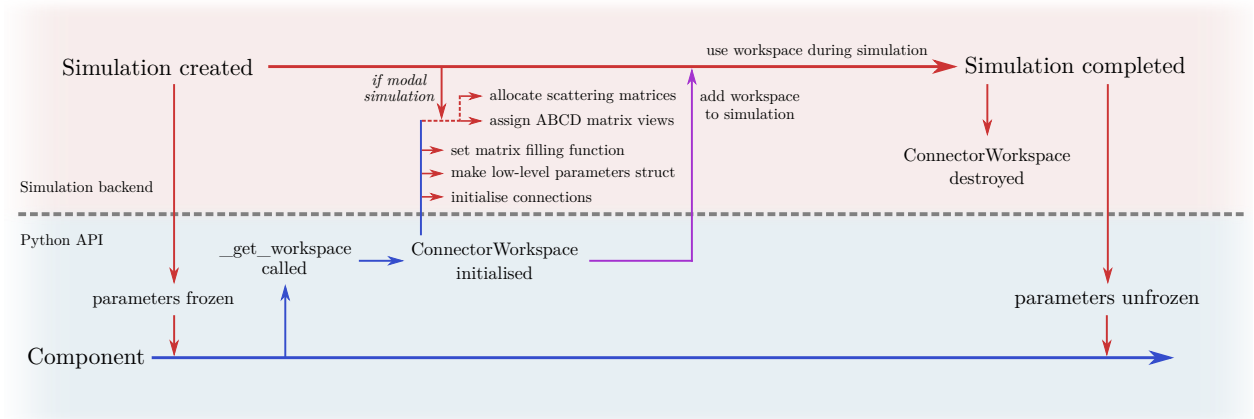


Figure 6.2: A representation of the workspace pattern for a typical component of a FINESSE 3 model and simulation. This highlights the process by which components are “converted” into optimised workspaces when a simulation is created, providing a separation of the Python facing connector object and intermediate-level Cython code associated with this connector. Note that this is quite a broad overview, there are other steps (such as determining which sub-matrices are changing, which scattering matrices need to be re-computed, etc.) taken in addition to those shown here.

component can do so by simply defining a `_get_workspace` method in this new object, as implied by Figure 6.2, in which the necessary attributes of a connector workspace are set. All of the lower-level handling is then carried out automatically, allowing the full power of the Cython extensions to be utilised for any such custom component. The process by which a user would implement a custom component is shown conceptually in Figure 6.3, from their initial idea through to the realisation of this via interfacing with the FINESSE 3 API. In this figure the user can code their new component entirely in Python, without needing to write any Cython code nor worry about the underlying Cython extensions. This is possible via various interface methods to the low-level code which initialise quantities such as scattering matrices and tell the connector workspace to use specific matrix filling methods; as highlighted in the steps 4 and 5 in the “implementation” box of this figure.

6.4 Program flow

Now that we have given an overview of the structure and some of the inner-workings of FINESSE 3, it makes sense to summarise these concepts through a logical flow-chart of the procedure undertaken by the program. Such a chart is displayed in Figure 6.4. In this figure the run-time “state” of the program is separated into four distinct categories — the *Model*, the *Simulation*, the *Output* and *Error-Handling*. The *Model* region describes all the interactions and processes followed within the main interface through which configurations are constructed. It is from this object that the *Simulation* is typically first created. As

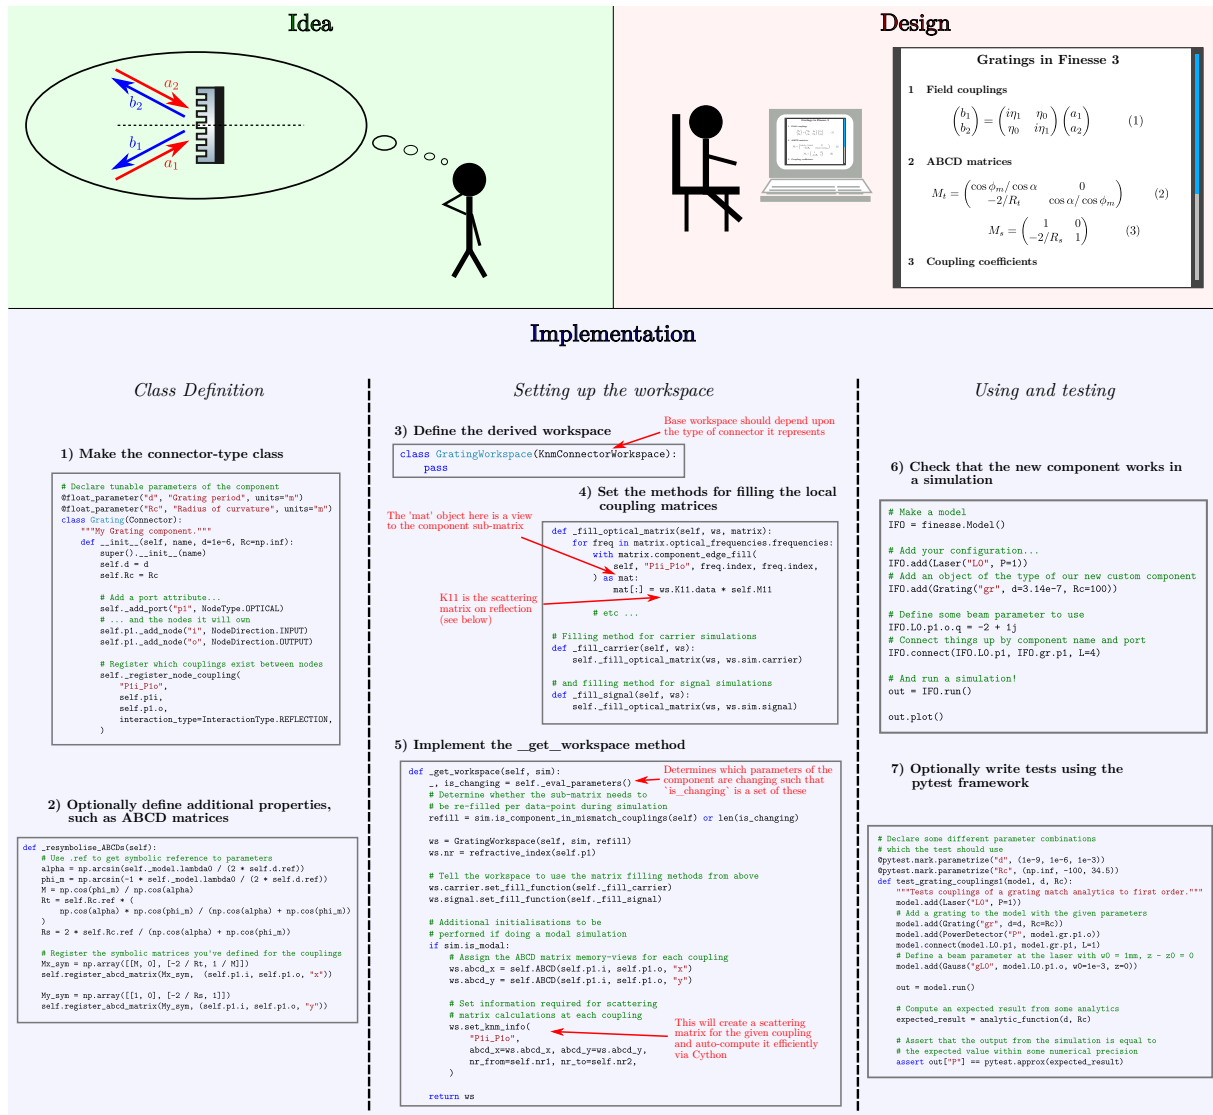


Figure 6.3: An overview of the idea to the implementation process for creating a custom component in FINESSE 3. This highlights the steps necessary for writing a custom connector type object, which can then be used in conjunction with simulations. Note that at no point during the implementation phase is a user required to write or touch any Cython code; various interfacing functions are available through the API to initialise, in a very general way, the necessary variables for executing the low-level code required for running simulations.

described earlier, in terms of the component workspaces (see §6.3), this realm is the low-level, heavily-optimised space in which the interferometer matrices are filled, scattering matrices are computed and beam traces are performed, as and when required over the axis scan(s) of the analysis being executed. When the analysis is completed, the simulation is destroyed (de-allocating the potentially memory-expensive attributes such as the interferometer matrix itself) and a solution object is handed to the user — as shown in the *Output* zone of Figure 6.4. This is where the user may interact with their solved system by accessing specific detector outputs, plotting the solution etc. Concurrent to all these domains is the *Error-Handling* section. This, as the name implies, handles and raises any exceptions which may be raised during the whole file-running process. The form in which errors are displayed to the user is dependent upon the environment in which they are running their FINESSE model.

In the *Simulation* domain there are a few distinct modes of operation — matrix-solver simulations (with the optional signal matrix solving) and so-called “matrix-less” simulations. Both of these derive from the main thread shown in this domain, where common tasks (such as initialising the workspaces) are performed. The latter type of simulation is typically for models containing detectors which only require knowledge of the modal basis, rather than needing any field amplitudes. We briefly covered such a type of simulation, and the benefits of separating it out, in this context in §2.2.4.2. It is important to note that the existing subtle separation of simulation types will become more distinguished in the future, stemming from the plans to implement such simulation types as, e.g., time-domain and FFT solvers.

6.5 Documentation

FINESSE 3 uses the Sphinx Python documentation generator [161] as the framework and tool for writing and producing comprehensive documentation — in both HTML (see [81] for the latest online version) and L^AT_EX formats. Both of these are produced automatically as part of the LIGO GitLab repository pipeline on the master branch, resulting in online documentation which is always consistent with the latest development version. I was the principal developer of the documentation — most features mentioned in this section were implemented by myself.

The documentation is separated into different sections — introduction (including installation instructions), the manual (containing detailed physics behind the features used in FINESSE 3), using FINESSE (examples of using the software for different tasks), the syntax reference (to refer to the kat-script syntax), API documentation (covering all accessible classes and functions exposed by the API), a developer guide (for information on how to contribute to the software) and other miscellaneous sections such as a bibliography and contributions list. Note that, as of writing, FINESSE 3 is in an alpha state of development and so the above sections are in no way complete at this stage. The framework for writing more content for the documentation is, however, in a stable state. This will make it much easier for current, and future, developers to contribute to the pages.

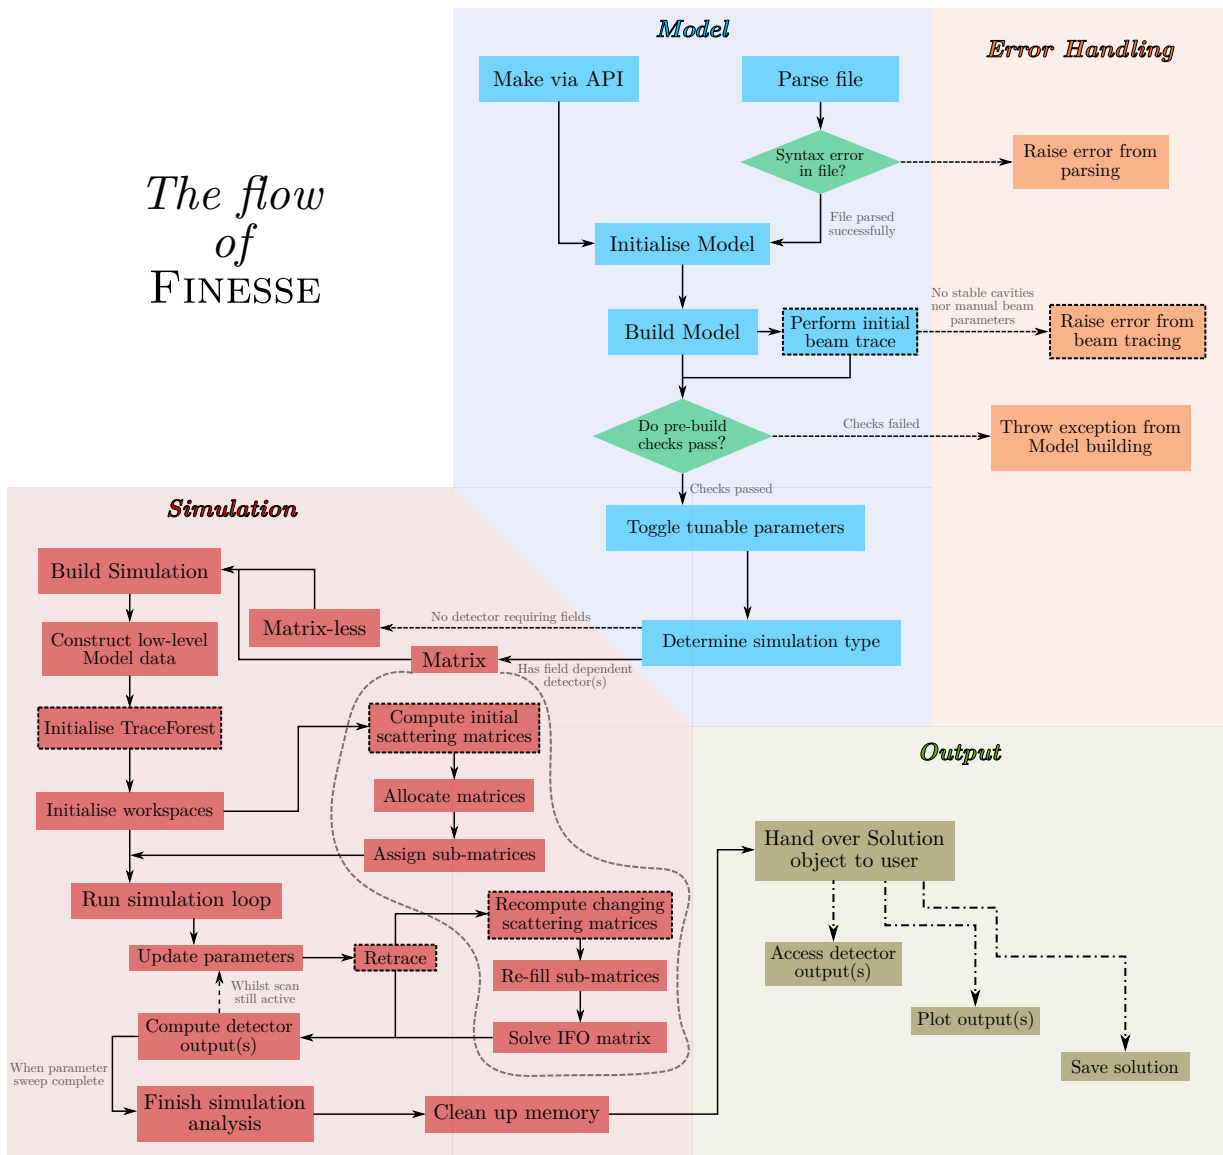


Figure 6.4: This figure displays the flow of the typical pipeline of a FINESSE modelling task — from parsing an input file to handling the output data. Each shaded region describes a distinct state of the program being executed. The *Model*, *Output* and *Error-Handling* realms serve as Python API points; i.e. the user can easily interact with these parts through standard Python object-oriented interfaces. Meanwhile, the *Simulation* zone is at a much lower-level for optimisation purposes; as described in §6.2 and expanded upon in §6.3. The boxes with dotted boundaries represent steps required only for modal simulations. Note that, in the interests of reducing figure complexity, steps associated with signal simulations are not shown here. These steps are similar to those shown in the matrix simulation area already, but acting on a separate “signal” matrix.

6.5.1 API documentation

The API pages of the documentation are generated automatically via an external tool I developed called `reslate` [162]. This tool ensures that all of the API pages are consistent with the current state of the source code. In this section of the documentation, the accessible parts of the FINESSE 3 API are listed and documented using the “doc-strings” of each of these exposed modules, classes and functions. One can use Sphinx to combine these native doc-strings with externally produced images — providing the potential for detailed API documentation pages. One such example of this is shown in Figure 6.5, where a screenshot of the current documentation page for the `Mirror.ABCD` method is displayed. In this figure one can see that we have included graphics to help guide the reader on the conventions used for this part of the code — a feature that is extendable to the doc-strings of any other object in the source code.

Mirror.ABCD(*from_node, to_node, direction='x', symbolic=False, copy=True, retboth=False*) [\[source\]](#)

Returns the ABCD matrix of the mirror for the specified coupling.

In both cases below, the sign of the radius is defined such that R_c is negative if the centre of the sphere is located in the direction of propagation.

Transmission

For transmission this is given by,

$$M_t = \begin{pmatrix} 1 & 0 \\ \frac{n_2 - n_1}{R_c} & 1 \end{pmatrix},$$

where n_2 and n_1 are the indices of refraction of the spaces connected to the mirror and R_c is the radius of curvature of the mirror.

The matrix for transmission in the opposite direction of propagation is identical.

Reflection

In the case of reflection the matrix is,

$$M_r = \begin{pmatrix} 1 & 0 \\ -\frac{2n_1}{R_c} & 1 \end{pmatrix}.$$

Reflection at the back surface can be described by the same type of matrix by setting the C element to $C = 2n_2/R_c$.

See `connector.ABCD()` for descriptions of parameters, return values and possible exceptions.

[◀ Previous](#) [Next ▶](#)

Figure 6.5: This figure shows an example from the API documentation, displaying the page for the ABCD method of the mirror class. Sphinx allows us to use both rendered equations and images in the doc-strings for any object in the source code, providing the ability to produce pages such as this.

As an aside, another aspect of the API pages that we can glean from Figure 6.5 is the presence of the navigation sidebar. Each sub-module is listed in this bar, grouped into the sub-packages in which they sit and ordered alphabetically for convenience. This view, along with the search function, provides easy access to all the API documentation and also provides

an insight into how the code is structured; as implied earlier from Figure 6.1.

6.6 Current status and future work

Considerable progress has already been made on FINESSE 3, with most of the core features having been successfully ported from the previous version of the software. These features have also been improved, particularly in the case of modal aspects of the code as was outlined throughout Chapter 2, and optimised for better performance. A completely new parser for KatScript, implemented by S. Leavey, has been built into the code and thoroughly tested. This new parser contains powerful features including support for symbolic expressions, keyword arguments in commands, better syntax for connecting components together and a clean separation of the configuration versus analyses to perform. As of writing there is currently a suite of over 2000 tests which all pass in the project pipeline; these include core physics validation tests which check a variety of systems, in both plane wave and modal cases, against analytic solutions. The project also contains automatically generated and built documentation pages for the API already, as outlined in §6.5.

With this current state of the code, we can run tasks using full detector files (such as the aLIGO design file). This includes being able to run complex models on the full interferometer (plus input & output optics) containing mode mismatches and misalignments, radiation pressure effects on surfaces and various types of signal injections including (but not limited to) — pitch and yaw motions of optics, GW strain inputs for free-spaces (based on [163]) and quantum noise couplings at open ports. At the end of April 2021, an initial “Alpha 1” version of FINESSE 3 was released. This is currently being tested by a small selection of expert users of previous versions of the software.

The work beyond Alpha 1 will be defined by the feedback obtained during this stage of development. Each subsequent pre-release version is intended for increasingly larger audiences — i.e. the alpha stage will involve expert users, whilst beta will include the GW detector community more broadly. In this way, the future development of FINESSE 3 can take place incrementally based on the feedback of an expanding user base. This cycle of development allows for “crowd-sourcing” of bug fixes (as our current test suite is comprehensive, but still cannot cover all eventualities at this stage due to the scope of the project) and a steady stream of feature requests which can then be implemented leading up to the release version of the package.

6.6.1 Longer-term plans and ideas

Beyond the release-requirements of FINESSE 3, several other features are planned for the longer-term. These include, but are not limited to, a time-domain solver, FFT propagation based simulation and the ability to model closed-loop control systems complete with electrical

couplings. With regards to the first two especially, hints at how these might be hooked into the existing code structure are given via Figure 6.1. In that figure, we previously saw that one such planned sub-module could be `components.fft` — existing alongside the current `components.modal` which provides workspaces and functions for solving the given system via the modal method (covered extensively throughout Chapter 2). This new FFT sub-package would then, correspondingly, provide a different set of workspaces required for executing a simulation through FFT propagation instead. The existing framework is built in such a way as to allow these different types of simulations, with corresponding workspace definitions, to be implemented relatively easily. This idea can be extended to other types of system solvers too, for example a time-domain simulation solver as mentioned earlier.

The extension of the software to provide alternative solvers would, in principle, not require any major changes to the core API of FINESSE 3. As we have covered previously in this chapter, the model object is the main interface for the user when it comes to interacting with their configuration. This object is simply a representation of the provided optical system and so would be identical whether looking at modal simulations, FFT, time-domain or some other simulation based solver. Further to this, the KatScript syntax and associated parser (mentioned earlier in this section) will also provide a consistent interface regardless of the simulation type required. In fact this new parser implementation is flexible and extensible, meaning that new commands (which could include different types of analyses) can be added to the specification in a much easier way than previous versions of FINESSE.

6.7 Summary

In this chapter I have introduced the overarching design of the new version of FINESSE, a software package to which my contributions have been vital (including the HOM modelling capabilities covered in Chapter 2) and broad (consisting of developing the documentation, contributing to most areas of the code and writing several tests within the framework developed principally by S. Leavey, as detailed in the documentation [81]).

The motivation for developing FINESSE 3 in the first place was discussed in §6.1. In summary of this, as well as technical reasons for implementing new features, one of the major points in favour of a new version is the ability to create a modular structure which is easy to work with and provides an interface through which new users can mould the code for their own requirements. This also links heavily into the future work (see §6.6), giving existing developers of FINESSE 3 the tools necessary to implement different simulation types — e.g. a time-domain solver — in the longer term.

The aforementioned modular structure of the code-base is shown in some detail in §6.1.1, with Figure 6.1 in particular displaying a reference network of the current state of the package. Through this figure, we explored the different domains / libraries that the software consists of; highlighting the effort that has been put into keeping logical separation of varying physics

topics associated with the elements of the code.

In §6.2 and §6.3, the requirement for highly-optimised back-end code whilst maintaining a clear and convenient Python API was detailed. The pattern of converting Python objects to low-level workspaces when simulations are triggered is key to this necessity, whilst simultaneously providing the means for new users and / or developers to “hook” into the fast code with custom components as explained via Figure 6.3. Tightly coupled to this, in §6.4 we detailed the procedure by which models in FINESSE 3 are executed — Figure 6.4 displays a flow-chart outlining this process.

Maintaining a consistent and useful set of documentation for such a large software package is vital for its longevity and, thus, in §6.5 a brief overview of the framework for the FINESSE 3 documentation was given. In short, the use of Sphinx [161] and internally developed tools provide us with the means to generate automatically up-to-date documentation of the API. These tools also allow us to craft well-designed pages which detail the recommended uses and applications of FINESSE 3; as well as providing a central location for describing the physics associated with each part of the code.

Finally, in §6.6, we discussed potential future lines of work which are planned for subsequent releases of FINESSE 3. Given the vast size of the field of interferometric modelling, in the contexts of both GW detectors and otherwise, there are many as-yet unexplored avenues of development for a package such as FINESSE 3. I believe that this new version of FINESSE provides an excellent platform from which to launch new features over the coming decade, as and when they are needed, whilst the design and construction of third generation gravitational wave detectors enters a new, and exciting, phase.

Chapter Seven

Summary and Conclusions

Understanding the effects that transverse beam shapes have on interferometric GW detectors is a broad topic, and one of increasing importance as the sensitivity of current and future detectors is increased even further. In this thesis I have explored the impact of larger beams in upcoming facilities, focusing particularly on ET, and proposed an innovative design of the core interferometer as pertaining to modal effects.

The impact of larger beams was discussed in an indirect sense in Chapter 3, where a key motivation for designing new mode matching systems comes from the greater beam sizes which will be used in future detectors for reducing coating thermal noise. As discussed in that chapter, larger beams require either a very large central beam splitter optic or beam expanders in the interferometer arms. We analysed Z-shaped telescopes for ET-LF and ET-HF, providing a nominal initial set of design values which result in a small beam at the beam splitter whilst maintaining stable recycling cavities. In this work it was also demonstrated that this initial design is flexible in terms of changes required by future trade-off studies. The contents of that chapter were also published in my paper *“Feasibility Study of Beam Expanding Telescopes in the Interferometer Arms for the Einstein Telescope”* [1].

A follow up to this work was presented in Chapter 4. This chapter contains an analysis of the possible actuation strategies for adaptive mode matching in ET-HF given the telescope design obtained in the previous chapter. Using this design, we determined that actuation on the RoC of ZM2 could provide the means to correct for small thermal lens distortions, whilst deformation of the ZM1 surface were shown to be relatively ineffectual in comparison. Furthermore, we demonstrated that adaptive displacement shifting of the telescope mirrors would also have an insignificant impact on real-time mode matching. An additional outcome of this chapter was a characterisation of the initial region of phase space that the SRC mode could hold, whilst retaining adaptive mode matching potential of at least 99.9% to the arm cavities. This area estimate was predicated upon actuators being present at the ZM2 and recycling mirror optics; with actuation ranges based on existing technology in aLIGO.

In Chapter 5 a more direct consequence of larger beams in future detectors was discussed, this time in the context of the impact on the alignment to longitudinal coupling. Here I used both simplified analytics and a more complicated aLIGO model to demonstrate that, nominally, the scaling of alignment noise coupling goes as approximately the beam size to the power of six. My modelling work in this chapter also provided an insight into the process of simulating these types of effects for a full detector file, whereby geometric changes in the arm cavity designs required careful mode matching of the recycling cavities to ensure that scattering into even order modes was minimised. As an additional investigation, I highlighted the nuances and caveats in this otherwise-startling scaling relationship. In particular, the final remarks in this chapter note that the larger beams will also result in stronger alignment signals and, thus, a reduction in the residual misalignment of the arm cavity mirrors.

Concurrently running throughout this thesis, I have also highlighted how the computational modelling of HOMs via modal simulations is vital for furthering knowledge in this domain. Implicit in this regard, I have also demonstrated the need for better modal modelling tools going forwards and, correspondingly, shown that my contributions to the new interferometer simulation tool FINESSE 3 are key in providing such tools.

As the design of third-generation GW detectors ramps up over the coming years, one area of increasing and immediate interest will be the formulation of new optical configurations based on different beam geometries (as compared to those in current detectors). The modelling efforts in relation to this topic will necessitate smarter tools for tracing beams — in particular for easily finding mode mismatches, calculating accumulated Gouy phases and plotting beam sizes. A large portion of Chapter 2, see §2.2.3 in particular, was dedicated to highlighting my work in this regard. The new and flexible beam propagation library I have implemented in FINESSE 3 will provide a good interface in the area of testing proposed interferometer designs for geometrically-related beam issues.

The performance of modal simulations is a key concern for modelling transverse spatial effects in GW detectors. The complexity of the calculations required combined with the “volume” of computations necessary for accurately simulating high order effects means that such simulations can take a significant time to execute; often due to requiring high mode orders to converge to a solution. My work on implementing very efficient modal calculations in FINESSE 3 was shown in Chapter 2 — and, specifically for coupling coefficients, in §2.3.2.3. The significant performance enhancement that these coupling computations provides over the previous versions of FINESSE will be of great utility for executing models in a much shorter time span; as well as allowing for larger and higher resolution parameter spaces to be explored in future modelling based research. This is of paramount importance when simulation results are required on short notice, e.g., for detector commissioning tasks.

Given that the development and use of FINESSE 3 was a primary theme throughout this thesis, and indeed throughout my PhD as a whole, I provided a broad overview of the design of this new interferometer modelling tool in Chapter 6. This new version of FINESSE is a complete re-development of the already popular software (used extensively in our field) and I have played a leading role in shaping the form of this tool such that it is “future-proofed”

for many years to come. My contributions to the code include (but are not limited to) development of all of the current modal modelling capabilities, as detailed in Chapter 2, integrating Cython extensions into the software to leverage low-level code performance and implementing the documentation. With the package modernised in this way, FINESSE 3 should provide an excellent tool for the optical modelling of interferometers in general — a task of great importance now that third generation GW detectors are on the horizon.

Appendix A

Beam Tracing with Trees

In this appendix we will explore the more technical aspects of the FINESSE 3 beam tracer in terms of its specific implementation, this should be read in the context of the concepts introduced already in §2.2.3 and §2.2.3.1. In addition to the core logic of the tracing library, we will also explore how this implementation is used for handling complicated simulations — for example, dealing with regions of parameter scans which result in unstable optical cavities.

The `TraceTree` object is central to the beam tracing algorithm implemented in FINESSE 3. This data structure represents a single optical node in the interferometer network, with links then to the other physically connected nodes (i.e. via transmission, reflection and free-space propagation couplings) and the ABCD matrices associated with such couplings. Each tree also stores a reference to the trace dependency object which it relies upon — these can be optical cavities or manually defined beam parameters (called Gauss objects in FINESSE 3 nomenclature). A simplified visual representation of the `TraceTree` structure is shown in Figure A.1. Note that if the node stored at a tree is an *output* node, then only the left sub-tree can exist — as one output node can only ever map to a single input node. This also means that the ABCD matrix to the left sub-tree will always correspond to propagation over a free-space in this case.

Executing a beam trace through a single `TraceTree`, from top to bottom, is a simple task using this data structure. The procedure undertaken is a case of recursing through the tree: getting the beam parameter from the parent tree node, applying the relevant ABCD matrix to this parameter then storing the result at the entry for the next tree node. This occurs until all branches of the tree are exhausted; i.e. when there are no more sub-trees for all recursed paths. A slightly simplified version of this is shown in Listing 2. This block of code represents the algorithm used in FINESSE 3 for recursively propagating a beam through a `TraceTree`. Note that this is implemented in Cython for performance reasons (see §6.2 for details on Cython extensions in FINESSE 3 in general). The simple and concise nature of this algorithm is beneficial in multiple ways. Firstly it offloads the complicated logic involved

with deciding how to piece together all the trace trees into a separate, initial stage. This provides a boost in terms of the performance of the code and allows for easier debugging and testing of the beam tracing. In relation to this, it also ensures that the act of beam tracing itself is robust — i.e. not prone to bugs. The main testing of the beam tracer can then be written in terms of planting of the `TraceForest`, rather than the propagation of the beam itself.

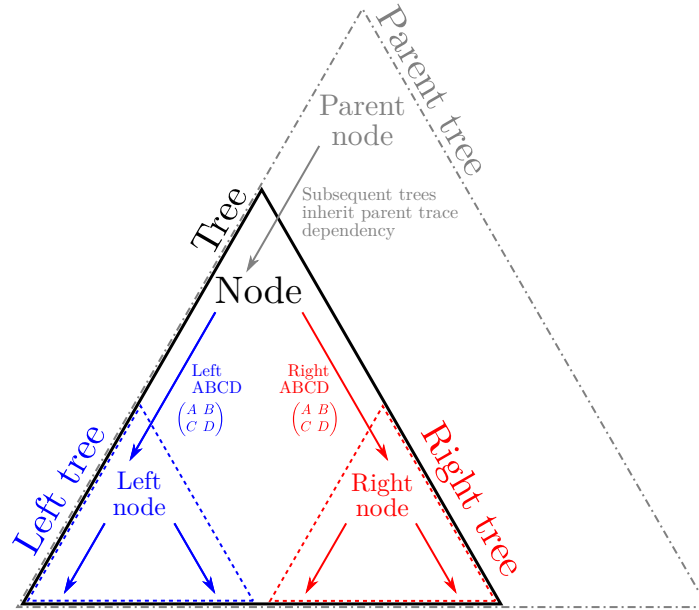


Figure A.1: This figure shows a simplified representation of the `TraceTree` structure implemented in FINESSE 3. This data structure is employed as the core object through which a beam trace from a given trace dependency is executed. Each `TraceTree` instance keeps references to its parent tree, left sub-tree and right sub-tree (if any of these exist) whilst the data it stores are the optical node it represents, a reference to the trace dependency it relies upon and the value of the refractive index of the associated medium. If the left sub-tree exists then it will also keep a view to the ABCD matrix corresponding to this coupling, and similarly for the right sub-tree.

As detailed in §2.2.3.1, the `TraceForest` structure serves as the container for all the different `TraceTree` instances associated with a given optical network. The order in which each tree appears in the forest is given by the priority value of the trace dependency referenced by that tree; a concept which was briefly covered in §2.2.3.5 and, again, in Figure 2.7 using the example of a Michelson interferometer with Fabry-Perot cavities in the arms. We will cover the internal implementation of this in the next section. Applying the algorithm detailed in Listing 2 to a full `TraceForest` is a simple case of iterating over each stored `TraceTree` in order and propagating the beam through each of these trees via this code.

```

1 cdef void propagate(TraceTree tree, dict trace, bint symmetric):
2     """Recursively trace a beam through a TraceTree."""
3     cdef TraceTree ltree = tree.left
4     cdef TraceTree rtree = tree.right
5
6     # Get the beam parameter at the current tree
7     qx1, qy1 = trace[tree.node]
8
9     # If the left tree exists...
10    if ltree is not None:
11        # ... transform q by applying ABCDs for node -> left node coupling
12        qx2 = transform_q(tree.left_abcd_x, qx1, tree.nr, ltree.nr)
13        qy2 = transform_q(tree.left_abcd_y, qy1, tree.nr, ltree.nr)
14
15        # and store the result for the left tree node
16        trace[ltree.node] = qx2, qy2
17        if symmetric:
18            trace[ltree.node.opposite] = qx2.reverse(), qy2.reverse()
19
20        # then continue propagating through the left tree
21        propagate(ltree, trace, symmetric)
22
23    # If the right tree exists...
24    if rtree is not None:
25        # ... transform q by applying ABCDs for node -> right node coupling
26        qx2 = transform_q(tree.right_abcd_x, qx1, tree.nr, rtree.nr)
27        qy2 = transform_q(tree.right_abcd_y, qy1, tree.nr, rtree.nr)
28
29        # and store the result for the right tree node
30        trace[rtree.node] = qx2, qy2
31        if symmetric:
32            trace[rtree.node.opposite] = qx2.reverse(), qy2.reverse()
33
34        # then continue propagating through the right tree
35        propagate(rtree, trace, symmetric)

```

Listing 2: Slightly simplified code corresponding to the recursive tracing algorithm implemented in FINESSE 3. This function is called on all branches through a `TraceTree` until all of these paths are exhausted. Note that the `transform_q` function shown here implements the ABCD law given by Equation (2.2) exactly.

A.1 Planting a forest

The planting of a `TraceForest` refers to the procedure carried out for constructing each `TraceTree` and piecing these together in the forest container by the order in which each should be traced when triggering a beam trace on the forest. This action occurs, whenever a

beam trace is called on a model, only if a fundamental change has been made to the model since the last time the forest was planted — i.e. if any optical components have been added or removed or if any trace dependency priorities were changed. Figure A.2 contains a visual aide for interpreting how trace forests are planted, in terms of an arbitrary list of ordered beam trace dependencies. For a practical example of this, refer back to Figure 2.7.

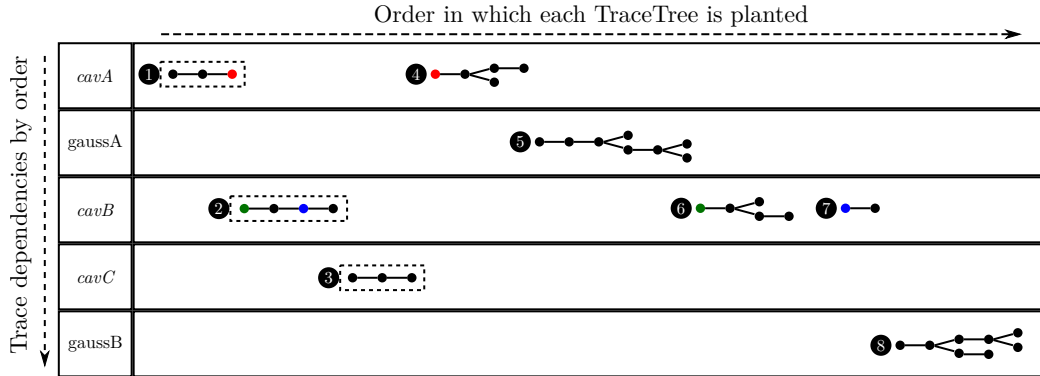


Figure A.2: This diagram gives a conceptual overview of how `TraceForest` structures are planted, from an arbitrary list of ordered dependencies, in `FINESSE 3`. The same type of procedure applies to any other ordered dependency list. The trees highlighted with dashed boxes represent the internal trace trees of cavity instances — hence why they decay to linked lists as shown. These types of trees are always planted first, in the order that each cavity appears in the trace order, to ensure that beam parameters at cavity nodes will be set according to the trace of the cavity eigenmode. After these trees are planted, all the other (“external”) trace trees are then constructed and set in the forest according to the ordering of their dependency object. The shaded nodes represent tree nodes which depend upon the node of an internal cavity tree. Note that a cavity dependency may have no external trace trees if these nodes were covered by previous trace trees; this is implied for the `cavC` in this example. Also note that the different sizes of the internal cavity trees here simply implies varying geometries of such cavities.

A.2 Constructing the changing trace forest

Whilst Appendix A.1 shows how the full `TraceForest` of a model is planted, this structure is not the one which gets used during a simulation in which a geometric parameter¹ is being swept. Instead, using the aforementioned full forest, a *changing* trace forest is constructed based on the trace dependencies and / or intermediate parameters (between dependencies) which will be changing during the simulation. This is an optimisation which provides two principle functions. Firstly, it ensures that only those beam parameters which can change

¹A geometric parameter, in this context, is defined as a parameter which is a dependency of any components’ ABCD matrix. This includes, for example, distances between components and radii of curvature of surfaces.

during a given simulation will get updated — all other beam parameters, at nodes which are not part of the changing forest, are fixed at their initial values. And secondly, it allows us to pre-determine which node couplings will be potentially mode mismatched during a simulation — a feature which will be detailed in Appendix A.3.

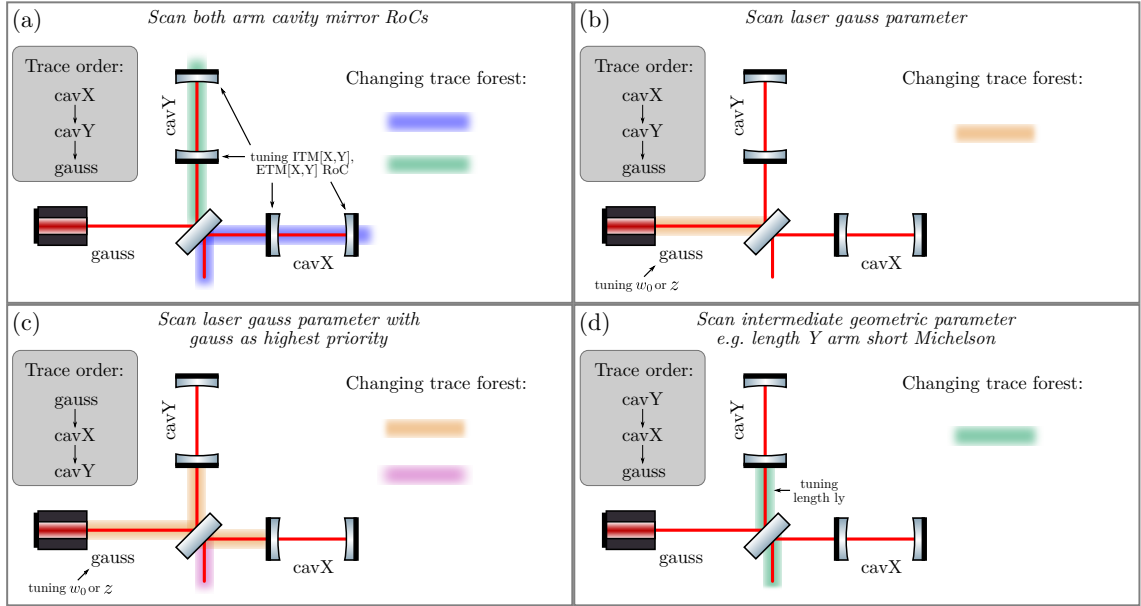


Figure A.3: Four examples of different changing trace forest structures based on a Fabry-Perot Michelson interferometer, where varied geometric parameters are marked as tunable. In (a) the arm cavity mirror RoCs are tuned, leading to the changing beam parameters consisting of both arm paths and the output, or anti-symmetric (AS), port of the central beam splitter. Boxes (b) and (c) consist of the cases where a parameter of the gauss object at the laser is varied (i.e. w_0 and $/$ or z); where box (c) shows the changing structure when this gauss is prioritised in the tracing order. Finally, (d) gives a scenario where an intermediate geometric parameter, in this case the length ly , is swept — such that the changing beam parameters are only from this point onwards (until other dependencies are encountered).

The easiest way to understand how changing forests are built is through examples, and, thus, in Figure A.3 we show several cases centred around the same type of configuration — namely a Michelson interferometer with Fabry-Perot cavities in the arms. Each box in this figure highlights a different forest of changing beam parameters based on varied geometric parameters being scanned at alternate points in the interferometer. These examples capture most of the types of changing forest structures which can occur, with differing trace dependency orders framing how this priority also affects the changing beam propagation during simulations. More complicated configurations, such as full detector files, will still follow the core logic given in these simpler examples.

A.3 Finding potentially varying mode mismatch couplings

Using the changing trace forest, we can begin performing some smart, pre-simulation loop, functions for optimisation. One such stage, carried out immediately after constructing this forest, is to find the optical node couplings which will be potentially mode mismatched (with changing mismatch values) during the simulation. For both symmetric and asymmetric beam traces, these couplings will be the intersection points between trees with different dependencies in the changing forest and the full model trace forest. In addition, for simulations which employ symmetric tracing, the node couplings corresponding to “self-reflections” (i.e. at mirror-type components) will also have possibly changing mode mismatches during the simulation — see §2.2.3.3 for details on symmetric vs asymmetric tracing.




Figure A.4: A zoomed-in view of the central beam splitter corresponding to box (a) shown on Figure A.3. This figure demonstrates the potentially mode mismatched node couplings in such a scenario, due to the three different trace trees intersecting at the beam splitter. The full (stored by the model) and changing (kept by the simulation) **TraceForest** structures in FINESSE 3 provide the ability to pre-determine all the node couplings which will exhibit changing mode mismatches during a simulation. This feature is useful for the user, for inspecting mismatch locations in their configuration, and also confers significant performance benefits to any modal simulation where geometric parameters are being swept.

We can utilise the configuration from Figure A.3, and specifically box (a) of this diagram, to give an example of how this works. This is shown in Figure A.4, where we have focused in on the central beam splitter to demonstrate the intersection of the three different trace trees at this component. Each port of the beam splitter is indicated by a shaded oval, where the colour corresponds to the trace dependency / tree that will set the beam parameters of the nodes at this port. The modal basis of the nodes at ports 3 and 4 are controlled by the trace from the cavity in the X-arm, as this cavity has the highest tracing priority (see Figure A.3 (a)). The black arrows on this figure then display the couplings which will have changing mode mismatch values during the simulation.

A.4 Dealing with unstable tracing regions

In simulations where geometric parameters are being scanned, regions of these parameter spaces can often yield unstable cavities — this is especially true for multi-dimensional scans over large ranges. The previous versions of FINESSE did not handle such regions and, instead, simply raised errors indicating that an unstable cavity was encountered. This is now handled in a smarter way in FINESSE 3, as we will detail in this section.

Central to the logic of dealing with unstable cavity regions is the “contingent” trace forest. As the changing forest structure is built using the specified trace dependencies, if one of these cavity-type dependencies becomes unstable at certain data points of a simulation (due to an associated geometric parameter being swept) then this structure cannot be applied, in general, for these points. Instead a new trace forest (termed a “contingent” trace forest), corresponding to the combination of unstable cavities at those points, is built and used as a substitute structure to trace the beam. This contingent forest is also cached in the simulation object such that it will be used for future data points where the exact same combination of unstable cavities occurs; preventing the need to build such structures multiple times.




Figure A.5: This figure gives an hypothetical 2D geometric parameter scan involving a model with two cavities, A and B, where at certain, arbitrary points these cavities become unstable. These highlighted regions then show how the beam tracing is performed, in a broad sense, at these regions of different cavity instabilities. In the blue region only cavity B is unstable and, hence, a contingent forest $\mathcal{F}\langle B \rangle$ is constructed and cached once this region of the simulation is encountered. Whilst, in the red region, both cavities A and B are unstable leading to a new contingent forest $\mathcal{F}\langle A, B \rangle$ in this zone of the parameter space. Note that, in this latter case, if no other stable tracing dependencies exist in the model then no contingent forest is built — details on what happens in such a scenario will be given in the next section.

Figure A.5 contains a conceptual example of how the contingent forest logic works, in the context of a simulation in which two geometric parameters (affecting two separate optical cavities, A and B) are scanned. One can observe from this figure that for the majority of data points the main changing forest is used to trace the beam — i.e. when both cavities are geometrically stable. In this example, cavity B becomes unstable in the blue region and so a new contingent forest with this cavity removed is constructed and then used for any data points in which this cavity is unstable. Similar behaviour then applies to other regions where different combinations of cavities are unstable — e.g. the red region. The dots on the boundaries of these two regions indicate where each contingent forest is built and cached for future use; assuming that the parameter scan propagates from the lower-left corner of this plot. Finally, the solid boundaries at the right edges of these zones correspond to “exit-points” from the unstable regime(s) in the simulation. These points are where the simulation switches back to the main changing forest.

The behaviour of unstable cavities in practice underpins the physical motivation for handling such simulation regimes in this way. A geometrically unstable cavity generally experiences large diffraction losses [100, 164, 165], especially for large order HOMs. In modal simulations this can be modelled as the optical cavity ceasing to act as a resonator for any light field, thereby losing power quickly. In the frequency-domain this will translate to sharp drops in the circulating cavity power for those data points which result in the cavity becoming geometrically unstable.

A.4.1 What happens when there are no stable trace dependencies?

In the previous section we detailed how contingent trace forests are built and used when unstable cavity regions are encountered. However, in some scenarios it may occur that, during a geometric parameter scan (or scans), all of the cavities in a model become unstable whilst no other trace dependencies are present. In such a case, the configuration as a whole becomes “undefined” from a modal basis perspective as there is then no way for any beam parameters to be propagated throughout the system. To deal with this situation in a user-friendly way, rather than simply raising an error and crashing out of the simulation, the simulation registers when this occurs and masks the corresponding data point(s) of the parameter scan(s) for any detector which requires a modal basis to be set in order to compute its output. The output data, exposed to the user, for such detectors are then masked NumPy arrays [107] which can be used in an established, and well-defined, way in mathematical operations and plotting.

To demonstrate the utility of such an approach, Figure A.6 gives an exemplary case of this simulation masking. In this example we scan the length of an aLIGO arm cavity $L \in [1, 5]$ km and detect the spot size on both cavity mirrors whilst also probing the g -factor of the cavity. The former detectors require the modal basis to be defined² whilst the latter does not. The

²This is because the beam size at a given node is computed via $w(q) = |q| \sqrt{\frac{\lambda}{\pi n_r \Im\{q\}}}$ in FINESSE 3.

effect of this is that the beam size detector outputs are masked at the points where the arm cavity becomes unstable — as there are no other trace dependencies present in this simple file. On the other hand, the g -factor detector output array is not masked at these points; providing us with the ability to see where these values breach the stable limits on the plot. This behaviour is automatic in FINESSE 3 in that only those detectors which require a defined modal basis will have their outputs masked in completely unstable trace regions.




Figure A.6: An example of detector output masking in simulations. This uses an aLIGO arm cavity, with an ITM RoC of 1934 m and an ETM RoC of 2245 m, where the length of the cavity is scanned. The detectors placed in the model were probes for the beam sizes at the mirrors (blue traces) and the cavity stability factor g (red trace). As the beam size detectors require the modal basis to be defined, the outputs of these are masked when the cavity becomes unstable — showing up as gaps in the data in the plot. The g -factor probe, on the other hand, only requires the cavity round-trip ABCD matrix and so the output data for this entry in the solution are not masked in the unstable regime. This also lets us observe where the g -factor enters unstable regions; as seen from the plot.

Appendix B

Analytic solution to WS phase space overlap integrals

In §4.1 we have seen that the overlap between any mode, represented in WS space as (W, S) , and a primary mode, (W_P, S_P) , can be computed by:

$$O(W, S) = \iint E(W, S) E(W_P, S_P)^* dx dy \times \iint E(W, S)^* E(W_P, S_P) dx dy, \quad (\text{B.1})$$

where the field $E(W, S)$ is given as

$$E(W, S) = \sqrt{\frac{2}{\pi}} \frac{1}{W} \exp\left(- (x^2 + y^2) \left(\frac{1}{W^2} + ik\frac{S}{2}\right)\right). \quad (\text{B.2})$$

In this appendix we will derive an analytic expression for Equation (B.1) by solving the overlap integrals therein. This will be performed for both an infinite plane (Appendix B.1) and also a finite circular aperture (Appendix B.2) for completeness.

First we will define some convenience symbols, namely we will let:

$$a = \sqrt{\frac{2}{\pi}} \frac{1}{W}, \quad a_p = \sqrt{\frac{2}{\pi}} \frac{1}{W_P}, \quad \alpha = \frac{1}{W^2} + ik\frac{S}{2}, \quad \alpha_p = \frac{1}{W_P^2} + ik\frac{S_P}{2}. \quad (\text{B.3})$$

We can then express the overlap integrands as,

$$\begin{aligned}
 E(W, S) E(W_P, S_P)^* &= aa_p \exp(-\alpha(x^2 + y^2)) \exp(-\alpha_p^*(x^2 + y^2)) \\
 &= aa_p \exp(-(\alpha + \alpha_p^*)(x^2 + y^2)),
 \end{aligned} \tag{B.4}$$

and

$$\begin{aligned}
 E(W, S)^* E(W_P, S_P) &= aa_p \exp(-\alpha^*(x^2 + y^2)) \exp(-\alpha_p(x^2 + y^2)) \\
 &= aa_p \exp(-(\alpha^* + \alpha_p)(x^2 + y^2)).
 \end{aligned} \tag{B.5}$$

It is useful to use a shorthand at this point of $\alpha + \alpha_p^* = \beta$, and so, commensurately, $\alpha^* + \alpha_p = \beta^*$. We can then write Equation (B.1) as

$$\begin{aligned}
 O(W, S) &= \iint aa_p \exp(-\beta(x^2 + y^2)) dx dy \times \iint aa_p \exp(-\beta^*(x^2 + y^2)) dx dy \\
 &= a^2 a_p^2 \iint \exp(-\beta x^2) \exp(-\beta y^2) dx dy \times \iint \exp(-\beta^* x^2) \exp(-\beta^* y^2) dx dy.
 \end{aligned} \tag{B.6}$$

B.1 Overlap on an infinite plane

Now, assuming that we are integrating over an infinite plane, we can note that

$$\int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \exp(-bx^2) \exp(-by^2) dx dy = \frac{\pi}{b}, \tag{B.7}$$

which, operating on Equation (B.6) in turn, yields

$$\begin{aligned}
 O(W, S) &= a^2 a_p^2 \frac{\pi}{\beta} \frac{\pi}{\beta^*} \\
 &= a^2 a_p^2 \frac{\pi^2}{|\beta|^2}.
 \end{aligned} \tag{B.8}$$

Using the expressions for a , a_p and β from earlier, and noting that

$$\beta = \alpha + \alpha_p^* = \frac{1}{W^2} + \frac{1}{W_P^2} + i\frac{k}{2}(S - S_P), \quad (\text{B.9})$$

we can expand the final expression out to

$$\begin{aligned} O(W, S) &= \frac{4}{\pi^2} \frac{1}{W^2 W_P^2} \frac{\pi^2}{\left| \left(\frac{1}{W^2} + \frac{1}{W_P^2} \right) + i\frac{k}{2}(S - S_P) \right|^2} \\ &= \frac{4}{W^2 W_P^2} \frac{1}{\left| \left(\frac{1}{W^2} + \frac{1}{W_P^2} \right) + i\frac{k}{2}(S - S_P) \right|^2} \\ \Rightarrow O(W, S) &= \frac{4}{W^2 W_P^2} \frac{1}{\left(\frac{1}{W^2} + \frac{1}{W_P^2} \right)^2 + \frac{k^2}{4}(S - S_P)^2}. \end{aligned} \quad (\text{B.10})$$

This exact form of the WS space overlap quantity is given in Equation (4.5) and used throughout Chapter 4 as the foundation from which each phase space plot is built.

B.2 Computing the overlap for a finite circular aperture

For completeness, we will now investigate how the overlap quantity changes when computing the overlap integrals in Equation (B.6) for a finite aperture (with a circular profile) rather than an infinite plane. We will label this ‘‘aperture overlap’’ now as $O_A(W, S)$.

Similarly to the previous section, we can note that integrating Gaussian profiles over a disk of radius R yields

$$\iint_{x^2+y^2 < R^2} \exp(-bx^2) \exp(-by^2) dx dy = \frac{\pi}{b} (1 - \exp(-bR^2)), \quad (\text{B.11})$$

which recovers back to the result in Equation (B.7) when $R \rightarrow \infty$, as expected. If we apply

this now to Equation (B.6), we obtain

$$\begin{aligned}
 O_A(W, S) &= a^2 a_p^2 \frac{\pi^2}{|\beta|^2} (1 - \exp(-\beta R^2)) (1 - \exp(-\beta^* R^2)) \\
 &= (1 - \exp(-\beta R^2)) (1 - \exp(-\beta^* R^2)) O(W, S) \\
 &= [1 + \exp(-2\Re\{\beta\}R^2) - (\exp(-\beta R^2) + \exp(-\beta^* R^2))] O(W, S),
 \end{aligned} \tag{B.12}$$

where $O(W, S)$ is the WS phase space overlap for an infinite plane, as given by Equation (4.5). At this point we can simplify Equation (B.12) by noting that

$$\begin{aligned}
 \exp(-\beta R^2) + \exp(-\beta^* R^2) &= \exp(-\Re(\beta)R^2) (\exp(i\Im(\beta)R^2) + \exp(-i\Im(\beta)R^2)) \\
 &= 2 \cos(\Im(\beta)R^2) \exp(-\Re(\beta)R^2).
 \end{aligned} \tag{B.13}$$

Plugging this into Equation (B.12) then yields,

$$\begin{aligned}
 O_A(W, S) &= [1 + \exp(-2\Re\{\beta\}R^2) - 2 \cos(\Im\{\beta\}R^2) \exp(-\Re\{\beta\}R^2)] O(W, S) \\
 &= [1 + (\exp(-\Re\{\beta\}R^2) - 2 \cos(\Im\{\beta\}R^2)) \exp(-\Re\{\beta\}R^2)] O(W, S).
 \end{aligned} \tag{B.14}$$

Now we can recall, using Equation (B.9), that

$$\Re\{\beta\} = \frac{1}{W^2} + \frac{1}{W_p^2}, \quad \text{and} \quad \Im\{\beta\} = \frac{k}{2} (S - S_p), \tag{B.15}$$

which allows us to finally express $O_A(W, S)$ as,

$$O_A(W, S) = \left[1 + \left(\exp(-(r^2 + r_p^2)) - 2 \cos\left(\frac{kR^2}{2} \Delta S\right) \right) \exp(-(r^2 + r_p^2)) \right] O(W, S), \tag{B.16}$$

where we have defined $r = R/W$, $r_p = R/W_p$ and $\Delta S = S - S_p$ for convenience. The former shorthand variables here are the ratios of the aperture radius to the given beam radii, whilst

the latter is simply the difference in defocus of the two compared modes. For a mode equal to the primary mode, i.e. $W = W_P$ and $S = S_P$, Equation (B.16) reduces to:

$$O_A(W_P, S_P) = 1 - 2 \exp(-2r_p^2) + \exp(-4r_p^2). \quad (\text{B.17})$$

Just for fun we can plot Equation (B.17), resulting in Figure B.1. As expected, for $r_p < 1$ we see a sharp fall in the overlap of the primary mode with itself, whilst for $r_p \gtrsim 2$ this quantity is approximately unity; as a tiny portion of the beam amplitude is clipped at this point.




Figure B.1: A simple plot showing the overlap of a primary mode (W_P, S_P) with itself for increased ratios of a circular aperture radius to the beam size. For large apertures one recovers the form of the overlap function on an infinite plane — leading to $O_A(W_P, S_P)$ tending towards unity.

Appendix C

Field couplings for a DC and AC misaligned Fabry-Perot cavity

In chapter 5, specifically §5.1, we saw that the first order alignment scaling for a DC plus AC misaligned cavity is (with only the upper sideband included in this equation),

$$P \sim a_{01}^*(\omega_0) a_{01}(\omega_0 + \omega_m), \quad (\text{C.1})$$

where ω_0 is the carrier field frequency and ω_m is the modulation frequency of the tilt of one of the mirrors of the cavity. Here P is simply the signal that would be detected on the photodiode corresponding to the configuration in Figure 5.1. Given that our goal was to find the alignment scaling with the beam size of the cavity mode in the aforementioned chapter, in this appendix we will thus derive an expression for Equation (C.1) in terms of this beam size.

To assist with this derivation, Figure C.1 gives a repeat view of the configuration under consideration — but now with each field quantity present. The fields highlighted in blue are the sideband fields around the carrier, created as a result of the modulated tilt of the end mirror. We will represent these tilt modulation couplings as “injections” of the field i_2 . Note that, whilst both upper and lower sidebands are shown on this figure, we will only consider the upper sideband fields when looking at the AC picture in this appendix. All couplings determined in this regime will equally apply to the lower sideband. For the static tilt at the input mirror, we will similarly represent this as an injection of the field i_1 upon reflection from this mirror. This tilt will couple to the first order mode at the carrier frequency.




Figure C.1: This shows the reference configuration for a DC plus AC misaligned Fabry-Perot cavity. The circulating carrier fields are given at each location. The $i_{1|2}$ terms represent “injections” of fields caused by the tilt couplings for the DC tilt at the input and AC tilt at the output mirror, respectively. The incident field is given by i_0 and is assumed to be a pure HG_{00} mode.

We will only consider *pitched* tilts of the mirrors for the maths that follows, allowing us to set the tangential mode index $n = 0$ in all cases. This means that whenever the symbols $a_{j,m}$ are used, these really expand to $a_{j,0m}$ — i.e. the field at position j (according to Figure C.1) in the mode HG_{0m} . An additional convention which we will use is to refer to coupling coefficients on reflection from the input mirror with k (e.g. k_{01} is the coefficient for the mode HG_0 to HG_1 for this coupling), and on reflection from the end mirror with κ (e.g. κ_{01} is the coefficient for the mode HG_0 to HG_1 , but this time for AC couplings).

For all of the following analyses we will assume that the cavity is tuned such that the HG_0 field at the carrier frequency is resonant. This effectively means that we can ignore fields in the HG_1 mode when considering HG_0 couplings; as $a_{j,0} \gg a_{j,1}$ under this assumption. Additionally, we will assume that the tilt angles of the mirrors are sufficiently small such that the coupling coefficients k_{00} , k_{11} and κ_{00} , κ_{11} are all approximately unity. Finally, we will assume an impedance matched cavity for convenience — i.e. the amplitude reflectivities and transmissivities shown on Figure C.1 obey $r = r_1 = r_2$ and $t = t_1 = t_2$.

When referring to fields at DC (i.e. at the carrier frequency ω_0) we will simply use $a_{j,m}$, whilst fields at AC (i.e. at the upper sideband frequency $\omega_0 + \omega_m$) will be labelled as $\bar{a}_{j,m}$ to distinguish these quantities.

For the analyses that follow, our overarching goal is to find expressions for the fields $a_{4,1}$ and $\bar{a}_{4,1}$ — i.e. the circulating HG_1 mode at both the carrier and signal sideband frequencies. These can then be applied directly via Equation (C.1) to obtain a scaling relation for this quantity as a function of the beam size. Note that, evidently, the signal at the photodiode will not be equal to the beating of the fields at this location, however the extra coupling terms after a_4 will be simple transmission (through the input mirror) and reflection (from the beam splitter) terms and so will not affect the scaling relation in this case.

C.1 HG₀₀ field couplings at DC

The simplest regime to consider is solving for the circulating field in the fundamental mode (HG₀₀) at the carrier frequency. In such a case, the injected field terms i_1 and i_2 are zero. The system of linear equations for our field terms in this regime are then

$$\begin{aligned}
 a_{1,0} &= r a_{4,0} + i t i_0 \\
 a_{2,0} &= \exp(-i(kL + \psi)) a_{1,0} \\
 a_{3,0} &= r a_{2,0} \\
 a_{4,0} &= \exp(-i(kL + \psi)) a_{3,0},
 \end{aligned} \tag{C.2}$$

where ψ is the Gouy phase accumulated over the intra-cavity space. Note that here we have used the FINESSE convention of a 90° phase-shift on transmission through a surface, hence the factor of i shown in the expression for $a_{1,0}$. Solving this set of equations, we find

$$a_{4,0} = \frac{i t r \exp(-i(2kL + \psi_{\text{RT}}))}{1 - R \exp(-i(2kL + \psi_{\text{RT}}))} i_0, \tag{C.3}$$

where we have noted that $2\psi = \psi_{\text{RT}}$ is the round-trip Gouy phase of the cavity. Again, as we are treating this cavity as being tuned to the resonance of the fundamental mode, it follows that $2kL + \psi_{\text{RT}} = 2\pi N$ with $N \in \mathbb{Z}$. This reduces Equation (C.3) to simply,

$$a_{4,0} = \frac{i t r}{1 - R} i_0. \tag{C.4}$$

By noting that $r = \sqrt{R}$ and $t = \sqrt{1 - R}$ (assuming loss-less mirrors), we can further simplify this to,

$$a_{4,0} = i \sqrt{\frac{R}{1 - R}} i_0. \tag{C.5}$$

C.2 HG₀₁ field couplings at DC

Now that we have an expression for the circulating field in the fundamental mode in Equation (C.5), we can find a complete equation for the circulating field in the HG₁ mode at the carrier frequency. This mode is excited via the static pitch misalignment of the input mirror,

as indicated by Figure C.1, and so we now have a non-zero i_1 as an injected coupling. The injected terms i_0 and i_2 are both zero in this case — the former because the incident beam is a pure HG_0 mode and the latter due to the AC tilt of the end mirror having no effect on the fields at this frequency.

Our set of linear equations for this mode and frequency is then

$$\begin{aligned}
 a_{1,1} &= r a_{4,1} + i_1 \\
 a_{2,1} &= \exp(-i(kL + 2\psi)) a_{1,1} \\
 a_{3,1} &= r a_{2,1} \\
 a_{4,1} &= \exp(-i(kL + 2\psi)) a_{3,1},
 \end{aligned} \tag{C.6}$$

where the phase accumulation over a single trip of the cavity is now 2ψ due to the accumulated Gouy phase for a mode being linearly dependent upon the mode order (see Equation (1.26)). We can solve for $a_{4,1}$ to yield

$$a_{4,1} = \frac{r \exp(-i(2kL + 2\psi_{\text{RT}})) i_1}{1 - R \exp(-i(2kL + 2\psi_{\text{RT}}))}. \tag{C.7}$$

As before, in Appendix C.1, we can use the assumption that the fundamental carrier mode is resonant (such that $\exp(-i(2kL + \psi_{\text{RT}})) = 1$) to further simplify this to

$$a_{4,1} = \frac{r \exp(-i\psi_{\text{RT}}) i_1}{1 - R \exp(-i\psi_{\text{RT}})}. \tag{C.8}$$

As the injected term in this case represents the static misalignment of the input mirror, we can thus write this term as $i_1 = r k_{01} a_{4,0}$. Using this, $a_{4,1}$ is then

$$a_{4,1} = \frac{R k_{01} \exp(-i\psi_{\text{RT}})}{1 - R \exp(-i\psi_{\text{RT}})} a_{4,0}. \tag{C.9}$$

We can now begin to express this in terms of the beam size w of the cavity mode at the mirrors. Firstly, we can recognise that the denominator of Equation (C.9) is equivalent to the reciprocal of Equation (5.16) with $n = 0$ and $m = 1$; we shall refer to this quantity as d_1 . From the Equation (5.24) in §5.2.2 we know that this term is given approximately by

$$d_1 = \frac{1}{2} \frac{\pi}{L\lambda} w^2 \exp(i\phi_{d_1}), \tag{C.10}$$

where L is the cavity length and ϕ_{d_1} is the phase of the complex resonance enhancement factor, and, as an aside, can be shown to be equivalent to

$$\phi_{d_1} = \arctan \left(\sin \left(\frac{\psi_{\text{RT}}}{2} \right) \tan \left(\frac{\psi_{\text{RT}}}{2} \right) \right). \quad (\text{C.11})$$

The other term in Equation (C.9) which we can substitute out is k_{01} — using the approximation given in Equation (5.13). Putting all this together, we obtain

$$\begin{aligned} a_{4,1} &\approx \frac{1}{2} \left(\frac{\pi}{L\lambda} \right) w^2 \exp(i\phi_{d_1}) R \frac{\pi}{\lambda} w \gamma_{\text{DC}} \exp(i\varphi) \exp(-i\psi_{\text{RT}}) a_{4,0} \\ &= \frac{R}{2} \left(\frac{\pi}{\lambda} \right)^2 \frac{w^3}{L} \gamma_{\text{DC}} \exp(i(\phi_{d_1} + \varphi - \psi_{\text{RT}})) a_{4,0}, \end{aligned} \quad (\text{C.12})$$

where, recalling from §5.2.1, $\varphi = \arctan(\cot(\psi_{\text{RT}}/4))$ meaning that all the phase terms here can be expressed in terms of only the round-trip Gouy phase. Now substituting in Equation (C.5) for $a_{4,0}$, and noting that $i = \exp(i\pi/2)$, yields

$$a_{4,1} \approx \sqrt{\frac{R^3}{1-R}} \left(\frac{\pi}{\lambda} \right)^2 \frac{w^3}{L} \beta_{\text{DC}} \exp \left(i \left(\frac{\pi}{2} + \phi_{d_1} + \varphi - \psi_{\text{RT}} \right) \right) i_0, \quad (\text{C.13})$$

where we have also noted that the static misalignment angle $\gamma_{\text{DC}} = 2\beta_{\text{DC}}$; the tilt angle β_{DC} is shown on Figure C.1. One can note immediately that this field amplitude has a w^3 dependence.

To check that Equation (C.13) is approximately correct for small static misalignment angles, we can compute this quantity as a function of β_{DC} and compare it to the results of a FINESSE 3 simulation. This is shown on Figure C.2, using an aLIGO-like symmetric and impedance matched arm cavity. The x -axis is given as a misalignment angle as a fraction of the cavity divergence, similarly to the $|k_{01}|$ comparison plot shown earlier in Figure 5.3. On the y -axis the amplitude of $a_{4,1}$ is shown. As expected, given the aforementioned results of Figure 5.3, the analytical approximation of Equation (C.13) is valid as long as the misalignment angle is small — i.e. $\gamma_{\text{DC}} \ll \bar{\gamma}_0$; where $\bar{\gamma}_0$ is the cavity divergence angle.

C.3 HG₀₁ field couplings at AC

Contrary to the carrier solutions, we do not require knowledge of the HG₀ mode field amplitudes at the signal sideband frequency in order to find an expression for $\bar{a}_{4,1}$ to first order. Given this, we can set the injected terms i_0 and i_1 to zero — where the former is true as




Figure C.2: This figure gives a comparison of the analytical approximation of $|a_{4,1}|$ against a FINESSE 3 simulation. The basis of these results is an aLIGO-style arm cavity but with symmetric cavity mirror RoCs and in an impedance matched regime. As expected, the approximation agrees well with the simulation for small tilt angles $\beta_{\text{DC}} = \gamma_{\text{DC}}/2$.

the incident beam is a pure fundamental mode at the carrier frequency. Thus, our set of equations in this regime is

$$\begin{aligned}
 \bar{a}_{1,1} &= r\bar{a}_{4,1} \\
 \bar{a}_{2,1} &= \exp(-i(kL + 2\psi))\bar{a}_{1,1} \\
 \bar{a}_{3,1} &= r\bar{a}_{2,1} + i_2 \\
 \bar{a}_{4,1} &= \exp(-i(kL + 2\psi))\bar{a}_{3,1},
 \end{aligned} \tag{C.14}$$

where we are making the assumption that the tilt modulation frequency of the end mirror is small enough such that $\bar{k} = k + k_m \approx k$; with $k_m = \omega_m/c$ as the wavenumber of the modulated field. Solving for $\bar{a}_{4,1}$ yields

$$\bar{a}_{4,1} = \frac{\exp(-i(kL + \psi_{\text{RT}})) i_2}{1 - R \exp(-i\psi_{\text{RT}})}, \tag{C.15}$$

in which we have again noted that $\exp(-i(2kL + \psi_{\text{RT}})) = 1$ given our assumption that the HG_0 mode is resonant, at the carrier frequency, inside the cavity. The injected term, i_2 , here

represents the modulated tilt of the end mirror and so we can write this as

$$\dot{\mathbf{i}}_2 = r\kappa_{01}a_{2,0}, \quad (\text{C.16})$$

which is, qualitatively, the coupling of the HG_0 mode at the *carrier* frequency into the HG_1 mode at the *signal* frequency due to this tilt modulation. We can now use the derivations in Appendix C.1 to express the field $a_{2,0}$ as

$$\begin{aligned} a_{2,0} &= \exp(-i(kL + \psi)) [ra_{4,0} + iti_0] \\ &= \exp(-i(kL + \psi)) \left[ir\sqrt{\frac{R}{1-R}} i_0 + iti_0 \right] \\ &= i\frac{1}{\sqrt{1-R}} \exp(-i(kL + \psi)) i_0. \end{aligned} \quad (\text{C.17})$$

Substituting these terms into Equation (C.15) then gives us

$$\begin{aligned} \bar{a}_{4,1} &= i\frac{r\kappa_{01} \exp(-i(kL + \psi_{\text{RT}}))}{1 - R \exp(-i\psi_{\text{RT}})} \frac{1}{\sqrt{1-R}} \exp(-i(kL + \psi)) i_0 \\ &= i\sqrt{\frac{R}{1-R}} \frac{\kappa_{01} \exp(-i\frac{\psi_{\text{RT}}}{2})}{1 - R \exp(-i\psi_{\text{RT}})} i_0. \end{aligned} \quad (\text{C.18})$$

Similarly to Appendix C.2, we can now express this in terms of the beam size w . The procedure is identical to the aforementioned section, with the only difference now being that the κ_{01} term is dependent upon the AC misalignment angle $\gamma_{\text{AC}} = 2\beta_{\text{AC}}$. After a short amount of algebra we obtain

$$\bar{a}_{4,1} \approx \frac{1}{2} \sqrt{\frac{R}{1-R}} \left(\frac{\pi}{\lambda}\right)^2 \frac{w^3}{L} \beta_{\text{AC}} \exp\left(i\left(\frac{\pi}{2} + \phi_{d_1} + \varphi - \frac{\psi_{\text{RT}}}{2}\right)\right) i_0. \quad (\text{C.19})$$

C.4 Computing the signal beat at the photodiode

As noted at the start of this appendix, the signal beat which we want to quantify is given by Equation (C.1) — thus we now need to compute $|a_{4,1}^* \bar{a}_{4,1}|$. This quantity can be quickly obtained from the expressions in Equation (C.13) and Equation (C.19). First we can note

that

$$a_{4,1}^* \bar{a}_{4,1} \approx \frac{1}{2} \frac{R^2}{1-R} \left(\frac{\pi}{\lambda}\right)^4 \frac{w^6}{L^2} \beta_{\text{DC}} \beta_{\text{AC}} \exp\left(i \frac{\psi_{\text{RT}}}{2}\right) i_0^2. \quad (\text{C.20})$$

Then the amplitude of this field is simply

$$|a_{4,1}^* \bar{a}_{4,1}| \approx \frac{1}{2} \frac{R^2}{1-R} \left(\frac{\pi}{\lambda}\right)^4 \frac{w^6}{L^2} |\beta_{\text{DC}} \beta_{\text{AC}}| |i_0|^2. \quad (\text{C.21})$$

Thus we have shown that, to first order, the alignment to longitudinal coupling scales as approximately the beam size to the power of six.

We noted earlier that this is the beating of the *circulating* fields in the the first order mode, rather than the actual beat on the photodiode upon reflection from the cavity. This is not an important point in this case as the main goal was to show the dependence upon the beam size — which will not be altered, to first order, by the ensuing couplings from this point. However, for completeness, if we label the reflected field in the mode HG_1 as $a_{\text{refl},1}$ then the corresponding beat at this “node” will be given by

$$P_{\text{refl},1} = |a_{\text{refl},1}^* \bar{a}_{\text{refl},1}| \approx \frac{R^2}{2} \left(\frac{\pi}{\lambda}\right)^4 \frac{w^6}{L^2} |\beta_{\text{DC}} \beta_{\text{AC}}| |i_0|^2, \quad (\text{C.22})$$

i.e. a factor of $t^2 = 1 - R$ difference due to the transmission through the input mirror for both $a_{4,1}$ and $\bar{a}_{4,1}$. This is, again, assuming an impedance matched cavity and assuming that the reflection coupling, into first order modes at the front surface of the statically pitched input mirror, is negligible.

C.4.1 Error estimation

We can use FINESSE to perform an estimation of the error on the derived approximation in Equation (C.22). To do so, a model was set-up using an aLIGO-style arm cavity with symmetric RoCs ($R_C = 2090$ m) in a loss-less, impedance matched regime. The beating of the carrier and signal sideband fields in the HG_{01} mode were detected upon reflection from the cavity, corresponding to the quantity given by Equation (C.22). Figure C.3 then shows the absolute difference between the two methods, computed as a function of the mirror tilt angles β_{DC} and β_{AC} . As one would expect, and as implied previously for the carrier-only regime by Figure C.2, the analytical approximation $P_{\text{refl},1}$ agrees well with the FINESSE simulation as long as the angles of misalignment are relatively small.




Figure C.3: The plot here gives an estimation on the error for the derived analytical approximation of the alignment signal in Equation (C.22), as a function of the tilt angles of the cavity mirrors. Plotted as contours are the absolute errors of this equation against the results of a FINESSE simulation — detecting the same quantity upon reflection from an aLIGO-style arm cavity with mirrors of equal reflectivities (and zero loss).

Bibliography

- [1] Samuel Rowlinson et al. “Feasibility study of beam-expanding telescopes in the interferometer arms for the Einstein Telescope”. In: *Phys. Rev. D* 103 (2 Jan. 2021), p. 023004. DOI: [10.1103/PhysRevD.103.023004](https://doi.org/10.1103/PhysRevD.103.023004). URL: <https://link.aps.org/doi/10.1103/PhysRevD.103.023004> (cit. on pp. iii, 23, 67, 137).
- [2] J. D. Hunter. “Matplotlib: A 2D graphics environment”. In: *Computing in Science & Engineering* 9.3 (2007), pp. 90–95. DOI: [10.1109/MCSE.2007.55](https://doi.org/10.1109/MCSE.2007.55) (cit. on p. vi).
- [3] Alexander Franzen. *ComponentLibrary: a vector graphics library for illustration of optics experiments*. <http://www.gwoptics.org/ComponentLibrary/>. 2006 (cit. on p. vi).
- [4] W. King. *Dragonslayer*. A Warhammer novel. Games Workshop Pub., 2000. ISBN: 9780743411578. URL: <https://books.google.co.uk/books?id=50UAAAAACAAJ> (cit. on p. vii).
- [5] B. P. Abbott et al. “Observation of Gravitational Waves from a Binary Black Hole Merger”. In: *Phys. Rev. Lett.* 116 (6 Feb. 2016), p. 061102. DOI: [10.1103/PhysRevLett.116.061102](https://doi.org/10.1103/PhysRevLett.116.061102). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.116.061102> (cit. on p. 1).
- [6] B. P. Abbott et al. “GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral”. In: *Phys. Rev. Lett.* 119 (16 Oct. 2017), p. 161101. DOI: [10.1103/PhysRevLett.119.161101](https://doi.org/10.1103/PhysRevLett.119.161101). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.119.161101> (cit. on p. 1).
- [7] B. P. Abbott et al. “Multi-messenger Observations of a Binary Neutron Star Merger”. In: *The Astrophysical Journal* 848.2 (Oct. 2017), p. L12. DOI: [10.3847/2041-8213/aa91c9](https://doi.org/10.3847/2041-8213/aa91c9). URL: <https://doi.org/10.3847/2041-8213/aa91c9> (cit. on p. 1).
- [8] F. J. Fattoyev, J. Piekarewicz, and C. J. Horowitz. “Neutron Skins and Neutron Stars in the Multimessenger Era”. In: *Phys. Rev. Lett.* 120 (17 Apr. 2018), p. 172702. DOI: [10.1103/PhysRevLett.120.172702](https://doi.org/10.1103/PhysRevLett.120.172702). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.120.172702> (cit. on p. 1).
- [9] Eemeli Annala et al. “Gravitational-Wave Constraints on the Neutron-Star-Matter Equation of State”. In: *Phys. Rev. Lett.* 120 (17 Apr. 2018), p. 172703. DOI: [10.1103/PhysRevLett.120.172703](https://doi.org/10.1103/PhysRevLett.120.172703). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.120.172703> (cit. on pp. 1, 6).

BIBLIOGRAPHY

- [10] LIGO. *LIGO Detection Companion Papers*. URL: <https://www.ligo.caltech.edu/page/detection-companion-papers> (cit. on p. 1).
- [11] B. P. Abbott et al. “GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs”. In: *Phys. Rev. X* 9 (3 Sept. 2019), p. 031040. DOI: [10.1103/PhysRevX.9.031040](https://doi.org/10.1103/PhysRevX.9.031040). URL: <https://link.aps.org/doi/10.1103/PhysRevX.9.031040> (cit. on p. 1).
- [12] *GraceDB — O3 gravitational-wave candidate event database*. <https://gracedb.ligo.org/superevents/public/O3/>. 2021 (cit. on p. 1).
- [13] J Aasi et al. “Advanced LIGO”. In: *Classical and Quantum Gravity* 32.7 (Mar. 2015), p. 074001. DOI: [10.1088/0264-9381/32/7/074001](https://doi.org/10.1088/0264-9381/32/7/074001) (cit. on pp. 1, 2, 3, 4, 67).
- [14] F Acernese et al. “Advanced Virgo: a second-generation interferometric gravitational wave detector”. In: *Classical and Quantum Gravity* 32.2 (Dec. 2014), p. 024001. DOI: [10.1088/0264-9381/32/2/024001](https://doi.org/10.1088/0264-9381/32/2/024001). URL: <https://doi.org/10.1088/0264-9381/32/2/024001> (cit. on pp. 1, 67).
- [15] T. Akutsu et al. “KAGRA: 2.5 generation interferometric gravitational wave detector”. In: *Nature Astronomy* 3.1 (Jan. 2019), pp. 35–40. ISSN: 2397-3366. DOI: [10.1038/s41550-018-0658-y](https://doi.org/10.1038/s41550-018-0658-y). URL: <https://doi.org/10.1038/s41550-018-0658-y> (cit. on pp. 1, 8).
- [16] B. P. Abbott et al. “Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA”. In: *Living Reviews in Relativity* 23.1 (Sept. 2020), p. 3. ISSN: 1433-8351. DOI: [10.1007/s41114-020-00026-9](https://doi.org/10.1007/s41114-020-00026-9). URL: <https://doi.org/10.1007/s41114-020-00026-9> (cit. on p. 1).
- [17] J Abadie et al. “Predictions for the rates of compact binary coalescences observable by ground-based gravitational-wave detectors”. In: *Classical and Quantum Gravity* 27.17 (2010), p. 173001. URL: <http://stacks.iop.org/0264-9381/27/i=17/a=173001> (cit. on p. 1).
- [18] B. S. Sathyaprakash and Bernard F. Schutz. “Physics, Astrophysics and Cosmology with Gravitational Waves”. In: *Living Reviews in Relativity* 12.1 (Dec. 2009), p. 2. ISSN: 1433-8351. DOI: [10.12942/lrr-2009-2](https://doi.org/10.12942/lrr-2009-2). URL: <https://doi.org/10.12942/lrr-2009-2> (cit. on pp. 1, 2).
- [19] Kiwamu Izumi and Daniel Sigg. “Advanced LIGO: length sensing and control in a dual recycled interferometric gravitational wave antenna”. In: *Classical and Quantum Gravity* 34.1 (Dec. 2016), p. 015001. DOI: [10.1088/0264-9381/34/1/015001](https://doi.org/10.1088/0264-9381/34/1/015001). URL: <https://doi.org/10.1088/0264-9381/34/1/015001> (cit. on p. 3).
- [20] Muzammil A. Arain and Guido Mueller. “Design of the Advanced LIGO recycling cavities”. In: *Opt. Express* 16.14 (July 2008), pp. 10018–10032. DOI: [10.1364/OE.16.010018](https://doi.org/10.1364/OE.16.010018). URL: <http://www.opticsexpress.org/abstract.cfm?URI=oe-16-14-10018> (cit. on pp. 3, 48, 68, 72).
- [21] Koji Arai et al. *Output Mode Cleaner (OMC) Design*. Tech. rep. T1000276-v5. 2011. URL: <https://dcc.ligo.org/LIGO-T1000276/public> (cit. on pp. 3, 30).

-
- [22] H. Wittel et al. “Matrix heater in the gravitational wave observatory GEO 600”. In: *Opt. Express* 26.18 (Sept. 2018), pp. 22687–22697. DOI: [10.1364/OE.26.022687](https://doi.org/10.1364/OE.26.022687). URL: <http://www.opticsexpress.org/abstract.cfm?URI=oe-26-18-22687> (cit. on pp. 4, 72).
- [23] Brian J. Meers. “Recycling in laser-interferometric gravitational-wave detectors”. In: *Phys. Rev. D* 38 (8 Oct. 1988), pp. 2317–2326. DOI: [10.1103/PhysRevD.38.2317](https://doi.org/10.1103/PhysRevD.38.2317). URL: <https://link.aps.org/doi/10.1103/PhysRevD.38.2317> (cit. on p. 4).
- [24] J. Mizuno et al. “Resonant sideband extraction: a new configuration for interferometric gravitational wave detectors”. In: *Physics Letters A* 175.5 (1993), pp. 273–276. ISSN: 0375-9601. DOI: [https://doi.org/10.1016/0375-9601\(93\)90620-F](https://doi.org/10.1016/0375-9601(93)90620-F). URL: <https://www.sciencedirect.com/science/article/pii/037596019390620F> (cit. on p. 4).
- [25] Chris Mueller. *Detector Characterization IMC Overview Presentation*. <https://dcc.ligo.org/LIGO-G1400096/public>. 2014. URL: <https://dcc.ligo.org/LIGO-G1400096/public> (cit. on p. 5).
- [26] A. Buikema et al. “Sensitivity and performance of the Advanced LIGO detectors in the third observing run”. In: *Phys. Rev. D* 102 (6 Sept. 2020), p. 062003. DOI: [10.1103/PhysRevD.102.062003](https://doi.org/10.1103/PhysRevD.102.062003). URL: <https://link.aps.org/doi/10.1103/PhysRevD.102.062003> (cit. on pp. 5, 6, 25, 101, 119).
- [27] LIGO Scientific Collaboration. *Instrument Science White Paper 2020*. Tech. rep. T2000407. 2020. URL: <https://dcc.ligo.org/LIGO-T2000407/public> (cit. on pp. 5, 102, 120, 121, 122).
- [28] J. G. Rollins, E. Hall, and C. Wipf. *pygwinc: Python Gravitational Wave Interferometer Noise Calculator*. <https://git.ligo.org/gwinc/pygwinc>. 2020 (cit. on p. 5).
- [29] J. G. Rollins and J. D. E. Creighton. *GW detector inspiral range calculation tools*. <https://git.ligo.org/gwinc/inspiral-range>. 2020 (cit. on p. 5).
- [30] Hsin-Yu Chen et al. “Distance measures in gravitational-wave astrophysics and cosmology”. In: *Classical and Quantum Gravity* 38.5 (Jan. 2021), p. 055010. DOI: [10.1088/1361-6382/abd594](https://doi.org/10.1088/1361-6382/abd594). URL: <https://doi.org/10.1088/1361-6382/abd594> (cit. on pp. 5, 6).
- [31] John Miller. *Inspiral range with cosmology*. Tech. rep. T1500491. 2015. URL: <https://dcc.ligo.org/LIGO-T1500491> (cit. on p. 5).
- [32] Lee Samuel Finn. “Binary inspiral, gravitational radiation, and cosmology”. In: *Phys. Rev. D* 53 (6 Mar. 1996), pp. 2878–2894. DOI: [10.1103/PhysRevD.53.2878](https://doi.org/10.1103/PhysRevD.53.2878). URL: <https://link.aps.org/doi/10.1103/PhysRevD.53.2878> (cit. on p. 6).
- [33] Alessandra Buonanno and Yanbei Chen. “Quantum noise in second generation, signal-recycled laser interferometric gravitational-wave detectors”. In: *Phys. Rev. D* 64 (4 July 2001), p. 042006. DOI: [10.1103/PhysRevD.64.042006](https://doi.org/10.1103/PhysRevD.64.042006). URL: <https://link.aps.org/doi/10.1103/PhysRevD.64.042006> (cit. on p. 6).
- [34] Carlton M. Caves. “Quantum-Mechanical Radiation-Pressure Fluctuations in an Interferometer”. In: *Phys. Rev. Lett.* 45 (2 July 1980), pp. 75–79. DOI: [10.1103/PhysRevLett.45.75](https://doi.org/10.1103/PhysRevLett.45.75). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.45.75> (cit. on p. 6).

BIBLIOGRAPHY

- [35] Vladimir B. Braginsky, Farid Ya Khalili, and Kip S. Thorne. *Quantum Measurement*. 1992 (cit. on p. 6).
- [36] M. Tse et al. “Quantum-Enhanced Advanced LIGO Detectors in the Era of Gravitational-Wave Astronomy”. In: *Phys. Rev. Lett.* 123 (23 Dec. 2019), p. 231107. DOI: [10.1103/PhysRevLett.123.231107](https://doi.org/10.1103/PhysRevLett.123.231107). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.123.231107> (cit. on p. 6).
- [37] Carlton M. Caves and Bonny L. Schumaker. “New formalism for two-photon quantum optics. I. Quadrature phases and squeezed states”. In: *Phys. Rev. A* 31 (5 May 1985), pp. 3068–3092. DOI: [10.1103/PhysRevA.31.3068](https://doi.org/10.1103/PhysRevA.31.3068). URL: <https://link.aps.org/doi/10.1103/PhysRevA.31.3068> (cit. on p. 6).
- [38] C Markakis et al. “Neutron star equation of state via gravitational wave observations”. In: *Journal of Physics: Conference Series* 189 (Oct. 2009), p. 012024. DOI: [10.1088/1742-6596/189/1/012024](https://doi.org/10.1088/1742-6596/189/1/012024). URL: <https://doi.org/10.1088/1742-6596/189/1/012024> (cit. on p. 6).
- [39] M. Evans et al. “Realistic filter cavities for advanced gravitational wave detectors”. In: *Phys. Rev. D* 88 (2 July 2013), p. 022002. DOI: [10.1103/PhysRevD.88.022002](https://doi.org/10.1103/PhysRevD.88.022002). URL: <https://link.aps.org/doi/10.1103/PhysRevD.88.022002> (cit. on p. 7).
- [40] F. Ya. Khalili. “Optimal configurations of filter cavity in future gravitational-wave detectors”. In: *Phys. Rev. D* 81 (12 June 2010), p. 122002. DOI: [10.1103/PhysRevD.81.122002](https://doi.org/10.1103/PhysRevD.81.122002). URL: <https://link.aps.org/doi/10.1103/PhysRevD.81.122002> (cit. on p. 7).
- [41] P. Kwee et al. “Decoherence and degradation of squeezed states in quantum filter cavities”. In: *Phys. Rev. D* 90 (6 Sept. 2014), p. 062006. DOI: [10.1103/PhysRevD.90.062006](https://doi.org/10.1103/PhysRevD.90.062006). URL: <https://link.aps.org/doi/10.1103/PhysRevD.90.062006> (cit. on p. 7).
- [42] David Reitze. *The Future of Ground-based Gravitational-wave Detectors*. Tech. rep. G1800292. 2018. URL: <https://dcc.ligo.org/LIGO-G1800292/public> (cit. on p. 7).
- [43] L. McCuller et al. “Frequency-Dependent Squeezing for Advanced LIGO”. In: *Phys. Rev. Lett.* 124 (17 Apr. 2020), p. 171102. DOI: [10.1103/PhysRevLett.124.171102](https://doi.org/10.1103/PhysRevLett.124.171102). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.124.171102> (cit. on p. 7).
- [44] ET Steering Committee Editorial Team. *Einstein Telescope design report update 2020*. Einstein Telescope Collaboration, Sept. 2020. URL: <https://apps.et-gw.eu/tds/ql/?c=15418> (cit. on pp. 7, 8, 67, 68, 69, 71, 83, 84, 92, 101, 102).
- [45] David Reitze et al. “Cosmic Explorer: The U.S. Contribution to Gravitational-Wave Astronomy beyond LIGO”. In: 51 (July 2019) (cit. on pp. 7, 102).
- [46] V.B. Braginsky and S.P. Vyatchanin. “Thermodynamical fluctuations in optical mirror coatings”. In: *Physics Letters A* 312.3 (2003), pp. 244–255. ISSN: 0375-9601. DOI: [https://doi.org/10.1016/S0375-9601\(03\)00473-0](https://doi.org/10.1016/S0375-9601(03)00473-0). URL: <https://www.sciencedirect.com/science/article/pii/S0375960103004730> (cit. on p. 7).

-
- [47] J. Steinlechner. “Development of mirror coatings for gravitational-wave detectors”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 376.2120 (2018), p. 20170282. DOI: [10.1098/rsta.2017.0282](https://doi.org/10.1098/rsta.2017.0282). eprint: <https://royalsocietypublishing.org/doi/pdf/10.1098/rsta.2017.0282>. URL: <https://royalsocietypublishing.org/doi/abs/10.1098/rsta.2017.0282> (cit. on pp. 7, 8).
- [48] William Yam, Slawek Gras, and Matthew Evans. “Multimaterial coatings with reduced thermal noise”. In: *Phys. Rev. D* 91 (4 Feb. 2015), p. 042002. DOI: [10.1103/PhysRevD.91.042002](https://doi.org/10.1103/PhysRevD.91.042002). URL: <https://link.aps.org/doi/10.1103/PhysRevD.91.042002> (cit. on p. 7).
- [49] Gregory M Harry et al. “Thermal noise in interferometric gravitational wave detectors due to dielectric optical coatings”. In: *Classical and Quantum Gravity* 19.5 (Feb. 2002), pp. 897–917. DOI: [10.1088/0264-9381/19/5/305](https://doi.org/10.1088/0264-9381/19/5/305). URL: <https://doi.org/10.1088/0264-9381/19/5/305> (cit. on p. 7).
- [50] S M Aston et al. “Update on quadruple suspension design for Advanced LIGO”. In: *Classical and Quantum Gravity* 29.23 (Oct. 2012), p. 235004. DOI: [10.1088/0264-9381/29/23/235004](https://doi.org/10.1088/0264-9381/29/23/235004). URL: <https://doi.org/10.1088/0264-9381/29/23/235004> (cit. on p. 8).
- [51] F Matichard et al. “Seismic isolation of Advanced LIGO: Review of strategy, instrumentation and performance”. In: *Classical and Quantum Gravity* 32.18 (Aug. 2015), p. 185003. DOI: [10.1088/0264-9381/32/18/185003](https://doi.org/10.1088/0264-9381/32/18/185003). URL: <https://doi.org/10.1088/0264-9381/32/18/185003> (cit. on p. 8).
- [52] C M Mow-Lowry and D Martynov. “A 6D interferometric inertial isolation system”. In: *Classical and Quantum Gravity* 36.24 (Nov. 2019), p. 245006. DOI: [10.1088/1361-6382/ab4e01](https://doi.org/10.1088/1361-6382/ab4e01). URL: <https://doi.org/10.1088/1361-6382/ab4e01> (cit. on pp. 8, 101).
- [53] Herbert B. Callen and Theodore A. Welton. “Irreversibility and Generalized Noise”. In: *Phys. Rev.* 83 (1 July 1951), pp. 34–40. DOI: [10.1103/PhysRev.83.34](https://doi.org/10.1103/PhysRev.83.34). URL: <https://link.aps.org/doi/10.1103/PhysRev.83.34> (cit. on p. 8).
- [54] Peter R. Saulson. “Thermal noise in mechanical experiments”. In: *Phys. Rev. D* 42 (8 Oct. 1990), pp. 2437–2445. DOI: [10.1103/PhysRevD.42.2437](https://doi.org/10.1103/PhysRevD.42.2437). URL: <https://link.aps.org/doi/10.1103/PhysRevD.42.2437> (cit. on p. 8).
- [55] T. Akutsu et al. “The status of KAGRA underground cryogenic gravitational wave telescope”. In: *Journal of Physics: Conference Series* 1342 (Jan. 2020), p. 012014. DOI: [10.1088/1742-6596/1342/1/012014](https://doi.org/10.1088/1742-6596/1342/1/012014). URL: <https://doi.org/10.1088/1742-6596/1342/1/012014> (cit. on p. 8).
- [56] C. Bond et al. “Interferometer techniques for gravitational-wave detection”. In: *Living Reviews in Relativity* 19 (Feb. 2017). DOI: [10.1007/s41114-016-0002-8](https://doi.org/10.1007/s41114-016-0002-8) (cit. on pp. 10, 11, 12, 15, 17, 20, 28, 29, 31, 41, 52, 58, 70).
- [57] H. Kogelnik and T. Li. “Laser beams and resonators”. In: *Proceedings of the IEEE* 54.10 (1966), pp. 1312–1329. DOI: [10.1109/PROC.1966.5119](https://doi.org/10.1109/PROC.1966.5119) (cit. on pp. 10, 14).

BIBLIOGRAPHY

- [58] A.E. Siegman. *Lasers*. University Science Books, 1986. ISBN: 9780935702118. URL: <https://books.google.co.uk/books?id=1BZVwUZLTkAC> (cit. on pp. 10, 11, 12, 14, 28, 29, 31, 33).
- [59] Herwig Kogelnik. “On the Propagation of Gaussian Beams of Light Through Lenslike Media Including those with a Loss or Gain Variation”. In: *Appl. Opt.* 4.12 (Dec. 1965), pp. 1562–1569. DOI: [10.1364/AO.4.001562](https://doi.org/10.1364/AO.4.001562). URL: <http://ao.osa.org/abstract.cfm?URI=ao-4-12-1562> (cit. on pp. 10, 16, 29, 70).
- [60] (Virgo Collaboration) J.Y. Vinet. *The Virgo Physics Book, Vol. II: Optics and Related Topics*. 2006 (cit. on pp. 10, 60).
- [61] Jean-Yves Vinet. “On Special Optical Modes and Thermal Issues in Advanced Gravitational Wave Interferometric Detectors”. In: *Living Reviews in Relativity* 12.1 (Dec. 2009), p. 5. ISSN: 1433-8351. DOI: [10.12942/lrr-2009-5](https://doi.org/10.12942/lrr-2009-5). URL: <https://doi.org/10.12942/lrr-2009-5> (cit. on p. 10).
- [62] E. Abramochkin and V. Volostnikov. “Beam transformations and nontransformed beams”. In: *Optics Communications* 83.1-2 (May 1991), pp. 123–135. DOI: [10.1016/0030-4018\(91\)90534-K](https://doi.org/10.1016/0030-4018(91)90534-K) (cit. on p. 10).
- [63] D. Speiser. *Discovering the Principles of Mechanics 1600-1800*. Birkhauser Basel, 2008. ISBN: 978-3-7643-8564-4. URL: <https://www.springer.com/gp/book/9783764385644> (cit. on p. 10).
- [64] Joseph W. Goodman. *Introduction to Fourier optics*. 1995 (cit. on p. 11).
- [65] Melvin Lax, William H. Louisell, and William B. McKnight. “From Maxwell to paraxial wave optics”. In: *Phys. Rev. A* 11 (4 Apr. 1975), pp. 1365–1370. DOI: [10.1103/PhysRevA.11.1365](https://doi.org/10.1103/PhysRevA.11.1365). URL: <https://link.aps.org/doi/10.1103/PhysRevA.11.1365> (cit. on p. 11).
- [66] Robert W. Boyd. “Intuitive explanation of the phase anomaly of focused light beams”. In: *J. Opt. Soc. Am.* 70.7 (July 1980), pp. 877–880. DOI: [10.1364/JOSA.70.000877](https://doi.org/10.1364/JOSA.70.000877). URL: <http://www.osapublishing.org/abstract.cfm?URI=josa-70-7-877> (cit. on p. 15).
- [67] P. Kwee et al. “Stabilized high-power laser system for the gravitational wave detector advanced LIGO”. In: *Opt. Express* 20.10 (May 2012), pp. 10617–10634. DOI: [10.1364/OE.20.010617](https://doi.org/10.1364/OE.20.010617). URL: <http://www.opticsexpress.org/abstract.cfm?URI=oe-20-10-10617> (cit. on p. 17).
- [68] Fabian Thies et al. “Nd:YVO4 high-power master oscillator power amplifier laser system for second-generation gravitational wave detectors”. In: *Opt. Lett.* 44.3 (Feb. 2019), pp. 719–722. DOI: [10.1364/OL.44.000719](https://doi.org/10.1364/OL.44.000719). URL: <http://ol.osa.org/abstract.cfm?URI=ol-44-3-719> (cit. on p. 17).
- [69] A. Householder. “Some Numerical Methods for Solving Systems of Linear Equations”. In: *American Mathematical Monthly* 57 (1950), p. 453 (cit. on p. 17).
- [70] G.H. Golub et al. *Matrix Computations*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, 1996. ISBN: 9780801854149. URL: <https://books.google.co.uk/books?id=mlOa7wPX6OYC> (cit. on p. 17).

-
- [71] L.N. Trefethen and D. Bau. *Numerical Linear Algebra*. Other Titles in Applied Mathematics. Society for Industrial and Applied Mathematics (SIAM, 3600 Market Street, Floor 6, Philadelphia, PA 19104), 1997. ISBN: 9780898719574. URL: <https://books.google.co.uk/books?id=JaPtxOytY7kC> (cit. on p. 17).
- [72] Timothy A. Davis. “Front Matter”. In: *Direct Methods for Sparse Linear Systems*, pp. i–xii. DOI: [10.1137/1.9780898718881.fm](https://doi.org/10.1137/1.9780898718881.fm). eprint: <https://epubs.siam.org/doi/pdf/10.1137/1.9780898718881.fm>. URL: <https://epubs.siam.org/doi/abs/10.1137/1.9780898718881.fm> (cit. on p. 17).
- [73] Daniel David Brown and Andreas Freise. *Finesse*. You can download the binaries and source code at <http://www.gwoptics.org/finesse>. May 2014. DOI: [10.5281/zenodo.821364](https://doi.org/10.5281/zenodo.821364). URL: <https://doi.org/10.5281/zenodo.821364> (cit. on pp. 17, 26, 64).
- [74] Daniel D. Brown et al. “Pykat: Python package for modelling precision optical interferometers”. In: *SoftwareX* 12 (2020), p. 100613. ISSN: 2352-7110. DOI: <https://doi.org/10.1016/j.softx.2020.100613>. URL: <http://www.sciencedirect.com/science/article/pii/S2352711020303265> (cit. on pp. 17, 26, 55, 56, 64, 123).
- [75] Daniel Brown et al. *Finesse 3*. <https://git.ligo.org/finesse/finesse3>. 2020 (cit. on pp. 17, 70, 123, 125).
- [76] Andreas Freise, Daniel Brown, and Charlotte Bond. *Finesse 2.0 User manual*. 2014 (cit. on pp. 20, 21, 37, 52, 59, 61).
- [77] Aydın Buluç et al. “Parallel sparse matrix-vector and matrix-transpose-vector multiplication using compressed sparse blocks”. In: *IN SPAA*. 2009, pp. 233–244 (cit. on p. 22).
- [78] I. S. Duff, Roger G. Grimes, and John G. Lewis. “Sparse Matrix Test Problems”. In: *ACM Trans. Math. Softw.* 15.1 (Mar. 1989), pp. 1–14. ISSN: 0098-3500. DOI: [10.1145/62038.62043](https://doi.org/10.1145/62038.62043). URL: <https://doi.org/10.1145/62038.62043> (cit. on p. 22).
- [79] Timothy A. Davis and Ekanathan Palamadai Natarajan. “Algorithm 907: KLU, A Direct Sparse Solver for Circuit Simulation Problems”. In: *ACM Trans. Math. Softw.* 37 (2010), 36:1–36:17 (cit. on pp. 22, 127).
- [80] D Brown. “Interactions of light and mirrors: Advanced techniques for modelling future gravitational wave detectors”. PhD thesis. University of Birmingham, 2015 (cit. on pp. 22, 128).
- [81] Samuel Rowlinson et al. *Finesse 3 Documentation*. <https://finesse.docs.ligo.org/finesse3/>. 2020 (cit. on pp. 25, 29, 89, 123, 124, 131, 135).
- [82] Hello, Patrice and Vinet, Jean-Yves. “Analytical models of thermal aberrations in massive mirrors heated by high power laser beams”. In: *J. Phys. France* 51.12 (1990), pp. 1267–1282. DOI: [10.1051/jphys:0199000510120126700](https://doi.org/10.1051/jphys:0199000510120126700). URL: <https://doi.org/10.1051/jphys:0199000510120126700> (cit. on pp. 25, 83).

BIBLIOGRAPHY

- [83] Antonello Cutolo, Paolo Gay, and Salvatore Solimeno. “Mirror Deformations and Wavefront Aberrations Caused by C.w. High Power Laser Beams”. In: *Optica Acta: International Journal of Optics* 27.8 (1980), pp. 1105–1116. DOI: [10.1080/713820360](https://doi.org/10.1080/713820360). eprint: <https://doi.org/10.1080/713820360>. URL: <https://doi.org/10.1080/713820360> (cit. on p. 25).
- [84] L Barsotti, M Evans, and P Fritschel. “Alignment sensing and control in advanced LIGO”. In: *Classical and Quantum Gravity* 27.8 (Apr. 2010), p. 084026. DOI: [10.1088/0264-9381/27/8/084026](https://doi.org/10.1088/0264-9381/27/8/084026). URL: <https://doi.org/10.1088/0264-9381/27/8/084026> (cit. on pp. 25, 101, 119).
- [85] M. Evans, L. Barsotti, and P. Fritschel. “A general approach to optomechanical parametric instabilities”. In: *Physics Letters A* 374.4 (2010), pp. 665–671. ISSN: 0375-9601. DOI: <https://doi.org/10.1016/j.physleta.2009.11.023>. URL: <https://www.sciencedirect.com/science/article/pii/S0375960109014558> (cit. on p. 25).
- [86] Matthew Evans et al. “Observation of Parametric Instability in Advanced LIGO”. In: *Phys. Rev. Lett.* 114 (16 Apr. 2015), p. 161102. DOI: [10.1103/PhysRevLett.114.161102](https://doi.org/10.1103/PhysRevLett.114.161102). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.114.161102> (cit. on p. 25).
- [87] A C Green et al. “The influence of dual-recycling on parametric instabilities at Advanced LIGO”. In: *Classical and Quantum Gravity* 34.20 (Sept. 2017), p. 205004. DOI: [10.1088/1361-6382/aa8af8](https://doi.org/10.1088/1361-6382/aa8af8). URL: <https://doi.org/10.1088/1361-6382/aa8af8> (cit. on p. 25).
- [88] Aidan F. Brooks et al. *Point absorbers in Advanced LIGO*. 2021. arXiv: [2101.05828](https://arxiv.org/abs/2101.05828) [[physics.ins-det](https://arxiv.org/abs/2101.05828)] (cit. on p. 25).
- [89] J. Aasi et al. “Enhanced sensitivity of the LIGO gravitational wave detector by using squeezed states of light”. In: *Nature Photonics* 7.8 (Aug. 2013), pp. 613–619. ISSN: 1749-4893. DOI: [10.1038/nphoton.2013.177](https://doi.org/10.1038/nphoton.2013.177). URL: <https://doi.org/10.1038/nphoton.2013.177> (cit. on p. 25).
- [90] E. Oelker et al. “Squeezed light for advanced gravitational wave detectors and beyond”. In: *Opt. Express* 22.17 (Aug. 2014), pp. 21106–21121. DOI: [10.1364/OE.22.021106](https://doi.org/10.1364/OE.22.021106). URL: <http://www.opticsexpress.org/abstract.cfm?URI=oe-22-17-21106> (cit. on p. 25).
- [91] Matthew Evans. *Optickle Technical Report*. Tech. rep. T070260-v1. 2007. URL: <https://dcc.ligo.org/T070260/public> (cit. on p. 26).
- [92] Matthew Evans and Nic Smith. *Optickle, Lentickle and Pickle*. URL: https://awiki.ligo-wa.caltech.edu/wiki/ISC_Modeling_Software (cit. on p. 26).
- [93] Hiro Yamamoto. *Stationary Interferometer Simulation*. Tech. rep. G0900912. 2009. URL: <https://dcc.ligo.org/LIGO-G0900912/public> (cit. on p. 26).
- [94] Hiro Yamamoto. *FOGPrime 13 Manual*. Tech. rep. T1300942. 2018. URL: <https://dcc.ligo.org/LIGO-T1300942> (cit. on p. 26).
- [95] Jérôme Degallaix. “OSCAR a Matlab based optical FFT code”. In: *Journal of Physics: Conference Series* 228 (May 2010), p. 012021. DOI: [10.1088/1742-6596/228/1/012021](https://doi.org/10.1088/1742-6596/228/1/012021). URL: <https://doi.org/10.1088/1742-6596/228/1/012021> (cit. on p. 26).

- [96] Pablo Barriga et al. “Numerical calculations of diffraction losses in advanced interferometric gravitational wave detectors”. In: *J. Opt. Soc. Am. A* 24.6 (June 2007), pp. 1731–1741. DOI: [10.1364/JOSAA.24.001731](https://doi.org/10.1364/JOSAA.24.001731). URL: <http://josaa.osa.org/abstract.cfm?URI=josaa-24-6-1731> (cit. on p. 26).
- [97] D Brown, R J E Smith, and A Freise. “Fast simulation of Gaussian-mode scattering for precision interferometry”. In: *Journal of Optics* 18.2 (Jan. 2016), p. 025604. DOI: [10.1088/2040-8978/18/2/025604](https://doi.org/10.1088/2040-8978/18/2/025604). URL: <https://doi.org/10.1088/2040-8978/18/2/025604> (cit. on p. 27).
- [98] Wen Qiao et al. “Simple method of optical ring cavity design and its applications”. In: *Opt. Express* 22.12 (June 2014), pp. 14782–14791. DOI: [10.1364/OE.22.014782](https://doi.org/10.1364/OE.22.014782). URL: <http://www.opticsexpress.org/abstract.cfm?URI=oe-22-12-14782> (cit. on p. 30).
- [99] Koji Arai. *On the accumulated round-trip Gouy phase shift for a general optical cavity*. <https://dcc.ligo.org/LIGO-T1300189/public>. 2015. URL: <https://dcc.ligo.org/LIGO-T1300189/public> (cit. on p. 32).
- [100] A. E. Siegman. “Unstable Optical Resonators”. In: *Appl. Opt.* 13.2 (Feb. 1974), pp. 353–367. DOI: [10.1364/AO.13.000353](https://doi.org/10.1364/AO.13.000353). URL: <http://ao.osa.org/abstract.cfm?URI=ao-13-2-353> (cit. on pp. 32, 148).
- [101] H Wang. “Beware of Warped Surfaces: Near-Unstable Cavities for Future Gravitational Wave Detectors”. PhD thesis. University of Birmingham, 2017 (cit. on p. 32).
- [102] “The Fabry Perot Resonator”. In: *Laser Resonators and Beam Propagation: Fundamentals, Advanced Concepts and Applications*. New York, NY: Springer New York, 2005, pp. 189–215. ISBN: 978-0-387-25110-3. DOI: [10.1007/0-387-25110-3_5](https://doi.org/10.1007/0-387-25110-3_5). URL: https://doi.org/10.1007/0-387-25110-3_5 (cit. on p. 33).
- [103] Aric A. Hagberg, Daniel A. Schult, and Pieter J. Swart. “Exploring Network Structure, Dynamics, and Function using NetworkX”. In: *Proceedings of the 7th Python in Science Conference*. Ed. by Gaël Varoquaux, Travis Vaught, and Jarrod Millman. Pasadena, CA USA, 2008, pp. 11–15 (cit. on pp. 36, 128).
- [104] Donald E. Knuth. *The Art of Computer Programming, Volume 1 (3rd Ed.): Fundamental Algorithms*. USA: Addison Wesley Longman Publishing Co., Inc., 1997. ISBN: 0201896834 (cit. on p. 36).
- [105] C Mueller. “Techniques for Resonant Optical Interferometry with Applications to the Advanced LIGO Gravitational Wave Detectors”. PhD thesis. University of Florida, 2014 (cit. on p. 39).
- [106] L. McCuller, S. Biscans, and L. Barsotti. *Frequency Dependent Squeezing Final Optical Layout*. Tech. rep. T1900649. 2019. URL: <https://dcc.ligo.org/LIGO-T1900649> (cit. on p. 48).
- [107] Charles R. Harris et al. “Array programming with NumPy”. In: *Nature* 585.7825 (Sept. 2020), pp. 357–362. DOI: [10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2). URL: <https://doi.org/10.1038/s41586-020-2649-2> (cit. on pp. 48, 148).

BIBLIOGRAPHY

- [108] Samuel Rowlinson. *Finesse 3 Documentation Beam propagation*. https://finesse.docs.ligo.org/finesse3/usage/homs/propagating_beams.html. 2020 (cit. on p. 48).
- [109] F. Bayer-Helms. “Coupling coefficients of an incident wave and the modes of a spherical optical resonator in the case of mismatching and misalignment”. In: *Appl. Opt.* 23.9 (May 1984), pp. 1369–1380. DOI: [10.1364/AO.23.001369](https://doi.org/10.1364/AO.23.001369). URL: <http://ao.osa.org/abstract.cfm?URI=ao-23-9-1369> (cit. on pp. 51, 52, 59, 60, 65, 105, 127).
- [110] Dana Z. Anderson. “Alignment of resonant optical cavities”. In: *Appl. Opt.* 23.17 (Sept. 1984), pp. 2944–2949. DOI: [10.1364/AO.23.002944](https://doi.org/10.1364/AO.23.002944). URL: <http://ao.osa.org/abstract.cfm?URI=ao-23-17-2944> (cit. on p. 58).
- [111] Mischa Sallé. *Shift and phase change from tilted beamsplitters - part 1: theory*. <https://logbooks.ifosim.org/pykat/blog/shift-and-phase-change-from-tilted-beamsplitters-part-1-theory/>. 2021 (cit. on p. 60).
- [112] Mischa Sallé. *Shift and phase change from tilted beamsplitters - part 2: finesse examples*. <https://logbooks.ifosim.org/pykat/blog/shift-and-phase-change-from-tilted-beamsplitters-part-2-finesse-examples/>. 2021 (cit. on p. 60).
- [113] Samuel Rowlinson. *Finesse 3 Documentation Animated Beam Images*. https://finesse.docs.ligo.org/finesse3/usage/homs/imaging_beams.html#producing-animated-beam-images. 2020 (cit. on p. 62).
- [114] Leonardo Dagum and Ramesh Menon. “OpenMP: An Industry-Standard API for Shared-Memory Programming”. In: *IEEE Comput. Sci. Eng.* 5.1 (Jan. 1998), pp. 46–55. ISSN: 1070-9924. DOI: [10.1109/99.660313](https://doi.org/10.1109/99.660313). URL: <https://doi.org/10.1109/99.660313> (cit. on p. 62).
- [115] Michele Maggiore et al. “Science case for the Einstein telescope”. In: *Journal of Cosmology and Astroparticle Physics* 2020.03 (Mar. 2020), pp. 050–050. DOI: [10.1088/1475-7516/2020/03/050](https://doi.org/10.1088/1475-7516/2020/03/050). URL: <https://doi.org/10.1088/1475-7516/2020/03/050> (cit. on p. 67).
- [116] Alessandra Buonanno and Yanbei Chen. “Scaling law in signal recycled laser interferometer gravitational-wave detectors”. In: *Phys. Rev. D* 67 (6 Mar. 2003), p. 062002. DOI: [10.1103/PhysRevD.67.062002](https://doi.org/10.1103/PhysRevD.67.062002). URL: <https://link.aps.org/doi/10.1103/PhysRevD.67.062002> (cit. on p. 67).
- [117] C Buy et al. “Design of a high-magnification and low-aberration compact catadioptric telescope for the Advanced Virgo gravitational-wave interferometric detector”. In: *Classical and Quantum Gravity* 34.9 (Apr. 2017), p. 095011. DOI: [10.1088/1361-6382/aa65e3](https://doi.org/10.1088/1361-6382/aa65e3). URL: <https://doi.org/10.1088/1361-6382/aa65e3> (cit. on p. 68).
- [118] Teng Zhang et al. “Effects of static and dynamic higher-order optical modes in balanced homodyne readout for future gravitational waves detectors”. In: *Phys. Rev. D* 95 (6 Mar. 2017), p. 062001. DOI: [10.1103/PhysRevD.95.062001](https://doi.org/10.1103/PhysRevD.95.062001). URL: <https://link.aps.org/doi/10.1103/PhysRevD.95.062001> (cit. on p. 69).

-
- [119] Riccardo DeSalvo. *Considerations on Michelson beam splitters for third generation Gravitational Wave Observatories*. <https://dcc.ligo.org/LIGO-G1900927/public>. 2020. URL: <https://dcc.ligo.org/LIGO-G1900927/public> (cit. on p. 69).
- [120] Riccardo DeSalvo et al. *Beam expander telescopes for the Michelson beam splitters in third generation Gravitational Wave Observatories*. paper in preparation. 2020 (cit. on p. 69).
- [121] A Dmitriev et al. *Optmisation of beam expansion telescopes for third-generation gravitational wave detectors*. (in preparation). 2020 (cit. on p. 70).
- [122] The ET Science Team. *Einstein gravitational wave Telescope conceptual design*. European Commission, June 2011. URL: <http://www.et-gw.eu/index.php/relevant-et-documents> (cit. on pp. 71, 72, 83).
- [123] Philip Jones et al. “Implications of the quantum noise target for the Einstein Telescope infrastructure design”. In: *Phys. Rev. D* 101 (8 Apr. 2020), p. 082002. DOI: [10.1103/PhysRevD.101.082002](https://doi.org/10.1103/PhysRevD.101.082002). URL: <https://link.aps.org/doi/10.1103/PhysRevD.101.082002> (cit. on pp. 71, 85).
- [124] M Granata et al. “Design of the Advanced Virgo non-degenerate recycling cavities”. In: *Journal of Physics: Conference Series* 228 (May 2010), p. 012016. DOI: [10.1088/1742-6596/228/1/012016](https://doi.org/10.1088/1742-6596/228/1/012016). URL: <https://doi.org/10.1088/1742-6596/228/1/012016> (cit. on p. 72).
- [125] H Lück et al. “The upgrade of GEO 600”. In: *Journal of Physics: Conference Series* 228 (May 2010), p. 012012. DOI: [10.1088/1742-6596/228/1/012012](https://doi.org/10.1088/1742-6596/228/1/012012). URL: <https://doi.org/10.1088/1742-6596/228/1/012012> (cit. on p. 72).
- [126] Aidan F. Brooks et al. “Overview of Advanced LIGO adaptive optics”. In: *Appl. Opt.* 55.29 (Oct. 2016), pp. 8256–8265. DOI: [10.1364/AO.55.008256](https://doi.org/10.1364/AO.55.008256). URL: <http://ao.osa.org/abstract.cfm?URI=ao-55-29-8256> (cit. on pp. 83, 87, 92, 93, 94).
- [127] A Rocchi et al. “Thermal effects and their compensation in Advanced Virgo”. In: *Journal of Physics: Conference Series* 363 (June 2012), p. 012016. DOI: [10.1088/1742-6596/363/1/012016](https://doi.org/10.1088/1742-6596/363/1/012016). URL: <https://doi.org/10.1088/1742-6596/363/1/012016> (cit. on p. 83).
- [128] A. F. Brooks et al. *Active wavefront control in and beyond Advanced LIGO*. Tech. rep. T1500188–v1. 2015. URL: <https://dcc.ligo.org/LIGO-T1500188/public> (cit. on pp. 87, 88).
- [129] Antonio Perreca et al. “Analysis and visualization of the output mode-matching requirements for squeezing in Advanced LIGO and future gravitational wave detectors”. In: *Phys. Rev. D* 101 (10 May 2020), p. 102005. DOI: [10.1103/PhysRevD.101.102005](https://doi.org/10.1103/PhysRevD.101.102005). URL: <https://link.aps.org/doi/10.1103/PhysRevD.101.102005> (cit. on pp. 87, 88, 89, 91, 93, 94, 97, 99).
- [130] Jérôme Degallaix. “Compensation of strong thermal lensing in advanced interferometric gravitational waves detectors”. PhD thesis. University of Western Australia, 2006 (cit. on p. 92).

BIBLIOGRAPHY

- [131] Jérôme Degallaix. Private Communication. 2021 (cit. on p. 92).
- [132] Crystran. *Silica Glass Product Data*. <https://www.crystran.co.uk/optical-materials/silica-glass-sio2>. 2012 (cit. on p. 92).
- [133] Aidan F. Brooks. *SR3 heater test*. LLO aLOG 27262. 2016. URL: <https://alog.ligo-la.caltech.edu/aLOG/index.php?callRep=27262> (cit. on p. 94).
- [134] Aidan Brooks. Private Communication. 2021 (cit. on p. 94).
- [135] Michael L. Waskom. “seaborn: statistical data visualization”. In: *Journal of Open Source Software* 6.60 (2021), p. 3021. DOI: [10.21105/joss.03021](https://doi.org/10.21105/joss.03021). URL: <https://doi.org/10.21105/joss.03021> (cit. on p. 98).
- [136] D. V. Martynov et al. “Sensitivity of the Advanced LIGO detectors at the beginning of gravitational wave astronomy”. In: *Phys. Rev. D* 93 (11 June 2016), p. 112004. DOI: [10.1103/PhysRevD.93.112004](https://doi.org/10.1103/PhysRevD.93.112004). URL: <https://link.aps.org/doi/10.1103/PhysRevD.93.112004> (cit. on p. 101).
- [137] Katherine L. Dooley et al. “Angular control of optical cavities in a radiation-pressure-dominated regime: the Enhanced LIGO case”. In: *J. Opt. Soc. Am. A* 30.12 (Dec. 2013), pp. 2618–2626. DOI: [10.1364/JOSAA.30.002618](https://doi.org/10.1364/JOSAA.30.002618). URL: <http://josaa.osa.org/abstract.cfm?URI=josaa-30-12-2618> (cit. on p. 101).
- [138] Eiichi Hirose et al. “Angular instability due to radiation pressure in the LIGO gravitational wave detector”. In: *Appl. Opt.* 49.18 (June 2010), pp. 3474–3484. DOI: [10.1364/AO.49.003474](https://doi.org/10.1364/AO.49.003474). URL: <http://ao.osa.org/abstract.cfm?URI=ao-49-18-3474> (cit. on p. 101).
- [139] John A. Sidles and Daniel Sigg. “Optical torques in suspended Fabry–Perot interferometers”. In: *Physics Letters A* 354.3 (2006), pp. 167–172. ISSN: 0375-9601. DOI: <https://doi.org/10.1016/j.physleta.2006.01.051>. URL: <https://www.sciencedirect.com/science/article/pii/S0375960106001381> (cit. on pp. 101, 121).
- [140] Brett Shapiro et al. “Cryogenically cooled ultra low vibration silicon mirrors for gravitational wave observatories”. In: *Cryogenics* 81 (2017), pp. 83–92. ISSN: 0011-2275. DOI: <https://doi.org/10.1016/j.cryogenics.2016.12.004>. URL: <https://www.sciencedirect.com/science/article/pii/S0011227516302818> (cit. on p. 102).
- [141] Raffaele Flaminio. “Status and plans of the Virgo gravitational wave detector”. In: *Ground-based and Airborne Telescopes VIII*. Ed. by Heather K. Marshall, Jason Spyromilio, and Tomonori Usuda. Vol. 11445. International Society for Optics and Photonics. SPIE, 2020, pp. 205–214. DOI: [10.1117/12.2565418](https://doi.org/10.1117/12.2565418). URL: <https://doi.org/10.1117/12.2565418> (cit. on p. 102).
- [142] K. Ackley et al. “Neutron Star Extreme Matter Observatory: A kilohertz-band gravitational wave detector in the global network”. In: *Publications of the Astronomical Society of Australia* 37 (2020), e047. DOI: [10.1017/pasa.2020.39](https://doi.org/10.1017/pasa.2020.39) (cit. on p. 102).
- [143] Andreas Freise. *Alignment coupling with large beams*. Tech. rep. Dec. 2017. URL: <https://dcc.ligo.org/LIGO-G1702391> (cit. on pp. 102, 104, 108, 119, 122).

-
- [144] A Freise, A Bunkowski, and R Schnabel. “Phase and alignment noise in grating interferometers”. In: *New Journal of Physics* 9.12 (Dec. 2007), pp. 433–433. DOI: [10.1088/1367-2630/9/12/433](https://doi.org/10.1088/1367-2630/9/12/433). URL: <https://doi.org/10.1088/1367-2630/9/12/433> (cit. on p. 103).
- [145] J Aasi et al. “Advanced LIGO”. In: *Classical and Quantum Gravity* 32.7 (Mar. 2015), p. 074001. DOI: [10.1088/0264-9381/32/7/074001](https://doi.org/10.1088/0264-9381/32/7/074001). URL: <https://doi.org/10.1088/0264-9381/32/7/074001> (cit. on pp. 110, 111).
- [146] Samuel Rowlinson. *cavcalc*. <https://cavcalc.readthedocs.io/en/latest/>. 2020 (cit. on pp. 111, 113).
- [147] Coralia Cartis et al. “Improving the Flexibility and Robustness of Model-Based Derivative Free Optimization Solvers”. In: *ACM Trans. Math. Softw.* 45.3 (Aug. 2019). ISSN: 0098-3500. DOI: [10.1145/3338517](https://doi.org/10.1145/3338517). URL: <https://doi.org/10.1145/3338517> (cit. on p. 115).
- [148] D Martynov. “Lock Acquisition and Sensitivity Analysis of Advanced LIGO Interferometers”. PhD thesis. California Institute of Technology, 2015 (cit. on p. 115).
- [149] Euan Morrison et al. “Automatic alignment of optical interferometers”. In: *Appl. Opt.* 33.22 (Aug. 1994), pp. 5041–5049. DOI: [10.1364/AO.33.005041](https://doi.org/10.1364/AO.33.005041). URL: <http://ao.osa.org/abstract.cfm?URI=ao-33-22-5041> (cit. on p. 119).
- [150] A. Freise. *Introduction to Finesse: Motivation and History*. <https://finesse.docs.ligo.org/finesse3/introduction/index.html#motivation-and-history>. 2021 (cit. on p. 123).
- [151] A. Freise. *History and impact of the finesse simulation tool*. <http://www.gwoptics.org/finesse/impact.php>. 2017 (cit. on p. 123).
- [152] A Freise et al. “Frequency-domain interferometer simulation with higher-order spatial modes”. In: *Classical and Quantum Gravity* 21.5 (2004). Finesse is available at <http://www.gwoptics.org/finesse>., S1067–S1074. URL: <http://stacks.iop.org/0264-9381/21/S1067> (cit. on p. 123).
- [153] *Finesse 3 Contributors Page*. <https://git.ligo.org/finesse/finesse3/-/graphs/master>. 2021 (cit. on p. 124).
- [154] Stefan Behnel et al. “Cython: The best of both worlds”. In: *Computing in Science & Engineering* 13.2 (2011), pp. 31–39 (cit. on pp. 124, 127).
- [155] Holger Krekel et al. *pytest 6.2*. 2004. URL: <https://github.com/pytest-dev/pytest> (cit. on p. 125).
- [156] Philip Jones, Samuel Rowlinson, and Andreas Freise. *Finesse 3 Presentation LVC March 2019*. LVK. Mar. 2019. URL: <https://dcc.ligo.org/LIGO-G1900515> (cit. on p. 125).
- [157] Haixing Miao et al. “Enhancing the Bandwidth of Gravitational-Wave Detectors with Unstable Optomechanical Filters”. In: *Phys. Rev. Lett.* 115 (21 Nov. 2015), p. 211104. DOI: [10.1103/PhysRevLett.115.211104](https://doi.org/10.1103/PhysRevLett.115.211104). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.115.211104> (cit. on p. 125).

BIBLIOGRAPHY

- [158] Y. Ma et al. “Proposal for gravitational-wave detection beyond the standard quantum limit through EPR entanglement.” In: *Nature Physics* 13 (8 Aug. 2017), pp. 776–780. DOI: [10.1038/nphys4118](https://doi.org/10.1038/nphys4118). URL: <https://doi.org/10.1038/nphys4118> (cit. on p. 125).
- [159] Paul Ross. *Musings on Cython*. <https://notes-on-cython.readthedocs.io/en/latest/index.html>. 2014 (cit. on p. 127).
- [160] Paul Ross. *The Performance of Python, Cython and C on a Vector*. https://notes-on-cython.readthedocs.io/en/latest/std_dev.html. 2014 (cit. on p. 128).
- [161] Georg Brandl. *Sphinx Python Documentation Generator*. <https://www.sphinx-doc.org/en/master/>. 2021 (cit. on pp. 131, 136).
- [162] Samuel Rowlinson. *Reslate*. <https://gitlab.com/sjrowlinson/reslate>. 2021 (cit. on p. 133).
- [163] Charlotte Bond, Daniel Brown, and Andreas Freise. “Interferometer responses to gravitational waves: Comparing FINESSE simulations and analytical solutions”. In: *arXiv e-prints* (June 2013). arXiv: [1306.6752](https://arxiv.org/abs/1306.6752) [[physics.optics](https://arxiv.org/abs/1306.6752)] (cit. on p. 134).
- [164] Walter K. Kahn. “Unstable Optical Resonators”. In: *Appl. Opt.* 5.3 (Mar. 1966), pp. 407–413. DOI: [10.1364/AO.5.000407](https://doi.org/10.1364/AO.5.000407). URL: <http://ao.osa.org/abstract.cfm?URI=ao-5-3-407> (cit. on p. 148).
- [165] Lee W. Casperson and Susan D. Lunnam. “Gaussian Modes in High Loss Laser Resonators”. In: *Appl. Opt.* 14.5 (May 1975), pp. 1193–1199. DOI: [10.1364/AO.14.001193](https://doi.org/10.1364/AO.14.001193). URL: <http://ao.osa.org/abstract.cfm?URI=ao-14-5-1193> (cit. on p. 148).