FACULTY OF ENGINEERING
AND ARCHITECTURE

Calibrating gravitational wave detectors using scattered light

Arthur Callens
Student number: 02105458

Supervisor: Prof. Dr. Archisman Ghosh, Dr. Daniela Pascucci

Master's dissertation submitted in order to obtain the academic degree of Master of Science in Engineering Physics

GHENT
UNIVERSITY

Faculty of Engineering and Architecture
Academic year 2025-2026

Acknowledgements

This thesis would not have been possible without the support and guidance of many people, to whom I am deeply grateful. What follows is a brief acknowledgment of a few individuals whose contributions were invaluable in shaping this project into what it is today.

First and foremost, I would like to express my sincere gratitude to my supervisor, Dr. Daniela Pascucci, whose day-to-day guidance, patience, and expertise were invaluable at every stage of this research. Her support and insight greatly enriched both the process and the outcome of this thesis. I would also like to thank my promotor, Prof. Dr. Archisman Ghosh, for welcoming me into the Ghent Gravity Group and for his academic leadership throughout this project.

I am especially grateful to Prof. Dr. Hans Vanhaevermaet for leading the project meetings and for his continued involvement and valuable insight into my specific research topic. I would also like to thank Tom Dooney and my fellow thesis student, Stef Bordo, for their collaboration and the many discussions we shared throughout the year.

My thanks also go to the broader Ghent Gravity Group for fostering such a stimulating and collegial environment. The weekly group meetings with Wim Devis, Flor Boussier, and Jasper Vanwildemeersch were always engaging and enjoyable. Each pursued their own fascinating research, and their enthusiasm made them a source of both inspiration and friendship throughout the year.

This thesis was carried out in the context of the LIGO-Virgo-KAGRA (LVK) Collaboration, and I feel privileged to have contributed, even in a small way, to such a remarkable scientific endeavour. Working within a large international collaboration broadened my perspective enormously and taught me lessons that extend far beyond the scope of this thesis itself.

I would like to extend my heartfelt thanks to my fellow Engineering Physics classmates: (Daan, Felix, Hendrik, Simon, Emma, Sander, Tine, Evy, Rohan and Stan) without whom this journey would have been far more difficult. Their camaraderie, encouragement, and the many moments of laughter we shared made even the most challenging periods of writing this thesis a pleasure to endure.

On a more personal note, I owe my deepest gratitude to my parents, Bruno and Sophie, my sister Margot and to my family for their unwavering support and encouragement throughout my studies. Their belief in me has been a constant source of strength, especially during the more challenging moments of this journey. A special thank you to Sara, for her patience, her support, and for always being there when it mattered most.

Admission to loan

Toegang tot bruikleen

Arthur Callens, May 2026

Use of artificial intelligence

In accordance with the Faculty of Engineering and Architecture regulations for the academic year 2025–2026, this section documents the use of generative artificial intelligence (AI) tools in the preparation of this master's dissertation. The use of AI was limited to supportive purposes and carried out in full respect of the regulations concerning plagiarism, fraud, and scientific integrity.

During the research and writing process, the following generative AI systems were consulted: ChatGPT (GPT, OpenAI), GitHub Copilot, Microsoft Copilot, Google Gemini, and Claude (Anthropic). These tools were used primarily as technical and efficiency-enhancing aids. In particular, they supported programming tasks in MATLAB and Python by providing assistance with code structuring, syntax clarification, debugging suggestions, and the generation of plots. All AI-generated suggestions were carefully reviewed, tested, and, where necessary, corrected before being incorporated into the work.

Generative AI was also used to assist with LaTeX typesetting, including the formulation of mathematical expressions, the structuring of tables, and the formatting of figures. In some cases, limited language refinement was performed with AI support. However, the conceptual development, methodological design, analysis of results, and scientific reasoning presented in this thesis were carried out independently.

Furthermore, AI tools were used as a complementary extension to traditional academic search engines to help identify relevant literature and refine search queries. All referenced sources were independently verified, consulted in their original form, and critically assessed prior to inclusion.

Throughout the process, awareness was maintained regarding the potential limitations and risks of generative AI, such as inaccuracies, bias, and reliability concerns. No confidential or sensitive data were shared with these systems. The responsibility for the content of this thesis rests entirely with the author, and the work reflects independent research and critical reflection.

Arthur Callens, May 2026

Calibrating gravitational wave detectors using scattered light

by

Arthur Callens

Master's dissertation submitted in order to obtain the academic degree of Master of Science in Engineering Physics

Supervisor: Prof. Dr. Archisman Ghosh, Dr. Daniela Pascucci

Faculty of Engineering and Architecture

Ghent University

Abstract - Accurate calibration of interferometric gravitational wave detectors is essential for extracting reliable astrophysical information from observed signals. The end-mirror power transmission T_{WE} of the Fabry–Perot arm cavities in the Advanced Virgo detector is a key parameter in the optical response model, and an independent measurement technique is desirable to complement existing methods based on finesse measurements and manufacturer specifications.

This thesis develops and validates a novel calibration method that exploits controlled backscattered light generated by a piezoelectric bender mounted on the detector end bench. A beam transmitted through the end mirror is reflected off the oscillating bender surface and re-injected into the arm cavity, producing a phase-modulated signal that is recorded at both the main output photodiode (B1) and the end-bench photodiode (B8). From the ratio of the signal amplitudes at these two photodiodes, T_{WE} can be extracted without requiring knowledge of the absolute laser power or the intracavity field amplitude.

A time-domain simulation framework was constructed using the OSCAR FFT-based optical code to model the full scattering process, including tilt-induced higher-order mode coupling, the frequency-dependent cavity transfer function, and the finite optical memory of the arm cavity. The simulation demonstrates that the method recovers the injected reference value $T_{\text{WE,ref}} = 3.39 \times 10^{-6}$ with a sub-percent relative accuracy across the full range of drive frequencies tested ($f_{\text{mod}} = 1$ mHz to 50 Hz). At low drive frequencies a magnitude-only cavity pole correction is sufficient, at higher drive frequencies, where the signal frequency approaches or exceeds the cavity pole at ~ 52 Hz, the cavity storage-time phase lag is removed by a time-domain deconvolution of the B1 signal prior to the sinusoidal fit, restoring sub-percent accuracy.

An independent measurement of the cavity field decay rate is obtained from the double-exponential ring-down following a transient scattered-light injection, yielding a value consistent with the known mirror reflectivities. The tilt-induced mode coupling is validated through a Hermite–Gaussian modal decomposition of the intracavity field, confirming that the calibration signal is concentrated in the TEM$_{00}$ component within a narrow angular window around the bender zero-crossing.

The experimental hardware for the calibration (including the piezoelectric bender and a vacuum-compatible linear translation stage) has been designed, characterised, and prepared for deployment on the Advanced Virgo end bench.

Keywords - Gravitational wave observation, detector calibration, signal injection, light scattering, Fabry–Perot cavity, piezoelectric bender, Advanced Virgo

Calibrating Gravitational Wave Detectors Using Scattered Light

Arthur Callens
Supervisor: Prof. Dr. Archisman Ghosh, Dr. Daniela Pascucci
Engineering Physics, Ghent University

***Abstract*—Accurate calibration of interferometric gravitational wave detectors is essential for extracting reliable astrophysical information from observed signals. This work develops and validates a novel calibration method that exploits controlled scattered light generated by a piezoelectric bender mounted on the Advanced Virgo end bench. A small fraction of the laser light transmitted through the arm cavity end mirror is reflected off the oscillating bender surface and re-injected into the cavity, producing a modulated signal that is recorded at two photodiodes. Because the signals at these two photodiodes depend on the end-mirror power transmission T_{WE} in different ways, their ratio allows T_{WE} to be extracted without knowledge of the absolute laser power. A time-domain simulation framework demonstrates that the method recovers the injected reference value with a sub-percent accuracy across the full range of drive frequencies tested (1 mHz to 50 Hz), after applying a time-domain deconvolution that removes the cavity storage-time phase lag from the output signal. When combined with an independent laboratory measurement of the mirror transmission, this provides a direct test of the absolute strain calibration of the detector, fully independent of both the photon calibrator and the Newtonian calibrator.**

***Index Terms*—Gravitational wave observation, detector calibration, signal injection, light scattering, Fabry–Perot cavity, piezoelectric bender, Advanced Virgo**

I. Introduction

The direct detection of gravitational waves by the LIGO and Virgo collaborations in 2015 [1] opened a new observational window onto the universe, enabling the study of astrophysical phenomena from binary black hole mergers to neutron star collisions. The scientific output of these observations depends critically on calibration: the process of converting the raw detector output into a calibrated gravitational wave strain time series $h(t)$. Any systematic error in this conversion propagates directly into the inferred source parameters: masses, spins, distances, and sky locations [2]. A $1\,\%$ bias in the calibration amplitude, for example, translates directly into a $1\,\%$ error in the luminosity distance, which in turn affects cosmological parameter estimates when gravitational wave events are used as standard sirens [3].

Advanced Virgo is a laser interferometer with two 3 km Fabry–Perot arm cavities (figure 1). The laser light bounces back and forth roughly 300 times inside each cavity, amplifying the tiny length changes caused by a passing gravitational wave. The amount of light that leaks out through the end mirror on each bounce is characterised by the end-mirror power transmission T_{WE}. This parameter plays a dual role: it determines the fraction of intracavity power reaching the monitoring photodiodes behind the end mirror, and it sets the coupling strength of any stray light re-entering the cavity from the end bench. An independent measurement of T_{WE} is therefore valuable both for detector characterisation and for modelling scattered-light noise, which is one of the most persistent noise sources in the low-frequency band of Advanced Virgo [4].

Was et al. [4] showed during the third observing run (O3) that this scattered-light coupling can be turned from a noise source into a calibration tool, using the natural seismic motion of the end bench as the scattering source. The present work extends this approach by developing a *dedicated*, controllable scattering actuator (a piezoelectric bender driven at a known frequency) and a simulation framework that models the full scattering process, including effects that arise specifically from the bender geometry.

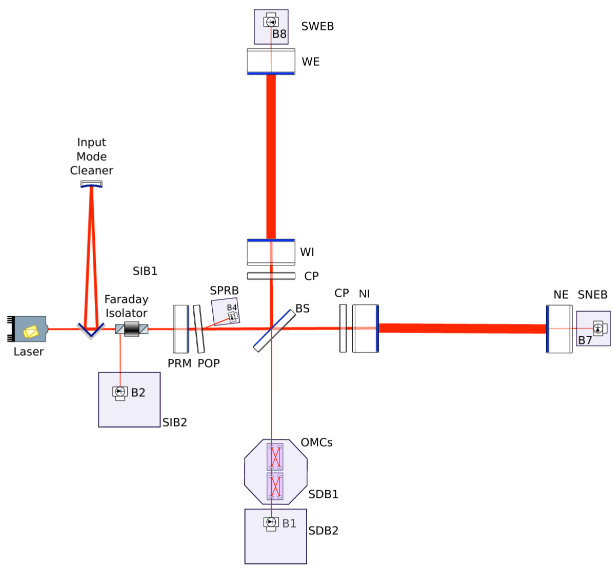

Fig. 1. Optical layout of Advanced Virgo showing the field paths relevant to the scattered-light calibration [4]. Light transmitted through the end mirror reaches the end bench, where a fraction reflects back into the arm cavity.

II. Scattered Light Model

A. Physical Principle

The calibration exploits the fact that the end mirror is not perfectly reflective: a tiny fraction $T_{\mathrm{WE}} \approx 3.4\,\mathrm{ppm}$ of the intracavity power passes through it on every round trip. This transmitted light reaches the end bench (figure 2), where a beam splitter directs part of it onto a photodiode (B8) for monitoring and part onto the piezoelectric bender. The bender surface reflects a small fraction back toward the cavity. Because this returning light has travelled an additional round-trip distance $2x(t)$ that varies with the bender displacement $x(t)$, it acquires a time-dependent phase shift $\varphi_{\rm sc}(t) = 4\pi x(t)/\lambda$ relative to the intracavity field.

When this phase-shifted light re-enters the cavity through the end mirror, it interferes with the circulating field. The interference produces power fluctuations that are recorded at both the main output photodiode (B1) at the interferometer's antisymmetric port and at B8 on the end bench. Crucially, the two signals depend on $T_{\rm WE}$ differently.

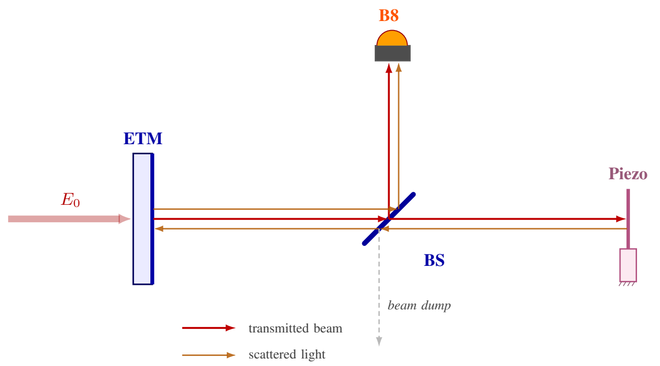

Fig. 2. Simplified diagram of the end-bench optical layout. The transmitted beam (red) passes through the end mirror (ETM) and is split by a beam splitter (BS): one path reaches the B8 photodiode, the other illuminates the piezoelectric bender. The scattered light (orange) returns along the same paths, with half re-entering the arm cavity through the ETM.

B. Signal at the Two Photodiodes

Following Was et al. [4], the scattered field returning to the cavity can be written as

$$E_{\rm sc}(t) \propto \sqrt{T_{\rm WE}}\ \sqrt{f_{r,m}}\ E_0\, e^{i\varphi_{\rm sc}(t)}, \tag{1}$$

where E_0 is the intracavity field amplitude and $f_{r,m} = f_r \cdot f_m$ is the product of the surface reflectivity f_r and the mode-matching efficiency f_m (the fraction of reflected light whose spatial profile matches the cavity's fundamental Gaussian mode). Only this mode-matched fraction builds up resonantly inside the cavity, all other spatial components are suppressed by the large transverse mode frequency splitting ($\Delta\nu_{10} \sim$ 10 kHz), which far exceeds the ~ 100 Hz cavity linewidth [5].

The scattered field re-entering the cavity is amplified by the frequency-dependent cavity gain [6], [7]

$$G(f) = \frac{G_{\rm max}}{1 + if/f_{\rm arm}}, \tag{2}$$

where $G_{\rm max} = 1/(1 - r_{\rm rt}) \simeq 151$ is the on-resonance amplitude gain, $r_{\rm rt}$ is the round-trip amplitude factor, and $f_{\rm arm} \approx 52$ Hz is the cavity pole frequency [5]. Using a signal with frequency below the pole, the gain is large and nearly constant. Above it, the gain falls as $1/f$ and acquires a phase lag because the field oscillates faster than the cavity can follow.

The interference between the scattered light and the circulating field produces measurable power fluctuations at both photodiodes. Advanced Virgo uses a detection scheme called DC readout, in which a small differential length offset ψ between the two arm cavities is introduced so that a small amount of carrier light leaks to the output port, serving as a local oscillator [6]. In this configuration, the power at the main output photodiode B1 is

$$P_{\rm B1} \approx \tfrac{1}{2} T_{\rm IM} |E_0|^2 \Big[\psi^2 + \psi\, |G|\, T_{\rm WE} \sqrt{f_{r,m}}\ \sin\varphi_{\rm sc}\Big], \tag{3}$$

where $T_{\rm IM}$ is the input mirror transmission. The first term is the constant baseline power set by the readout offset. The second is the oscillating scattered-light signal, whose amplitude depends on the product $G \cdot T_{\rm WE} \cdot \sqrt{f_{r,m}}$.

The power at the end-bench photodiode B8, which sits outside the cavity, is

$$P_{\rm B8} \approx \tfrac{1}{2} T_{\rm WE} |E_0|^2 \Big[1 + \sqrt{f_{r,m}}\ \cos\varphi_{\rm sc}\Big]. \tag{4}$$

Here the oscillating term depends on $\sqrt{f_{r,m}}$ alone, the cavity gain and the readout offset do not appear. This asymmetry is the key to the calibration: dividing signal amplitudes by baseline powers causes the intracavity power $|E_0|^2$ to cancel in both equations, and combining the two ratios isolates $T_{\rm WE}$.

C. Equivalent Strain

The scattered-light signal at B1 mimics a gravitational wave: both produce differential phase shifts between the two arm cavities that are converted into power fluctuations at the antisymmetric port. In the low-frequency limit ($f \ll f_{\rm arm}$), the scattered-light contribution can therefore be expressed as an equivalent gravitational wave strain [4]:

$$h_{\rm sc}(t) = \frac{T_{\rm WE}\sqrt{f_{r,m}}}{2L}\ \frac{\lambda}{4\pi}\ \sin\varphi_{\rm sc}(t). \tag{5}$$

This expression is independent of both the cavity gain G and the readout offset ψ, because the same detector response that amplifies a real gravitational wave equally amplifies the scattered-light signal: the two factors cancel when converting power fluctuations back to strain. This cancellation is the physical reason why the method can measure $T_{\rm WE}$ without independently knowing the cavity gain or the operating point.

If a laboratory measurement $T_{\rm WE}^{\rm LMA}$ of the mirror transmission is available, comparing the reconstructed strain to the absolutely calibrated scattered-light signal directly tests the accuracy of the detector calibration. During O3, Was et al. found this ratio to be 1.009 ± 0.05 for the west end mirror [4], consistent with unity. The dominant uncertainty arose from the 2–5 % precision of the laboratory coating measurement, not from the scattered-light technique itself.

III. Tilt-Induced Mode Coupling

An important complication arises from the bender geometry. The piezoelectric bender is a cantilever of length $\ell = 28$ mm, so a tip displacement x_{tip} is necessarily accompanied by a tilt $\theta \approx x_{\text{tip}}/\ell$ that deflects the reflected beam away from the cavity axis. This tilt imposes a linear transverse phase ramp on the scattered field [8], coupling power out of the cavity's fundamental Gaussian mode (TEM_{00}) and into higher-order Hermite–Gaussian modes [9], [10].

The fraction of scattered power that remains in TEM_{00}, and therefore builds up resonantly in the cavity, follows a Gaussian envelope:

$$f_m(\theta) = \exp\big[-(kw\theta)^2\big], \tag{6}$$

where $k = 2\pi/\lambda$ and $w \approx 5.5$ cm is the beam radius at the end mirror [11]. This coupling efficiency equals unity when the bender is perpendicular to the beam ($\theta = 0$) and drops rapidly with increasing tilt. The characteristic angular scale is $\theta_c = 1/(kw) \approx 3.1$ μrad: beyond this angle, essentially no scattered light couples back into the cavity mode.

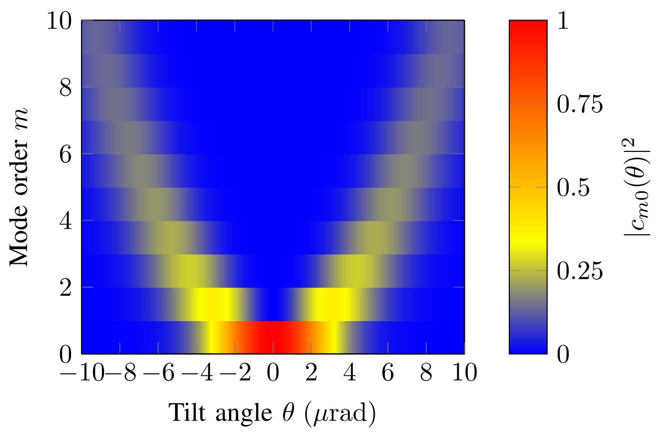

Fig. 3. Power coupled into each Hermite–Gaussian mode TEM_{m0} as a function of bender tilt angle θ (equation (6) generalised to all mode orders), for Virgo arm cavity parameters. At $\theta = 0$ all scattered power resides in TEM_{00} ($m = 0$). As $|\theta|$ increases, the distribution shifts to higher orders, and the TEM_{00} fraction drops rapidly.

During a sinusoidal bender oscillation, the tilt passes through zero twice per cycle. The calibration signal is therefore confined to narrow time windows around these zero-crossings, where $f_m \approx 1$ and the scattered field couples efficiently into the cavity (figure 3). For the realistic parameters $kw\Theta \approx 115$ and $f_{\text{mod}} = 10$ mHz, the condition $f_m > 0.99$ is satisfied within $|\delta t| < 14$ ms of the true zero-crossing, making the method robust to millisecond-level timing errors.

IV. Simulation Framework

A. OSCAR Optical Code

The simulations use OSCAR [12], [13], a MATLAB-based FFT code that represents the electromagnetic field on a two-dimensional spatial grid. Mirror interactions are modelled as spatially varying phase masks, and free-space propagation is computed in the Fourier domain. A simplified model of the Advanced Virgo west arm cavity was implemented with design parameters ($T_{\text{IM}} = 1.3\,\%$, $T_{\text{EM}} = 3.39$ ppm, $L = 3$ km) [5], [14], [15], including a 50/50 beam splitter and DC readout offset. Power recycling, signal recycling, thermal effects, and quantum noise are omitted, consistent with [4].

B. Time-Domain Cavity Dynamics

A standard frequency-domain treatment of the cavity response is impractical here because the bender tip displacement of ~ 10 μm at $\lambda = 1064$ nm produces a modulation depth $\beta = 4\pi x_{\text{tip}}/\lambda \approx 118$ [10], meaning the scattered-light signal distributes its power across hundreds of frequency sidebands. Combined with the tilt-induced spatial distortion, no small number of modes suffices.

The cavity dynamics are instead formulated as a discrete-time update [16], [7], [17]: at each time step, the intracavity field is a weighted sum of the previous field (attenuated by the round-trip losses) and a new contribution coupled through the input mirror. For compatibility with OSCAR's steady-state solver, this is cast as an effective input field:

$$E_{\text{eff}}^{(n+1)} = f_{\text{old}}\, E_{\text{eff}}^{(n)} + (1 - f_{\text{old}})\, E_{\text{new}}^{(n+1)}, \tag{7}$$

where $f_{\text{old}} = e^{-\gamma\Delta t}$ is the fraction of intracavity energy retained after one time step Δt, and $\gamma = 2\pi f_{\text{arm}} \approx 327$ s^{-1} is the cavity field decay rate. The corresponding storage time is $\tau_{\text{cav}} = 1/\gamma \simeq 3$ ms. At each step, $E_{\text{eff}}^{(n+1)}$ is passed to OSCAR as a static input, and the solver returns the intracavity field. This formulation is the main methodological contribution on the simulation side: it enables time-domain modelling of the cavity response to arbitrarily modulated scattered fields without modifying the OSCAR code.

V. Calibration Procedure

Near a zero-crossing, the baseline-subtracted B1 and B8 signals are locally sinusoidal. Their amplitudes can be written as

$$A_{\text{B1}} = \tfrac{1}{2} T_{\text{IM}} |E_0|^2\, \psi\, G\, T_{\text{WE}} \sqrt{f_{r,m}}, \tag{8}$$

$$A_{\text{B8}} = \tfrac{1}{2} T_{\text{WE}} |E_0|^2 \sqrt{f_{r,m}}. \tag{9}$$

The procedure to extract T_{WE} has four steps.

Step 1: determine $f_{r,m}$ from B8. Dividing A_{B8} by the B8 baseline power causes T_{WE} and $|E_0|^2$ to cancel:

$$f_{r,m} = \left(A_{\text{B8}}/P_{\text{B8,base}}\right)^2. \tag{10}$$

Step 2: verify ψ. The DC readout offset is cross-checked from the B1 baseline power as a consistency check on the operating point.

Step 3: recover T_{WE}. Dividing A_{B1} by the B1 baseline power similarly cancels $|E_0|^2$, and substituting $f_{r,m}$ from Step 1 isolates the mirror transmission:

$$T_{\text{WE}} = \frac{\psi}{\sqrt{f_{r,m}}\, G} \cdot \frac{A_{\text{B1}}}{P_{\text{B1,base}}}. \tag{11}$$

Step 4: averaging. Averaging the recovered T_{WE} over multiple zero-crossings across several bender oscillation cycles reduces the statistical uncertainty.

Zero-crossing identification. In practice, the zero-crossing times are identified from the data using three complementary approaches: synchronisation with the recorded drive voltage, the envelope of the B8 power burst (which peaks when $f_m \approx 1$), and the corresponding B1 envelope as a consistency check.

VI. Simulation Results

Synthetic data were generated with $T_{\mathrm{WE,ref}} = 3.39 \times 10^{-6}$, $f_r = 10^{-6}$, $x_{\mathrm{tip}} = 10\ \mu\mathrm{m}$, and drive frequencies spanning $f_{\mathrm{mod}} \in [10^{-3}, 50]$ Hz.

A. Zero-Crossing Fits

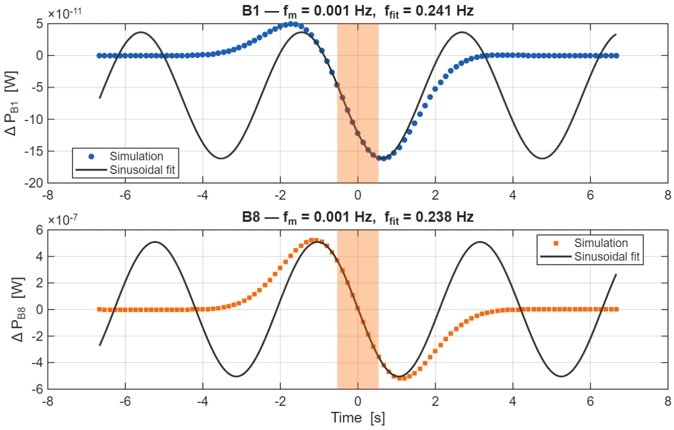

Fig. 4. Baseline-subtracted ΔP_{B1} (top) and ΔP_{B8} (bottom) at $f_{\mathrm{mod}} = 1$ mHz, plotted against bender displacement. The sinusoidal fit (grey) matches the simulation inside the zero-crossing window (shaded) and diverges outside, where the mode-matching efficiency drops.

A single-frequency sinusoidal fit reproduces the simulated photocurrents inside the nine-sample zero-crossing window to numerical precision (figure 4). Outside the window the fit diverges, as expected from the rapid drop in mode-matching efficiency. The fitted frequency scales as $f_{\mathrm{fit}} \simeq 241\, f_{\mathrm{mod}}$, exceeding $\beta \cdot f_{\mathrm{mod}} \approx 118\, f_{\mathrm{mod}}$. This is a geometric effect of the finite bender length (the tilt adds a transverse phase contribution), and does not affect the fitted amplitude entering the calibration.

B. Frequency Dependence of Cavity Gain

The cavity amplitude gain $G(\theta)$ was extracted as a function of bender tilt for different drive frequencies. At low f_{mod} the gain traces a sharply peaked curve centred on $\theta = 0$, reaching the DC value $G_{\max} \simeq 151.5$. As f_{mod} increases, the peak is progressively reduced (figure 5) because the bender sweeps through the zero-crossing faster than the cavity storage time $\tau_{\mathrm{cav}} \simeq 2.9$ ms: the intracavity field cannot fully build up before the bender has moved on.

C. Cavity Decay Rate from Ring-Down

After the bender sweeps past a zero-crossing, the mode-matching efficiency drops to $f_m \approx 0$ and the intracavity perturbation decays freely. Fitting a double-exponential model

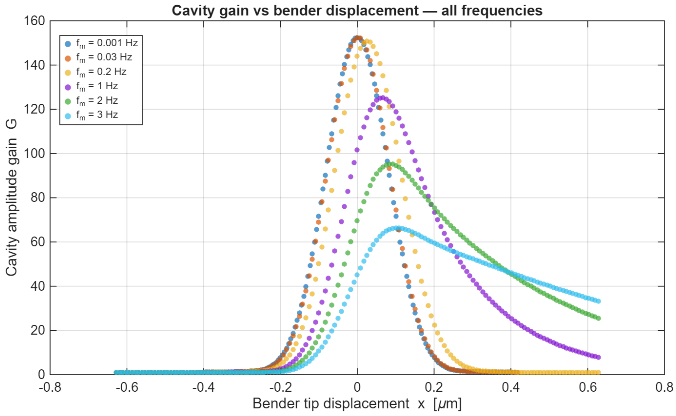

Fig. 5. Cavity amplitude gain as a function of bender tilt angle, for different drive frequencies. At low f_{mod} the gain peaks at $G_{\max} \simeq 151.5$. At higher frequencies the bender sweeps through faster than the cavity storage time, reducing the effective gain.

to the tail of the TEM$_{00}$ power yields $\gamma_{\mathrm{fit}} = 325.9\ \mathrm{s}^{-1}$, in close agreement with $\gamma_{\mathrm{known}} = 326.1\ \mathrm{s}^{-1}$ computed from the mirror reflectivities (figure 6). This validates the cavity memory model of equation (7) and provides an independent measurement of the cavity pole: $f_{\mathrm{arm,fit}} = \gamma_{\mathrm{fit}}/(2\pi) = 51.9$ Hz, consistent with the design value [11].

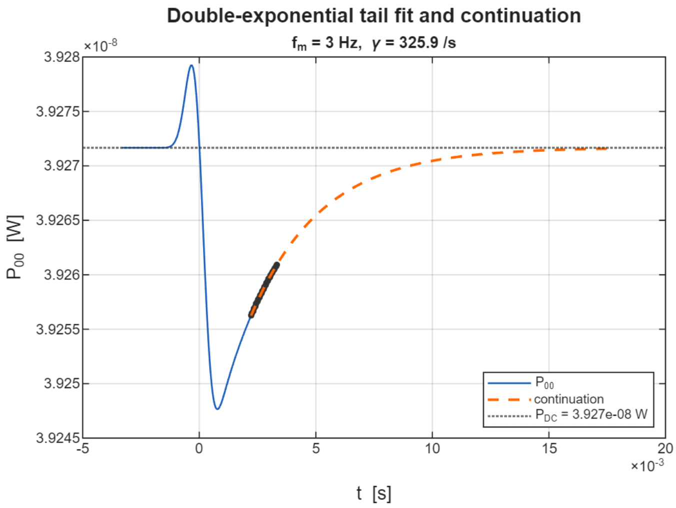

Fig. 6. TEM$_{00}$ power for $f_{\mathrm{mod}} = 3$ Hz (blue) with the double-exponential ring-down fit (orange dashed). The fitted decay rate $\gamma_{\mathrm{fit}} = 325.9\ \mathrm{s}^{-1}$ corresponds to $f_{\mathrm{arm}} = 51.9$ Hz.

D. End-Mirror Transmission Recovery

The main calibration result is shown in figure 7 and table I. Two variants of the cavity gain correction are compared.

With a magnitude-only correction, $G_{\mathrm{corr}} = G_{\max}/\sqrt{1 + (f_{\mathrm{fit}}/f_{\mathrm{arm}})^2}$, the recovered T_{WE} agrees with the reference to within $0.2\,\%$ for $f_{\mathrm{mod}} = 1$ mHz to 50 mHz, but develops a systematic upward bias at higher frequencies, reaching $T_{\mathrm{WE}}/T_{\mathrm{WE,ref}} = 2.09$ at $f_{\mathrm{mod}} = 3$ Hz. The origin of

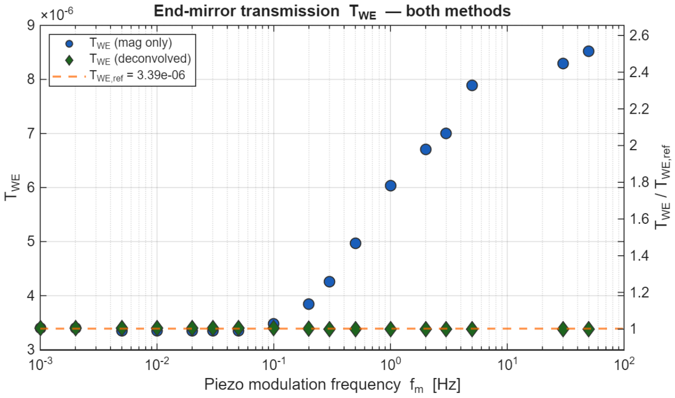

Fig. 7. Recovered $T_{\mathrm{WE}}/T_{\mathrm{WE,ref}}$ as a function of drive frequency. Blue circles: magnitude-only pole correction. Green diamonds: full complex correction via time-domain deconvolution (equation (12)). With the deconvolution, sub-percent accuracy is maintained across the full range.

this bias is the neglected phase lag $\varphi_s = \arctan(f_{\mathrm{fit}}/f_{\mathrm{arm}})$: the magnitude-only correction accounts for the amplitude roll-off above the cavity pole but not for the phase rotation, which distorts the transient signal shape at the zero-crossing and biases the fitted amplitude.

To understand why the phase lag biases a transient fit, consider what the cavity does to the scattered-light burst that occurs at each zero-crossing. When the bender passes through $\theta = 0$, a short pulse of mode-matched light enters the cavity and rings up the intracavity field. At low drive frequencies this build-up is nearly instantaneous compared to the bender motion, so the intracavity field tracks the input faithfully and the signal shape at B1 is a clean sinusoid. At high drive frequencies, however, the cavity responds with a delay of order τ_{cav}: the field peak is shifted in time relative to the input, and the leading and trailing edges of the burst are asymmetrically distorted. A symmetric sinusoidal fit to this asymmetric waveform systematically overestimates the amplitude, producing the observed upward bias in T_{WE}.

The bias is eliminated by deconvolving the cavity pole from the B1 signal before fitting. The arm cavity acts as a single-pole low-pass filter, whose inverse in the time domain is

$$\Delta P_{B1}^{(\mathrm{deconv})}(t) = \Delta P_{B1}(t) + \tau_{\mathrm{arm}} \frac{\mathrm{d}\,\Delta P_{B1}}{\mathrm{d}t}(t), \qquad (12)$$

where $\tau_{\mathrm{arm}} = 1/(2\pi f_{\mathrm{arm}})$ is the cavity storage time. This identity follows directly from rearranging the first-order ordinary differential equation that defines the cavity impulse response [18]: $x = y + \tau_{\mathrm{arm}}\,\dot{y}$. It is exact for arbitrary time dependence, including transient signals, and requires no steady-state assumption.

The sinusoidal fit is then performed on the deconvolved signal, and $G_{\max}$ replaces G_{corr} in equation (11). With this full complex correction, sub-percent accuracy is maintained across the entire range of drive frequencies tested, from 1 mHz to 50 Hz (figure 7).

TABLE I
RECOVERED T_{WE} FOR SELECTED DRIVE FREQUENCIES, COMPARING THE MAGNITUDE-ONLY AND DECONVOLVED CORRECTIONS.

f_{mod} [Hz]	f_{fit} [Hz]	$\frac{T_{\mathrm{WE}}}{T_{\mathrm{WE,ref}}}$ (mag.)	$\frac{T_{\mathrm{WE}}}{T_{\mathrm{WE,ref}}}$ (deconv.)
0.001	0.24	1.0016	1.0016
0.01	2.38	0.9874	1.0027
0.05	11.92	0.9893	1.0029
0.1	23.85	1.0251	1.0015
0.2	47.69	1.1340	0.9994
0.3	71.54	1.2554	0.9983
0.5	119.23	1.4645	0.9973
1.0	238.46	1.7801	0.9969
2.0	476.92	1.9785	0.9969
3.0	715.38	2.0656	0.9969
5.0	1207.36	2.3261	0.9982
30	7244.15	2.4473	0.9980
50	12073.58	2.5129	0.9977

E. Influence of Bender Length

The ratio $f_{\mathrm{fit}}/f_{\mathrm{mod}} \simeq 241$ exceeds $\beta \approx 118$ because the finite bender length couples tilt and displacement. Repeating the simulation for several bender lengths while keeping $x_{\mathrm{tip}} = 10\ \mu\mathrm{m}$ fixed confirms that f_{fit} converges to $\beta \cdot f_{\mathrm{mod}}$ as $\ell \to \infty$. The fitted amplitude, which is the quantity entering the calibration, is independent of bender length.

Practical operating envelope. The minimum practical tip displacement is approximately $x_{\mathrm{tip,min}} \approx 0.15\ \mu\mathrm{m}$, below which too small a fraction of an optical fringe lies within the zero-crossing window for a reliable sinusoidal fit. Operating at higher f_{mod} with proportionally reduced x_{tip} shortens the measurement time and reduces susceptibility to slow drifts in the detector operating point, while the deconvolution of equation (12) ensures that the cavity phase lag does not limit the accessible frequency range.

VII. Experimental Hardware

The experimental setup for generating controlled scattered light on the Advanced Virgo end bench consists of two main components.

Piezoelectric bender. A Thorlabs PB4NB2S bimorph bender provides up to $\pm 450\ \mu\mathrm{m}$ tip displacement with a resonant frequency of 370 Hz and is vacuum-compatible to 10^{-7} Torr ($\approx 1.3 \times 10^{-7}$ mbar). It is driven differentially by a Thorlabs MDT694B high-voltage controller. Two 10 kΩ protection resistors limit the system bandwidth to ≈ 29 Hz.

Linear translation stage. A Newport Agilis AG-LS25-27 piezo stage (27 mm travel, 100 nm minimum incremental motion) positions the bender into or out of the transmitted beam path, controlled remotely via USB.

Initial laboratory tests revealed that the bare ceramic surface of the bender absorbs most of the incident laser power (~ 2 W), causing rapid heating that risks damaging the piezoelectric ceramic and causing outgassing under vacuum. On the Advanced Virgo end bench, the transmitted power reaching the bender position is expected to reach approximately 1.6 W during O5 [19], focused to a ~ 1 mm spot. The baseline solution is to apply a thin metallic coating to the bender surface, increasing the reflectivity above $90\,\%$ at 1064 nm [20] and reducing the absorbed power by roughly an order of magnitude. A reflective coating also increases f_r, directly improving the calibration signal-to-noise ratio.

VIII. Conclusion

A novel method for measuring the end-mirror transmission of the Advanced Virgo arm cavities has been developed, exploiting controlled scattered light from a piezoelectric bender. The key findings are:

1) A time-domain simulation framework correctly models the full scattering process, including tilt-induced higher-order mode coupling and the frequency-dependent cavity response.
2) T_{WE} is recovered with sub-percent accuracy across the full range of drive frequencies tested (1 mHz to 50 Hz), using only measured power ratios. At low drive frequencies a magnitude-only cavity pole correction suffices, at higher frequencies, a time-domain deconvolution removes the cavity phase lag from the B1 signal before fitting, extending sub-percent accuracy to the full range. The method is self-calibrating: it requires no knowledge of the absolute laser power.
3) The cavity decay rate measured from the ring-down ($\gamma_{\mathrm{fit}} = 325.9\ \mathrm{s}^{-1}$) matches the known value ($326.1\ \mathrm{s}^{-1}$), validating the cavity memory model.
4) The TEM_{00} mode-matching follows $\exp[-(kw\theta)^2]$, confirming that the calibration window is correctly restricted to the bender zero-crossings.

When combined with a laboratory measurement of T_{WE}, the method provides an independent test of the absolute calibration of the gravitational wave detector, complementing existing techniques such as the photon calibrator [21], [22] and the Newtonian calibrator [23]. The three methods cover complementary frequency ranges: the Newtonian calibrator operates at a few hertz, the photon calibrator at tens to hundreds of hertz, and the scattered-light method at a few Hz to a few kHz set by the bender drive.

IX. Outlook

The remaining steps toward an in-situ measurement are: applying a thin coating (e.g. ~ 100 nm aluminium [20]) or a spectralon reflector [19] to the bender surface to reduce absorbed power and increase the scattered-light signal, calibrating the bender displacement using an independent sensor and verifying linearity well below the $370\,$Hz mechanical resonance, measuring the surface reflectivity at 1064 nm to determine the scattering parameter f_r, and verifying vacuum compatibility of the full assembly under representative thermal load before installation on the suspended end bench.

Once installed, a commissioning phase will confirm that the expected signals are produced and that no excess noise is introduced into the detector output. The bender can be retracted when not in use, leaving no residual effect on detector sensitivity. If the experimental results confirm the sub-percent accuracy demonstrated by the simulation, the scattered-light method could become a valuable addition to the calibration toolkit of current and future gravitational wave observatories [21].

References

[1] B. P. Abbott *et al.*, "Observation of gravitational waves from a binary black hole merger," *Phys. Rev. Lett.*, vol. 116, p. 061102, 2016.

[2] F. Acernese *et al.*, "Interferometer sensing and control for the Advanced Virgo experiment in the O3 scientific run," *Galaxies*, vol. 8, no. 4, p. 85, 2020.

[3] B. P. Abbott *et al.*, "A gravitational-wave standard siren measurement of the Hubble constant," *Nature*, vol. 551, pp. 85–88, 2017.

[4] M. Was, R. Gouaty, and R. Bonnand, "End benches scattered light modelling and subtraction in Advanced Virgo," *Class. Quantum Grav.*, vol. 38, p. 075020, 2021.

[5] F. Acernese *et al.*, "Advanced Virgo: a second-generation interferometric gravitational wave detector," *Class. Quantum Grav.*, vol. 32, p. 024001, 2015.

[6] A. Freise and K. Strain, "Interferometer techniques for gravitational-wave detection," *Living Rev. Relativ.*, vol. 13, 2010.

[7] M. J. Lawrence *et al.*, "Dynamic response of a Fabry–Perot interferometer," *J. Opt. Soc. Am. B*, vol. 16, pp. 523–532, 1999.

[8] D. Z. Anderson, "Alignment of resonant optical cavities," *Appl. Opt.*, vol. 23, no. 17, pp. 2944–2949, 1984.

[9] F. Bayer-Helms, "Coupling coefficients of an incident wave and the modes of a spherical optical resonator," *Appl. Opt.*, vol. 23, no. 9, pp. 1369–1380, 1984.

[10] B. E. A. Saleh and M. C. Teich, *Fundamentals of Photonics*, 2nd ed. Wiley, 2007.

[11] C. Jacquet *et al.*, "Arm cavity characterization," Virgo Collaboration, Tech. Rep. VIR-0540A-24, 2024.

[12] J. Degallaix, "Oscar: A matlab based package to simulate realistic optical cavities," *SoftwareX*, vol. 12, p. 100587, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2352711020303009

[13] ——, "Oscar a matlab based optical fft code," *Journal of Physics: Conference Series*, vol. 228, no. 1, p. 012021, may 2010. [Online]. Available: https://doi.org/10.1088/1742-6596/228/1/012021

[14] L. Pinard, "Advanced Virgo end mirror characterization report – EM03," Virgo Collaboration, Tech. Rep. VIR-0270A-15, 2015.

[15] ——, "Advanced Virgo input mirror characterization report IM03 (coatings C118046+C18013) spare IM," Virgo Collaboration, Specification/Characterisation VIR-0326A-23, Apr. 2023, internal Virgo document. https://tds.virgo-gw.eu/ql/?c=19196.

[16] M. Rakhmanov, "Doppler-induced dynamics of fields in Fabry–Perot cavities with suspended mirrors," *Appl. Opt.*, vol. 40, pp. 1942–1949, 2001.

[17] M. Rakhmanov *et al.*, "Dynamic resonance of light in Fabry–Perot cavities," *Phys. Lett. A*, vol. 305, pp. 239–244, 2002.

[18] A. V. Oppenheim, A. S. Willsky, and S. H. Nawab, *Signals and Systems*, 2nd ed. Prentice Hall, 1997.

[19] Virgo Collaboration, "Advanced Virgo Plus for O5 – technical design report," Virgo Collaboration, Tech. Rep. VIR-0499A-25, 2025.

[20] R. Lugolole and S. K. Obwoya, "The effect of thickness of aluminium films on optical reflectance," *Journal of Ceramics*, vol. 2015, no. 1, p. 213635, 2015. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1155/2015/213635

[21] F. Acernese *et al.*, "Calibration of Advanced Virgo and reconstruction of the detector strain $h(t)$ during the observing run O3," *Class. Quantum Grav.*, vol. 39, p. 045006, 2022.

[22] D. Estevez *et al.*, "The Advanced Virgo photon calibrators," *Class. Quantum Grav.*, vol. 38, p. 075007, 2021.

[23] D. Estevez, B. Mours, and T. Pradier, "Newtonian calibrator tests during the Virgo O3 data taking," *Class. Quantum Grav.*, vol. 38, p. 075012, 2021.

Table of Contents

List of Figures

List of Tables

1 Introduction

The direct detection of gravitational waves by the LIGO and Virgo collaborations in 2015 [2] opened a new observational window onto the universe, enabling the study of astrophysical phenomena (from binary black hole mergers to neutron star collisions) that are invisible to electromagnetic telescopes. The scientific output of this endeavour depends critically on the accuracy with which the raw detector output can be converted into a calibrated strain time series $h(t)$. Any systematic error in the calibration propagates directly into the inferred source parameters: masses, spins, distances, and sky locations of detected events all carry uncertainties that are bounded from below by the calibration accuracy. For example, a 1 % bias in the calibration amplitude translates directly into a 1 % error in the luminosity distance to a source, which propagates into cosmological parameter estimates when gravitational wave events are used as standard sirens. [3, 4].

At the heart of the calibration lies the optical response of the Fabry–Perot arm cavities that form the two arms of the Advanced Virgo interferometer. These cavities amplify the tiny differential displacement induced by a passing gravitational wave, but the amplification factor depends on the cavity finesse and, in particular, on the power transmission coefficients of the input and end test masses. Among these, the end-mirror transmission T_{WE} plays a dual role: it determines the fraction of intracavity power that reaches the end-bench monitoring photodiodes, and it sets the coupling strength of any scattered light that re-enters the cavity from the end bench. An accurate, independent measurement of T_{WE} is therefore valuable both for verifying the absolute calibration of the detector's strain response and for modelling scattered-light noise.

Scattered light is one of the most persistent noise sources in the low-frequency band of Advanced Virgo. When light transmitted through the end mirror reflects off optical components on the end bench and re-enters the arm cavity, it accumulates a phase that depends on the distance to the scattering surface. If this distance varies in time, for example, due to seismic motion of the bench, the resulting phase modulation generates a characteristic signal at the output photodiode. Considerable effort during the commissioning of Advanced Virgo has been devoted to identifying, characterising, and mitigating these scattering sources [5].

Was et al. [5] demonstrated during the third observing run that the scattered-light coupling mechanism can be turned from a noise source into a calibration tool, using the natural seismic motion of the end bench as the scattering source. The present thesis extends this approach by developing a dedicated controllable scattering actuator, a piezoelectric bender placed on the end bench and driven at a known frequency, and by constructing a simulation framework that models the full scattering process including effects that are absent in the bench-motion case. By placing a piezoelectric bender on the end bench and driving it at a known frequency, a controlled and reproducible scattered field is generated. The oscillating bender surface modulates both the phase and the propagation angle of the reflected beam, producing a well-characterised signal at the output photodiode (B1) and the end-bench photodiode (B8). Because the signal amplitudes at B1 and B8 depend on T_{WE} in different ways, the end-mirror transmission can be extracted from the ratio of these amplitudes without requiring knowledge of the absolute laser power.

The use of a piezoelectric bender, rather than uncontrolled bench motion, introduces several challenges that were not present in the original method. The bender tip displacement of $\sim 10\,\mu$m at a wavelength of 1064 nm produces a modulation depth of $\beta \approx 118$, far beyond the small-modulation regime where only two sidebands are relevant. The frequency-domain description of the cavity response becomes impractical in this regime, motivating a time-domain reformulation of the cavity dynamics. Furthermore, the finite length of the bender means that a tip displacement is necessarily accompanied by a tilt, which couples the reflected beam into higher-order Hermite–Gaussian modes that do not resonate in the arm cavity. The calibration signal is therefore concentrated in short time windows around the bender zero-crossings, where the tilt vanishes and the mode-matching efficiency peaks.

This thesis is structured as follows.

Chapter 2 describes the Advanced Virgo detector, following the laser beam from the source through the input mode cleaner, power recycling cavity, arm cavities, signal recycling, and output mode cleaner to the B1 photodiode. The chapter also covers seismic isolation, sensing and control of longitudinal degrees of freedom, and the quantum and thermal noise sources that set the detector sensitivity. Lastly, it covers the current methods for calibrating gravitational wave detectors.

Chapter 3 develops the mathematical foundations needed for the scattered-light model: Gaussian beam optics and Hermite–Gaussian modes, the Fabry–Perot cavity transfer function, phase modulation and sideband generation, and the overlap integral that quantifies mode matching.

Chapter 4 presents the theoretical model for scattered light coupling in the Virgo end benches, derives the equivalent gravitational wave strain produced by scattering, and analyses controlled injection measurements performed during the commissioning period preceding the fourth observing run (O4).

Chapter 5 describes the simulation framework built around the Oscar FFT code, derives the time-domain cavity update equation and its equivalence with the frequency-domain response, analyses tilt-induced mode coupling, and presents the four-step calibration procedure for recovering T_{WE} from the B1 and B8 signals.

Chapter 6 applies the calibration procedure to synthetic data generated by the simulation, validating the method across the full range of drive frequencies tested and demonstrating that a time-domain deconvolution of the cavity transfer function eliminates the systematic bias that arises above the cavity pole when using a magnitude-only correction.

Chapter 7 presents the experimental hardware: the piezoelectric bender, the linear translation stage, and the laser source, together with the outlook for deployment on the Advanced Virgo end bench.

Chapter 8 summarises the principal results, discusses their implications for the Advanced Virgo calibration programme, and identifies the remaining steps toward an in-situ measurement.

2

Advanced Virgo

2.1 Overview and scientific context

Advanced Virgo (AdV) is a second-generation laser interferometric gravitational wave detector located at the European Gravitational Observatory (EGO) near Cascina, Pisa, Italy. It is operated by the Virgo Collaboration, a large international group of research institutions across Europe. Together with the two Advanced LIGO detectors in the United States and the KAGRA detector in Japan, Advanced Virgo forms the core of the current global network of ground-based gravitational wave observatories. This network has enabled the detection of gravitational wave signals from compact binary coalescences, neutron star mergers, and other astrophysical events since the first direct detection, GW150914, in 2015 [2].

Gravitational waves are ripples in the fabric of spacetime, predicted by Albert Einstein's general theory of relativity in 1916. They are generated by the acceleration of massive objects, most notably by the inspiral and merger of compact binary systems such as pairs of black holes or neutron stars. As these objects orbit each other, they lose energy through gravitational radiation, causing their orbits to shrink until they eventually coalesce. The emitted waves propagate outward at the speed of light, stretching and compressing spacetime transversely as they pass through it.

The effect of a gravitational wave on a ring of freely falling test masses is a quadrupolar distortion: the ring is alternately stretched along one axis and compressed along the perpendicular axis, with the two polarisation states (h_+ and $h_\times$) rotated by 45° relative to each other (figures 2.1a and 2.1b). It is this differential length change that a Michelson-type interferometer is designed to detect.

The fundamental physical quantity that Advanced Virgo is designed to measure is the gravitational wave strain $h(t)$, defined as the dimensionless ratio of the differential length change δL induced by a passing gravitational wave to the arm length L [6]:

$$h(t) = \frac{2\,\delta L(t)}{L}. \tag{2.1}$$

For astrophysical sources within the detection band, h is typically of order 10^{-21} or smaller around 100 Hz, corresponding to a differential displacement of order 10^{-18} m for the 3 km Virgo arms (a fraction of the diameter of a proton). Achieving sensitivity at this level over a detection band spanning roughly 10 Hz to several kilohertz requires suppressing classical and quantum noise sources to extraordinary levels. The entire optical and mechanical design of Advanced Virgo is organised around this challenge.

The conceptual basis for the measurement is the Michelson interferometer: a laser beam is split into two orthogonal arms, and a gravitational wave passing through the detector differentially modulates the optical path lengths, causing a measurable change in the interference pattern at the output port. The mathematical foundation for how a Fabry–Perot cavity in a cavity-enhanced Michelson interferometer responds to such a phase signal is developed in chapter 3, particularly in sections 3.2 and 3.3. The present chapter translates that foundation into the physical and engineering reality of the Advanced Virgo instrument, following the path of light from the laser source through the main optical subsystems to the final detection, providing the context needed to understand what the scattered light must traverse and why generating a clean scattered light calibration signal is challenging. An overview of the detector layout is shown in figure 2.2.

2.2 Optical layout

2.2.1 Main laser and pre-stabilisation

The primary light source of Advanced Virgo is a continuous-wave Nd:YAG laser operating at a wavelength of $\lambda = 1064$ nm, delivering an output power of 17 W after the pre-stabilisation stage [9, 10]. The choice of wavelength reflects a combination of factors: the availability

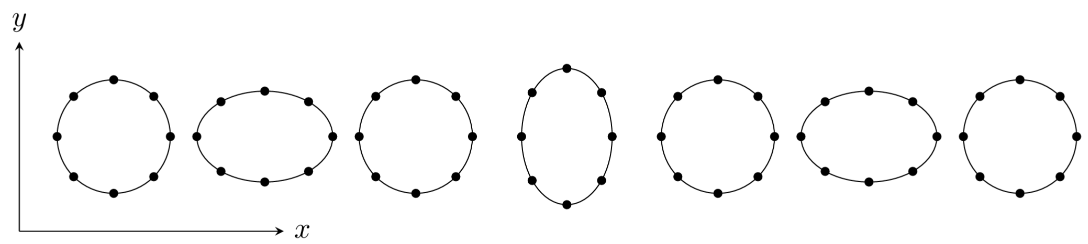

(a) + polarization

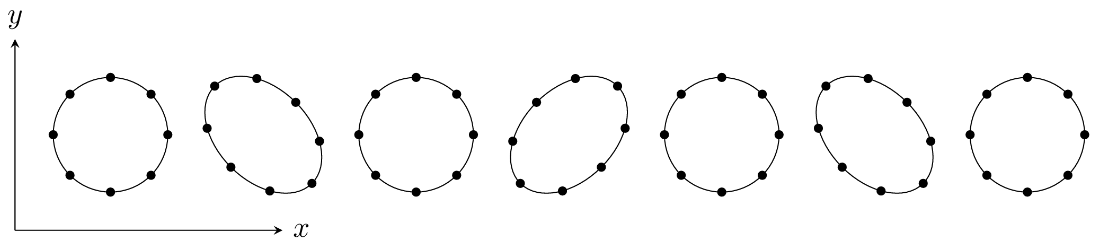

(b) × polarization

Figure 2.1: Effect of the two gravitational-wave polarization states on a ring of freely falling test particles. The $+$ polarization (top) periodically stretches space along one transverse axis while compressing it along the perpendicular axis. The $\times$ polarization (bottom) produces the same deformation pattern rotated by $45°$ [7].

of high-quality optical coatings with extremely low absorption and scattering at this wavelength, the low absorption in fused silica mirror substrates, and the compatibility with mature and reliable Nd:YAG laser technology.

The raw laser output contains frequency noise, intensity noise, and pointing fluctuations that would affect the gravitational wave signal if injected directly into the interferometer. The pre-stabilisation system (PSL) therefore reduces these noise contributions before the beam enters the input mode cleaner [10].

Frequency pre-stabilisation is achieved by locking the laser to a short, rigid reference cavity using the Pound–Drever–Hall (PDH) technique described in section 3.3.4. The laser frequency is servo-controlled to track the resonance of the reference cavity, and residual frequency noise at the PDH error point is fed back to the laser frequency actuator, reducing the laser linewidth to the kilohertz level or below [11].

2.2.2 Input mode cleaner

After pre-stabilisation, the beam enters the input mode cleaner (IMC), a three-mirror triangular ring cavity with a round-trip length of approximately 143 m. The IMC serves three distinct purposes simultaneously.

Spatial filtering. Any higher-order Hermite–Gaussian modes present in the input beam, whether from imperfect laser output or residual pointing fluctuations, are not resonant in the IMC. As described by the transverse mode frequency splitting of equation (3.25), each mode order $m + n$ resonates at a frequency shifted from the TEM_{00} resonance. For the IMC geometry, this splitting is large compared to the cavity linewidth, so only the TEM_{00} mode is resonantly transmitted. This ensures that the beam entering the main interferometer has a clean, well-defined Gaussian spatial profile, essential for the mode-matching conditions discussed in section 3.4.1 [12].

Further frequency stabilisation. The IMC is a high-finesse cavity and therefore provides a precise optical frequency reference. The laser is locked to this cavity via a PDH error signal, so that residual frequency fluctuations above the IMC linewidth are attenuated by the single-pole cavity roll-off described in section 3.2.2 [12].

Intensity and pointing noise reduction. Analogously to the frequency filtering, the IMC attenuates residual intensity and pointing noise at frequencies above its linewidth [12].

After transmission through the IMC, the beam has reduced spatial mode content, improved frequency stability, and lower intensity noise compared to the pre-stabilised laser output, and is ready for injection into the main interferometer [12].

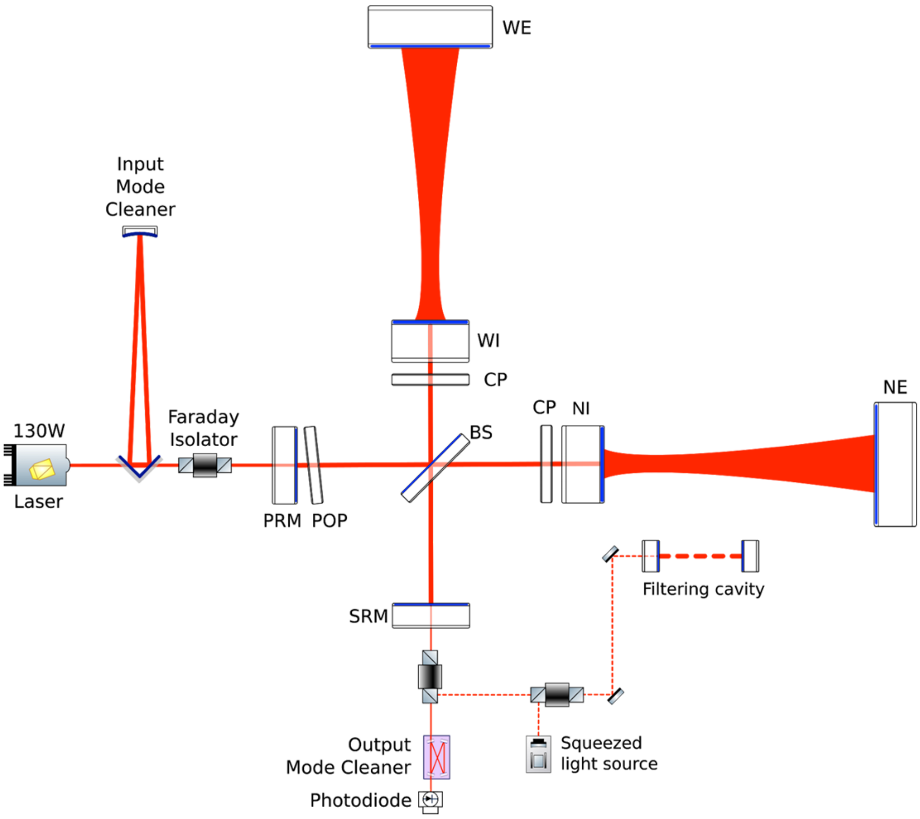

Figure 2.2: Simplified schematic optical layout of the Advanced Virgo interferometer. Light travels from the laser through the input mode cleaner (IMC) and power recycling cavity (PRC) before entering the two 3 km Fabry–Perot arm cavities. The gravitational wave signal is read out at the antisymmetric port via the output mode cleaner (OMC) and the B1 photodiode [8].

2.2.3 Power recycling mirror

The beam transmitted through the IMC is injected into the main interferometer through the power recycling mirror (PRM), which forms one end of the power recycling cavity (PRC).

When the interferometer is operated at or near the dark fringe (meaning almost all of the carrier light is reflected back toward the laser rather than transmitted to the output port) the reflected carrier constructively interferes with the incoming laser beam if the PRC is resonant. Light that would otherwise be lost is thereby redirected back into the interferometer, increasing the circulating power [13].

The PRC is a three-mirror cavity formed by the PRM and the two input test masses (ITMs), with the central beam splitter effectively acting as part of the recycling cavity structure. For Advanced Virgo, the power recycling gain amplifies the input laser power by a factor of approximately 40, raising the power at the beam splitter to roughly 1 kW at design sensitivity [14]. This increased circulating power is essential because the shot-noise-limited sensitivity of the interferometer improves as $1/\sqrt{P_{\text{in}}}$, where P_{in} is the power entering each arm cavity [15].

The PRC must be held on resonance with active length control, achieved using an auxiliary PDH locking scheme with radio-frequency (RF) sidebands that are resonant in the PRC but not in the arm cavities. This frequency selectivity separates the length sensing degrees of freedom of the two cavity systems and allows independent servo loops to control the PRC and the arm cavities simultaneously [16].

2.2.4 Central beam splitter

After entering the power recycling cavity, the beam arrives at the central beam splitter (BS), a fused silica optic with a 50/50 power splitting ratio that divides the incoming light equally into the two arms of the Michelson interferometer [17, 18].

Under perfectly asymmetric conditions, the Michelson interferometer destructively interferes all carrier light at the output (antisymmetric) port, reflecting it back toward the PRM. This is the dark fringe condition, and it is what makes power recycling possible.

In practice, small differences between the two arms and the deliberate introduction of a DC differential arm offset (described in section 2.4.2) mean that a small but controlled amount of carrier power is always present at the output port. The management of this carrier leakage is central to the DC readout technique [19].

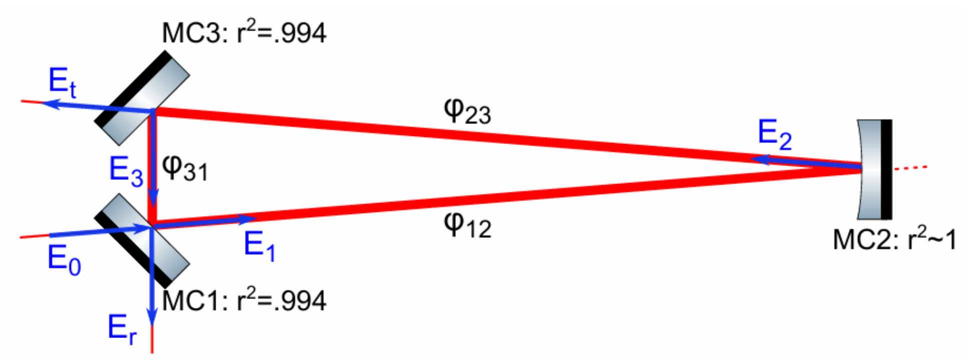

Figure 2.3: Schematic diagram of the input mode cleaner (IMC). The triangular ring cavity of length approximately 143 m transmits only the TEM$_{00}$ mode to the main interferometer [12].

2.2.5 Fabry–Perot arm cavities

Optical layout and parameters

Each arm of the Advanced Virgo interferometer contains a Fabry–Perot cavity formed by an input test mass (ITM) and an end test mass (ETM). For the west arm, these are the WI (west input) and WE (west end) mirrors. For the north arm, they are the NI and NE mirrors. Each arm cavity has a length of $L = 3$ km, and the mirrors are suspended by multi-stage seismic isolation systems described in section 2.3.

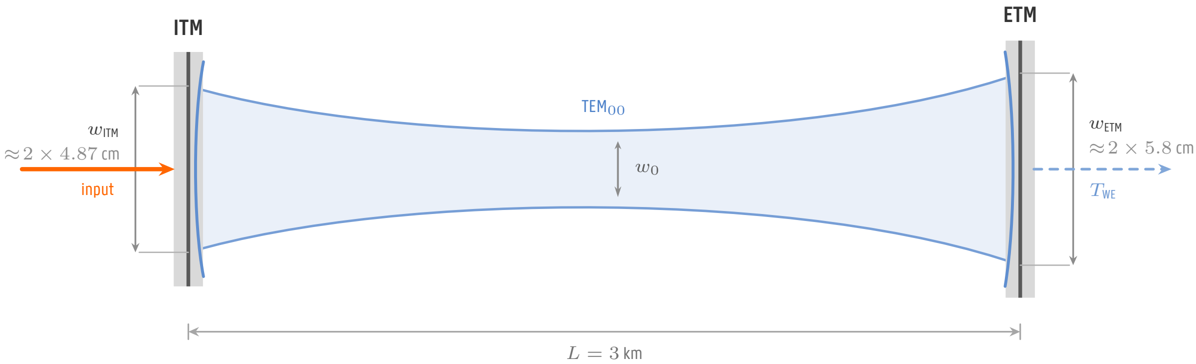

Figure 2.4: Schematic of a single Advanced Virgo Fabry–Perot arm cavity. The cavity is formed by two suspended mirrors separated by $L = 3$ km. The beam radii at the input and end mirrors are $w_{\text{ITM}} \approx 4.87$ cm and $w_{\text{ETM}} \approx 5.8$ cm, respectively [20].

The mirror geometry is designed to maximise the beam size on the test mass surfaces, thereby minimising thermal noise, which scales inversely with the illuminated area (section 2.5.4). The arm cavity mirrors have radii of curvature $R_{\text{ITM}} = 1420$ m and $R_{\text{ETM}} = 1683$ m [14], chosen such that the cavity is geometrically stable and the beam radii on the mirror surfaces are maximised. The resulting design beam radii are $w_{\text{ITM}} = 48.7$ mm and $w_{\text{ETM}} = 58$ mm [14], with the measured values during O4 reported in [21]. The stability criterion and beam radius calculation are derived in section 3.2.

The cavity geometry also determines the frequency spacing between transverse modes. For the Virgo arm cavities, the nearest higher-order mode resonance lies at a frequency offset of order $\Delta\nu_{10} \sim 10$ kHz from the carrier, far exceeding the cavity linewidth of ~ 100 Hz, ensuring that higher-order mode contributions are strongly off-resonance. The mathematical framework underlying this mode structure is developed in sections 3.1 and 3.2. The input mirror transmission is $T_{\text{ITM}} = 1.3$ % [22], while the end mirror is near-perfectly reflective with $T_{\text{ETM}} \approx 3.39$ ppm [23].

Resonance condition and length control

The arm cavities must be held on resonance with extreme precision: a differential mirror displacement of $\delta x \sim \lambda/(2\mathcal{F}) \sim 1$ nm would shift the cavity from resonance to the half-power point. Active length control achieves mirror position stability at the sub-nanometre level using the PDH error signal described in section 3.3.4. The error signal is generated using RF sidebands that are transmitted through the PRC but lie above the arm cavity linewidth, so they reflect from the ITMs and serve as a clean phase reference [21].

The arm cavity resonance condition also sets the operating frequency for the carrier laser, which is locked to the CARM (common arm length) degree of freedom through the control hierarchy described in section 2.4.

2.2.6 Signal recycling

Physical mechanism

In the ideal Michelson interferometer operated at the dark fringe, the carrier and the gravitational wave signal sidebands separate at the central beam splitter and exit through different ports. The signal sidebands are not injected from outside but are generated inside the interferometer by a passing gravitational wave, which transfers energy from the carrier into sideband fields at the gravitational wave frequency [24]. The signal recycling cavity (SRC) is formed between the signal recycling mirror (SRM) and the ITMs. Placing a highly reflective mirror in front of the output photodiode increases the signal power seen by that photodiode.

Tuned and detuned configurations

The response of the SRC is governed by the microscopic tuning of the SRM position, i.e. the round-trip phase accumulated by signal sidebands in the cavity. This is illustrated in figure 2.5, computed for a signal recycling mirror reflectivity of $R = 0.9$ [24].

In the tuned configuration, the SRC is resonant for the carrier, which maximises the enhancement of signal sidebands at frequencies near DC. The result is a flat plateau of sideband amplitude from very low frequencies up to the effective bandwidth set by the coupled system (arm cavity + SRC), followed by a roll-off at higher frequencies (red curve in figure 2.5). The detector sensitivity is therefore broadband but concentrated at low frequencies [24].

Introducing a microscopic offset to the SRM position detunes the cavity resonance away from DC to a non-zero gravitational wave frequency f_{det}. The sideband amplitude then exhibits a peaked response centred on f_{det}, visible in figure 2.5 for detuning values of 150 Hz, 500 Hz, and 5 kHz. As the detuning frequency increases, the peak narrows and the low-frequency plateau amplitude decreases while the peak amplitude grows. At extreme detuning, the detector becomes essentially a narrowband filter centred on a single frequency [24].

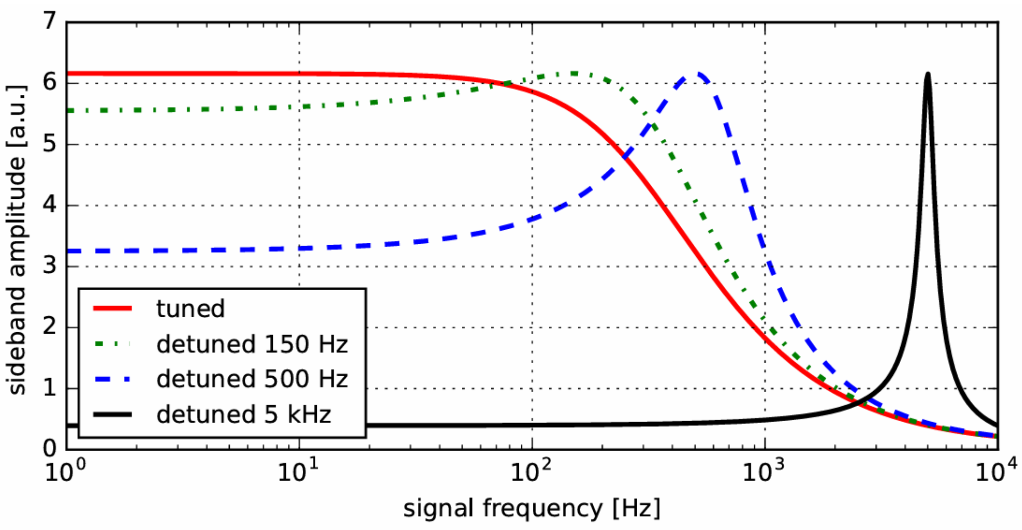

Figure 2.5: Signal sideband amplitude at the output port as a function of gravitational wave frequency for a Michelson interferometer with signal recycling mirror reflectivity $R = 0.9$. The tuned configuration (red) yields a flat, broadband response peaked near DC. Detuned configurations shift the peak to the detuning frequency (150 Hz, 500 Hz, 5 kHz), narrowing the bandwidth as the detuning increases [24].

Conservation of integrated gain

A fundamental property of the signal recycling technique is that the total integrated gain is the same for all tuning configurations [24]. The areas under the curves in figure 2.5 are constant regardless of detuning. Signal recycling therefore does not increase the overall sensitivity of the detector in a frequency-averaged sense. Rather, it shapes the response function, redistributing sensitivity from frequencies where it is not needed to those where it is. The choice of tuning is therefore an astrophysical one: broadband sensitivity for unmodelled transients, or narrowband enhancement targeted at a known source frequency.

Resonant sideband extraction and dual recycling

When arm cavities of very high finesse are combined with a signal recycling mirror tuned to or near the anti-resonant operating point, the configuration is known as resonant sideband extraction (RSE) [25]. In RSE, the effective bandwidth of the detector is increased relative to the bare arm cavities, since the anti-resonant SRC allows signal sidebands to escape more rapidly. An analysis of the relationship between signal recycling and RSE across the full space of configurations can be found in [26]. The simultaneous use of power recycling to enhance the circulating carrier power and signal recycling to control the signal bandwidth is known as dual recycling [14, 8].

2.2.7 Output mode cleaner

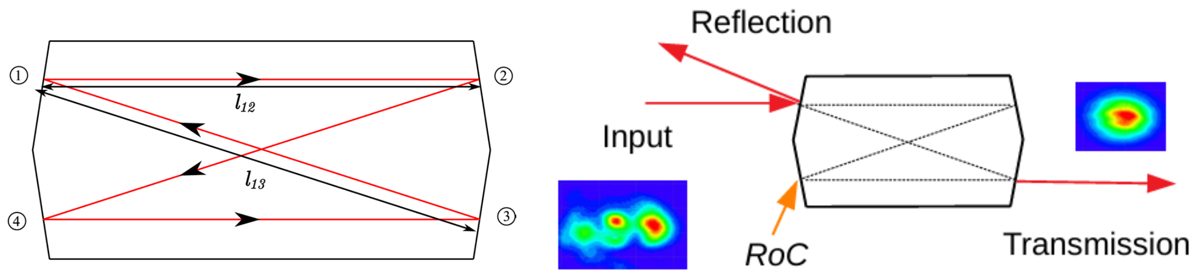

Figure 2.6: Layout (left) and working principle (right) of the output mode cleaner [27, 28].

The signal-recycled light transmitted through the antisymmetric port passes through the output mode cleaner (OMC) before reaching the B1 photodiode. The OMC is a four-mirror bow-tie cavity with a round-trip length of approximately 57 cm and a free spectral range of approximately 267 MHz [29]. It serves two purposes.

Spatial mode filtering. The OMC transmits only the TEM_{00} mode, reflecting all higher-order spatial mode content. Misalignments, mirror figure errors, thermal distortions, and audio-band beam jitter all generate power in higher-order modes at the output port. Since these modes carry no gravitational wave signal, transmitting them to the photodiode would add shot noise without adding signal. The OMC therefore improves the signal-to-shot-noise ratio by ensuring that only the signal-carrying TEM_{00} field reaches the B1 photodiode [29].

Frequency filtering. The OMC round-trip length and finesse are chosen so that neither the RF modulation sidebands used for interferometer length control nor the audio-band alignment sidebands in higher-order modes are resonant in the cavity. Both are reflected rather than transmitted [24]. Since neither contributes to the gravitational wave signal, their suppression reduces the shot noise at B1 without any loss of sensitivity.

After the OMC, the gravitational wave signal is detected by the B1 InGaAs photodiode, which converts the incident optical power to a photocurrent. This photocurrent is amplified, digitised, and processed by the Advanced Virgo data acquisition chain to produce the calibrated strain time series $h(t)$.

2.3 Seismic isolation

Ground seismic noise displaces the Earth's surface by amplitudes orders of magnitude larger than the gravitational wave signal at frequencies below a few hundred hertz. Without isolation, this motion would couple directly to the test masses and overwhelm any gravitational wave signal. The seismic noise floor sets the lower limit of the detection band at roughly 10 Hz for Advanced Virgo. Below this frequency, residual seismic coupling dominates the signal even with isolation [30].

Advanced Virgo employs the Superattenuator system (visible in figure 2.7), a multi-stage inverted pendulum and geometric anti-spring (GAS) filter tower that provides more than 10^{12} attenuation of seismic noise above a few hertz in both horizontal and vertical directions [32]. The underlying principle is that below the mechanical resonance frequency f_i of each pendulum stage, the suspended mass is mechanically decoupled from ground vibrations and behaves as a free-falling mass. For a single stage with resonance frequency f_i, the seismic attenuation at $f \gg f_i$ scales as $(f_i/f)^2$. Cascading N stages provides multiplicative attenuation, scaling as $(f_0/f)^{2N}$ at high frequencies [31, 32].

The test masses are suspended as the final stage from thin fused silica fibres, chosen for their extremely low mechanical loss angle, which minimises thermal noise from the suspension [33]. The entire Superattenuator tower is mounted on an active pre-isolation platform equipped with inertial sensors and feedback actuators that reduce low-frequency ground motion before it enters the pendulum chain [34, 35].

2.4 Sensing and control of longitudinal degrees of freedom

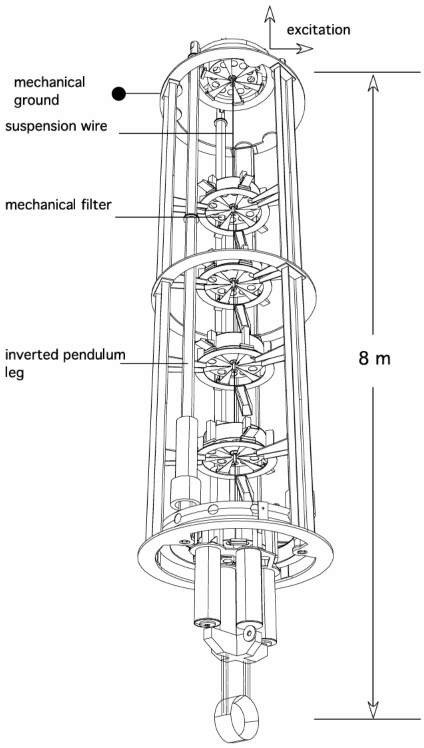

Figure 2.7: Schematic of the Advanced Virgo Superattenuator seismic isolation system. The multi-stage pendulum chain provides more than 10^{12} attenuation of horizontal seismic noise above a few hertz [31, 32].

2.4.1 Degrees of freedom

With multiple coupled cavities (the PRC, the two arm cavities, and the SRC), the Advanced Virgo interferometer has five longitudinal degrees of freedom that must each be independently controlled. These are conventionally decomposed into orthogonal combinations of mirror positions [24]:

- **CARM** (common arm length): the average of the two arm lengths, controlling the resonance of the carrier in the arm cavities. This is the most critical degree of freedom for sensitivity.
- **DARM** (differential arm length): the difference of the two arm lengths. This is the degree of freedom to which a gravitational wave signal couples directly, and it is therefore the primary science channel of the detector.
- **PRCL** (power recycling cavity length): the length of the power recycling cavity, controlling the PRC resonance.
- **SRCL** (signal recycling cavity length): the length of the signal recycling cavity.
- **MICH** (Michelson differential): the differential length of the short Michelson formed by the two ITMs and the beam splitter.

Each degree of freedom must be held at its operating point by an independent servo loop. The error signals for each loop are derived using the PDH technique described in section 3.3.4, with different RF sideband frequencies that resonate selectively in different parts of the optical layout. This selectivity allows a control matrix to be designed such that each error signal is predominantly sensitive to a single degree of freedom [3, 36].

The CARM loop keeps the arm cavities jointly resonant with the carrier, maintaining the full power build-up factor B of equation (3.15). The DARM loop keeps the differential arm length at the DC readout operating point (section 2.4.2). The PRCL and SRCL loops keep the recycling cavities on resonance.

2.4.2 DC readout

Necessity of a local oscillator

A gravitational wave differentially modulates the arm lengths, generating signal sidebands at f_{GW} that exit at the antisymmetric port. Since a photodiode responds to optical power rather than field amplitude, these sidebands must beat against a local oscillator (LO) field at the carrier frequency to produce a photocurrent modulation at f_{GW}.

In first-generation detectors, the LO was provided by RF modulation sidebands injected from the laser input (heterodyne readout). This scheme introduced several problems, including phase noise on the RF oscillator, laser relative intensity noise coupling through sideband amplitude imbalance, and additional shot noise arising because both RF sidebands contribute vacuum fluctuations while only one carries the gravitational wave signal [37, 19].

The DC readout scheme

In DC readout (a homodyne detection scheme), the LO is provided by the carrier itself [38, 19]. A small differential arm length offset ψ is introduced to leak a controlled fraction of the carrier power to the antisymmetric port. The leaked carrier and the gravitational wave sidebands co-propagate through the OMC to B1, where they produce the signal. Because the LO has passed through the full interferometer optical chain, it is automatically mode-matched to the signal field, spatially filtered, and free of the technical noise sources that contaminate an externally generated RF sideband [19].

Optimisation of the offset

The output power at B1 is given by equation (2.2), which will be explained in more depth in chapter 4:

$$P_{\text{B1}}(t) \approx \frac{T_{\text{IM}}}{2}\,|E_0|^2\Big(\psi^2 + 2\psi\,\frac{4\pi GL}{\lambda}\,h(t)\Big). \tag{2.2}$$

Here, T_{IM} is the effective power transmission from the cavity to the antisymmetric port, $|E_0|^2$ is the intracavity power, ψ is the DC readout offset, G is the arm cavity gain, and L the arm length. The gravitational wave signal term scales linearly with ψ, as does the shot noise amplitude $\sqrt{P_{\text{B1}}} \propto \psi$, so the shot-noise-limited signal-to-noise ratio is independent of ψ in the regime $\psi \gg \epsilon$, where ϵ is the residual carrier leakage due to contrast defect. The offset therefore does not affect the fundamental shot-noise limit but determines which noise regime the detector operates in.

If ψ is too small, the gravitational wave signal falls below classical technical noise sources such as laser amplitude and frequency noise, which couple to DARM independently of ψ. A minimum offset is therefore required to keep the detector shot-noise-limited [38].

If ψ is too large, intensity noise at the dark port contains terms scaling as ψ^2 while the signal scales only as ψ, so the intensity-noise-to-signal ratio grows linearly with the offset [39]. In addition, the DC power on B1 scales as ψ^2, eventually risking photodiode saturation, and the additional carrier power exerts growing radiation pressure back-action on the test masses. The operating value of ψ is therefore set during commissioning by balancing these competing effects [39, 14, 10].

2.5 Quantum noise and squeezing

The sensitivity of Advanced Virgo is ultimately limited by quantum noise, which arises from the discrete nature of photons and the Heisenberg uncertainty principle. Quantum noise manifests in two complementary forms: shot noise, which dominates at high frequencies and is caused by the statistical fluctuations in photon arrival times at the photodiode, and radiation pressure noise, which dominates at low frequencies and results from the random momentum kicks imparted to the suspended mirrors by the fluctuating photon flux. These two contributions cannot be simultaneously minimised for a given laser power, establishing a fundamental sensitivity floor known as the standard quantum limit. This section describes both noise sources, their frequency dependence, and the squeezed-light injection technique used in Advanced Virgo to surpass the standard quantum limit.

2.5.1 Shot noise

In the quantum picture, shot noise arises from vacuum fluctuations entering the interferometer through the open output port. Because the interferometer is operated at the dark fringe, input laser noise is reflected back toward the laser, while vacuum noise entering through the antisymmetric port reaches the output photodiode as the dominant noise contribution [40].

For a Michelson interferometer with DC readout, the shot noise power spectral density at the output is frequency-independent and proportional to the carrier power P_0 [24]:

$$S_P^{\text{shot}} = 2P_0\,\hbar\omega_0. \tag{2.3}$$

The noise-to-signal ratio (NSR), expressed as an equivalent gravitational wave strain, is obtained by dividing the shot noise by the signal transfer function $T_{\text{gw}\to P}(f)$ from strain to output power. For a simple Michelson interferometer with arm length L, this gives [24]

$$\text{NSR}(f) = \sqrt{\frac{2\hbar}{P_0\omega_0}}\cdot\frac{2\pi f}{\sin(2\pi f\,L/c)}\,\frac{h}{\sqrt{Hz}}. \tag{2.4}$$

The sensitivity improves as $1/\sqrt{P_0}$: doubling the laser power reduces the shot noise floor by $\sqrt{2}$. The displacement sensitivity does not depend on the DC offset and can be improved by increasing the laser power or by using high-finesse Fabry–Perot cavities.

2.5.2 Radiation pressure noise

At low frequencies, the sensitivity is instead limited by radiation pressure noise. Amplitude fluctuations of the vacuum field entering the dark port exert a stochastic force on the suspended mirrors [24]. The radiation pressure force from a beam of power $P(\Omega)$ acting on a mirror of mass M is

$$F_{\text{rp}}(\Omega) = \frac{2P(\Omega)}{c}. \tag{2.5}$$

Above the suspension resonance frequencies, the mirror behaves as a free mass with mechanical susceptibility $H(\Omega) = -1/(M\Omega^2)$, so the vacuum-driven force produces a displacement noise that scales as $1/\Omega^2$ [24]. The resulting phase noise power spectral density at the output for a single suspended mirror is

$$S_\phi(\Omega) = 2P_0\hbar\omega_0 + \frac{8\hbar P_0^3\omega_0 k_0^2}{M^2c^2\Omega^4}. \tag{2.6}$$

The first term is the shot noise floor and the second is the radiation pressure contribution. The radiation pressure noise spectral density grows as f^{-4} in strain, or equivalently as f^{-2} in strain amplitude spectral density $\sqrt{S_h(f)}$, which is the quantity plotted in sensitivity curves such as figure 2.8. It also increases with laser power, in direct opposition to shot noise, which decreases with power. This trade-off means that no choice of laser power can simultaneously suppress both contributions, establishing a lower bound on the total quantum noise known as the standard quantum limit (SQL) [40, 41]:

$$S_h^{\text{SQL}}(f) = \frac{8\hbar}{M(2\pi f)^2L^2}. \tag{2.7}$$

This equation shows that the Standard Quantum Limit decreases with increasing test-mass and interferometer arm length, and with increasing frequency.

2.5.3 Squeezed light injection

Shot noise originates from vacuum fluctuations entering the interferometer through the open output port. In a coherent state, quantum noise is distributed equally between the amplitude and phase quadratures, with their product bounded by the Heisenberg uncertainty relation. Squeezed states of light redistribute this noise asymmetrically: the uncertainty in one quadrature is reduced below the vacuum level at the expense of increased uncertainty in the conjugate quadrature [42].

By injecting squeezed vacuum states into the output port of the interferometer, the squeezed field spatially overlaps with the circulating carrier at the beam splitter. If the squeezed quadrature is aligned with the readout quadrature of the output field, the photon-counting noise at the B1 photodiode is reduced below the shot noise level without any increase in laser power [42]. The first demonstration of this technique in an operating gravitational wave observatory was achieved at GEO600, where a broadband noise reduction of up to 3.5 dB was observed in the shot-noise-limited frequency band [42].

Advanced Virgo injects squeezed vacuum states generated by parametric down-conversion in an optical parametric oscillator (OPO). During O4, frequency-dependent squeezing is implemented using a 285 m filter cavity that rotates the squeezing angle as a function of frequency, simultaneously reducing shot noise at high frequencies and radiation pressure noise at low frequencies [14, 8].

2.5.4 Thermal noise

In the intermediate frequency band between approximately 30 Hz and 300 Hz, the dominant noise source is thermal noise. By the fluctuation-dissipation theorem, any mechanical system in thermal equilibrium at temperature T exhibits stochastic displacement fluctuations whose spectral density is directly proportional to the mechanical dissipation of the system.

Coating Brownian noise

The most significant contribution is coating Brownian noise, arising from thermally driven fluctuations in the dielectric mirror coating layers. The highly reflective coatings of the Advanced Virgo test masses consist of alternating layers of titania-doped tantala (Ta_2O_5-TiO_2, high refractive index) and silica (SiO_2, low refractive index). The tantala layers dominate the mechanical loss of the coating stack and therefore set the Brownian noise floor [43]. Titania doping reduces the mechanical loss of tantala by approximately 40 % compared to the undoped material [44]. The displacement noise spectral density from coating Brownian noise is approximately white within the detection band and scales inversely with the beam radius on the mirror surface, which is one of the primary motivations for maximising the beam size on the test masses [43, 14].

Suspension thermal noise

A second contribution is suspension thermal noise, arising from mechanical dissipation in the fused silica fibres of the monolithic last-stage suspension [33]. The displacement noise from this source scales as $1/f^2$ and falls below the coating Brownian noise floor above approximately 30 Hz for the Advanced Virgo configuration [14, 45].

2.5.5 Noise budget and sensitivity

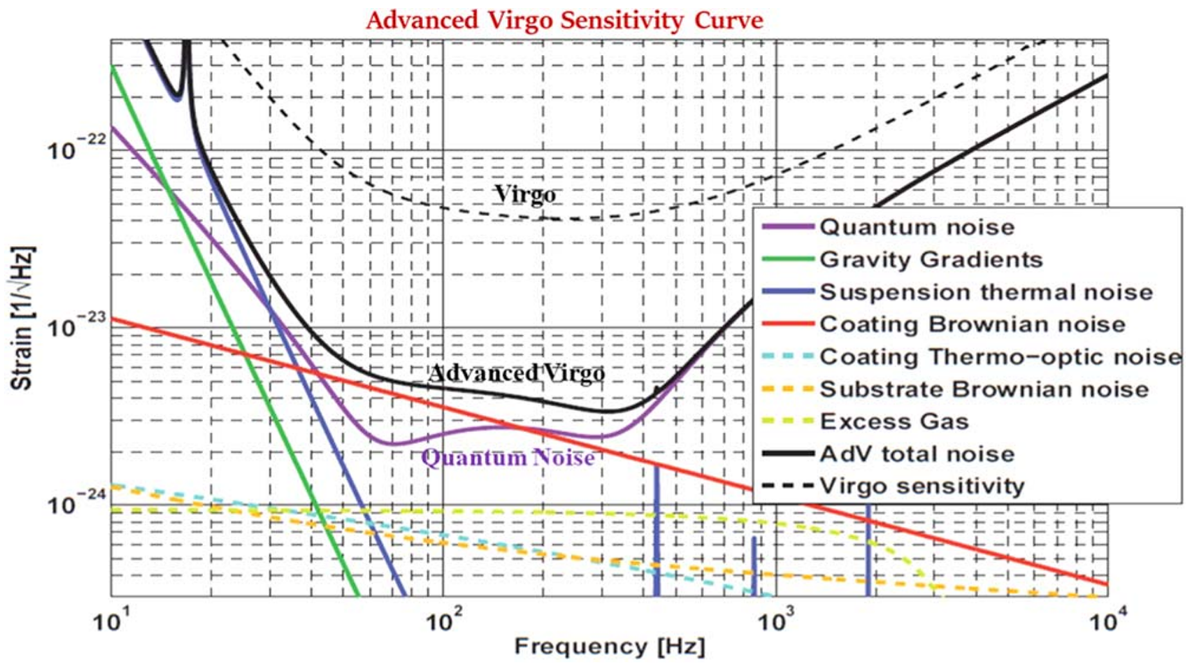

Figure 2.8: Noise budget of Advanced Virgo during the fourth observing run (O4). The total measured sensitivity (black) is compared to the individual noise model contributions [46].

The combined effect of all noise sources produces a strain amplitude spectral density $\sqrt{S_h(f)}$ that defines the detector sensitivity as a function of frequency. At design sensitivity, Advanced Virgo achieves a peak strain sensitivity of approximately 2.2×10^{-24} Hz$^{-1/2}$ near 300 Hz, with performance degrading toward lower frequencies due to seismic and thermal noise and toward higher frequencies due to shot noise.

2.6 Current calibration methods

The conversion of the raw detector output into a calibrated strain time series $h(t)$ requires an accurate knowledge of the detector's response function. This response must be measured independently of the gravitational wave signal itself, which means that known, well-characterised displacements of the test masses must be injected and compared against the detector output. Two primary calibration methods are currently employed in the global gravitational wave detector network: the photon calibrator, which uses radiation pressure from an auxiliary laser, and the Newtonian calibrator, which exploits the gravitational attraction of rotating masses. Both methods are described below, as they provide the context against which the scattered-light calibration technique developed in this thesis should be evaluated.

2.6.1 Photon calibration

The photon calibrator (PCal) is the principal calibration tool used across the LIGO–Virgo–KAGRA detector network [47, 1]. The technique exploits the radiation pressure exerted by an auxiliary, power-modulated laser beam reflecting off one of the suspended end test masses.

When a laser beam of power P reflects from a mirror, it transfers momentum at a rate $2P/c$, exerting a force

$$F_{\text{PCal}}(t) = \frac{2P(t)\cos\theta}{c}, \tag{2.8}$$

where θ is the angle of incidence and c is the speed of light [47]. The resulting displacement of the suspended mirror is

$$x_{\mathrm{PCal}}(f) = \frac{F_{\mathrm{PCal}}(f)}{M(2\pi f)^2}, \tag{2.9}$$

where M is the mirror mass and f is the modulation frequency, valid above the suspension resonance frequencies where the mirror behaves as a free mass [1].

In Advanced Virgo, one PCal system is installed at each end station (NE and WE). Each system consists of an auxiliary laser whose power is modulated at specific calibration frequencies, directed onto the high-reflectivity surface of the end mirror. The modulated laser power reflecting from the mirror is measured by calibrated photodiodes, providing an absolute reference for the injected displacement. Two beams are used, symmetrically displaced about the centre of the mirror face, to minimise the excitation of rotational modes of the suspended test mass and to reduce errors from localised elastic deformation of the mirror surface [47]. In Advanced Virgo, the PCal laser beam is generated with a diode laser source of 1047 nm as shown in the optical sketch of the injection and reflection benches in figure 2.9.

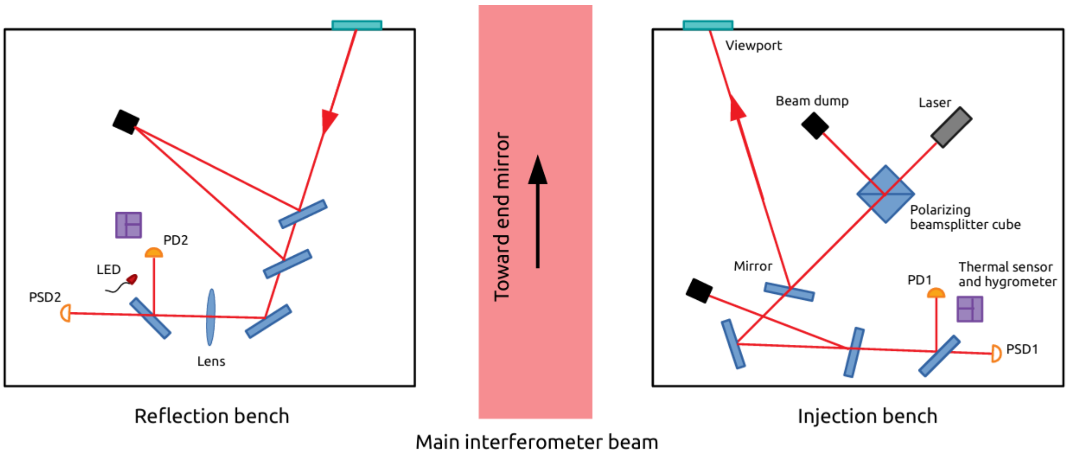

Figure 2.9: Detailed schematic of an Advanced Virgo photon calibrator viewed from the top [1].

The PCal was first used as the primary calibration reference for Virgo during the third observing run (O3) [1, 48]. Prior to O3, the reference method was the free-swinging Michelson technique, in which the interferometer mirrors are allowed to swing freely through optical fringes and the known fringe spacing provides the displacement calibration. The adoption of the PCal as the primary reference brought Virgo in line with the calibration methodology used by the LIGO detectors, enabling an intercalibration of the global network based on a common absolute reference. This intercalibration, based on a transfer standard known as the Gold Standard (a portable, NIST-calibrated power sensor that is circulated between the detector sites), revealed and corrected a systematic bias of 3.92 % that would otherwise have been present in the Virgo calibration [1, 49].

The dominant sources of uncertainty in the PCal measurement are the absolute calibration of the laser power sensors, the knowledge of the beam position on the mirror surface, and the elastic deformation (body modes) of the mirror substrate induced by the localised radiation pressure force. For Advanced Virgo during O3, the total uncertainty budget on the PCal-induced mirror displacement was estimated at 1.36 % for O3a and 1.40 % for O3b [1]. By the start of O4b, improvements in the power calibration chain reduced the Virgo PCal uncertainty to approximately 0.56 % [50].

A key advantage of the PCal is that it operates while the detector is in its most sensitive observing configuration, providing continuous calibration information throughout data-taking periods. Dedicated calibration lines (continuous sinusoidal excitations at fixed frequencies) are injected by the PCal during observation, enabling real-time monitoring of the detector response [48]. However, the PCal's accuracy is fundamentally limited by the precision with which the absolute laser power can be measured and traced to a national metrology standard. Discrepancies between national laser power standards at different metrology institutes (such as NIST in the United States and PTB in Germany) have historically been at the level of several percent, posing a challenge for achieving sub-percent calibration accuracy across the global network [49].

2.6.2 Newtonian calibration

The Newtonian calibrator (NCal) provides a fundamentally different approach to detector calibration. Rather than using electromagnetic radiation pressure, it exploits the Newtonian gravitational attraction between rotating masses and the suspended test mass of the interferometer [51, 50]. This approach has the advantage that the injected signal depends only on Newton's gravitational constant G, the geometry and masses of the rotating bodies, and the distance to the test mass, all of which can in principle be measured with very high precision using purely mechanical and metrological techniques.

The basic concept of using a modulated gravitational field for calibration was proposed in the early days of gravitational wave detector development. The Virgo implementation uses a set of rotor assemblies positioned outside the vacuum chamber near the north end mirror. Each rotor consists of a cylinder divided into two sectors of different densities (or equivalently, two sectors of a denser material separated by two sectors of a lighter material), creating a mass quadrupole. When the rotor spins at angular frequency ω_{rot}, the gravitational force on the nearby test mass is modulated at twice the rotation frequency, $f_{\text{NCal}} = 2f_{\text{rot}}$, because the quadrupole pattern repeats every half turn [51].

The gravitational acceleration induced by a single rotor on the test mass can be computed from Newtonian gravity. For an idealised point-mass quadrupole at distance d from the test mass, the acceleration amplitude scales as $1/d^3$, making the signal strongly dependent on the distance. A realistic calculation must account for the extended geometry of both the rotor and the cylindrical test mass, which is performed using both analytical models and finite-element numerical simulations with a dedicated code called FROMAGE (FRee Object Model for Advanced Gravitational Experiments) [51, 50].

For the fourth observing run (O4), the Virgo NCal system was significantly upgraded from the prototype tested during O3. The O4 configuration consists of two triplets of rotor actuators positioned around the north end mirror of the interferometer [50]. One triplet, called "Far", is placed at approximately 2.1 m from the mirror and uses aluminium rotors. The second triplet, called "Near", is placed at approximately 1.7 m from the mirror and uses polyvinyl chloride (PVC) rotors to minimise coupling from the magnetic fields generated by the spinning metallic rotors [50]. Each triplet includes rotors on either side of the mirror (designated North and South) and a third rotor (East) placed symmetrically to the North rotor relative to the beam axis. This redundant configuration provides independent measurements and allows the mirror position relative to the rotor array to be determined from the data itself, significantly reducing the dominant geometric uncertainty [51, 50].

At the start of O4b, the NCal calibration uncertainty was estimated at 0.12 %, substantially better than both the Virgo PCal uncertainty of 0.56 % and the LIGO PCal uncertainty of 0.29 % [50]. This improvement made the NCal the new absolute calibration reference for the Virgo detector, with the PCal serving as an independent cross-check. The NCal's superior accuracy arises because its systematic uncertainties are dominated by purely geometric quantities (rotor dimensions, distances, and alignment) that can be measured with high precision, rather than by laser power calibration, which is the limiting factor for the PCal. A top view of the system used in O4b is shown in figure 2.10.

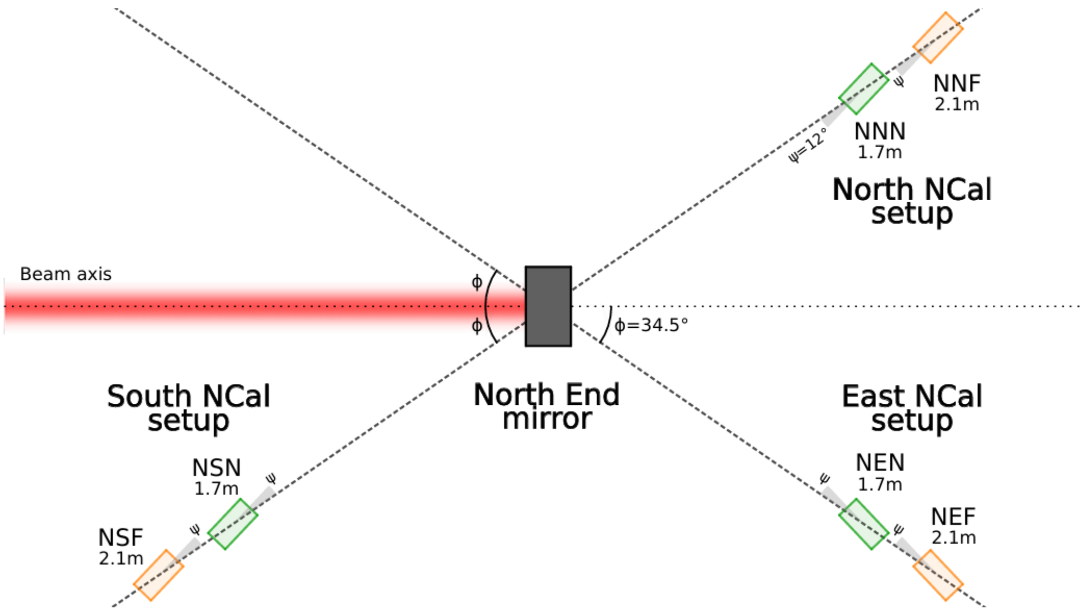

Figure 2.10: Newtonian calibrator system of the Virgo interferometer at the start of O4b. Each rotor is identified by three letters: the beam direction (N for North), the NCal position relative to the mirror (N for North, S for South and E for East), and the triplet identifier (N for Near and F for Far). The green rotors (Near) are made of PVC, while the orange rotors (Far) are in aluminium [51].

The NCal has two practical limitations. First, the signal amplitude is small: the gravitational acceleration is of order 10^{-13} m s^{-2}, requiring long integration times to achieve sufficient signal-to-noise ratio. Second, the NCal excitation frequency is fixed by the rotor speed, which is typically set to produce signals at frequencies of a few hertz to a few tens of hertz, within the sensitive band of the detector but limited to a narrow frequency range [50]. Despite these limitations, the NCal has established itself as a powerful calibration tool that is complementary to the PCal and provides a fundamentally independent cross-check of the absolute detector calibration.

3

Mathematical background

3.1 Gaussian beams and Hermite–Gaussian modes

The laser light used in gravitational wave detectors such as Advanced Virgo is not a plane wave. Rather, it is a Gaussian beam: a solution to the paraxial wave equation (an approximation of a propagating electromagnetic field) that remains Gaussian in cross-section at every point along its propagation axis. Understanding the spatial structure of these beams is essential for describing how light couples into an optical cavity and for quantifying the mode-matching condition that underlies the scattered light model developed in section 4.2.

3.1.1 The fundamental Gaussian beam

A monochromatic laser field propagating along the z-axis can be written in the paraxial approximation as

$$E(x, y, z, t) = E_0\, u(x, y, z)\, e^{i(\omega t - kz)}, \tag{3.1}$$

where $u(x, y, z)$ is a slowly varying complex envelope, ω the angular frequency and k the wave vector of the propagating field [6]. The simplest solution is the fundamental Gaussian mode TEM$_{00}$:

$$u_{00}(x, y, z) = \frac{w_0}{w(z)} \exp\left(-\frac{x^2 + y^2}{w(z)^2}\right) \exp\left(-i\frac{k(x^2 + y^2)}{2R(z)} + i\zeta(z)\right). \tag{3.2}$$

Three z-dependent quantities fully characterise the beam [52]:

- The **beam radius** $w(z)$, which defines the radial distance at which the field amplitude falls to $1/e$ of its on-axis value:

$$w(z) = w_0\sqrt{1 + \left(\frac{z}{z_R}\right)^2}. \tag{3.3}$$

  Here w_0 is the beam waist, defined as the minimum beam radius (by convention located at $z = 0$), and $z_R = \pi w_0^2/\lambda$ is the Rayleigh range, the propagation distance over which the beam area doubles.

- The **radius of curvature** of the wavefront:

$$R(z) = z\left(1 + \left(\frac{z_R}{z}\right)^2\right). \tag{3.4}$$

  At the waist, the wavefront is flat ($R \to \infty$). Far from the waist, the beam resembles a spherical wave ($R \approx z$).

- The **Gouy phase**:

$$\zeta(z) = \arctan\left(\frac{z}{z_R}\right), \tag{3.5}$$

  an additional longitudinal phase accumulated beyond the ordinary kz propagation phase. It ranges from $-\pi/2$ to $+\pi/2$ as the beam passes through the focal point.

It is clear by looking at equation (3.3) that the far field, where $z \gg z_R$, the beam radius grows linearly with propagation distance and the beam expands like a cone. The half-angle of this cone is the divergence angle, defined as

$$\theta = \frac{w_0}{z_R} = \frac{\lambda}{\pi w_0}. \tag{3.6}$$

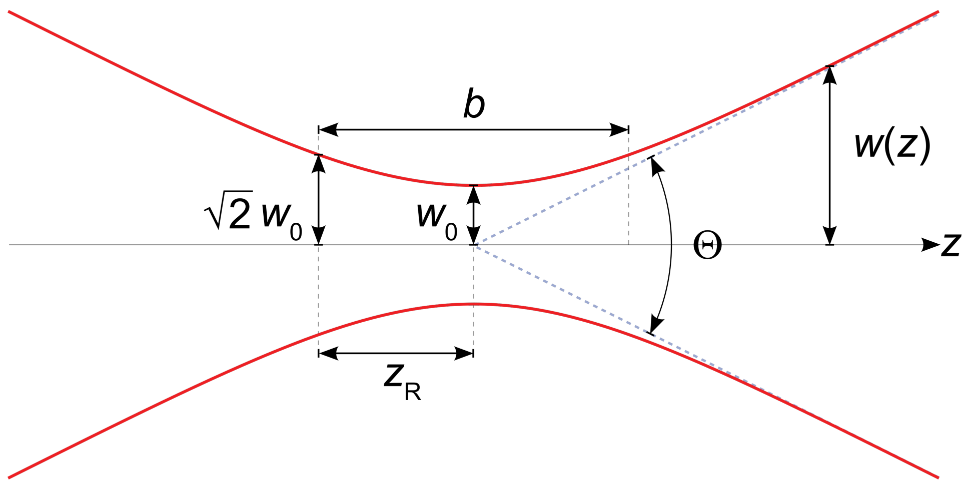

Figure 3.1: Diagram of a Gaussian beam, illustrating the beam waist w_0, the beam radius $w(z)$, the Rayleigh range z_R, and the divergence angle θ [53].

This angle characterises how rapidly the beam spreads beyond the waist and reflects a fundamental consequence of diffraction: a more tightly focused beam (smaller w_0) diverges more rapidly. The product $\theta \cdot w_0$ is therefore a constant set only by the wavelength. A visualisation of a Gaussian beam is shown in figure 3.1.

A particularly compact description of a Gaussian beam is provided by the complex beam parameter $q(z)$:

$$q(z) = z + iz_R = \left(\frac{1}{R(z)} - \frac{i\lambda}{\pi w(z)^2} \right)^{-1}. \tag{3.7}$$

This parameter encodes both the wavefront curvature and the beam radius in a single complex number [24].

The propagation of paraxial rays through optical elements can be described compactly using 2×2 ray-transfer matrices, commonly known as ABCD matrices [24]. In this formalism, a ray at any plane is characterised by its transverse displacement y from the optical axis and its slope $\theta = dy/dz$. An optical element (a lens, mirror, free-space propagation, or interface) maps the input ray (y_1, θ_1) to the output ray (y_2, θ_2) via

$$\begin{pmatrix} y_2 \\ \theta_2 \end{pmatrix} = \begin{pmatrix} A & B \\ C & D \end{pmatrix} \begin{pmatrix} y_1 \\ \theta_1 \end{pmatrix}. \tag{3.8}$$

For example, free-space propagation over a distance d corresponds to $A = 1, B = d, C = 0, D = 1$, while a thin lens of focal length f has $A = 1, B = 0, C = -1/f, D = 1$. Cascading multiple elements is achieved by multiplying their matrices in reverse order [6]. The ABCD formalism extends naturally to Gaussian beams: rather than transforming the ray coordinates, one transforms the complex beam parameter $q(z)$ according to

$$q' = \frac{Aq + B}{Cq + D}. \tag{3.9}$$

The q-parameter is therefore a powerful tool for designing optical systems and verifying mode-matching conditions analytically, and is used extensively for describing and shaping scattered light in sections 4.2 and 5.4.

3.1.2 Hermite-Gaussian modes

The TEM$_{00}$ mode is the lowest-order member of a complete, orthogonal family of solutions to the paraxial wave equation. The Hermite-Gaussian (HG) modes TEM$_{mn}$ are obtained by multiplying the Gaussian envelope by Hermite polynomials H_m and H_n in the x and y directions, respectively [54]:

$$u_{mn}(x,y,z) = \frac{w_0}{w(z)} H_m\left(\frac{\sqrt{2}\,x}{w(z)}\right) H_n\left(\frac{\sqrt{2}\,y}{w(z)}\right) \exp\left(-\frac{x^2+y^2}{w(z)^2}\right) \exp\left(-i\frac{k(x^2+y^2)}{2R(z)}\right) e^{i(m+n+1)\zeta(z)}. \tag{3.10}$$

The integers m and n are the transverse mode orders in x and y, and $N = m + n$ is the total mode order. Three properties of the HG modes are particularly important [6]:

- **Intensity pattern:** The intensity $|u_{mn}|^2$ displays $(m+1)(n+1)$ lobes arranged in a rectangular grid. TEM$_{00}$ has a single central lobe, TEM$_{10}$ has two horizontally separated lobes, and TEM$_{20}$ has three lobes, as illustrated in figure 3.2.

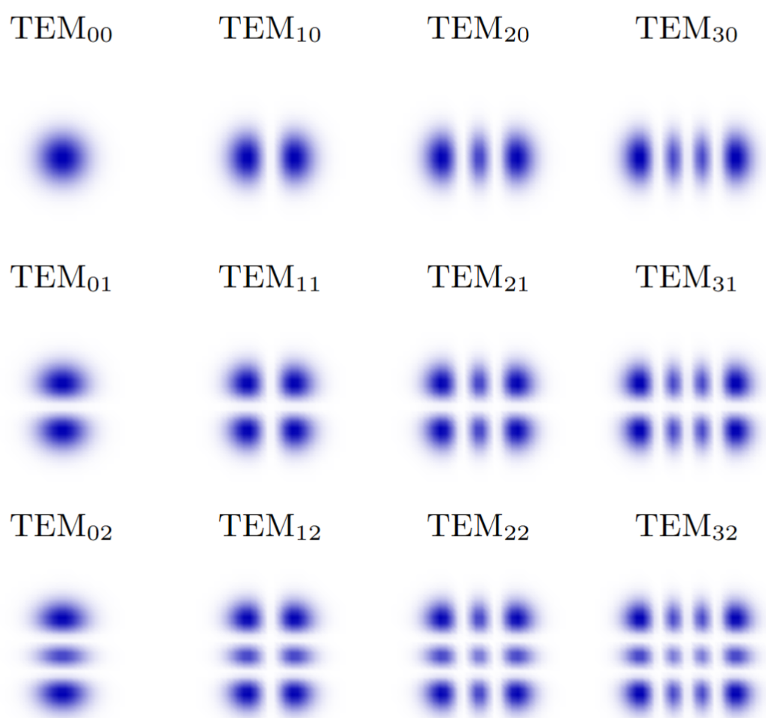

Figure 3.2: Normalised intensity patterns $|u_{mn}|^2$ of the Hermite-Gaussian modes TEM$_{mn}$ for $m \in \{0, 1, 2, 3\}$ and $n \in \{0, 1, 2\}$. Moving right increases the number of lobes in the x-direction ($m + 1$ lobes), moving down increases the number of lobes in the y-direction ($n + 1$ lobes).

- **Gouy phase:** Each mode TEM$_{mn}$ accumulates an additional phase shift of $(m+n+1)\zeta(z)$, where $\zeta(z)$ is the Gouy phase of the TEM$_{00}$ mode defined in equation 3.5. Since this phase shift depends on the mode order, higher-order modes acquire different total phases upon propagation. As will be discussed in Section 3.2, this leads to a frequency separation between transverse modes in optical cavities. This effect is exploited in arm cavity scans to identify and measure individual higher-order mode contributions [21].

- **Orthogonality and completeness:** The HG modes form a complete orthonormal basis for paraxial optical fields. Any field sharing the same beam axis can be decomposed as [24]

$$E(x, y, z) = \sum_{m,n} c_{mn}\, u_{mn}(x, y, z), \tag{3.11}$$

  where the coefficients are given by the overlap integral

$$c_{mn} = \iint u^*_{mn}(x, y, z)\, E(x, y, z)\, \mathrm{d}x\, \mathrm{d}y. \tag{3.12}$$

  The power fraction carried by mode TEM$_{mn}$ is $|c_{mn}|^2$. This decomposition is the mathematical basis for quantifying both mode mismatch and misalignment, and is used in section 5.4.1.

3.2 Fabry–Perot cavity: fields and gain

The arm cavities of Advanced Virgo are Fabry–Perot resonators, formed by two facing mirrors that trap light by allowing it to bounce back and forth many times before leaking out. This multi-pass storage amplifies the interaction between the light and any physical displacement of the mirrors, which is precisely what makes interferometric gravitational wave detection possible. This section derives the key parameters of a Fabry–Perot cavity that are used throughout the scattered light model of section 4.2.

3.2.1 Round-trip field factor and finesse

Consider a linear cavity formed by two mirrors with amplitude reflectivities r_1 and r_2 and power transmissions $T_1 = 1 - r_1^2$ and $T_2 = 1 - r_2^2$, separated by a length L. An input field E_{in} enters through the input mirror. After each round trip of length $2L$, the

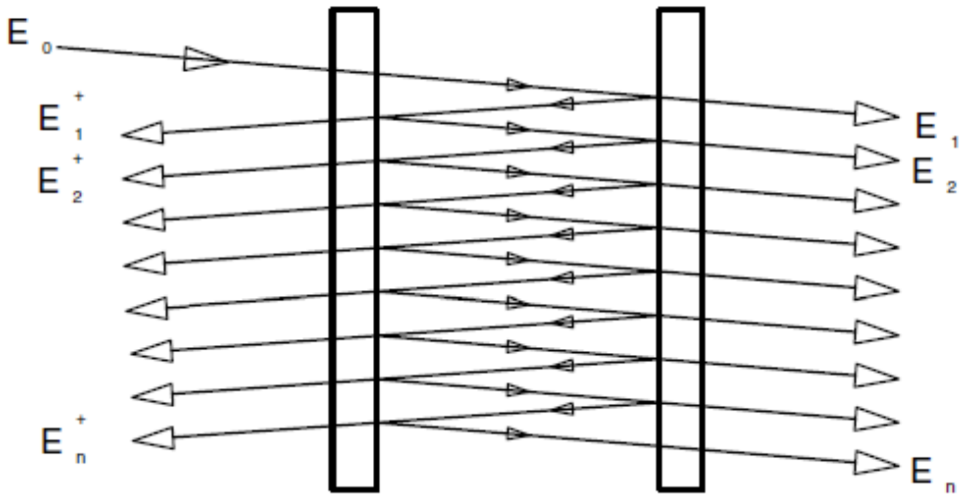

Figure 3.3: Working principle of a Fabry–Perot cavity showing that the intracavity field is the coherent sum of all multiply-reflected contributions [56].

field accumulates a phase $\phi = \omega \cdot 2L/c$ and is attenuated by the round-trip amplitude factor

$$r_{\text{rt}} = r_1 r_2 \sqrt{1 - \Lambda_{\text{arm}}}, \tag{3.13}$$

where Λ_{arm} models any additional round-trip power loss (absorption, scattering, clipping). The intracavity field is the coherent sum of all multiply-reflected contributions (as seen in figure 3.3), forming a geometric series that converges to the steady-state cavity field [55]:

$$E_{\text{cav}} = \frac{i t_1}{1 - r_{\text{rt}}\, e^{i\phi}}\, E_{\text{in}}. \tag{3.14}$$

Here, $t_1 = \sqrt{T_1}$ is the input mirror amplitude transmission. The intracavity power $P_{\text{cav}} = |E_{\text{cav}}|^2$ is maximised at resonance ($\phi = 2\pi n$), where the denominator of equation (3.14) is minimised. The ratio of the circulating power at resonance to the input power defines the power build-up factor:

$$B = \left. \frac{P_{\text{cav}}}{P_{\text{in}}} \right|_{\text{res}} = \frac{T_1}{(1 - r_{\text{rt}})^2}. \tag{3.15}$$

The sharpness of the resonance is characterised by the finesse $\mathcal{F}$, defined as the ratio of the free spectral range (FSR) to the resonance linewidth (full width at half maximum, FWHM) [6]:

$$\mathcal{F} = \frac{\nu_{\text{FSR}}}{\delta\nu_{\text{FWHM}}} = \frac{\pi\sqrt{r_{\text{rt}}}}{1 - r_{\text{rt}}}. \tag{3.16}$$

The free spectral range $\nu_{\text{FSR}} = c/2L$ is the frequency spacing between consecutive resonances and equals approximately 50 kHz for the 3 km Virgo arm cavities. The resonance linewidth is verified to be approximately 100 Hz [21, 57]. A plot of the cavity power enhancement is shown in figure 3.4.

3.2.2 Frequency-domain cavity response

When a field at frequency $\omega_0 + 2\pi f$ (offset by f from the carrier resonance) is injected into the cavity, it experiences a frequency-dependent gain. Starting from the steady-state cavity field of equation (3.14), the exact cavity gain at frequency offset f from resonance is

$$G_{\text{exact}}(f) = \frac{1}{1 - r_{\text{rt}}\, e^{-i2\pi f\, t_{\text{RT}}}}, \tag{3.17}$$

where the carrier resonance condition $e^{-i\omega_0 t_{\text{RT}}} = 1$ has been used to absorb the carrier phase. The value t_{RT} represents the time a photon needs to make one round trip in the cavity. This is simply equal to $2L/c$. In the limit $f \ll \nu_{\text{FSR}} = 1/t_{\text{RT}}$, the exponential can be expanded to first order:

$$e^{-i2\pi f\, t_{\text{RT}}} \approx 1 - i2\pi f\, t_{\text{RT}}, \tag{3.18}$$

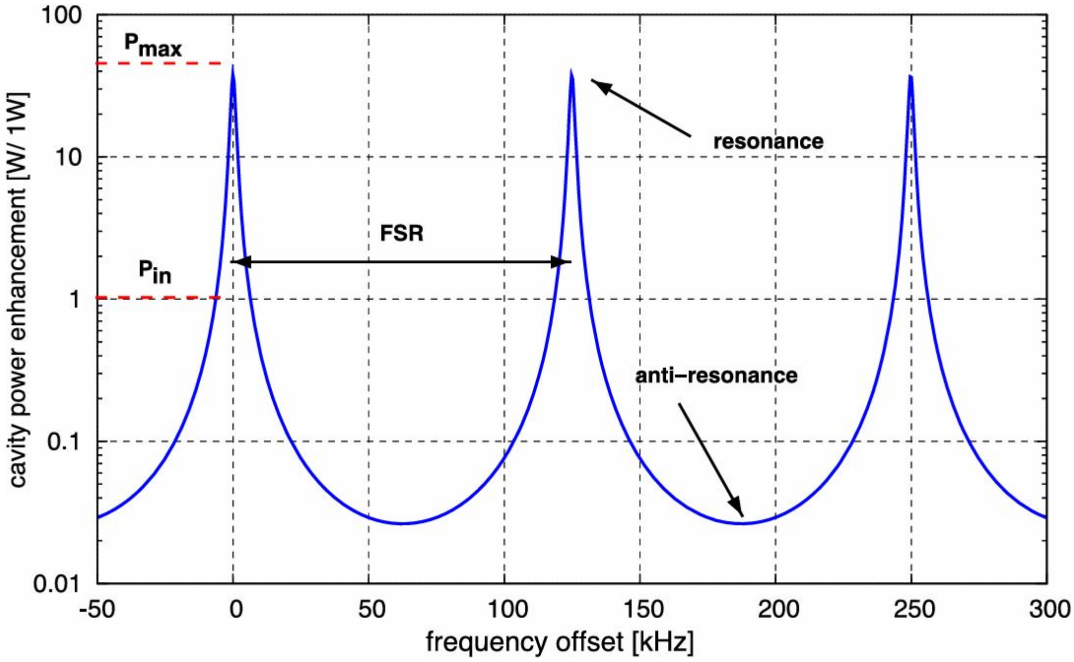

Figure 3.4: Power enhancement in a two-mirror cavity as a function of the offset of the laser-light frequency with respect to the nominal value. The peaks mark the resonances of the cavity, at which the injected light is resonantly enhanced. The frequency spacing between consecutive peaks is the free spectral range (FSR) [6].

so that the denominator of equation (3.17) becomes

$$1 - r_{\text{rt}}\, e^{-i2\pi f\, t_{\text{RT}}} \approx 1 - r_{\text{rt}} + i2\pi f\, r_{\text{rt}}\, t_{\text{RT}}. \tag{3.19}$$

Factoring out $(1 - r_{\text{rt}})$ and using $r_{\text{rt}} \approx 1$ in the high-finesse limit gives

$$1 - r_{\text{rt}}\, e^{-i2\pi f\, t_{\text{RT}}} \approx (1 - r_{\text{rt}}) \left(1 + i\frac{f}{f_{\text{arm}}}\right), \tag{3.20}$$

where the cavity pole frequency is identified as

$$f_{\text{arm}} = \frac{1 - r_{\text{rt}}}{2\pi\, t_{\text{RT}}} = \frac{(1 - r_{\text{rt}})\, c}{4\pi L} = \frac{\nu_{\text{FSR}}}{2\mathcal{F}}. \tag{3.21}$$

Substituting equation (3.20) into equation (3.17) gives the single-pole approximation of the cavity gain:

$$G(f) = \frac{1}{1 - r_{\text{rt}}} \cdot \frac{1}{1 + if/f_{\text{arm}}}, \tag{3.22}$$

where $1/(1 - r_{\text{rt}})$ is the on-resonance field build-up factor [24].

The cavity pole sets the bandwidth of the arm cavity, which acts as a low-pass filter for the injected signal. Signals at frequencies $f \ll f_{\text{arm}}$ are amplified by the full on-resonance gain $1/(1 - r_{\text{rt}}) \approx \mathcal{F}/\pi$. Signals above f_{arm} are increasingly suppressed as $1/f$. This behaviour is illustrated in figure 3.5 and plays a central role in the calibration method developed in section 5.5.2.

The Gouy phase accumulated per round trip determines the frequencies at which higher-order modes resonate. For a stable resonator, the round-trip Gouy phase is determined by the cavity geometry [24]:

$$\Delta\zeta_{\text{RT}} = 2\,\text{arccos}(\sqrt{g_1 g_2})\,. \tag{3.23}$$

Here, $g_i = 1 - L/R_i$ are the stability parameters of the two mirrors with radii of curvature R_i. Converting this phase condition into a frequency condition using the free spectral range ($\nu_{\text{FSR}} = c/(2L)$) yields the resonance frequencies of the transverse modes:

$$\nu_{qmn} = \nu_{\text{FSR}} \left[q + \frac{(m + n + 1)}{\pi}\, \text{arccos}\sqrt{g_1 g_2}\right]. \tag{3.24}$$

The TEM$_{mn}$ mode therefore resonates at a frequency shifted from the TEM$_{00}$ resonance by

$$\Delta\nu_{mn} = (m + n)\, \frac{\text{arccos}\sqrt{g_1 g_2}}{\pi}\, \nu_{\text{FSR}}. \tag{3.25}$$

This frequency splitting makes arm cavity scans a powerful tool for measuring the mirror radii of curvature [58].

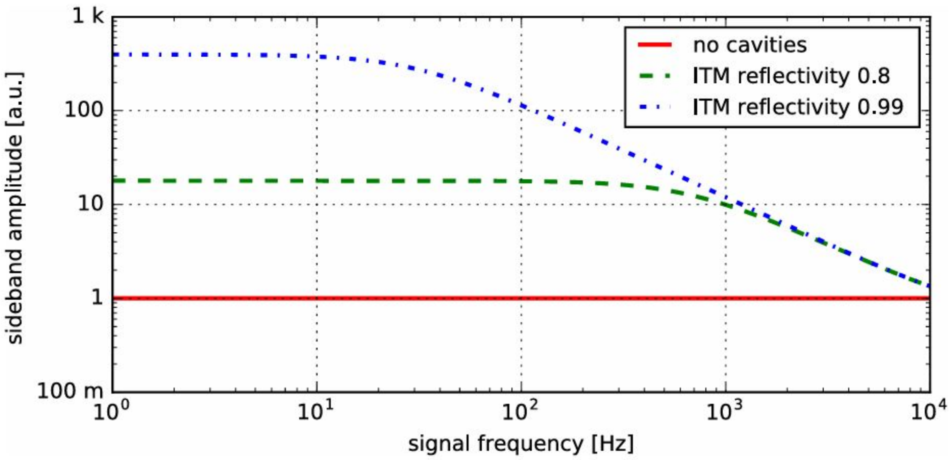

Figure 3.5: Signal sideband amplitude for a Fabry–Perot–Michelson interferometer. The signal is a differential arm length change detected at the antisymmetric output port, shown as a function of signal frequency. The solid red trace at an amplitude of 1 corresponds to the case without arm cavities. The other two traces show the increased amplitude for different input mirror reflectivities [6].

3.3 Phase modulation and sideband generation

Many of the sensing and control systems in Advanced Virgo rely on introducing additional frequency components into the laser field. These components, called sidebands, are generated by phase-modulating the carrier before it enters the interferometer. This section derives the sideband structure that results from phase modulation, forming the mathematical foundation for the cavity response to an oscillating scattering element developed in section 5.2.3.

3.3.1 Phase modulation of a monochromatic field

Consider a monochromatic laser field with angular frequency ω_0, whose electric field is

$$E_0(t) = A\, e^{i\omega_0 t}. \tag{3.26}$$

If this field reflects off a surface that oscillates sinusoidally at angular frequency Ω with displacement amplitude x_0, the round-trip optical path length changes by $2x_0\,\sin(\Omega t)$, and the reflected field accumulates an additional phase [24]:

$$E(t) = A\, e^{i\omega_0 t}\, e^{i\beta\,\sin(\Omega t)}, \tag{3.27}$$

where the modulation depth β is defined as

$$\beta = \frac{2\pi \cdot 2x_0}{\lambda} = \frac{4\pi x_0}{\lambda}. \tag{3.28}$$

This quantity gives the peak phase excursion in radians. For the small displacements typical in gravitational wave detectors, $\beta \ll 1$.

3.3.2 Jacobi–Anger expansion

The complex exponential $e^{i\beta\,\sin(\Omega t)}$ is periodic and can be expanded exactly in a Fourier series using the Jacobi–Anger identity [59]:

$$e^{i\beta\,\sin(\Omega t)} = \sum_{n=-\infty}^{+\infty} J_n(\beta)\, e^{in\Omega t}, \tag{3.29}$$

where $J_n(\beta)$ is the Bessel function of the first kind of order n. Substituting into equation (3.27) gives

$$E(t) = A \sum_{n=-\infty}^{+\infty} J_n(\beta)\, e^{i(\omega_0 + n\Omega)t}. \tag{3.30}$$

This result shows that a phase-modulated field is equivalent to a superposition of discrete frequency components: the carrier at ω_0 and an infinite set of sidebands at frequencies $\omega_0 \pm \Omega$, $\omega_0 \pm 2\Omega$, and so on, symmetrically spaced around the carrier by integer multiples of the modulation frequency Ω.

3.3.3 Sideband amplitudes

The amplitude of each frequency component is determined by the corresponding Bessel function $J_n(\beta)$. For small modulation depth $\beta \ll 1$, the Bessel functions simplify considerably using the small-argument approximations:

$$J_0(\beta) \approx 1 - \frac{\beta^2}{4} \approx 1, \tag{3.31}$$

$$J_1(\beta) \approx \frac{\beta}{2}, \tag{3.32}$$

$$J_n(\beta) \approx \frac{1}{n!}\left(\frac{\beta}{2}\right)^n \approx 0 \quad \text{for } n \geq 2. \tag{3.33}$$

$$J_{-n}(\beta) = (-1)^n J_n(\beta), \quad \text{for } n > 0. \tag{3.34}$$

In this limit, only the carrier and the first-order sidebands at $\omega_0 \pm \Omega$ carry significant power, and the field reduces to

$$E(t) \approx A\left[J_0(\beta)\, e^{i\omega_0 t} + J_1(\beta)\, e^{i(\omega_0+\Omega)t} + J_{-1}(\beta)\, e^{i(\omega_0-\Omega)t}\right]. \tag{3.35}$$

Using $J_{-1}(\beta) = -J_1(\beta)$, this can be written as

$$E(t) \approx A\, e^{i\omega_0 t}\left[1 + \frac{\beta}{2}\, e^{i\Omega t} - \frac{\beta}{2}\, e^{-i\Omega t}\right] = A\, e^{i\omega_0 t}\left[1 + i\beta \sin(\Omega t)\right], \tag{3.36}$$

recovering the first-order Taylor expansion of equation (3.27) directly. The power fraction carried by each first-order sideband relative to the total is $(\beta/2)^2$.

3.3.4 Pound–Drever–Hall sensing

The Pound–Drever–Hall (PDH) technique generates the error signal used to keep the arm cavities on resonance [60, 61]. It exploits the phase-modulated sideband structure discussed in the previous section, and its operating principle is shown in Figure 3.6.

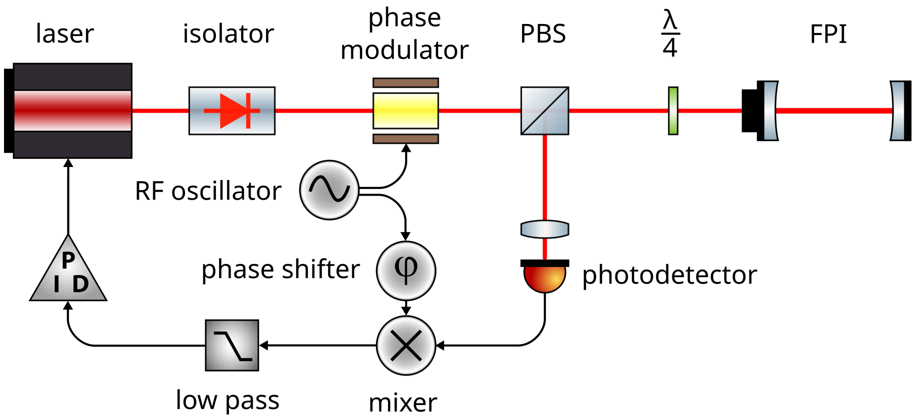

Figure 3.6: Schematic of a Pound–Drever–Hall locking setup. Phase-modulated light is sent into the cavity and the reflected beam is detected by a fast photodiode and mixed with the original modulation signal. After low-pass filtering, the resulting error signal is fed back to an actuator that controls the cavity length. Adapted from [60].

Phase-modulated light with carrier frequency ω_0 and sidebands at $\omega_0 \pm \Omega$ is directed onto a Fabry–Perot cavity whose complex reflection coefficient is

$$R(\omega) = \frac{r_1 - r_2\, e^{2i\omega L/c}}{1 - r_1 r_2\, e^{2i\omega L/c}}, \tag{3.37}$$

with r_1, r_2 the mirror amplitude reflectivities and L the cavity length. Because the cavity is a linear time-invariant system, each frequency component reflects independently and the reflected field is

$$E_{\text{ref}}(t) = E_0\, e^{i\omega_0 t}\left[R(\omega_0)\, J_0(\beta) + R(\omega_0+\Omega)\, J_1(\beta)\, e^{i\Omega t} - R(\omega_0-\Omega)\, J_1(\beta)\, e^{-i\Omega t}\right]. \tag{3.38}$$

The modulation frequency Ω is chosen far above the cavity linewidth, so the sidebands reflect promptly with $R(\omega_0 \pm \Omega) \approx -1$, while the carrier experiences the full frequency-dependent cavity response.

The reflected power $P_{\text{ref}} = |E_{\text{ref}}|^2$ contains terms at DC, at Ω, and at 2Ω. The terms oscillating at Ω arise from the beat between the reflected carrier and the promptly reflected sidebands and carry the phase information of $R(\omega_0)$. To extract this information, the photodiode signal is mixed with a phase-shifted copy of the modulation signal and sent through a low-pass filter. In the regime $\Omega \gg \delta\nu_{\text{FWHM}}$, the resulting PDH error signal reduces to [60]

$$\epsilon_{\text{PDH}} \propto \text{Im}[R(\omega_0)]\,. \tag{3.39}$$

This error signal has two properties that make it ideally suited for cavity locking. First, it is antisymmetric in the detuning $\delta = \omega_0 - \omega_{\text{res}}$: it changes sign across resonance, telling the servo loop not only how far the laser is from resonance but also in which direction to correct. Second, it is linear near resonance:

$$\epsilon_{\text{PDH}} \approx -\frac{8\,J_0(\beta)\,J_1(\beta)}{\delta\nu_{\text{FWHM}}}\,\delta, \tag{3.40}$$

so for small detunings the signal is directly proportional to the frequency offset. Additionally, because the error signal is derived from a phase measurement, it responds to the laser frequency independently of intensity fluctuations, which is a decisive advantage over simpler methods such as a side-of-fringe lock [60].

The zero crossing at resonance has a clear physical origin. At exact resonance the carrier is fully coupled into the cavity and its reflected amplitude vanishes, $R(\omega_0) = 0$. The two sidebands still reflect, but they can only beat against each other at frequency 2Ω, which is rejected by the low-pass filter. A signal at Ω requires a beat between the carrier and a sideband. Without a reflected carrier this beat is absent and the error signal is exactly zero. Just off resonance a small carrier amplitude reappears with $\text{Im}[R(\omega_0)] \neq 0$, restoring the beat at Ω and producing an error signal whose sign indicates which side of resonance the cavity is on. As shown in Figure 3.6, this signal is fed back through a servo loop to an actuator that adjusts the cavity length, closing the feedback loop.

In Advanced Virgo, the PDH technique is applied with modulation frequencies far above the arm cavity linewidth [6], ensuring prompt sideband reflection and a clean phase reference. The resulting error signals drive the length control loops that maintain the cavities within a tiny fraction of a linewidth of resonance, preserving the power build-up factor on which detector sensitivity depends.

3.3.5 Breakdown of the small-modulation approximation

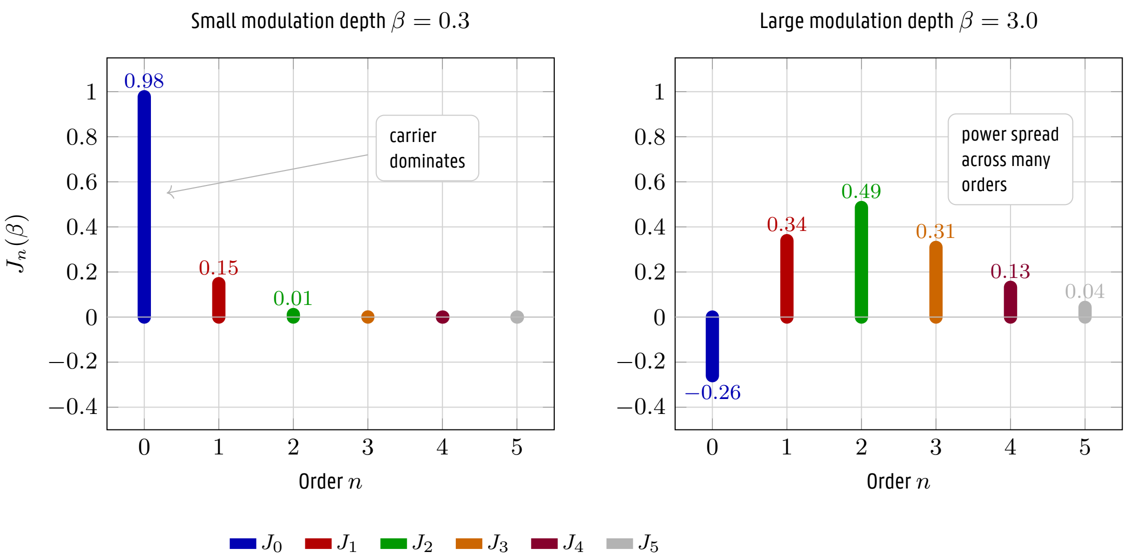

Figure 3.7: Bessel function amplitudes $J_n(\beta)$ for small and large modulation depth. (a) For $\beta = 0.3$, almost all power resides in the carrier (J_0) and the first sideband (J_1), justifying the small-modulation approximation. (b) For $\beta = 3.0$, power is distributed across many orders and the small-modulation approximation breaks down.

The sideband decomposition of equation (3.30) is exact for any modulation depth β. The practical utility of the frequency-domain description, however, depends on how many sideband orders carry significant power. In the small-modulation regime $\beta \ll 1$, only the first-order sidebands matter and the three-component approximation of equation (3.36) is sufficient. This is the regime relevant to the modulation sidebands used for PDH sensing in Advanced Virgo, where $\beta \sim 0.1$–0.3.

In this thesis, this approximation is not valid because the displacement amplitudes evaluated here are $x_B \sim 10\,\mu$m, giving a modulation depth of

$$\beta_B = \frac{4\pi x_B}{\lambda} \approx 118. \tag{3.41}$$

At this modulation depth, the Bessel function coefficients $J_n(\beta)$ remain significant for orders up to $n \sim \beta$, as illustrated in figure 3.7. The phase-modulated field is therefore a superposition of hundreds of sidebands symmetrically spaced around the carrier, each at a frequency $\omega_0 + n\Omega_B$ with n ranging from roughly $-\beta$ to $+\beta$.

3.4 Mode matching and the overlap integral

When a beam is injected into an optical cavity, only the fraction of the incident field that spatially coincides with a cavity eigenmode can build up resonantly. Any mismatch between the injected beam and the cavity mode couples into higher-order modes that either do not resonate in the cavity or resonate at different frequencies. Quantifying this coupling is essential for the scattered light model developed in section 4.2.

3.4.1 The overlap integral

The fraction of an incident field E_{inc} that couples into a target mode u_{mn} is given by the overlap integral

$$c_{mn} = \langle u_{mn} | E_{\text{inc}} \rangle = \iint_{-\infty}^{+\infty} u_{mn}^*(x, y)\, E_{\text{inc}}(x, y)\, \mathrm{d}x\, \mathrm{d}y, \tag{3.42}$$

where both fields are evaluated in the same transverse plane and normalised such that $\langle u_{mn}|u_{mn}\rangle = 1$. The power fraction coupled into mode TEM$_{mn}$ is $|c_{mn}|^2$, and the total coupled power satisfies

$$\sum_{m,n} |c_{mn}|^2 = 1, \tag{3.43}$$

3.4.2 Gaussian beam overlap

For two Gaussian beams with complex beam parameters q_1 and q_2 evaluated in the same plane, the overlap integral can be computed analytically. Writing both beams in terms of their waist sizes w_1, w_2 and waist positions z_1, z_2 along the optical axis, the mode-matching efficiency is [55]

$$\eta = |c_{00}|^2 = \frac{4}{\left(\dfrac{w_1}{w_2} + \dfrac{w_2}{w_1}\right)^2 + \left(\dfrac{\lambda\, \Delta z}{\pi w_1 w_2}\right)^2}, \tag{3.44}$$

where $\Delta z = z_1 - z_2$ is the distance between the waist positions. Perfect mode matching ($\eta = 1$) requires both conditions to be satisfied simultaneously:

$$w_1 = w_2 \qquad \text{(waist size match)}, \tag{3.45}$$

$$z_1 = z_2 \qquad \text{(waist position match)}. \tag{3.46}$$

Any deviation from either condition reduces η below unity. Equation (3.44) shows that the two contributions are independent: a waist size mismatch degrades η even when the waist positions coincide, and vice versa.

3.4.3 Effect of beam tilt

If the incident beam arrives at a small transverse angle α with respect to the cavity axis, the field acquires a linear phase ramp across its transverse profile. Taking the tilt to lie in the x-direction, the field at the end mirror plane becomes

$$E_{\text{sc}}(x, y) = E_{\text{Gaussian}}(x, y)\, e^{ikx \sin\alpha} \approx E_{\text{Gaussian}}(x, y)\, e^{ikx\alpha}, \tag{3.47}$$

where the small-angle approximation $\sin\alpha \approx \alpha$ has been used. This phase ramp redistributes power from the fundamental mode into higher-order transverse modes. Since the tilt lies entirely in the x-direction, only modes with $n = 0$ are excited. As derived in appendix A, the power coupled into TEM$_{m0}$ follows a Poisson distribution in the mode index m:

$$|c_{m0}|^2 = \frac{\mu^m}{m!} e^{-\mu}, \quad \mu = \left(\frac{\pi w \alpha}{\lambda}\right)^2, \tag{3.48}$$

with mean $\langle m \rangle = \mu$. For small tilts ($\mu \ll 1$), the dominant higher-order contribution is the TEM$_{10}$ mode with $|c_{10}|^2 \approx \mu$, confirming that beam tilt couples primarily into the first-order mode along the tilt direction.

Setting $m = 0$ gives the fraction of scattered power that remains in the fundamental mode after the tilt:

$$\eta(\alpha) = \eta_0 \exp\left(-\frac{(k\, w(z_{\text{EM}})\, \alpha)^2}{4}\right), \tag{3.49}$$

where η_0 accounts for any residual on-axis mismatch due to a beam parameter difference between the returning beam and the cavity eigenmode (section 5.4.3). This is the only fraction of scattered light that is both resonantly amplified by the arm cavity and spatially matched to the local oscillator at the dark port, and therefore the only fraction that contributes to the calibration signal.

The mathematical tools developed in this chapter (Gaussian beam propagation, the Fabry–Perot cavity transfer function, phase modulation and Bessel-function sideband decomposition, and the mode-matching overlap integral) provide the theoretical framework needed to model the scattered light coupling in the Virgo end benches. Chapter 4 applies these tools to derive the scattered-light field model and the equivalent gravitational wave strain, while Chapter 5 extends the treatment to the time domain to handle the large modulation depths encountered in practice.

4

Scattered light in Advanced Virgo

When light is transmitted out of the arm cavity through the end mirror, it enters the end bench, where it encounters a variety of optical elements. A small fraction of this transmitted light is reflected back toward the cavity and re-enters the interferometer. This scattered light accumulates a round-trip phase that depends on the distance to the scattering surface, and when this phase varies in time it mimics the signature of a gravitational wave signal at the B1 photodiode.

This chapter explains the physical origin of scattered light, develops the theoretical model for scattered light coupling, derives the equivalent gravitational wave strain it produces, and presents an analysis of controlled injection measurements performed on the Virgo end benches during the commissioning period preceding O4. These measurements serve both as a validation of the model and as the observational context that motivates the calibration method developed in section 5.5.2.

4.1 End bench scattering

4.1.1 Transmitted power and end bench layout

A small fraction of the power circulating inside each arm cavity is transmitted through the end mirror (ETM) onto the end bench. For the west arm this is the Suspended West End Bench (SWEB), and for the north arm it is the Suspended North End Bench (SNEB). The transmitted power is

$$P_{\text{trans}} = T_{\text{WE}} \cdot P_{\text{circ}}, \tag{4.1}$$

where T_{WE} is the power transmission coefficient of the west end mirror and P_{circ} is the circulating arm power.

The end benches house a variety of optical elements, including photodiodes used for beam alignment and angular sensing (the B7 photodiode on SNEB and the B8 photodiode on SWEB, visible in figure 4.2), as well as mode-matching telescopes, auxiliary pick-off optics, lenses, viewports, and diagnostic photodiodes. The B7 and B8 photodiodes are of particular importance for this thesis, as they monitor the light leaving the arm cavity through the end mirror and are therefore sensitive to the scattered-light signals described in Chapter 4.

4.1.2 Scattered light from the end benches

The end benches are suspended on vibration isolation platforms, though with less aggressive seismic attenuation than the test mass Superattenuators, since the optical elements on the bench are not part of the precision measurement chain. Nevertheless, small mechanical motions of the end bench can produce scattered light that re-enters the arm cavity through the end mirror. A fraction of the light transmitted onto the bench is reflected back by optical elements such as quadrant photodiodes, viewports, lenses, and other optical surfaces. Part of this reflected light is mode-matched to the arm cavity TEM$_{00}$ mode and re-couples into the interferometer [10].

The scattered light accumulates a round-trip phase proportional to the distance between the ETM and the scattering surface. When this distance varies over time due to bench motion, the resulting phase modulation can mimic the signature of a gravitational-wave signal at the B1 detector. Scattered light is therefore one of the most persistent noise sources in the low-frequency band (up to 100 Hz) of the Virgo detector, particularly below a few tens of hertz [10].

During multiple observing runs, scattered light produced characteristic features in the detector output, appearing as arches in spectrograms (see section 4.3). Most of these features were observed and characterised during O3 and the commissioning period leading up to O4. The end benches were consequently identified as critical locations for scattering investigations, and the study of their behaviour under controlled conditions became an important part of the Virgo noise-hunting strategy.

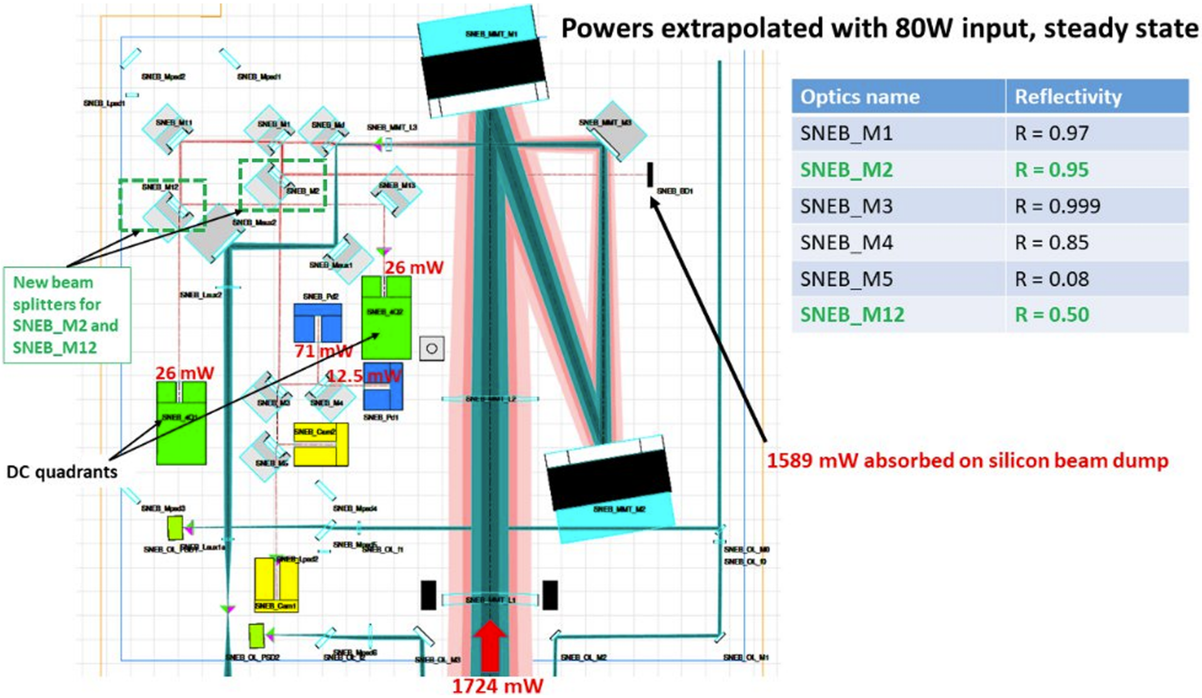

Figure 4.1: Optical layout of the SWEB and SNEB end benches as planned for O5. [10]

The DET (Detector group) technical noise investigation reports document several aspects of these studies, including the role of beam dumps, bidirectional reflectance distribution function (BRDF) measurements, and the identification of relevant scattering surfaces such as TGG (Terbium Gallium Garnet) crystals, meniscus lenses, output mode cleaner mirrors, and photodiodes. These investigations demonstrated that scattered light is a multi-faceted problem present throughout the detector infrastructure.

4.2 Field model: transmitted and scattered fields

The derivation presented in this section follows the reference work of Was et al. [5], extended to include a more accurate beam splitter model [6].

Throughout this chapter, optical fields are normalised so that the detected power scales as $|E|^2$. All constant proportionality factors are omitted, since only relative power changes are relevant for the analysis. All symbols are defined in table 4.1.

4.2.1 Transmitted and returning fields

The calibration process exploits the small fraction of light transmitted through the end mirror onto the bench. This light is reinjected in a controlled way, and the effect of the injection is observed at the main photodiodes. The derivation starts from the intracavity field E_0 at the WE mirror. This field is transmitted through the WE mirror, reflected off a scattering element, and returned to the mirror, accumulating a total round-trip phase $\phi_{\text{sc}}(t)$. The schematic overview that forms the basis of this derivation is shown in figure 4.2.

The returning field at the WE mirror is [5]

$$E_{\text{sc}}(t) = \sqrt{T_{\text{WE}}\, f_{r,m}}\; E_0(t)\, e^{\,i\phi_{\text{sc}}(t)}, \qquad \phi_{\text{sc}}(t) = \frac{4\pi}{\lambda}\, x(t), \qquad f_{r,m} = f_r\, f_m. \tag{4.2}$$

A key parameter in this expression is the reflected light power fraction $f_{r,m}$, which represents the fraction of light that reflects back toward the cavity, remains in the fundamental TEM$_{00}$ mode, and is mode-matched to the arm cavity eigenmode. It consists of two factors: f_r, the fraction of light that is reflected by the scattering element, and f_m, the fraction of the reflected light that is mode-matched to the cavity. The distinction between these two factors is important because the light reaching B8 does not necessarily satisfy both conditions, and special care must be taken when interpreting the measured signal (see section 5.5.2). When the scattering element oscillates, f_m becomes time-dependent through the instantaneous tilt angle $\theta(t)$ of the scattering surface, as analysed in detail in Section 5.4.2. In the case considered here, $f_{r,m}$ is treated as a constant.

The quantity $x(t)$ is the distance between the HR surface of the WE mirror and the scattering element, which determines the accumulated round-trip phase.

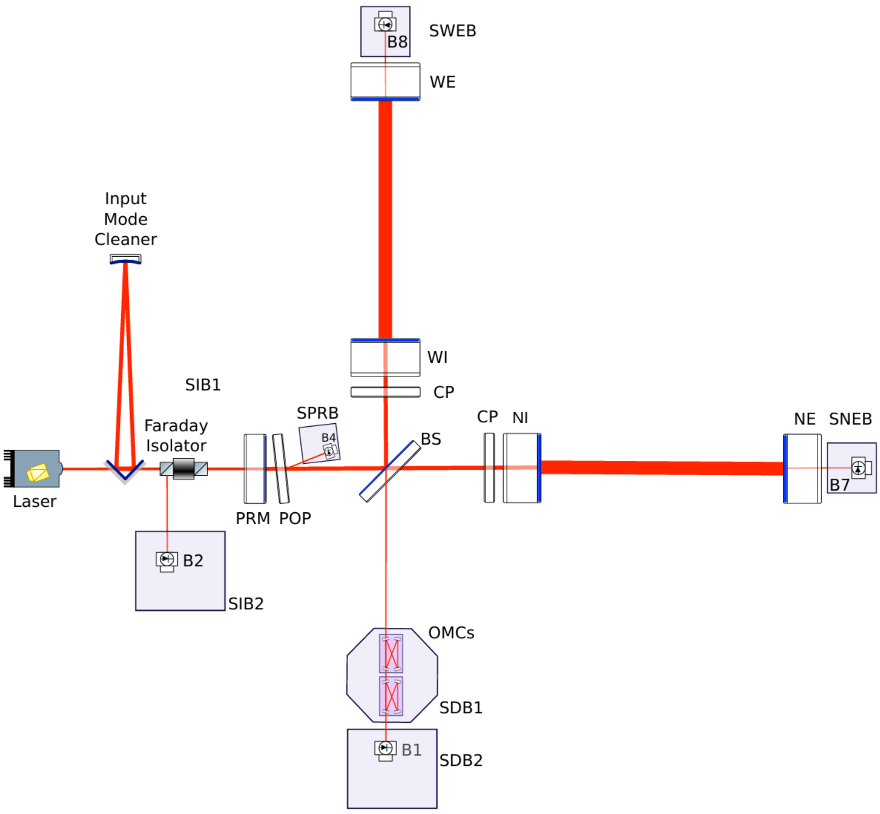

Figure 4.2: Overview of the most important optical elements of the Virgo detector [5].

The total field reaching the B8 photodiode on the bench is the sum of the field transmitted out of the cavity and the returning scattered field $E_{\text{sc}}(t)$. The exact form of this expression depends on the optical layout of the end bench. A more realistic treatment including the beam splitter geometry is presented in section 5.5.2. In the simplified model, the field and power at B8 are [5]

$$E_{\text{B8}}(t) = \sqrt{\alpha\, T_{\text{WE}}}\, E_0(t)\big(1 + \sqrt{f_{r,m}}\, e^{i\phi_{\text{sc}}(t)}\big), \tag{4.3}$$

$$P_{\text{B8}}(t) = |E_{\text{B8}}(t)|^2 = \alpha\, T_{\text{WE}}|E_0|^2\, \left|1 + \sqrt{f_{r,m}}\, e^{i\phi_{\text{sc}}(t)}\right|^2, \tag{4.4}$$

where the factor α accounts for optical losses between the WE mirror and the B8 photodiode. This factor does not affect the calibration method and is included only for completeness.

The quantity used in practice is the relative intensity noise (RIN), defined as the fractional power change at B8 due to the scattering relative to the baseline power:

$$\frac{\delta P_{\text{B8}}(t)}{P_{\text{B8}}} \;=\; 2\sqrt{f_{r,m}}\cos\phi_{\text{sc}}(t) \;+\; \mathcal{O}(f_{r,m}). \tag{4.5}$$

This expression provides a means of determining the amount of scattered light from the B8 data. Fitting to $\cos 2\phi_{\text{sc}}(t)$ rather than $\cos\phi_{\text{sc}}(t)$ can reveal second-order scattering, in which the light performs an additional round trip before re-entering the cavity.

If the light performs two round trips, the amplitude contribution is $\propto f_r^2$ and appears as $\cos 2\phi_{\text{sc}}$ in the B8 RIN. More generally, if the field at B8 is

$$E_{\text{B8}} \propto 1 + a_1 e^{i\phi_{\text{sc}}} + a_2 e^{i2\phi_{\text{sc}}} + \cdots, \tag{4.6}$$

then

$$\frac{\delta P_{\text{B8}}}{P_{\text{B8}}} \approx 2a_1\cos\phi_{\text{sc}} + 2a_2\cos 2\phi_{\text{sc}} + \cdots \tag{4.7}$$

This higher-order contribution is not used in the calibration described here, owing to $f_{r,m} \ll 1$. Here a_1 and a_2 are general coefficients representing the first- and second-order scattering amplitudes.

4.2.2 Output power at the antisymmetric port

The scattered light that re-enters the cavity interacts with the intracavity field and is amplified by the frequency-dependent cavity gain $G(f)$. In its simplest form, this gain is given by the single-pole model of equation (3.22):

$$G(f) = \frac{1}{1 - r_{rt}} \cdot \frac{1}{1 + i\dfrac{f}{f_{\text{arm}}}}, \qquad r_{\text{rt}} = r_1 r_2 \sqrt{1 - \Lambda_{\text{arm}}}, \tag{4.8}$$

where Λ_{arm} is the round-trip loss and $f_{\text{arm}} \approx (1 - r_{rt})\, c/(4\pi L)$ is the cavity pole. The prefactor $1/(1 - r_{rt})$ is the on-resonance field build-up factor (see equation (3.22)).

The cavity gain $G(f)$ applies to the scattered light only insofar as it is a perfect TEM_{00} mode that is mode-matched to the cavity. If this condition is not met, the effective gain is reduced and the analysis must be modified accordingly. The total intracavity field is the sum of the unperturbed field and the amplified scattered contribution from equation (4.2):

$$E_{\text{cavity}}(t) = E_0(t)\big(1 + G\, T_{\text{WE}} \sqrt{f_{r,m}}\, e^{\,i\phi_{\text{sc}}(t)}\big). \tag{4.9}$$

In this expression, G represents the cavity gain at the frequency of the relevant sideband component. Since the scattered field contains many sideband orders (section 3.3.5), a full treatment requires applying $G(n\Omega)$ to each order n individually. This is handled by the time-domain simulation developed in chapter 5.

Equation (4.9) uses the same transmission coefficient T_{WE} for both the outgoing and returning fields. This assumption is justified by optical reciprocity: the transmission amplitude through a dielectric multilayer coating is identical in both propagation directions, even though the reflectivities from the HR and AR sides differ. An elaboration on this is given in Appendix B.

The scattered light therefore modifies the intracavity field and consequently the power at the output photodiode B1. To interpret this effect, it is useful to compare it to the phase shift produced by a real gravitational wave, which induces a round-trip differential phase of [5]

$$\Delta\phi_{\text{GW}}(t) = \frac{4\pi G L}{\lambda}\, h(t). \tag{4.10}$$

The Virgo detector operates in DC readout, in which a small differential offset ψ from the dark fringe is introduced (section 2.4.2). Including both the gravitational wave phase and the scattered light perturbation, the fields returning to the beam splitter are [37]

$$E_{\text{cavity,west}}(t) = \sqrt{T_{\text{IM}}}\, E_0(t)\, e^{i\left(\frac{\psi}{2} + \frac{2\pi G L}{\lambda} h(t)\right)} \left[1 + G\, T_{\text{WE}} \sqrt{f_{r,m}}\, e^{i\phi_{\text{sc}}(t)}\right], \tag{4.11}$$

$$E_{\text{cavity,north}}(t) = \sqrt{T_{\text{IM}}}\, E_0(t)\, e^{-i\left(\frac{\psi}{2} + \frac{2\pi G L}{\lambda} h(t)\right)}. \tag{4.12}$$

The field and power reaching the B1 photodiode are

$$E_{\text{out}}(t) = \frac{1}{\sqrt{2}}\Big(E_{\text{cavity,west}}(t) - (1 - \epsilon)\, E_{\text{cavity,north}}(t)\Big), \tag{4.13}$$

$$P_{\text{B1}}(t) = \frac{1}{2}\left|E_{\text{cavity,west}}(t) - (1 - \epsilon)\, E_{\text{cavity,north}}(t)\right|^2, \tag{4.14}$$

where ϵ accounts for asymmetries between the two arms and the factor $1/\sqrt{2}$ arises from the beam splitter transmission. Expanding to leading order in ψ, h, and $f_{r,m}$ (see appendix C for the full derivation), the power at B1 becomes

$$\begin{aligned} P_{\text{B1}}(t) \approx \frac{T_{\text{IM}}}{2}\, |E_0|^2 \Big[&\psi^2 + \epsilon^2 + 2\psi\, \frac{4\pi G L}{\lambda}\, h(t) \\ &+ 2\psi\, G\, T_{\text{WE}} \sqrt{f_{r,m}} \sin\phi_{\text{sc}}(t) + 2\epsilon\, G\, T_{\text{WE}} \sqrt{f_{r,m}} \cos\phi_{\text{sc}}(t)\Big]. \end{aligned} \tag{4.15}$$

Three observations follow from this expression. The term $2\psi\,(4\pi G L/\lambda)\, h(t)$ is the gravitational wave signal, linear in h. The term proportional to $\sin\phi_{\text{sc}}(t)$ behaves as an additional phase signal which could be compared to a gravitational wave signal. The term proportional to $\cos\phi_{\text{sc}}(t)$ represents an amplitude-type leakage that couples through the contrast defect ϵ.

A more complete treatment retaining the complex cavity gain (appendix C) gives, in the absence of a gravitational wave signal and arm asymmetries,

$$P_{\text{B1}}(t) = \frac{T_{\text{IM}}}{2} |E_0|^2 \left[\underbrace{\psi^2}_{\text{DC power}} + \underbrace{2\psi\, |G(f)|\, T_{\text{WE}} \sqrt{f_{r,m}}\, \sin(\phi_{\text{sc}} - \varphi)}_{\text{scattered light — phase coupling}}\right], \tag{4.16}$$

with

$$|G(f)| = \frac{1}{1-r_{\text{rt}}} \cdot \frac{1}{\sqrt{1+(f/f_{\text{arm}})^2}}, \tag{4.17}$$

$$\varphi(f) = \arctan\left(\frac{f}{f_{\text{arm}}}\right). \tag{4.18}$$

Symbol	Meaning	Units
E_0	Carrier field amplitude inside the west Fabry–Perot cavity at the HR surface of the WE mirror	$\text{W}^{1/2}$ or arb.
$T_{\text{WE}}, T_{\text{WI}}$	Power transmission coefficients of WE, WI mirrors	—
$f_{r,m}$	Fraction of light scattered back into TEM_{00} mode and mode-matched to the cavity	—
$\phi_{\text{sc}}(t)$	Scattered-light round-trip phase $= 4\pi x(t)/\lambda$	rad
$x(t)$	Distance between HR surface of WE mirror and scattering surface	m
λ	Laser wavelength (1064 nm for Virgo)	m
L	Arm cavity length (3 km for Virgo)	m
$G(f)$	Frequency-dependent cavity gain	—
ψ	DC readout differential phase offset	rad
ϵ	Contrast defect parameter	—
T_{IM}	Input mirror power transmission	—
B8	Photodiode on the end bench monitoring scattered light	—
$h(t)$	Gravitational wave strain	—

Table 4.1: Definitions of symbols and quantities used in the calibration calculation.

4.2.3 Equivalent strain from scattering

For the calibration procedure, it is useful to express the scattered light contribution to B1 as an equivalent gravitational wave strain. This is the strain that, if real, would produce the same photodiode signal. This is obtained by dividing the scattered light terms in equation (4.15) by the gravitational wave scaling factor $2\psi\,(4\pi GL/\lambda)$ [5].

From the sin ϕ_{sc} (phase coupling) term:

$$h_{\text{sc,phase}}(t) \equiv \frac{1}{L}\, T_{\text{WE}} \sqrt{f_{r,m}}\, \frac{\lambda}{4\pi}\, \sin\phi_{\text{sc}}(t). \tag{4.19}$$

From the cos φ_{sc} (amplitude coupling) term:

$$h_{\text{sc,amplitude}}(t) \equiv \frac{\epsilon}{\psi} \cdot \frac{1}{L}\, T_{\text{WE}} \sqrt{f_{r,m}}\, \frac{\lambda}{4\pi}\, \cos\phi_{\text{sc}}(t). \tag{4.20}$$

These expressions are valid in the low-frequency limit $f \ll f_{\text{arm}}$, where the cavity gain is real and positive. The generalisation to arbitrary signal frequencies, in which the complex cavity gain introduces a phase lag $\phi_s = \arctan(f/f_{\text{arm}})$ (appendix C), equation (C.4), is treated in section 5.5.2 and validated numerically in section 6.6.

Two important observations follow. First, the phase coupling term (4.19) is independent of both ψ and G. Its amplitude is set entirely by T_{WE}, $f_{r,m}$, and the wavelength. This independence is the basis of the calibration procedure in chapter 5.5.2. Second, the amplitude coupling term (4.20) is suppressed relative to the phase coupling by the factor ϵ/ψ. Since Advanced Virgo operates with $\epsilon \ll \psi$, the amplitude coupling is a small correction and is neglected in the leading-order calibration analysis.

Other coupling paths are present but are irrelevant for the calibration considered in this thesis. Numerical simulations and a detailed discussion of additional coupling mechanisms are provided in [5].

For the calibration method in section 5.5.2, the more general equation (4.16) is used. All quantities and their definitions are summarised in table 4.1.

The expressions derived above assume a direct optical path between the end mirror and the scattering element. In the chapter 5, these are generalised to account for a beam splitter present on the end bench (figure 5.3), which modifies the numerical prefactors in the B1 and B8 signal amplitudes while leaving the structure of the calibration procedure unchanged.

In summary, the equivalent strain formulation shows that scattered light produces a signal at B1 whose amplitude depends on the product $T_{\text{WE}}\sqrt{f_{r,m}}$. By independently measuring the scattered-light signal at B8 (which depends on $\sqrt{f_{r,m}}$ alone), the end-mirror transmission T_{WE} can in principle be extracted from the ratio of the two measurements. This observation is the basis for the calibration procedure developed in chapter 5.

4.2.4 From noise source to calibration tool

The equivalent strain expressions derived in the previous section show that backscattered light produces a signal at B1 whose amplitude is set by the product $T_{\text{WE}}\sqrt{f_{r,m}}$, while the corresponding signal at B8 depends on $\sqrt{f_{r,m}}$ alone. This asymmetry is the key observation that makes scattered light useful for calibration: by measuring the signal amplitudes at both photodiodes independently and taking their ratio, the end-mirror transmission T_{WE} can be extracted without requiring knowledge of the absolute laser power, the intracavity field amplitude, or the optical losses between the end mirror and the B8 photodiode.

If, in addition, an independent laboratory measurement of T_{WE} is available (for example, from the coating manufacturer), the ratio of the reconstructed strain $\hat{h}_{\text{rec}}$ to a reference strain $\hat{h}_{\text{signal}}$ constructed from this laboratory value directly tests the absolute accuracy of the detector's strain calibration, as will be shown in section 5.5. This provides a calibration cross-check that is fully independent of both the photon calibrator and the Newtonian calibrator described in section 2.6.

Exploiting this opportunity requires a scattering source that is controlled, reproducible, and well characterised. Seismic motion of the end bench, while responsible for the scattered light noise observed during normal operations, is neither predictable nor reproducible at the level needed for a precision measurement. A dedicated actuator is therefore needed to generate the scattered field on demand. The piezoelectric bender, described in detail in section 7.1, fulfils this role: driven at a known frequency and amplitude, it produces a controlled oscillation of the scattering surface that generates a well-defined scattered-light signal at both B1 and B8. The remainder of this thesis develops the simulation framework needed to validate this calibration concept (chapter 5), presents the simulation results (chapter 6), and describes the experimental hardware that has been prepared for deployment on the Advanced Virgo end bench (chapter 7).

4.3 Scattered light behaviour in Virgo

During the commissioning period preceding the fourth observing run (O4), controlled scattered-light injections were performed on both the SWEB and SNEB end benches by deliberately oscillating the bench suspension at low frequencies. These measurements serve two purposes: they validate the theoretical model of section 4.2 against real detector data, and they characterise the scattering properties of each bench to inform the choice of installation site for the piezoelectric bender. The analysis presented below was performed by Stef Bordo [62]. The results are included here to provide the observational context that motivates the calibration method developed in the following chapters.

4.3.1 Spectral signature and arch morphology

A scattered light signal can be observed by oscillating the entire end bench at a low frequency (below 1 Hz). The resulting signals at various photodiodes around the interferometer provide insight into the amount of scattered light produced and can reveal differences or asymmetries between the two arms. One such asymmetry was observed when comparing scattering with open and closed quadrant shutters: SWEB exhibited a much more predictable response, while closing the quadrant shutters on SNEB produced an increase in scattering power [5].

Investigations of scattering from TGG crystals and OMC optics, as well as from lenses, were also performed. These studies revealed that misaligned optics can scatter light over a wide angular range, with only a small fraction needing to recouple into the arm cavity to produce a measurable signal [21].

The scattered light signals appear as arches in the output spectrogram at low frequencies, arising when bench motion upconverts low-frequency seismic fluctuations. The shape and amplitude of these arches are directly related to the instantaneous bench velocity: larger bench motion produces higher instantaneous frequencies, resulting in a broader spectrogram pattern [5].

A consistent observation is the presence of higher-order scattering processes, particularly at higher injection amplitudes. These correspond to scenarios in which light undergoes more than one scattering event before re-entering the interferometer, as described mathematically in section 4.2.

4.3.2 Data analysis of injected signals

The injections analysed in this section targeted the SNEB and SWEB with a dimensionless amplitude value A equal to 140. This value represents a normalised actuation amplitude and does not correspond directly to a physical displacement. The conversion to physical displacement amplitude A_{physical} (in micrometres) is linear:

$$\frac{A_{\text{physical}}}{A} \approx 0.14\,\mu m. \tag{4.21}$$

Each unit increase in A therefore corresponds to an increase of approximately $0.14\,\mu$m in physical displacement.

These longitudinal excitations were chosen to drive the benches into regimes where the differential phase accumulated by the backscattered light undergoes large excursions, producing strong and clearly visible scattering signals. Each injection consists of a sinusoidal bench motion with an amplitude sufficient to wrap through multiple optical fringes, which occurs when the displacement exceeds the carrier wavelength. The data recorded during each injection include the bench motion via LC (Local Control) channels (LC_X, LC_Y, LC_Z), the responses of the B7 and B8 photodiodes on SNEB and SWEB, respectively, and the reconstructed strain from the B1 photodiode.

Plateau structure and frequency cutoffs

A clear feature visible in the data is a plateau-like behaviour in the low-frequency band of the B8 signal. At higher injection amplitudes, a secondary plateau becomes prominent, corresponding to the primary and secondary scattering processes described in section 4.2. These plateaus can be seen in figure 4.3.

The cutoff frequency of the primary plateau is directly related to the maximum instantaneous velocity of the bench during the injection. For a sinusoidal injection at frequency f_{mod} and physical amplitude A_{physical}, the maximum instantaneous velocity is

$$v_{\text{max}} = 2\pi\, A_{\text{physical}}\, f_{\text{mod}}, \tag{4.22}$$

and the maximum frequency of the primary plateau is

$$f_{\text{max,1}} = \frac{2\, v_{\text{max}}}{\lambda}. \tag{4.23}$$

The maximum frequency of the secondary scattering process is always $f_{\text{max,2}} = 2\, f_{\text{max,1}}$.

These cut-off frequencies can also be identified in the output spectrograms as the turning points of the characteristic arch pattern (figure 4.4). The maximum frequencies observed in the spectrograms agree well with equations (4.22) and (4.23).

SWEB versus SNEB

The bench motion of both benches is driven by a clean sinusoidal input signal (top subplots of figure 4.3). The SWEB results are cleaner and more reproducible than the SNEB results. The underlying scattering processes are identical in both benches, as can be verified from the measured instantaneous displacement. The less stable SNEB response was assumed to be an issue with the suspension system of that bench. The difference in plateau clarity between the two benches is illustrated in figure 4.4.

The spectrograms from the SWEB and SNEB injections are shown in figure 4.4 for an injection amplitude of 140 and injection frequency of 0.2 Hz. The less stable injections at SNEB are apparent from the broader and less regular arch morphology.

Directional dependence

Injections produce signals in all three translational directions (x, y, z), and the results are compared in figure 4.5. The longitudinal direction (z) produces the strongest signal, as expected from the model: longitudinal motion directly modulates the round-trip phase φ_{sc}. Transverse direction signals appear weaker through a combination of beam clipping and the angular sensitivity of $f_{r,m}$.

Rotational injections were also performed by rotating the bench about the x, y, and z axes (figure 4.6). These rotational injections are of particular relevance to the present work, as they resemble the injection configuration that will be produced by the piezoelectric bender. The scattered light coupling depends not only on the magnitude of longitudinal motion but also on angular changes that alter the scattering angle and modify the overlap of the scattered field with the cavity eigenmode. This angular sensitivity is characterised in detail in section 3.4.

4.3.3 Summary

The injection measurements confirm the theoretical model of Section 4.2: longitudinal bench motion produces phase-modulated backscattered light whose spectral signature (plateau structure, arch morphology, and cutoff frequencies) matches the predictions of equations (4.22) and (4.23). Rotational injections demonstrate that angular coupling is also present, motivating the tilt analysis developed in Chapter 5. These observations establish the experimental basis for the controlled scattering calibration method.

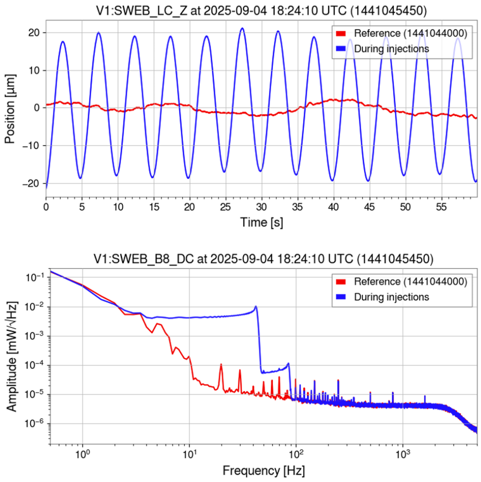

(a) SWEB injection B8 results at 0.2 Hz for actuation amplitude $A = 140$

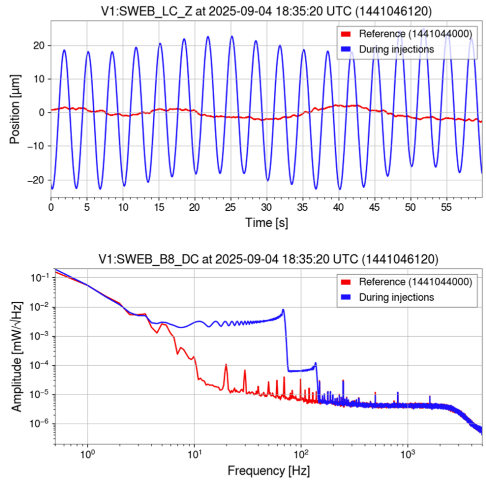

(b) SWEB injection B8 results at 0.3 Hz for actuation amplitude $A = 140$

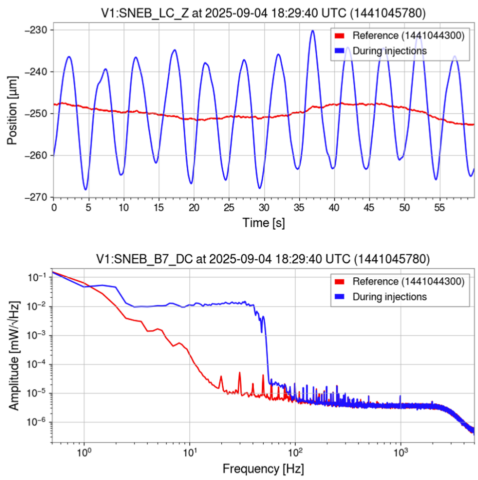

(c) SNEB injection B8 results at 0.2 Hz for actuation amplitude $A = 140$

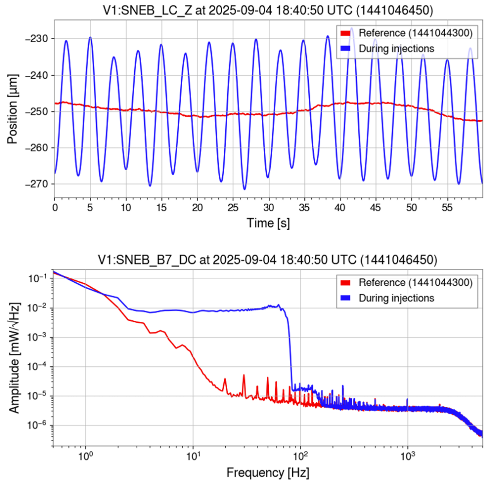

(d) SNEB injection B8 results at 0.3 Hz for actuation amplitude $A = 140$

Figure 4.3: Comparison of B8 measurements for SWEB and SNEB injection measurements with actuation amplitude $A = 140$.

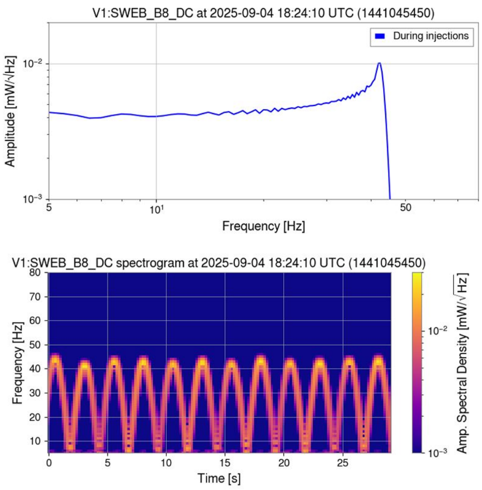

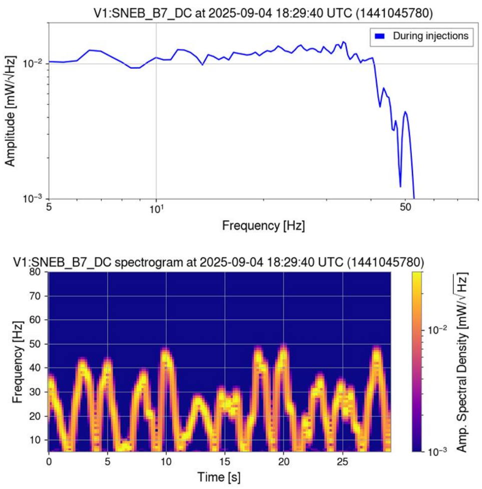

Figure 4.4: Comparison of the primary plateau shape and corresponding spectrogram for SWEB (left) and SNEB (right). Injections performed at actuation amplitude $A = 140$ and frequency 0.2 Hz [62].

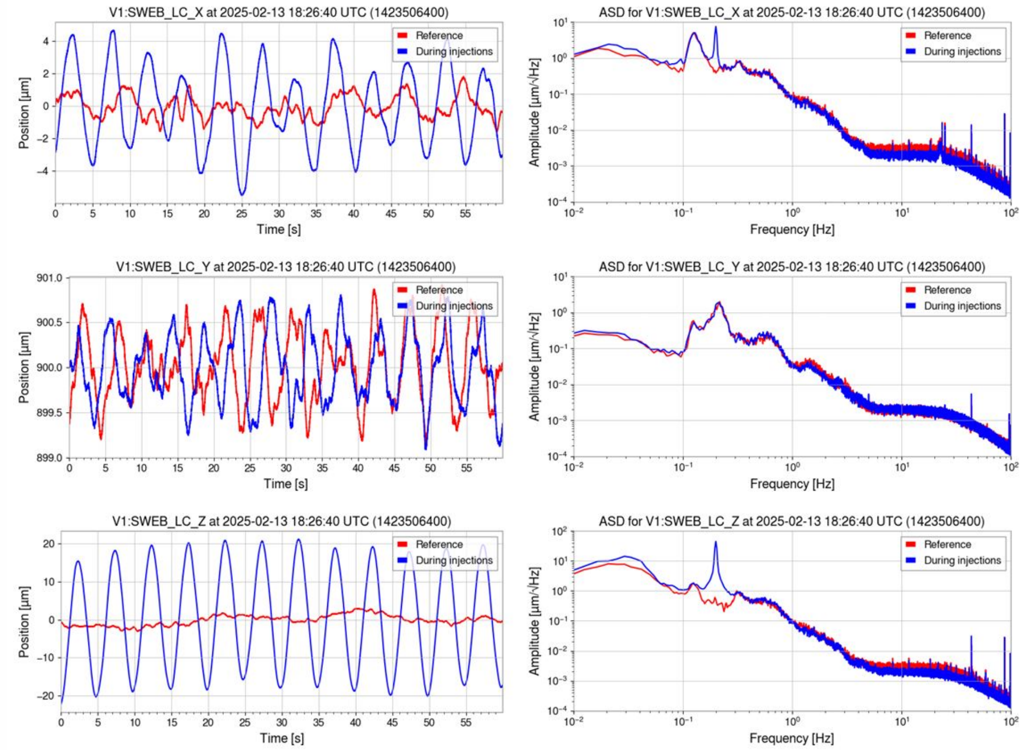

Figure 4.5: Comparison of scattered light injections produced by oscillating the SWEB in the x, y, and z directions.

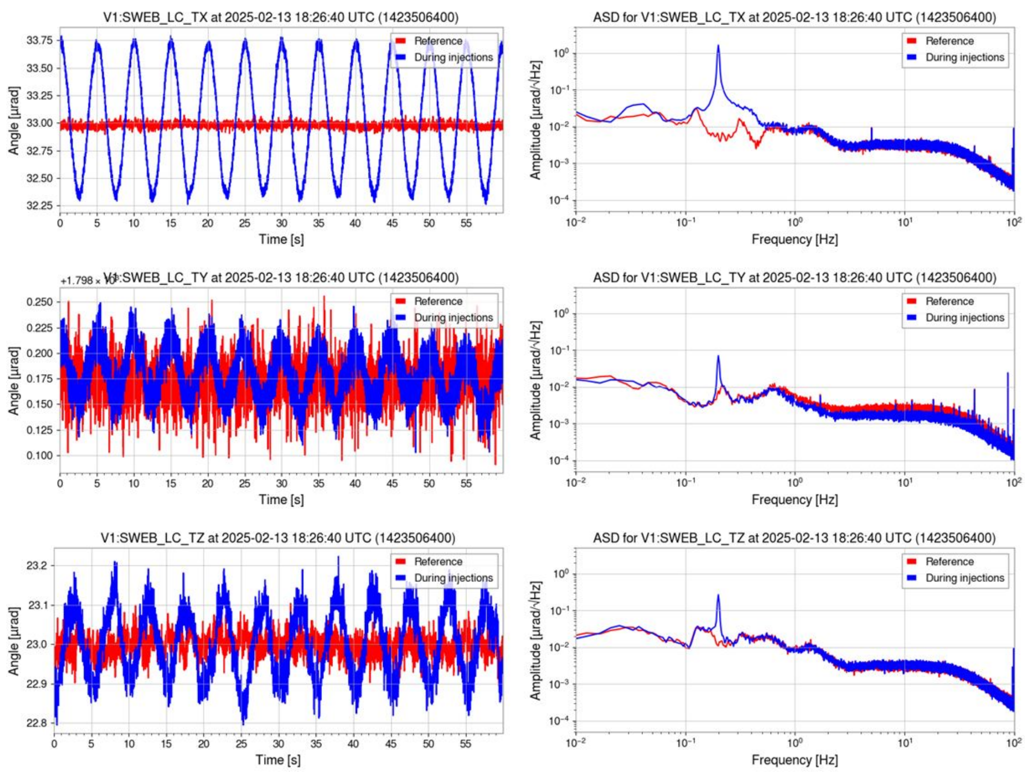

Figure 4.6: Comparison of scattered light injections produced by rotating the SWEB about the x, y, and z axes [62].

5

Modelling of scattered light from a piezoelectric bender

The PCal and NCal methods described in section 2.6 are well-established and provide calibration accuracies at the percent or sub-percent level. However, both rely on injecting a known mechanical displacement of the test mass and comparing the detector's response. The scattered-light calibration method [5] described in this thesis operates on a different principle: it measures an optical parameter of the detector (the end-mirror transmission T_{WE}) using only power ratios at two photodiodes, without injecting any mechanical force on the test mass. As will be discussed in section 4.2.3, the ratio of the reconstructed strain to a reference strain constructed from an independent laboratory measurement of T_{WE} provides a direct test of the absolute detector calibration that is fully independent of both the PCal and the NCal. The three methods also operate in complementary frequency ranges: the NCal injects signals at a few hertz (set by the rotor speed), the PCal covers tens to hundreds of hertz, and the scattered-light method operates at frequencies of 50 Hz to 2 kHz set by the bender drive [8]. Together, with different dominant systematics and different frequency coverage, the three methods form a set of complementary cross-checks that strengthen confidence in the overall calibration of the gravitational wave detector network.

In this chapter, the simulation and calibration methodology is explained for the scattered light calibration generated by an oscillating and bending surface. This oscillating surface corresponds to the bending surface of a piezoelectric bender which will be used in Advanced Virgo [10].

5.1 Simulation framework

5.1.1 OSCAR: field representation and cavity computation

The optical simulations presented in this chapter were performed using OSCAR (Optical Simulation Containing Ansys Results), a Matlab-based FFT code developed by Jérôme Degallaix [63]. The name reflects the software's origins: it was first created to simulate optical cavities subject to thermal distortions, with non-spherical surface deformations computed using the finite-element package Ansys [63]. The first version was written in Igor in 2005, after which it was translated to Matlab and has since been continuously extended to cover an increasingly broad range of cavity configurations [64]. Since its first public release in 2008, OSCAR has been downloaded more than 4000 times and has been used extensively in the design and commissioning of the Advanced Virgo gravitational wave detector [63].

Field representation and free-space propagation

In OSCAR, the electromagnetic field is represented as a two-dimensional array of complex numbers sampled on a uniform Cartesian grid, corresponding to the transverse distribution of the electric field amplitude. A fundamental Gaussian beam of beam radius w, wavefront radius of curvature R, and amplitude A is discretised as

$$E(x,y) = A\, e^{i\varphi}\, \text{exp}\left(-\frac{x^2+y^2}{w^2}\right)\, \text{exp}\left(-i\frac{2\pi}{\lambda}\frac{x^2+y^2}{2R}\right), \tag{5.1}$$

where φ accounts for an overall phase shift, including the Gouy phase [63]. Arbitrary higher-order modes and non-Gaussian beam profiles can be represented on the same grid without any modification to the propagation kernel.

Free-space propagation is implemented using the angular spectrum method under the paraxial approximation [64, 63]. The propagation of a field over a distance d proceeds in three steps:

1. The transverse field distribution is decomposed into a sum of plane waves using a two-dimensional FFT.

2. Each plane-wave component at transverse spatial frequency (f_x, f_y) acquires the propagation phase $e^{ik_z d}$, where $k_z = \sqrt{k^2 - (2\pi f_x)^2 - (2\pi f_y)^2}$ and $k = 2\pi/\lambda$.
3. The propagated field is recovered by a two-dimensional inverse FFT.

All operations other than propagation are performed directly on the spatial-domain grid and require no Fourier transform.

Mirror interaction and wavefront distortions

The interaction of the field with an optical element is described as a spatially varying phase shift. An input field $E_i(x, y)$ passing through an element with optical path difference $\Delta\text{OPL}(x, y)$ and finite aperture $\mathcal{A}(x, y)$ becomes [64]

$$E_o(x, y) = E_i(x, y)\ \exp\big(-ik\,\Delta\text{OPL}(x, y)\big) \times \mathcal{A}(x, y), \tag{5.2}$$

where $\mathcal{A}(x, y)$ is a binary mask equal to one within the mirror aperture and zero outside. For a spherical mirror with radius of curvature R_C, the optical path difference is twice the surface height (sagitta) as seen in figure 5.1:

$$\Delta\text{OPL}(x, y) = 2\left(R_C - \sqrt{R_C^2 - x^2 - y^2}\right). \tag{5.3}$$

Arbitrary surface maps measured by wavefront sensors can be loaded directly and added to the nominal spherical surface, making it straightforward to assess the effect of realistic mirror polishing defects on cavity performance [63].

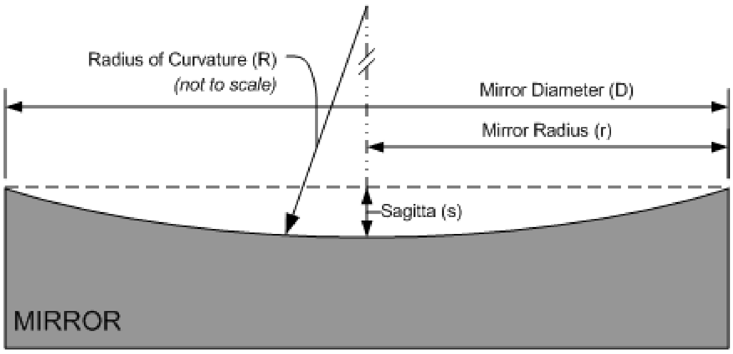

Figure 5.1: Diagram of a mirror with defining parameters [65].

Cavity field calculation

The steady-state circulating field in a Fabry–Perot cavity is computed by iterating the round-trip map and summing the successive contributions at a fixed reference plane [64], as illustrated schematically in figure 5.2. Starting from the field E_1 obtained by propagating the input beam E_{in} through the input mirror (IM), each round trip consists of four operations: propagation across the cavity length by FFT, reflection off the end mirror (EM) surface via equation (5.2), propagation back to the IM, and reflection off the IM surface. The total circulating field is the coherent sum of all contributions:

$$E_{\text{circ}} = \sum_{n=1}^{\infty} E_n. \tag{5.4}$$

The number of terms required for convergence is set by the cavity finesse. OSCAR is designed to simulate the steady-state classical optical field inside a single Fabry–Perot cavity and does not account for radiation pressure or quantum effects [64].

5.1.2 Simulated Virgo model

To simulate the scattered light calibration procedure, a simplified model of the Advanced Virgo west arm cavity was implemented in OSCAR. This section describes which components of the full detector were included, which were omitted, and how the piezoelectric bender on the end bench was modelled.

Optical layout

The simulation represents a Michelson interferometer with two Fabry–Perot arm cavities using the Advanced Virgo design parameters. Each arm cavity is modelled as a pair of mirror objects, each consisting of an HR and an AR coating interface separated by a 200 mm

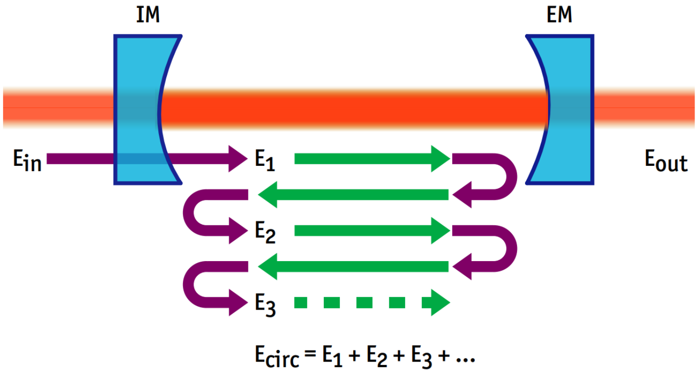

Figure 5.2: Schematic view of the algorithm used in OSCAR to compute the circulating field in a Fabry–Perot cavity [64].

fused silica substrate of refractive index $n = 1.45$. The input mirror (ITM) has a radius of curvature of 1420 m, a power transmissivity $T_{\text{IM}} = 1.3$ %, and a loss of 2.8 ppm. The end mirror (ETM) has a radius of curvature of 1683 m, a transmissivity $T_{\text{EM}} = 3.39$ ppm, and a loss of 3.2 ppm. The cavity length is $L = 3000$ m, and the beam splitter is modelled as a 50/50 power splitter [18, 22, 23].

The input field is a Gaussian beam with beam radius $w = 4.91$ cm and wavefront radius of curvature $R = -1431$ m at the ITM, corresponding to the Advanced Virgo input mode-matching condition. The circulating power in both cavities is set to $P_{\text{circ}} = 151$ W, which corresponds to the intracavity power of the simplified model without power recycling and a laser source with a power of 1 W. Since the calibration relies entirely on power ratios (equation (5.71)), the absolute value of P_{in} cancels and does not affect the recovered T_{WE}.

A small DC readout offset of $\psi = 10^{-4}$ rad is applied differentially to the arm cavities to displace the Michelson operating point from the dark fringe. With the symmetric beam splitter convention [24]

$$S = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & i \\ i & 1 \end{pmatrix}, \tag{5.5}$$

each reflection acquires a factor of i. Both the north and west return paths undergo one reflection and one transmission at the beam splitter, accumulating the same factor irt. The dark-port field E_{B1} is therefore

$$E_{\text{B1}} = i\,r\,t\left[E_{\text{north}} + \left(E_{\text{west}} + E_{\text{west,scatter}}\right)\right], \tag{5.6}$$

with $t = r = 1/\sqrt{2}$, where $E_{\text{west,scatter}}$ is the scattered-light field transmitted back through the west arm cavity. The dark fringe is produced by a macroscopic $\lambda/2$ path length difference between the two arms, which introduces a relative π phase shift between E_{north} and E_{west} so that $E_{\text{north}} = -E_{\text{west}}$ at the operating point. The differential offset ψ then produces a residual DC readout power

$$P_{\text{DC}} = |E_{\text{B1}}|^2 \propto \sin^2\!\left(\tfrac{\psi}{2}\right) \approx \frac{\psi^2}{4}\,. \tag{5.7}$$

The end-bench signal E_{B8} is formed as a linear combination of the ETM-transmitted field and the scattered field returning from the bender. A simplified diagram of the optical setup on the end bench is shown in figure 5.3.

In this layout, half of the incoming field amplitude is directed straight to the photodiode. The remaining field is reflected by the bender and transmitted through the beamsplitter a second time, introducing a factor of $1/\sqrt{2}$ in field amplitude for each beamsplitter interaction. After reflection from the ETM, which is approximated as perfectly reflective, the field returns to the beamsplitter, where another factor of $1/\sqrt{2}$ directs it toward the photodiode. Consequently, the field arriving at the photodiode through this path has an overall amplitude factor of

$$\left(\frac{1}{\sqrt{2}}\right)^3 = \frac{1}{2\sqrt{2}}$$

relative to the incident field.

Components not included

Several subsystems of the full Advanced Virgo detector are omitted to keep the model tractable: there is no power recycling cavity, no signal recycling cavity, no input mode cleaner, and no output mode cleaner. Thermal effects, realistic two-dimensional mirror surface maps, radiation pressure, quantum noise, and technical noise are all excluded. This is consistent with the approach taken in [5].

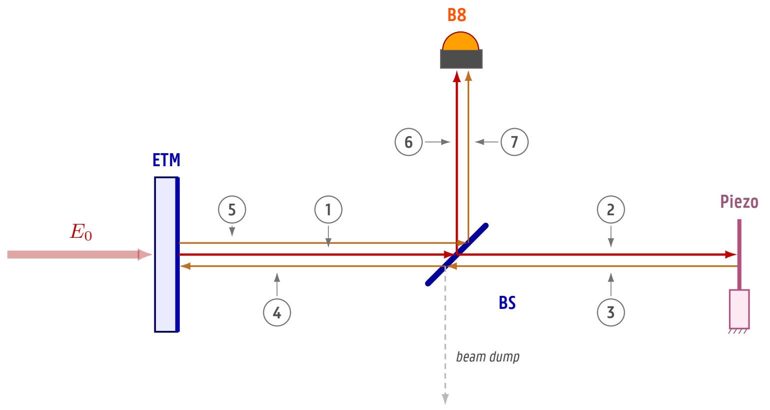

$$
\begin{aligned}
① &= \sqrt{T_{\text{WE}}}\,E_0\\
② &= \tfrac{1}{\sqrt{2}}\sqrt{T_{\text{WE}}}\,E_0\\
③ &= \tfrac{1}{\sqrt{2}}\sqrt{f_{r,m}}\sqrt{T_{\text{WE}}}\,E_0\\
④ &= \tfrac{1}{2}\sqrt{f_{r,m}}\sqrt{T_{\text{WE}}}\,E_0\\
⑤ &= \tfrac{1}{2}\sqrt{f_{r,m}}\sqrt{T_{\text{WE}}}\,E_0\\
⑥ &= \tfrac{1}{\sqrt{2}}\sqrt{T_{\text{WE}}}\,E_0\\
⑦ &= \tfrac{1}{2\sqrt{2}}\sqrt{f_{r,m}}\sqrt{T_{\text{WE}}}\,E_0
\end{aligned}
$$

$$
\begin{aligned}
E_{\text{B8}} &= ⑥ + ⑦\\
&= \tfrac{1}{\sqrt{2}}\sqrt{T_{\text{WE}}}\,E_0\,e^{i\phi_1}\\
&\quad + \tfrac{1}{2\sqrt{2}}\sqrt{f_{r,m}}\sqrt{T_{\text{WE}}}\,E_0\,e^{i\phi_2}
\end{aligned}
$$

Figure 5.3: Simplified diagram of the end bench optical layout with the corresponding electric field amplitudes at each beam path segment.

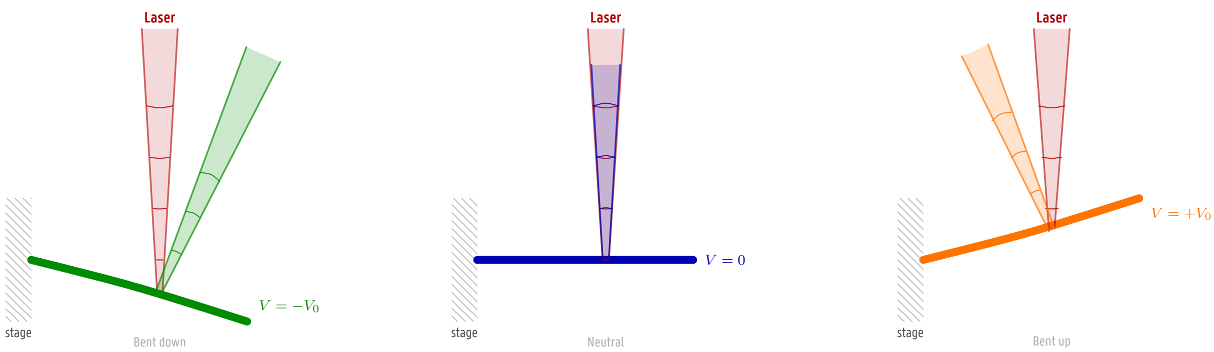

Figure 5.4: Deflection of a piezoelectric bender under applied voltage and the resulting change in reflection angle of an incident laser beam.

Piezo bender model

The piezoelectric bender is modelled as a rigid cantilever of length $\ell = 28\,\text{mm}$ driven sinusoidally at frequency f_{mod} with a tip displacement amplitude of $x_{\text{max}} = 10\,\mu\text{m}$. The corresponding maximum bending angle is

$$\theta_{\text{max}} = \arctan\left(\frac{x_{\text{max}}}{\ell}\right) \approx 357\,\mu\text{rad}. \tag{5.8}$$

At time t, the instantaneous bender angle and tip displacement are

$$\theta(t) = \theta_{\text{max}} \sin(\omega_m t), \tag{5.9}$$

$$x(t) = \ell \sin\big(\theta(t)\big) \approx x_{\text{max}} \sin(\omega_m t), \tag{5.10}$$

where $\omega_m = 2\pi f_{\text{mod}}$. The scattering element is placed at a distance $d \approx 1$ m from the ETM. The law of reflection dictates that the reflected beam is deflected by twice the bender tilt angle:

$$\alpha(t) = 2\theta(t). \tag{5.11}$$

This tilt is applied in OSCAR using the `Add_Tilt` method, which multiplies the transverse field by the linear phase factor $e^{ik\alpha x_\perp}$. The scattered field amplitude at the bender is scaled by the reflection coefficient $\sqrt{f_r} = 10^{-3}$. The deflection positions of the bender are illustrated in figure 5.4.

5.2 Time-domain cavity dynamics

5.2.1 Motivation for a time-domain approach

The frequency-domain cavity transfer function derived in section 3.2 provides an exact description of how a Fabry–Perot cavity filters an injected field: each frequency component at $\omega_0 + n\Omega$ of the injected phase-modulated field acquires an independent complex gain $G(\omega_0 + n\Omega)$, and the intracavity field is the superposition of all filtered components [6, 55]. This description is practical as long as the injected field contains only a small number of significant frequency components.

For the piezoelectric bender considered in this thesis, the conditions that make the frequency-domain approach convenient break down simultaneously:

- **Large modulation depth.** The bender tip displacement $x_B \sim 10\,\mu$m at $\lambda = 1064$ nm produces a modulation depth

$$\beta_B = \frac{4\pi x_B}{\lambda} \approx 118. \tag{5.12}$$

  The Bessel coefficients $J_n(\beta)$ remain significant for $|n| \sim \beta$, meaning several hundred sidebands carry non-negligible power. Tracking the cavity gain at each of these frequencies individually is both computationally wasteful and numerically fragile.

- **Spatial distortion of the reflected field.** When the bender surface is tilted by an instantaneous angle $\theta(t)$, the reflected field acquires a transverse phase ramp $e^{ikx\theta(t)}$ (equation (3.47)). The reflected beam is therefore no longer a pure TEM_{00} mode but a superposition of higher-order Hermite–Gaussian modes, each of which resonates at a different frequency inside the cavity. A scalar frequency-domain treatment cannot account for this spatial coupling without a full modal decomposition at every sideband frequency, which is impractical.

Both difficulties are resolved by reformulating the cavity dynamics in the time domain, where the full two-dimensional transverse field is propagated from one time step to the next through a single recursive update that handles spatial and spectral complexity simultaneously. The following sections derive this formulation from first principles and establish its exact equivalence with the frequency-domain result [66, 67, 68, 69].

5.2.2 Discrete-time cavity update equation

Consider a linear Fabry–Perot cavity formed by an input mirror with amplitude transmissivity $t_{\text{in}} = \sqrt{T_{\text{in}}}$ and reflectivity $r_1 = \sqrt{1 - T_{\text{in}}}$, and an end mirror with reflectivity r_2. The mirrors are separated by a length L, so that one round trip takes

$$t_{\text{RT}} = \frac{2L}{c}. \tag{5.13}$$

Any additional round-trip power loss (absorption, clipping, or scattering) is modelled by a single amplitude factor $\sqrt{1 - \Lambda_{\text{arm}}}$, giving the net round-trip amplitude factor

$$r_{\text{rt}} = r_1 r_2 \sqrt{1 - \Lambda_{\text{arm}}}, \qquad |r_{\text{rt}}| < 1. \tag{5.14}$$

Sampling the intracavity field once per round trip and defining $E_n \equiv E_{\text{cav}}(n\, t_{\text{RT}})$, two contributions add coherently at step $n+1$: the stored field E_n, attenuated by r_{rt} after one additional round trip, and a new contribution coupled in through the input mirror with amplitude t_{in}. This gives the discrete-time cavity update equation:

$$E_{n+1} = r_{\text{rt}}\, E_n + t_{\text{in}}\, E_{\text{in}}(n\, t_{\text{RT}}). \tag{5.15}$$

The condition $|r_{\text{rt}}| < 1$ guarantees that the stored field decays over successive round trips and the recursion is stable [55].

General solution

Starting from an initially empty cavity ($E_0 = 0$), the general intracavity field after n round trips can be obtained by iterating equation (5.15) through repeated back-substitution:

$$\begin{aligned}
E_n &= r_{\text{rt}}\, E_{n-1} + t_{\text{in}}\, E_{\text{in}}\big((n-1)\, t_{\text{RT}}\big) \\
&= r_{\text{rt}} \left[r_{\text{rt}}\, E_{n-2} + t_{\text{in}}\, E_{\text{in}}\big((n-2)\, t_{\text{RT}}\big)\right] + t_{\text{in}}\, E_{\text{in}}\big((n-1)\, t_{\text{RT}}\big) \\
&= r_{\text{rt}}^2\, E_{n-2} + t_{\text{in}}\, r_{\text{rt}}\, E_{\text{in}}\big((n-2)\, t_{\text{RT}}\big) + t_{\text{in}}\, E_{\text{in}}\big((n-1)\, t_{\text{RT}}\big) \\
&= r_{\text{rt}}^3\, E_{n-3} + t_{\text{in}}\, r_{\text{rt}}^2\, E_{\text{in}}\big((n-3)\, t_{\text{RT}}\big) + t_{\text{in}}\, r_{\text{rt}}\, E_{\text{in}}\big((n-2)\, t_{\text{RT}}\big) + t_{\text{in}}\, E_{\text{in}}\big((n-1)\, t_{\text{RT}}\big) \\
&\;\;\vdots \\
&= r_{\text{rt}}^n \underbrace{E_0}_{=0} + t_{\text{in}} \sum_{k=0}^{n-1} r_{\text{rt}}^k\, E_{\text{in}}\big((n-1-k)\, t_{\text{RT}}\big).
\end{aligned} \tag{5.16}$$

Since the cavity is initially empty, the first term vanishes and the intracavity field is given by

$$E_n = t_{\text{in}} \sum_{k=0}^{n-1} r_{\text{rt}}^k\, E_{\text{in}}\big((n-1-k)\, t_{\text{RT}}\big). \tag{5.17}$$

Equation (5.17) shows that the intracavity field is a weighted sum of past input fields with exponentially decaying weights r_{rt}^k. The cavity therefore has a finite optical memory: inputs from the distant past contribute negligibly, while recent inputs contribute with weight close to unity. The characteristic storage time is

$$\tau_{\text{cav}} = -\frac{t_{\text{RT}}}{\ln|r_{\text{rt}}|} \approx \frac{t_{\text{RT}}}{1-|r_{\text{rt}}|} = \frac{1}{2\pi f_{\text{arm}}}, \tag{5.18}$$

where the approximation holds in the high-finesse limit and the last equality uses the cavity pole definition of section 3.2.2.

5.2.3 Equivalence with the frequency-domain response

The discrete-time update equation must reproduce the frequency-domain cavity transfer function of equation (3.22) exactly. This equivalence is established below for two cases: a monochromatic input and a fully phase-modulated input.

Monochromatic input

Let the input be a monochromatic field at angular frequency ω:

$$E_{\text{in}}(t) = E_0\, e^{i\omega t}. \tag{5.19}$$

Substituting into equation (5.16), factoring out the k-independent terms, and defining $q \equiv r_{\text{rt}}\, e^{-i\omega t_{\text{RT}}}$ with $|q| < 1$, the geometric series converges absolutely as $n \to \infty$:

$$\sum_{k=0}^{n-1} q^k \xrightarrow{n\to\infty} \frac{1}{1-q} = \frac{1}{1 - r_{\text{rt}}\, e^{-i\omega t_{\text{RT}}}}. \tag{5.20}$$

Passing to the steady state and converting to continuous time gives a steady-state cavity field:

$$E_{\text{cav}}^{\text{ss}}(t) = \frac{t_{\text{in}}}{1 - r_{\text{rt}}(\omega)}\, E_{\text{in}}(t) \equiv t_{\text{in}}\, G(\omega)\, E_{\text{in}}(t), \tag{5.21}$$

where $r_{\text{rt}}(\omega) \equiv r_{\text{rt}}\, e^{-i\omega t_{\text{RT}}}$ and $G(\omega)$ is exactly the frequency-domain transfer function of equation (3.22). The time-domain update equation thus reproduces the frequency-domain result exactly for a monochromatic input.

Phase-modulated input

Now consider the phase-modulated input

$$E_{\text{in}}(t) = E_0\, e^{i\omega_0 t}\, e^{i\beta\,\sin(\Omega t)}. \tag{5.22}$$

Applying the Jacobi–Anger expansion (section 3.3.2) and substituting into equation (5.16), each sideband at $\omega_0 + m\Omega$ generates an independent geometric series of the same form as equation (5.20). In the steady-state limit:

$$E_{\text{cav}}^{\text{ss}}(t) = E_0 \sum_{m=-\infty}^{+\infty} J_m(\beta)\, t_{\text{in}}\, G(\omega_0 + m\Omega)\, e^{i(\omega_0+m\Omega)t}. \tag{5.23}$$

This is term-by-term identical to applying the frequency-domain transfer function $G(\omega_0+m\Omega)$ independently to each sideband. This establishes the central result of this section:

> *The discrete-time update equation (5.15) and the frequency-domain transfer function produce identical steady-state intracavity fields for any phase-modulated input field, regardless of the modulation depth β.*

Two limiting regimes follow from equation (5.23). For slow modulations ($\Omega \ll \gamma$), all sidebands experience essentially the same gain as the carrier and the cavity responds adiabatically. For fast modulations ($\Omega \gg \gamma$), successive round trips add fields with rapidly oscillating phases, leading to destructive interference and strong suppression of the sidebands. The cavity thus acts as a low-pass filter on the modulation spectrum, with a -3 dB bandwidth equal to the cavity pole frequency $f_{\text{arm}} = \gamma/(2\pi)$ [67, 66].

5.2.4 Generalisation to spatially varying fields

The scalar update equation (5.15) treats the intracavity field as a single complex number, valid only when the input field perfectly matches the fundamental TEM$_{00}$ cavity eigenmode. When the piezoelectric bender is tilted by a time-varying angle $\theta(t)$, the reflected field acquires a spatially dependent phase ramp (section 3.4.3):

$$E_{\text{in}}(\mathbf{r}_\perp, t) = E_0(\mathbf{r}_\perp)\, e^{ikx\,\sin\theta(t)} \approx E_0(\mathbf{r}_\perp)\, e^{ikx\theta(t)}, \tag{5.24}$$

where $\mathbf{r}_\perp = (x, y)$ denotes the transverse coordinates and the small-angle approximation $\sin\theta \approx \theta$ is valid for the tilt angles considered here. This field is no longer a pure TEM$_{00}$ mode. By the modal decomposition of equation (3.11), the linear phase ramp generates coupling into all odd-order Hermite–Gaussian modes, with the TEM$_{10}$ mode receiving the largest contribution [55].

Because different transverse modes accumulate different Gouy phases per round trip, the round-trip operation can no longer be represented by a single scalar factor r_{rt}. It must be replaced by a round-trip propagation operator $\hat{U}_{\text{RT}}$ acting on the full two-dimensional field profile. The update equation (equation 5.15) generalises to

$$E_{n+1}(\mathbf{r}_\perp) = \hat{U}_{\text{RT}}[E_n(\mathbf{r}_\perp)] + t_{\text{in}}\, E_{\text{in}}^{(n+1)}(\mathbf{r}_\perp), \tag{5.25}$$

where $\hat{U}_{\text{RT}}$ encodes free-space propagation over $2L$, phase imprints from the mirror surface profiles and curvatures, diffraction and aperture clipping effects, and the mode-order-dependent Gouy phase $(m+n+1)\Delta\zeta_{\text{RT}}$ accumulated per round trip (section 3.1.2).

Because $\hat{U}_{\text{RT}}$ is mode-order-dependent, a tilted input that carries power in TEM$_{10}$ finds that mode shifted in frequency by $\Delta\nu_{10} \sim 10$ kHz from the TEM$_{00}$ resonance, far exceeding the cavity linewidth of ~ 100 Hz. Unless $\Delta\nu_{10}$ coincidentally falls on a cavity resonance, TEM$_{10}$ does not build up resonantly and its contribution to the intracavity field is strongly suppressed.

The operator form (5.25) is the foundation of the Oscar simulations described in section 5.1.1. Rather than decomposing the field into HG modes explicitly, Oscar represents $E_n(\mathbf{r}_\perp)$ on a two-dimensional discrete spatial grid and implements $\hat{U}_{\text{RT}}$ via FFT-based propagation and mirror phase masks. The iterative solution of equation (5.25) until convergence constitutes the Oscar steady-state algorithm as described in section 5.1.1.

5.3 OSCAR-compatible iterative implementation

The time-domain update equation (5.25) describes the cavity dynamics in terms of the intracavity field E_n at each time step. Implementing this directly would require propagating the full two-dimensional field $E_n(\mathbf{r}_\perp)$ through the cavity at every iteration, which is precisely what OSCAR does in its steady-state solver. However, OSCAR is designed to compute the steady-state intracavity field for a given fixed input field, not to track a time-varying intracavity field explicitly. This section derives a reformulation of equation (5.25) that is fully compatible with OSCAR's steady-state approach by moving the cavity memory from the intracavity field to the input field. The effective input-field formulation presented in this section is the main methodological contribution of this thesis on the simulation side, enabling the use of OSCAR's existing steady-state solver for time-varying scattered fields without modifying the code itself.

5.3.1 Cavity decay rate and storage time

The starting point is the discrete-time update equation (5.15). For a continuous-time description, it is more natural to characterise the decay by a rate rather than a per-round-trip factor. The field decay rate γ is defined by requiring that the intracavity field decays as $e^{-\gamma t}$ in the absence of any input. After one round trip of duration t_{RT} the field must be attenuated by r_{rt}, so

$$e^{-\gamma t_{\text{RT}}} = r_{\text{rt}} \implies \gamma = -\frac{\ln r_{\text{rt}}}{t_{\text{RT}}} \approx \frac{1 - r_{\text{rt}}}{t_{\text{RT}}} = 2\pi f_{\text{arm}}, \tag{5.26}$$

where the approximation uses $\ln r_{\text{rt}} \approx -(1 - r_{\text{rt}})$ for r_{rt} close to unity, and the last equality uses the definition of the cavity pole frequency from equation (3.21). The field decay rate γ and the cavity pole f_{arm} are therefore the same physical quantity expressed in different units.

The associated cavity storage time is

$$\tau_{\text{cav}} = \frac{1}{\gamma} = \frac{t_{\text{RT}}}{1 - r_{\text{rt}}} \approx \frac{\mathcal{F}\, t_{\text{RT}}}{\pi} = \frac{1}{2\pi f_{\text{arm}}}, \tag{5.27}$$

where the relation $\mathcal{F} \approx \pi/(1 - r_{\text{rt}})$ for a high-finesse cavity has been used. The per-round-trip survival factor and its generalisation to an arbitrary time step Δt are

$$f_{\text{old}} \equiv e^{-\gamma \Delta t} = r_{\text{rt}}^{\Delta t / t_{\text{RT}}}. \tag{5.28}$$

This generalisation is important for the numerical implementation, where the time step Δt is set by the desired temporal resolution of the simulation rather than by t_{RT}.

5.3.2 Effective input-field formulation

OSCAR computes the exact steady-state intracavity field $E_{\text{cav}} = \mathcal{C}[E_{\text{in}}]$ for a given input field E_{in}, where $\mathcal{C}[\cdot]$ denotes the steady-state cavity response operator. The goal is to reformulate the time-domain update so that at each iteration n only a single steady-state calculation is required, driven by an effective input field that already encodes the cavity memory.

Suppose the cavity is in steady state with input $E_{\text{in}}^{(n)}$ and intracavity field $E_{\text{cav}}^{(n)} = \mathcal{C}[E_{\text{in}}^{(n)}]$. At the next time step, the input changes to $E_{\text{new}}^{(n+1)}$. The effective input field $E_{\text{eff}}^{(n+1)}$ is constructed such that, when passed to OSCAR's steady-state solver, it produces the same result as the exact update (5.25). Since $\mathcal{C}[\cdot]$ is a linear operator, this is satisfied by

$$E_{\text{eff}}^{(n+1)} = f_{\text{old}}\, E_{\text{in}}^{(n)} + f_{\text{new}}\, E_{\text{new}}^{(n+1)}, \tag{5.29}$$

where

$$f_{\text{old}} = e^{-\gamma \Delta t}, \tag{5.30}$$

$$f_{\text{new}} = 1 - f_{\text{old}}. \tag{5.31}$$

The choice $f_{\text{new}} = 1 - f_{\text{old}}$ ensures that in steady state ($E_{\text{new}}^{(n+1)} = E_{\text{in}}^{(n)} \equiv E_{\text{in}}$), the effective input reduces to the true input:

$$E_{\text{eff}} = f_{\text{old}}\, E_{\text{in}} + (1 - f_{\text{old}})\, E_{\text{in}} = E_{\text{in}}. \tag{5.32}$$

The full iterative scheme is therefore

$$E_{\text{eff}}^{(n+1)} = f_{\text{old}}\, E_{\text{in}}^{(n)} + (1 - f_{\text{old}})\, E_{\text{new}}^{(n+1)}, \qquad E_{\text{cav}}^{(n+1)} = \mathcal{C}\left[E_{\text{eff}}^{(n+1)}\right]. \tag{5.33}$$

At each iteration, $E_{\text{eff}}^{(n+1)}$ is constructed from the previous effective input and the newly generated scattered field, then passed to OSCAR as if it were a static input. The cavity memory is handled entirely at the level of the input field through the weights f_{old} and f_{new}.

The physical interpretation of equation (5.33) is straightforward: $f_{\text{old}} = e^{-\gamma \Delta t}$ is the fraction of the previously stored field that survives one time step. For $\Delta t \ll \tau_{\text{cav}}$ the cavity retains essentially perfect memory. For $\Delta t \gg \tau_{\text{cav}}$ it has completely forgotten the previous state. The two weights sum to unity, ensuring that the total field amplitude is conserved in steady state. The formulation (5.33) is mathematically equivalent to the original update equation (5.15) under the assumption that the input field varies slowly compared to the round-trip time t_{RT}, i.e. that $\Delta t \gg t_{\text{RT}}$.

5.3.3 Round-trip phase and additional phase contributions

In section 5.2.3, the frequency-dependent round-trip factor was written as $r_{\text{rt}}(\omega) = r_{\text{rt}}\, e^{-i\omega t_{\text{RT}}}$, encoding only the free-propagation phase. In a real cavity, additional phase contributions are present: the Gouy phase accumulated per round trip, phase shifts from the mirror coatings, and geometric phases from the beam path. Collecting all such contributions into a single term ϕ_{extra}, the full round-trip factor becomes

$$r_{\text{rt}}(\omega) = r_{\text{rt}}\, e^{-i\omega t_{\text{RT}}}\, e^{-i\phi_{\text{extra}}(\omega)}. \tag{5.34}$$

In the scalar, single-mode derivation of section 5.2.3, ϕ_{extra} was set to zero without loss of generality. In the OSCAR simulations, ϕ_{extra} does not need to be treated separately. The steady-state operator $\mathcal{C}[\cdot]$ propagates the full two-dimensional field through one complete round trip, including all phase contributions: free-space propagation, mirror surface profiles, diffraction, and the mode-order-dependent Gouy phase. The resonance condition is satisfied self-consistently by OSCAR's iterative solver. Any phase errors introduced by a tilted input beam are therefore automatically included in the steady-state solution, and the time-dependent part of the problem reduces to constructing the correct sequence of effective input fields $E_{\text{eff}}^{(n)}$.

5.4 Tilt-induced mode coupling

When the piezoelectric bender surface is perfectly perpendicular to the incident beam, the reflected field achieves maximum coupling to the cavity. For the calibration, it is essential to use data taken near this perpendicular position. This section derives how a tilt θ generates higher-order Hermite–Gaussian modes in the reflected field, quantifies the power coupled into each mode, analyses the consequences for the intracavity field and the calibration signal, explains how mode matching is achieved in both the physical end bench and the Oscar simulation, and derives the grid sampling constraints imposed by the tilt.

5.4.1 Coupling coefficients for a reflected beam

The tilt-induced mode coupling was introduced in section 3.4.3, where it was shown that a scattered field arriving at a small angle with respect to the cavity axis acquires a linear transverse phase ramp (equation (3.47)). Upon reflection off a surface tilted by θ about the y-axis, the reflected beam is deflected by 2θ, so the scattered field at the end mirror plane is [54]

$$E_{\text{sc}}(x, y) = \sqrt{T_{\text{WE}}}\, u_{00}(x, y, z_{\text{EM}})\, e^{i2kx\theta}, \tag{5.35}$$

where $k = 2\pi/\lambda$, T_{WE} is the power transmission of the west end mirror, z_{EM} is the axial position of the end mirror, and $u_{00}(x, y, z_{\text{EM}})$ is the cavity TEM$_{00}$ mode evaluated at that plane. The coupling coefficient into mode TEM$_{mn}$ follows from the overlap integral of equation (3.42):

$$c_{mn}(\theta) = \iint u_{mn}^*(x, y, z_{\text{EM}})\, u_{00}(x, y, z_{\text{EM}})\, e^{i2kx\theta}\, \mathrm{d}x\, \mathrm{d}y. \tag{5.36}$$

Because the phase ramp depends only on x, and because the HG modes factorise as $u_{mn}(x, y, z) = u_m(x, z)\, u_n(y, z)$ (equation (3.10)), the y-integral gives δ_{n0}: a tilt about the y-axis generates only modes with $n = 0$. Evaluating the x-integral by substituting the explicit HG mode forms and completing the square in the exponent yields [55, 70]

$$c_{m0}(\theta) = \frac{(ikw\theta)^m}{\sqrt{m!}}\, e^{-\frac{(kw\theta)^2}{2}}, \tag{5.37}$$

giving a power fraction

$$|c_{m0}(\theta)|^2 = \frac{(kw\theta)^{2m}}{m!}\, e^{-(kw\theta)^2}. \tag{5.38}$$

This is a Poisson distribution in the mode order m with mean $\mu = (kw\theta)^2$, where $w = w(z_{\text{EM}})$ is the beam radius at the end mirror surface. The distribution is shown in figure 5.5. It is clear to see that at $\theta = 0$ all scattered power resides in TEM$_{00}$. As $|\theta|$ increases, the distribution shifts to higher orders, with the Poisson mean $\mu = (kw\theta)^2$ following the ridge of maximum intensity.

For the Virgo arm cavities, the beam radius at the end mirror is $w(z_{\text{EM}}) \approx 5.5$ cm [21], giving a characteristic angular scale

$$\theta_c = \frac{1}{k\, w(z_{\text{EM}})} = \frac{\lambda}{2\pi\, w(z_{\text{EM}})} \approx 3.1\, \mu\text{rad}, \tag{5.39}$$

which is the angle at which the Poisson mean reaches unity ($\mu = 1$). For tilts $\theta \ll \theta_c$, nearly all power remains in the fundamental mode and the coupling is dominated by the first-order mode ($m = 1$). For $\theta \gtrsim \theta_c$, many higher-order modes become significantly excited. Geometrically, θ_c corresponds to the angular displacement at which the tilted beam has shifted by roughly one beam radius in phase space, so that the overlap with the original fundamental mode starts to drop substantially.

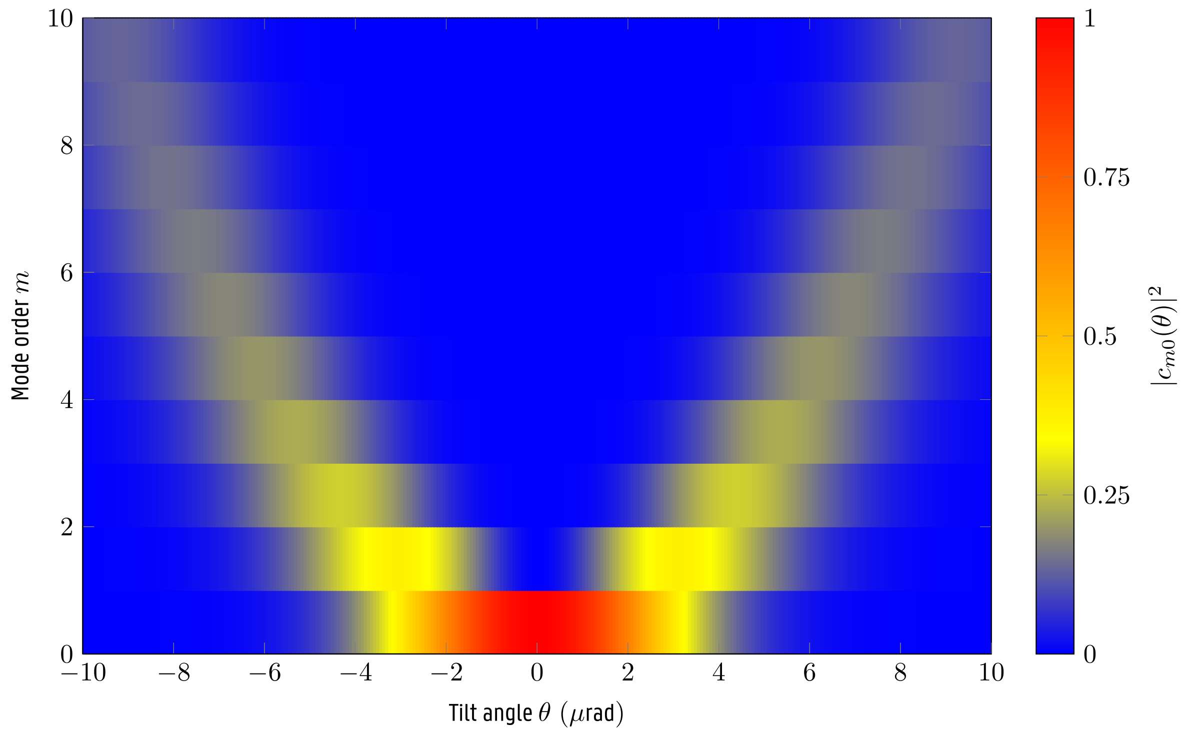

Figure 5.5: Power $|c_{m0}(\theta)|^2$ coupled into TEM$_{m0}$ as a function of bender tilt angle θ (equation (5.38)), for Virgo arm cavity parameters ($w(z_{\text{EM}}) \approx 5.5$ cm, $\lambda = 1064$ nm). At $\theta = 0$ all scattered power resides in TEM$_{00}$ ($m = 0$, maximum value 1). As $|\theta|$ increases, the distribution shifts to higher orders, with the Poisson mean $\mu = (kw\theta)^2$ following the ridge of maximum intensity.

5.4.2 Coupling behaviour during bender oscillation

The piezoelectric bender oscillates sinusoidally with angular frequency Ω_B and angular tip amplitude Θ:

$$\theta(t) = \Theta \sin(\Omega_B t). \tag{5.40}$$

The instantaneous mode-matching efficiency from equation (5.38) is

$$|c_{00}(t)| = \exp\left[-\frac{\left(k\, w(z_{\text{EM}})\, \Theta\right)^2}{2} \sin^2(\Omega_B t)\right], \tag{5.41}$$

and the half-width of the coupling window for large-amplitude oscillations ($k\, w(z_{\text{EM}})\, \Theta \gg 1$) is

$$\Delta t_{1/e} \approx \frac{1}{\Omega_B\, k\, w(z_{\text{EM}})\, \Theta}. \tag{5.42}$$

The coupling window therefore narrows as the oscillation amplitude Θ or frequency Ω_B increases. This has two important consequences for the calibration. First, the calibration signal at B1 is concentrated in short bursts near the zero-crossings of $\theta(t)$, so the data analysis must identify and select these windows. Second, the maximum of f_m always occurs at the same phase of the bender oscillation, providing a natural phase reference for fitting the sinusoidal signals at B1 and B8 described in section 5.5.2. The behaviour of $f_m(t)$ over one full oscillation period is illustrated in figure 5.6.

The higher-order mode components generated by the tilt are strongly suppressed in the detected signal, through a combination of cavity dynamics, spatial filtering, and the output mode cleaner [27, 29].

Off-resonance suppression. As derived in section 3.2.2, each transverse mode TEM$_{mn}$ resonates at a frequency shifted from the TEM$_{00}$ resonance by

$$\Delta\nu_{mn} = (m+n)\, \frac{\arccos\sqrt{g_1 g_2}}{\pi}\, \nu_{\text{FSR}}. \tag{5.43}$$

For the Virgo arm cavities, this Gouy-phase splitting is of order $\Delta\nu_{10} \sim 10$ kHz, far exceeding the cavity linewidth of ~ 100 Hz. Since the carrier is locked to the TEM$_{00}$ resonance, most higher-order mode components of the scattered field sit far off-resonance and experience a strongly reduced cavity gain:

$$|G_{\text{exact}}(\Delta\nu_{mn})| \ll G_{\text{exact}}(0) \quad \text{for most } m+n > 0. \tag{5.44}$$

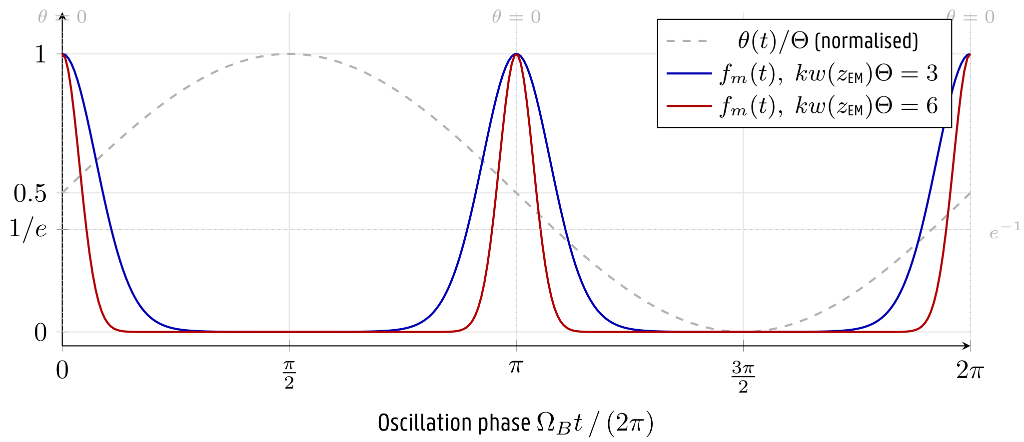

Figure 5.6: Mode-matching efficiency $f_m(t) = \exp\left[-(kw(z_{\mathrm{EM}})\theta(t))^2\right]$ as a function of oscillation phase for two values of the dimensionless tilt amplitude $kw(z_{\mathrm{EM}})\Theta$. The dashed grey curve shows the normalised tilt angle $\theta(t)/\Theta$ shifted to $[0, 1]$ for visual clarity. The coupling is sharply peaked at the zero-crossings $\theta(t) = 0$ and drops rapidly to zero elsewhere. A larger amplitude Θ narrows the coupling window, as seen by comparing the blue ($kw_0\Theta = 3$) and red ($kw_0\Theta = 6$) curves. The dotted horizontal line marks the $1/e$ level.

In practice, however, mirror surface imperfections and thermal lensing can distort the intracavity eigenmode, shifting certain higher-order mode resonances closer to the TEM$_{00}$ frequency. These near-degenerate modes can experience a non-negligible cavity gain and are not fully suppressed by the Gouy-phase splitting alone.

Output mode cleaner. To prevent such near-degenerate higher-order modes from reaching the detection photodiode, the output beam passes through an output mode cleaner (OMC) before detection at B1. The OMC is a short, independently locked cavity whose free spectral range and linewidth are chosen to transmit only the TEM$_{00}$ component, rejecting residual higher-order modes as well as the RF control sidebands.

Spatial mismatch with the local oscillator. Any residual higher-order mode content that is transmitted by the OMC is further suppressed at detection. The B1 photodiode measures the beat between the intracavity field and the local oscillator provided by the DC offset. Since the local oscillator is predominantly TEM$_{00}$, contributions from higher-order modes are suppressed by the orthogonality of the transverse mode basis:

$$\langle u_{00} | u_{mn} \rangle = \delta_{m0}\,\delta_{n0} = 0 \quad \text{for } m + n > 0. \tag{5.45}$$

For the purposes of the calibration, only the TEM$_{00}$ component of the scattered field contributes meaningfully to the signal at B1. The parameter $f_m(\theta)$ defined in equation (5.38) correctly captures this by giving the fraction of scattered power that is both spatially matched to the cavity mode and resonantly amplified.

5.4.3 Mode matching by beam reshaping

The analysis of the preceding subsections assumes that at the zero-crossing ($\theta = 0$), the returning scattered field is a perfect TEM$_{00}$ mode with $f_m = 1$. However, even at zero tilt there is a second independent source of coupling loss: the beam transmitted through the end mirror and the arm cavity TEM$_{00}$ mode do not, in general, share the same complex beam parameter q at the end mirror plane.

The arm cavity eigenmode is characterised by a specific complex beam parameter q_{cav} at the end mirror. The field transmitted through the end mirror propagates through the end bench optics (which include lenses, a beamsplitter, and steering mirrors [20, 5]) before reaching the scattering element. These optical elements transform the beam parameter along the path. After reflection off the scattering element and propagation back through the same optics to the end mirror, the returning beam arrives with a complex beam parameter $q_{\mathrm{sc}} \neq q_{\mathrm{cav}}$ in general. From equation (3.44), the resulting mode-matching efficiency is

$$\eta_q = \frac{4}{\left(\dfrac{w_{\mathrm{sc}}}{w_{\mathrm{cav}}} + \dfrac{w_{\mathrm{cav}}}{w_{\mathrm{sc}}}\right)^2 + \left(\dfrac{\lambda\,\Delta z}{\pi\, w_{\mathrm{sc}}\, w_{\mathrm{cav}}}\right)^2}, \tag{5.46}$$

where w_{sc} and w_{cav} are the beam radii at the end mirror plane and Δz is the difference in waist positions. The total coupling fraction, including both tilt and q-mismatch, is

$$f_m(\theta) = \eta_q \cdot \exp\left[-(kw(z_{\mathrm{EM}})\theta)^2\right]. \tag{5.47}$$

The value of η_q therefore depends on the specific optical configuration of the end bench. In the Virgo suspended end benches (SNEB and SWEB), the transmitted beam passes through a set of lenses that reduce the beam size for the B7/B8 photodiodes and alignment sensors [20]. These same optics also determine the beam parameter of any light reflected back toward the end mirror. In general, the round-trip transformation through the bench optics does not perfectly restore q_{cav}, so $\eta_q < 1$. The degree of mode mismatch depends on the lens positions, focal lengths, and the distance to the scattering element, and can in principle be reduced by careful optical design of the bench layout.

In the Oscar simulation, perfect mode matching ($\eta_q = 1$) is achieved by construction through time-reversal of the transmitted field after reflection off the scattering element. The time-reversed field is obtained by conjugating both the complex spatial field and the stored complex beam parameter:

$$E_{\text{tr}}(\mathbf{r}_\perp) = E^*_{\text{trans}}(\mathbf{r}_\perp), \qquad q_{\text{tr}} = q^*_{\text{trans}}. \tag{5.48}$$

Conjugating q reverses the sign of the wavefront curvature: a diverging beam with $\text{Re}(1/q) > 0$ becomes a converging beam with $\text{Re}(1/q_{\text{tr}}) < 0$. After propagating the return path back to the end mirror, the curvature is restored to its original value and the returning beam arrives with $q = q_{\text{cav}}$, giving $\eta_q = 1$ by construction. The tilt is applied to the time-reversed field before the return propagation, so the tilt-dependent factor $e^{-(kw(z_{\text{EM}})\theta)^2}$ remains in the final result.

The time-reversal approach yields $\eta_q = 1$ exactly because the return propagation is assumed to perfectly undo the outgoing propagation. In the real Virgo end bench, the beam may encounter slight astigmatism from imperfect optics, and the lens system introduces its own alignment and positioning errors [5]. Any reduction in the observed f_m relative to the simulated prediction should therefore be attributed at least in part to imperfect mode matching ($\eta_q < 1$) in the end bench rather than to the scattered light model itself.

5.4.4 Grid sampling constraints and numerical complexity

The Oscar simulation represents the two-dimensional transverse field $E(\mathbf{r}_\perp)$ on a discrete spatial grid with spacing Δx. This discretisation introduces a fundamental limitation on the spatial frequencies that can be faithfully represented, directly constraining the range of tilt angles θ that can be simulated without numerical artefacts.

When the bender surface is tilted by an angle θ, the reflected field acquires a linear transverse phase ramp (equation (5.35)), corresponding to a plane-wave component with transverse spatial frequency

$$f_x = \frac{2k\theta}{2\pi} = \frac{2\theta}{\lambda}. \tag{5.49}$$

On a discrete grid with spacing Δx, the Nyquist frequency[1] is $f_N = 1/(2\Delta x)$. Any spatial frequency exceeding f_N is aliased back into the representable interval $[-f_N, +f_N]$. Requiring $|f_x| < f_N$ gives the sampling criterion

$$|\theta| < \frac{\lambda}{4\Delta x}. \tag{5.50}$$

Violating this condition causes the phase ramp to wrap at the grid boundary, a purely numerical artefact with no physical meaning.

The sampling criterion (5.50) constrains Δx for a given maximum tilt angle θ_{max}. Simultaneously, the grid must be large enough to contain the full beam profile without truncation: for a Gaussian beam with radius $w(z_{\text{EM}})$ at the end mirror, the grid half-width must satisfy $L_{\text{grid}}/2 > 3w$ to capture more than $99\,\%$ of the beam power. Together, these two requirements determine the minimum number of grid points N along each transverse direction:

$$N > \frac{L_{\text{grid}}}{\Delta x} > \frac{6w(z_{\text{EM}})}{\Delta x} > \frac{24\,w(z_{\text{EM}})\,\theta_{\text{max}}}{\lambda}. \tag{5.51}$$

For the Virgo arm cavities with $w(z_{\text{EM}}) \approx 5.5$ cm and $\lambda = 1064$ nm, using an effective maximum tilt of $\Theta = 50\,\mu$rad beyond which f_m is negligible, equation (5.51) requires $N > 62$.

The simulations in this work use a grid of $N = 256$ points with a half-width $L_{\text{grid}}/2 = 0.25$ m, giving $\Delta x \approx 1.95$ mm. This comfortably satisfies the Nyquist criterion across the physically relevant range of bender tilts. The maximum tilt before aliasing occurs is

$$\theta_{\text{max}} = \frac{\lambda}{4\Delta x} = \frac{1064 \times 10^{-9}}{4 \times 1.95 \times 10^{-3}} \approx 136\,\mu\text{rad}, \tag{5.52}$$

which exceeds the effective maximum tilt by a factor of $136/50 \approx 2.7$. With $\Delta x \approx 1.95$ mm and $w(z_{\text{EM}}) \approx 5.5$ cm, there are approximately $w(z_{\text{EM}})/\Delta x \approx 28$ grid points across one beam radius, sufficient to represent the field profile accurately.

The dominant computational cost per iteration is the two-dimensional FFT used to propagate the field through the cavity. For a grid of $N \times N$ points, each FFT costs $\mathcal{O}(N^2 \log N)$ operations.

[1] The Nyquist–Shannon sampling theorem states that a signal sampled at intervals Δx can faithfully represent spatial frequencies up to $f_N = 1/(2\Delta x)$. Frequencies above this limit are indistinguishable from lower-frequency components, an effect known as aliasing.

5.5 Calibration

The preceding sections have established all the ingredients needed for the calibration. The scattered light model of section 4.2 gives the power at the B1 and B8 photodiodes as a function of the scattering parameters. The mode-matching analysis of section 5.4 shows that the coupling f_m (and hence $f_{r,m}$) is maximised and well-defined at the zero-crossings of the bender oscillation. This section combines these results into a concrete procedure for recovering the end mirror transmission T_{WE}, preceded by a ring-down method for measuring the cavity decay rate γ required by the iterative implementation of section 5.3.

5.5.1 Field decay rate from ring-down

Before the scattered light calibration can be applied, the cavity decay rate γ must be known accurately, since it enters the effective input-field formulation of equation (5.33) directly through the memory factor $f_{\mathsf{old}} = e^{-\gamma \Delta t}$. This subsection derives an analytical expression for the power at B1 following a brief injection of scattered light and shows how fitting this expression to the measured ring-down signal yields γ and thereby the value r_{rt} (equation 5.26). The scattering element in these calculations is located on the SWEB which results in a signal originating from the west arm.

Ring-down signal at the output port

Consider the situation immediately after a brief injection of scattered light has ended. The north arm remains in steady state throughout, while the west arm contains a transient excitation that decays exponentially at rate γ:

$$E_N(t) = E_N^{(0)}, \tag{5.53}$$

$$E_W(t) = E_W^{(0)} + E_1\, e^{-\gamma t}, \tag{5.54}$$

where $E_N^{(0)}$ and $E_W^{(0)}$ are the steady-state cavity fields and E_1 is a complex amplitude determined by the strength and phase of the injected excitation. The transient begins at $t = 0$, the moment the injection ends.

Including the small DC offset ψ that keeps the interferometer slightly off the dark fringe, the field at the antisymmetric port is

$$E_{\mathsf{out}}(t) = \frac{1}{\sqrt{2}} \Big[E_N^{(0)}\, e^{-i\psi/2} - \Big(E_W^{(0)} + E_1\, e^{-\gamma t} \Big)\, e^{+i\psi/2} \Big] \equiv E_\infty + E_T\, e^{-\gamma t}, \tag{5.55}$$

where E_∞ is the steady-state output field due to the DC offset and $E_T = (1/\sqrt{2})\, E_1\, e^{+i\psi/2}$ is the complex amplitude of the decaying transient.

Detected power and fitting formula

The power detected at B1 is $P(t) = T_{\mathsf{IM}} |E_{\mathsf{out}}(t)|^2$. Expanding equation (5.55):

$$P(t) = P_1 + P_2\, e^{-\gamma t} + P_3\, e^{-2\gamma t}, \tag{5.56}$$

with real coefficients

$$P_1 = T_{\mathsf{IM}} |E_\infty|^2, \tag{5.57}$$

$$P_2 = 2 T_{\mathsf{IM}}\ \mathsf{Re}(E_\infty E_T^*), \tag{5.58}$$

$$P_3 = T_{\mathsf{IM}} |E_T|^2. \tag{5.59}$$

The three terms have clear physical interpretations. The constant P_1 is the DC power at B1, set by the steady-state operating point. The term $P_2\, e^{-\gamma t}$ arises from the beat between the steady-state field E_∞ and the decaying transient E_T. Its sign depends on the relative phase of these two fields, and it decays at rate γ. The term $P_3\, e^{-2\gamma t}$ is the power stored in the transient field itself, decaying at twice the field decay rate since power scales as the square of the field amplitude.

Fitting equation (5.56) to the measured B1 signal after a brief scattered light injection yields γ directly. From γ, the product of the mirror reflectivities follows via

$$r_{\mathsf{ITM}}\, r_{\mathsf{ETM}} = e^{-\gamma t_{\mathsf{RT}}} = \exp\left(-\frac{2\gamma L}{c} \right). \tag{5.60}$$

Relation to the cavity pole and finesse

The decay rate γ is directly related to the cavity pole frequency f_{arm} and finesse $\mathcal{F}$ introduced in section 3.2.2. From equations (5.26) and (3.21):

$$\gamma = 2\pi f_{\mathrm{arm}} = \frac{(1 - r_{\mathrm{rt}})\, c}{2L} = \frac{\pi c}{2L\mathcal{F}}. \tag{5.61}$$

For the Virgo arm cavities with $L = 3$ km and $\mathcal{F} \approx 480$:

$$\gamma = \frac{\pi \times 3 \times 10^8}{2 \times 3000 \times 480} \approx 327\,\mathrm{s}^{-1}, \tag{5.62}$$

corresponding to a storage time $\tau_{\mathrm{cav}} = 1/\gamma \approx 3$ ms and a cavity pole $f_{\mathrm{arm}} = \gamma/(2\pi) \approx 52.1$ Hz. The ring-down signal following a brief injection therefore decays on a millisecond timescale, well resolved by the B1 photodiode at typical sampling rates of ~ 10 kHz. The ring-down fit (5.56) thus serves a dual purpose: it provides an independent measurement of γ that can be compared to the value inferred from the finesse measurement [21], and it validates the cavity memory model of section 5.3 before applying it to the full scattered light calibration.

5.5.2 Calibration procedure

The scattered light signal from the end-bench bender can be used not only to characterise the scattering properties of the bench, but also to provide an independent, absolute calibration of the detector's strain response. The idea, first demonstrated during the third observing run of Advanced Virgo [5], is that the scattered light enters the arm cavity through the same optical path as a gravitational wave, yet its amplitude can be determined independently from the B8 photodiode on the end bench, without involving the detector's strain calibration. Comparing what the detector reports ($\hat{h}_{\mathrm{rec}}$) with what the signal is known to be ($\hat{h}_{\mathrm{signal}}$) then directly reveals any calibration error.

This section describes the full procedure: the signal model at B1 and B8, the extraction of the relevant observables from the data, the calibration steps themselves, and the practical identification of zero-crossings required for the measurement.

The four-step procedure below follows the structure of [5], adapted here for the end-bench geometry of Figure 5.3 in which a beam splitter modifies the numerical prefactors.

Signal model at B1 and B8

From the scattered light model derived in section 4.2.1, applied to the end-bench geometry of figure 5.3 in which the B8 photodiode monitors the field via a beam splitter, the powers at the two photodiodes near a zero-crossing of $\theta(t)$ are

$$P_{\mathrm{B1}}(t) \approx \frac{1}{2}\, T_{\mathrm{IM}} |E_0|^2 \left[\psi^2 + 2\psi\, \frac{4\pi G L}{\lambda}\, h(t) + \psi\, G\, T_{\mathrm{WE}} \sqrt{f_{r,m}}\, \sin\!\big(\phi_{\mathrm{sc}}(t) + \phi_{\mathrm{B1}}\big)\right], \tag{5.63}$$

$$P_{\mathrm{B8}}(t) \approx \frac{\alpha}{2}\, T_{\mathrm{WE}} |E_0|^2 \left[1 + \sqrt{f_{r,m}}\, \cos\!\big(\phi_{\mathrm{sc}}(t) + \phi_{\mathrm{B8}}\big)\right], \tag{5.64}$$

where $\phi_{\mathrm{sc}}(t) = 4\pi x(t)/\lambda$ is the round-trip scattered light phase, ϕ_{B1} and ϕ_{B8} are static phase offsets arising from the finite mirror thickness and the additional optical path to B8, α accounts for optical losses between the end mirror and B8, and $G = G_{\mathrm{exact}}(f_{\mathrm{signal}})$ is the cavity gain at the signal frequency. The scattered light is assumed here to experience the same cavity gain as the circulating light. Compared to the simplified model of equation (4.15), the beam splitter on the end bench introduces an additional factor of $1/2$ in the scattered-field amplitude reaching the cavity, which cancels the factor of 2 in the cross-term. These expressions are valid under the two conditions established in sections 5.4.2 and 5.4.3: that the bender is near a zero-crossing ($f_m \approx 1$) and that the returning beam is mode-matched to the cavity ($\eta_q \approx 1$).

It should be noted that the real optical layout of the Virgo end benches is considerably more complex than the simplified geometry used here (as can be seen from figure 4.1). The expressions above capture the essential physics of the calibration method, but the numerical prefactors would need to be revised when applying the method to the actual Advanced Virgo end bench configuration.

Observable extraction

The calibration requires four measured quantities, extracted from the B1 and B8 time series in two stages.

Baseline powers. When the bender is far from a zero-crossing, $f_m \approx 0$ and the oscillating terms in equations (5.63) and (5.64) vanish. The signals reduce to their baseline values:

$$P_{\text{B1,base}} = \frac{T_{\text{IM}}}{2} |E_0|^2 \, \psi^2, \tag{5.65}$$

$$P_{\text{B8,base}} = \frac{\alpha \, T_{\text{WE}}}{2} |E_0|^2. \tag{5.66}$$

These are extracted by averaging over time intervals where the bender is well away from perpendicular. They provide measurements of the products $T_{\text{IM}}|E_0|^2\psi^2$ and $\alpha \, T_{\text{WE}}|E_0|^2$, respectively, without requiring knowledge of the individual factors.

Oscillation amplitudes. Near a zero-crossing, the bender passes through perpendicular at angular velocity $\dot{\theta} = \Omega_B \Theta$, producing a locally sinusoidal scattered light phase $\phi_{\text{sc}}(t) \approx \phi_0 + 2\pi v_{\text{inst}} \, t$, where v_{inst} is the instantaneous tip velocity. Fitting a sinusoid to each baseline-subtracted signal near the zero-crossing yields the amplitudes

$$A_{\text{B1}} = \frac{T_{\text{IM}}}{2} |E_0|^2 \, \psi \, G \, T_{\text{WE}} \sqrt{f_{r,m}}, \tag{5.67}$$

$$A_{\text{B8}} = \frac{\alpha \, T_{\text{WE}}}{2} |E_0|^2 \sqrt{f_{r,m}}. \tag{5.68}$$

The four observables $P_{\text{B1,base}}$, $P_{\text{B8,base}}$, A_{B1}, and A_{B8} form a system from which the scattering fraction and end mirror transmission can be extracted, as described in the following section.

Calibration steps

We define the reconstructed strain amplitude $\hat{h}_{\text{rec}}$ as the gravitational wave amplitude that the detector's calibration pipeline assigns to a given power fluctuation at B1. If the calibration is perfect, $\hat{h}_{\text{rec}}$ equals the true strain equivalent of the scattered light signal. Any miscalibration appears as a systematic offset in $\hat{h}_{\text{rec}}$.

Step 1: determine $f_{r,m}$ from B8. Dividing the oscillation amplitude A_{B8} by the baseline power $P_{\text{B8,base}}$:

$$\frac{A_{\text{B8}}}{P_{\text{B8,base}}} = \frac{\alpha T_{\text{WE}}/2 \cdot |E_0|^2 \sqrt{f_{r,m}}}{\alpha T_{\text{WE}}/2 \cdot |E_0|^2} = \sqrt{f_{r,m}} \quad \Rightarrow \quad f_{r,m} = \left(\frac{A_{\text{B8}}}{P_{\text{B8,base}}} \right)^2. \tag{5.69}$$

The factors α, T_{WE}, and $|E_0|^2$ cancel in the ratio, so $f_{r,m}$ depends only on measured quantities at B8.

Step 2: verify the DC offset ψ. As a consistency check, the DC offset can be inferred from the B1 baseline power:

$$\psi = \sqrt{\frac{2 \, P_{\text{B1,base}}}{T_{\text{IM}} |E_0|^2}}. \tag{5.70}$$

Since T_{IM} and $|E_0|^2$ are known independently, this can be compared to the value set by the control system.

Step 3: determine T_{WE} from B1. With $f_{r,m}$ known from step 1, the end mirror transmission follows from the B1 amplitude-to-baseline ratio:

$$\frac{A_{\text{B1}}}{P_{\text{B1,base}}} = \frac{G \, T_{\text{WE}} \sqrt{f_{r,m}}}{\psi} \quad \Rightarrow \quad T_{\text{WE}} = \frac{\psi}{\sqrt{f_{r,m}} \, G} \cdot \frac{A_{\text{B1}}}{P_{\text{B1,base}}}. \tag{5.71}$$

Again T_{IM} and $|E_0|^2$ cancel. This result can be expressed in terms of the reconstructed strain by noting that, by definition of $\hat{h}_{\text{rec}}$,

$$A_{\text{B1}} = T_{\text{IM}} |E_0|^2 \, \psi \, G \, \frac{4\pi L}{\lambda} \, \hat{h}_{\text{rec}}. \tag{5.72}$$

Dividing by $P_{\text{B1,base}}$ and substituting into equation (5.71), the factors G and ψ cancel:

$$T_{\text{WE}} = \frac{8\pi \, L}{\sqrt{f_{r,m}} \, \lambda} \, \hat{h}_{\text{rec}}. \tag{5.73}$$

The factor of 8π instead of 4π compared to equation (4.19) arises from the beam splitter in the end-bench layout (figure 5.3). The cancellation of G and ψ reflects the fact that converting the B1 power fluctuation into an equivalent strain divides out the same detector response that amplified the scattered light signal in the first place, leaving only the intrinsic scattering parameters and known constants [5].

Step 4: absolute calibration from the strain ratio. Equation (5.73) provides a measurement of T_{WE} that depends on the detector's strain calibration through $\hat{h}_{\text{rec}}$. If an independent laboratory measurement $T_{\text{WE}}^{\text{meas}}$ of the end mirror transmission is available, one can construct an absolutely calibrated reference strain by inverting equation (5.73):

$$\hat{h}_{\text{signal}} = \frac{T_{\text{WE}}^{\text{meas}} \sqrt{f_{r,m}}\, \lambda}{8\pi\, L}. \tag{5.74}$$

This reference strain depends only on $T_{\text{WE}}^{\text{meas}}$, the scattering fraction $f_{r,m}$ from step 1, and the precisely known constants L and λ. None of these involve the detector's strain calibration. The ratio of the reconstructed strain to this reference then directly probes the calibration accuracy:

$$\frac{\hat{h}_{\text{rec}}}{\hat{h}_{\text{signal}}} = \frac{T_{\text{WE}}}{T_{\text{WE}}^{\text{meas}}}. \tag{5.75}$$

If the calibration is perfect this ratio equals unity. Any deviation indicates a systematic error in the detector's strain response. During O3, this ratio was measured to be 1.009 ± 0.05 for the west end mirror and 1.014 ± 0.025 for the north end mirror [5], consistent with unity in both cases. The dominant uncertainty arises from the 2–5 % precision of the laboratory coating transmission measurement. The statistical error from the scattered light measurement itself is of order 1 % and could be reduced with longer dedicated injections.

This calibration method is fully independent of the photon calibrator [47, 71], which relies on an auxiliary laser pushing on the test masses through radiation pressure and is currently limited by disagreements between national laser power standards at the level of several percent. The scattered light method requires only a relative power measurement at B8 and a laboratory measurement of the mirror coating transmission, neither of which involves absolute laser power calibration.

Averaging and uncertainty propagation. The sinusoidal fits are performed over a short window around a single zero-crossing, so the statistical precision of A_{B1} and A_{B8} is limited by the number of samples in each window. Averaging the recovered T_{WE} over multiple zero-crossings across several bender oscillation cycles reduces this uncertainty and suppresses systematic effects from residual seismic motion or slow drifts in the operating point. The dominant remaining sources of uncertainty are the fit amplitudes A_{B1} and A_{B8} (the latter entering through $f_{r,m}$), the cavity gain G at the signal frequency, and the DC offset ψ. These are propagated in quadrature through equations (5.69) and (5.71) to obtain the combined relative uncertainty on T_{WE}.

Cavity phase correction

The calibration formula (5.71) uses the cavity amplitude gain G evaluated at the fitted signal frequency f_{fit}. In equation (6.1) this gain is approximated by its magnitude,

$$G_{\text{corr}} = \frac{G_{\text{max}}}{\sqrt{1 + (f_{\text{fit}}/f_{\text{arm}})^2}}, \tag{5.76}$$

which neglects the phase lag $\phi_s = \arctan(f_{\text{fit}}/f_{\text{arm}})$ introduced by the finite cavity storage time (appendix C, equation (C.4). For signal frequencies well below the cavity pole this approximation is excellent, but once $f_{\text{fit}} \gtrsim f_{\text{arm}}$ the phase lag distorts the shape of the transient scattered-light signal at the zero-crossing and biases the sinusoidal fit amplitude A_{B1}.

The bias can be eliminated by removing the cavity pole from the ΔP_{B1} time series *before* the sinusoidal fit. The arm cavity acts as a single-pole low-pass filter with transfer function

$$H(f) = \frac{1}{1 + i\, f/f_{\text{arm}}}. \tag{5.77}$$

Inverting this filter in the time domain amounts to applying the inverse of the first-order ordinary differential equation that defines the cavity impulse response [5]:

$$\Delta P_{B1}^{(\text{deconv})}(t) = \Delta P_{B1}(t) + \tau_{\text{arm}} \frac{\mathrm{d}\,\Delta P_{B1}}{\mathrm{d}t}(t), \tag{5.78}$$

where $\tau_{\text{arm}} = 1/(2\pi f_{\text{arm}})$ is the cavity storage time. Because equation (5.78) is a purely algebraic operation on the time series and its derivative, it is valid for transient signals and does not assume steady-state oscillation. The derivative is evaluated numerically using centred finite differences on the full simulation time series before extracting the zero-crossing window.

Derivation. The normalised cavity transfer function $H(f) = 1/(1 + if/f_{\text{arm}})$ is equivalent in the time domain to the first-order linear ODE

$$\tau_{\text{arm}} \frac{\mathrm{d}y}{\mathrm{d}t}(t) + y(t) = x(t), \tag{5.79}$$

where $x(t)$ is the input signal, $y(t)$ is the filtered output, and $\tau_{\text{arm}} = 1/(2\pi f_{\text{arm}})$ is the cavity storage time. This can be verified by Fourier transformation: replacing $\mathrm{d}/\mathrm{d}t \to i\omega$ yields $\hat{Y}(\omega)\,(1 + i\omega\tau_{\text{arm}}) = \hat{X}(\omega)$, hence $\hat{Y}/\hat{X} = H(f)$. Rearranging equation (5.79) for the input gives the exact inverse filter

$$x(t) = y(t) + \tau_{\text{arm}} \frac{\mathrm{d}y}{\mathrm{d}t}(t)\,, \tag{5.80}$$

which is equation (5.78). No approximation is involved: the identity holds for arbitrary time dependence of $y(t)$, including transient signals [72].

After deconvolution, the cavity pole has been removed from $\Delta P_{B1}^{(\text{deconv})}$: the signal behaves as if the cavity had a flat, frequency-independent gain G_{max}. Equation (5.71) is then applied with $G = G_{\text{max}}$ in place of G_{corr}:

$$T_{\text{WE}} = \frac{\psi}{\sqrt{f_{r,m}}\, G_{\text{max}}} \cdot \frac{A_{B1}^{(\text{deconv})}}{P_{B1,\text{base}}}\,, \tag{5.81}$$

where $A_{B1}^{(\text{deconv})}$ is the amplitude obtained by fitting a sinusoid to the deconvolved signal inside the zero-crossing window. The B8 signal requires no deconvolution because it is measured outside the arm cavity and is not filtered by the cavity transfer function.

A practical consideration is that the time derivative in equation (5.78) amplifies high-frequency noise. For the simulation data used in this work this is negligible, but in an experimental implementation a low-pass differentiator (for example a Savitzky–Golay filter [73]) may be needed to suppress sensor noise above the frequency band of interest.

Zero-crossing identification

The calibration relies on fitting the oscillation amplitudes near the bender zero-crossing, where the tilt vanishes and the mode-matching efficiency peaks at $f_m \approx 1$. Identifying the zero-crossing time t_0 accurately is therefore essential: any timing error δt shifts the fit window away from perpendicular, reducing the effective mode-matching and biasing the recovered T_{WE}.

Simulation. In the simulation, the zero-crossing times are known exactly from the analytic drive signal $\theta(t) = \Theta \sin(\Omega_B t)$: they occur at $t_0 = n\pi/\Omega_B$ for integer n, and the fit window is centred on these times by construction.

Experimental identification. In an experimental implementation, t_0 must be inferred from the data. Three complementary approaches are available, listed in order of decreasing reliance on external information.

1. Drive signal synchronisation. The voltage waveform sent to the piezo controller is recorded alongside the photodiode signals. Since the bender displacement is proportional to the applied voltage at frequencies well below both the mechanical resonance and the RC bandwidth of section 7.1.5, the zero-crossings of the drive voltage directly mark the zero-crossings of the bender angle. The dominant timing error comes from the phase lag of the RC filter formed by the protection resistors and the piezo capacitance. At a drive frequency of ~ 10 mHz, this lag is

$$\varphi = \arctan\left(\frac{f_{\text{mod}}}{f_{-3\,\text{dB}}}\right) \approx \arctan\left(\frac{0.01}{29}\right) \approx 0.02^\circ, \tag{5.82}$$

corresponding to a timing offset of $\delta t_{\text{RC}} = \varphi/(2\pi f_{\text{mod}}) \approx 0.5$ ms, which is small compared to the duration of the zero-crossing window.

2. B8 power envelope. The B8 signal contains an oscillating term whose amplitude is modulated by the mode-matching efficiency $f_m(t)$. Near a zero-crossing, f_m peaks sharply (figure 5.6), producing a burst of high-frequency oscillation in P_{B8}. The zero-crossing time can be estimated as the centroid of the envelope of this burst. This approach is self-contained: it uses only the photodiode data and does not require knowledge of the drive signal timing.

3. B1 signal amplitude. An analogous envelope peak appears in the baseline-subtracted signal ΔP_{B1}. Because the B1 signal passes through the full interferometer transfer function, its signal-to-noise ratio may differ from that of B8, but the timing information is redundant and serves as a consistency check.

Sensitivity to timing errors. A timing offset δt from the true zero-crossing shifts the fit window to a region where the instantaneous tilt is $\theta(\delta t) \approx \Omega_B \Theta\, \delta t$. The mode-matching efficiency at this offset is

$$f_m(\delta t) = \exp\Big[-\big(k\, w(z_{\text{EM}})\, \Omega_B\, \Theta\, \delta t\big)^2\Big]. \tag{5.83}$$

For the parameters of this work ($k\, w(z_{\text{EM}})\, \Theta \approx 115$, $\Omega_B = 2\pi \times 0.01$ rad s^{-1}), the condition $f_m > 0.99$ is satisfied for $|\delta t| < 14$ ms. The method is therefore robust to timing errors at the millisecond level, well within the precision achievable by all three identification approaches.

6

Simulation results

This chapter applies the calibration procedure developed in chapter 5 to synthetic data generated with the Oscar simulation, in order to validate the method under controlled conditions where all physical parameters are known exactly. The simulation uses an end-mirror transmission $T_{\text{WE,ref}} = 3.39 \times 10^{-6}$ and a scattering fraction $f_r = 10^{-6}$, drives the piezoelectric bender at modulation frequencies $f_{\text{mod}} \in [10^{-3},\ 50]$ Hz with a fixed tip amplitude $x_{\text{tip}} = 10\ \mu$m, and records the powers $P_{\text{B1}}(t)$ and $P_{\text{B8}}(t)$. The fitted amplitudes A_{B1} and A_{B8}, extracted over a nine-sample window centred on the bender zero-crossing, are combined with the baseline powers and the cavity gain to recover T_{WE} through equation (5.71).

6.1 Zero-crossing waveforms and fit quality

Before extracting any calibration parameters, it is important to verify that the scattered-light signal is well described by a single-frequency sinusoid within the zero-crossing window. This section examines the waveforms at B1 and B8 and assesses the quality of the sinusoidal fit that underpins the entire procedure.

Figure 6.1 shows the baseline-subtracted photocurrents $\Delta P_{\text{B1}}(t)$ and $\Delta P_{\text{B8}}(t)$ for the lowest drive frequency $f_{\text{mod}} = 1$ mHz, together with the sinusoidal fit obtained from the nine samples inside the zero-crossing window (shaded orange region).

Inside the window, the fit reproduces the simulation data to within the plotting resolution, confirming that the scattered-light signal is well approximated by a single-frequency sinusoid over the short interval where the mode-matching efficiency $f_m \approx 1$. Outside the window, the fit diverges rapidly from the data because f_m drops below unity and the true signal is a nonlinear function of the bender phase $\varphi_{\text{sc}}(t)$ whose instantaneous frequency varies with the bender velocity. This contrast justifies restricting the calibration to the zero-crossing window established in section 5.5.2.

The fitted frequency $f_{\text{fit}} = 0.241$ Hz is notably higher than the value $\beta \cdot f_{\text{mod}} \approx 0.118$ Hz (equation (3.41)) that would be expected from the modulation depth alone. This discrepancy arises because the bender is a finite-length cantilever: as it deflects, it introduces a tilt as well as a displacement, causing the waveform to deviate from a pure sinusoid near the zero-crossing. Simulations investigating the dependence of the instantaneous frequency on bender length show that for longer benders (where the tilt per unit displacement is smaller), the fitted frequency converges asymptotically to $\beta \cdot f_{\text{mod}}$, confirming that the discrepancy is a geometric effect of the finite bender length rather than a modelling error. This dependence is analysed in detail in section 6.7. All simulations reported in the following sections use a bender length of 28 mm, matching the physical device (section 7.1).

Importantly, this frequency shift does not affect the calibration, since the fitted amplitude, which is the quantity that enters the T_{WE} recovery remains unchanged with varying bender length.

6.2 Fitted signal frequency and bender geometry

The frequency of the scattered-light signal at the zero-crossing is set by the instantaneous velocity of the bender tip and determines where the signal sits relative to the cavity pole. This section examines how the fitted signal frequency scales with the drive frequency across the full range of datasets.

Figure 6.2 plots the fitted signal frequency f_{fit} against the drive frequency f_{mod} for all datasets. The points follow a power law with unit slope on log–log axes, confirming the expected linear scaling. A linear fit yields the proportionality constant $f_{\text{fit}}/f_{\text{mod}} \simeq 241$, which exceeds the modulation depth $\beta \approx 118$ for the reasons discussed in section 6.1.

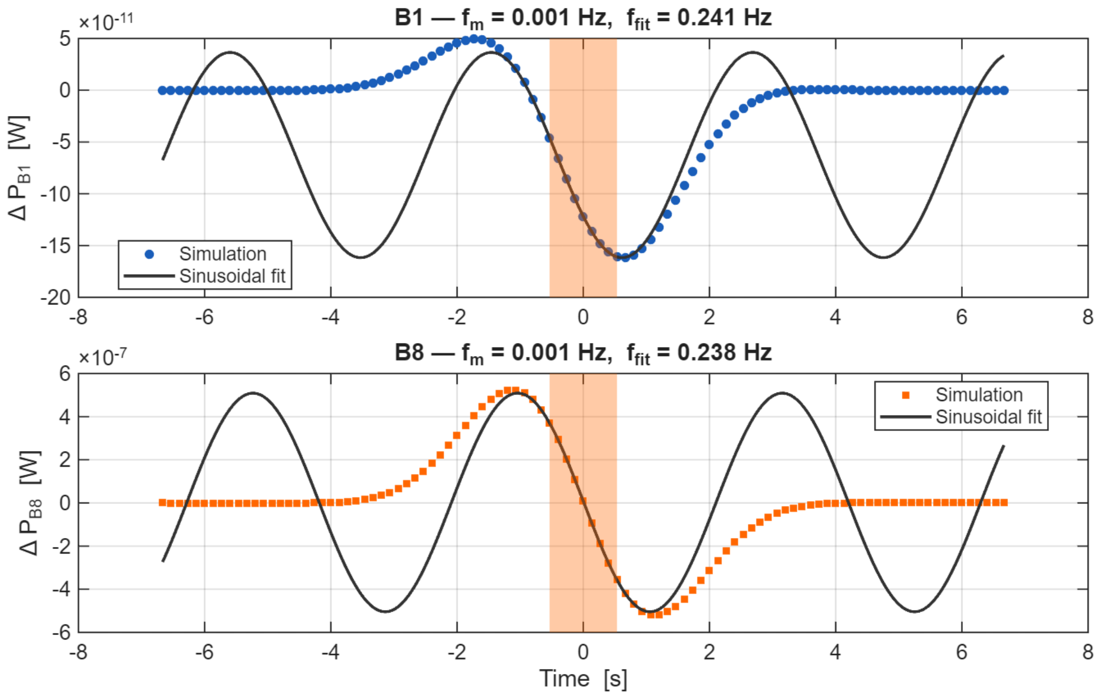

Figure 6.1: Baseline-subtracted photocurrents ΔP_{B1} (top) and ΔP_{B8} (bottom) at $f_{\mathrm{mod}} = 1$ mHz, shown over the entire simulated interval. The solid grey curve is the single-frequency sinusoid obtained by fitting only the nine samples within the shaded window. Inside the window the fit reproduces the data to within the plotting resolution, outside the window the fit and the simulation diverge because the true signal is not a single-frequency sinusoid. The fitted frequency is $f_{\mathrm{fit}} = 0.2415$ Hz, compared with the drive frequency $f_{\mathrm{mod}} = 1$ mHz.

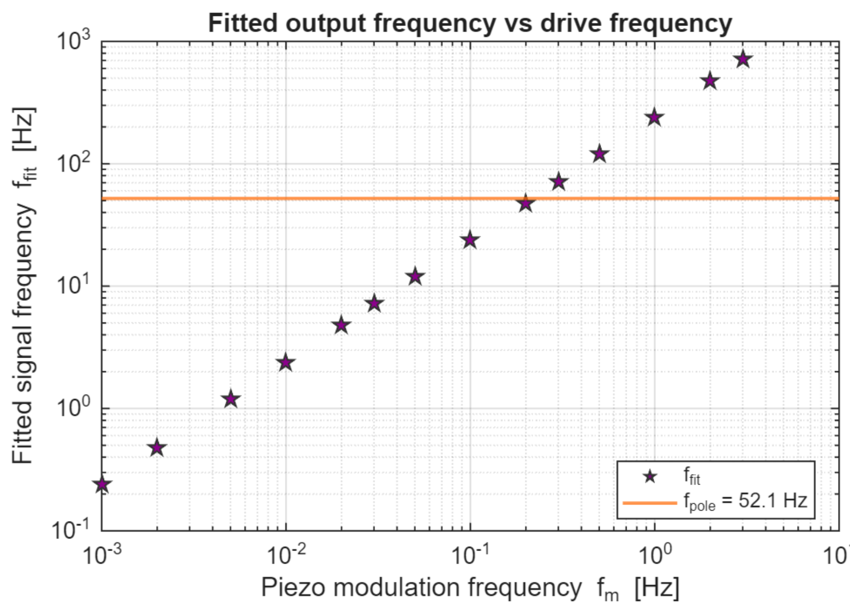

Figure 6.2: Fitted signal frequency f_{fit} as a function of the piezo drive frequency f_{mod}, on a log–log scale. The data follow $f_{\mathrm{fit}} \simeq 241\, f_{\mathrm{mod}}$. The horizontal line marks the cavity pole $f_{\mathrm{arm}} = 52.1$ Hz. Drive frequencies above $f_{\mathrm{mod}} \simeq 0.2$ Hz place the scattered-light signal above the pole, where the cavity transfer function begins to attenuate the signal.

The horizontal line in figure 6.2 marks the cavity pole $f_{\text{arm}} = 52.1$ Hz. The data cross this line near $f_{\text{mod}} \simeq 0.2$ Hz, dividing the parameter space into two regimes: below this threshold the cavity gain is effectively frequency-independent, while above it the pole correction of equation (6.1) becomes significant. This boundary reappears in sections 6.3 and 6.6.

6.3 Cavity gain

The cavity gain at the signal frequency is the final ingredient needed to convert the fitted B1 amplitude into a measurement of T_{WE}. This section examines how the gain depends on the bender dynamics and compares the instantaneous gain $G(\theta)$ with the pole-corrected value G_{corr} used in the calibration.

Figure 6.3 shows the cavity amplitude gain $G(\theta)$ as a function of bender tilt angle for different drive frequencies. At low f_{mod}, the gain traces a sharply peaked curve centred on $\theta = 0$, reaching the DC peak value $G_{\text{max}} = 1/(1 - r_{\text{rt}}) \simeq 151.5$. As f_{mod} increases, the peak is progressively reduced and the curves broaden. This behaviour reflects the finite cavity storage time $\tau_{\text{cav}} = 1/\gamma \simeq 2.9$ ms: at higher drive frequencies the bender sweeps through the zero-crossing faster than τ_{cav}, and the intracavity field cannot fully build up before the coupling window closes.

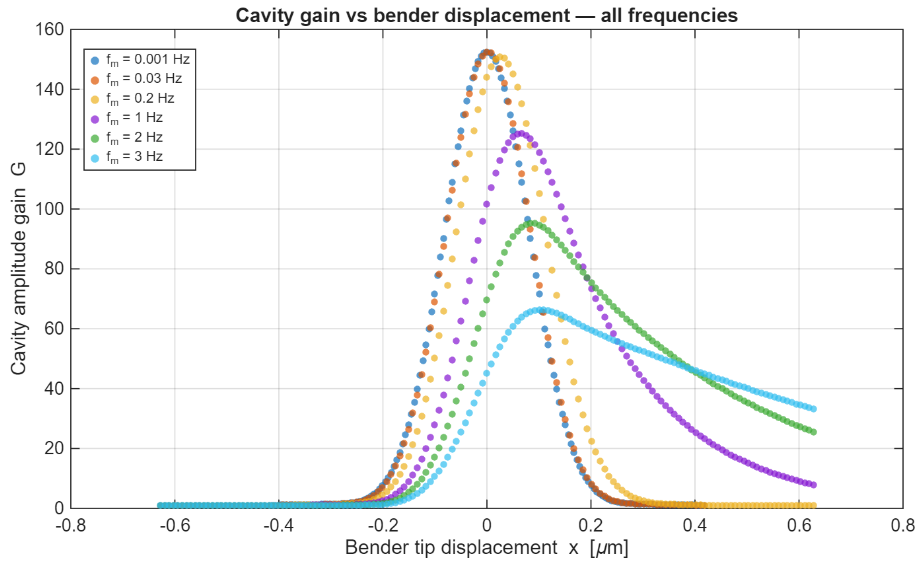

Figure 6.3: Cavity amplitude gain G as a function of bender tilt angle θ, with different drive frequencies overlaid. At low f_{mod} the gain is sharply peaked at $\theta = 0$. As f_{mod} increases, the peak value decreases and the curves broaden because the bender sweeps through the zero-crossing faster than the cavity storage time $\tau_{\text{cav}} \simeq 2.9$ ms.

Some insights of the cavity gain can be obtained by comparing the corrected cavity gain (equation 6.1) with the maximum cavity gain.

$$G_{\text{corr}} = \frac{G_{\text{max}}}{\sqrt{1 + (f_{\text{fit}}/f_{\text{arm}})^2}}. \tag{6.1}$$

Figure 6.4 compares G_{max} and G_{corr} as a function of f_{mod}. Below $f_{\text{mod}} \simeq 0.02$ Hz the two are indistinguishable: the signal frequency lies well below the pole and the correction factor is negligible. Above $f_{\text{mod}} \simeq 0.2$ Hz, where f_{fit} approaches and exceeds f_{arm}, the corrected gain drops rapidly. At $f_{\text{mod}} = 1$ Hz it has fallen to $G_{\text{corr}} \simeq 32$, and at $f_{\text{mod}} = 3$ Hz to $G_{\text{corr}} \simeq 11$, roughly a factor of fourteen below the DC value. This reduction in G_{corr} relative to G_{max} is the origin of the systematic bias observed when using the magnitude-only correction at high drive frequencies. However, as described in section 5.5.2, the bias is eliminated by deconvolving the cavity pole from the B1 time series using equation (5.78) before the sinusoidal fit, after which G_{max} is used directly. The simulation results confirming this are presented in section 6.6.

6.4 Higher-order mode content during bender oscillation

The tilt-induced mode coupling derived in section 5.4 predicts that the power scattered into the transverse mode TEM$_{m0}$ follows a Poisson distribution in the mode order m, with mean $\mu = (kw\theta)^2$ set by the instantaneous bender tilt angle θ (equation (5.38)). This section compares that prediction with the modal decomposition extracted directly from the OSCAR simulation.

At each time step, OSCAR decomposes the injected field into Hermite–Gaussian modes up to order $N_{\text{max}} = 20$ using the overlap integral of equation (5.36). The fractional power $|c_{m0}|^2$ in each mode is recorded as a function of $\theta(t)$, producing a two-dimensional map of

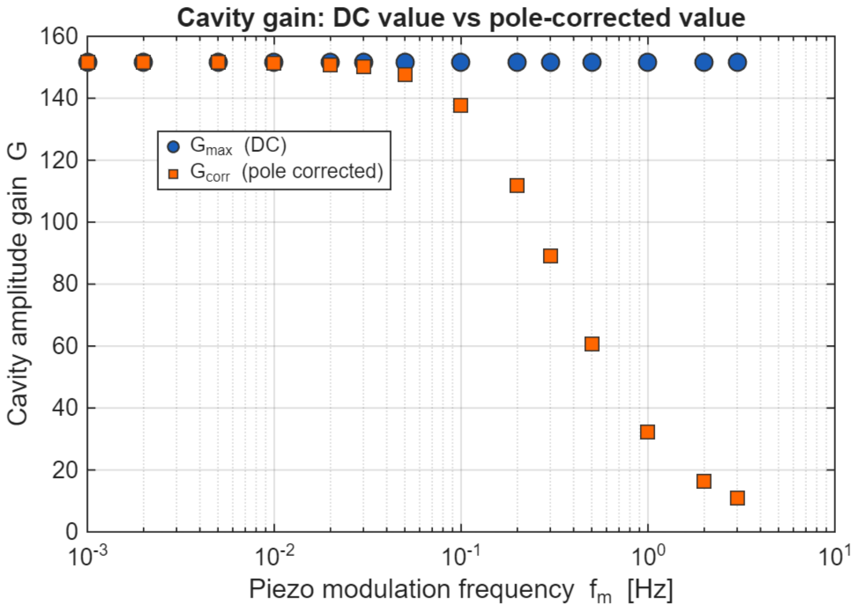

Figure 6.4: Cavity amplitude gain as a function of drive frequency. Blue circles: DC value G_{max} extracted from the peak of $G(\theta)$. Orange squares: pole-corrected gain G_{corr} from equation (6.1). The two coincide at low f_{mod} and diverge above the cavity-pole crossing at $f_{\text{mod}} \simeq 0.2$ Hz.

higher-order mode content versus tilt angle. Figure 6.5 shows this map for the lowest drive frequency $f_{\text{mod}} = 1$ mHz, where the bender moves quasi-statically and the cavity reaches steady state at every tilt angle.

Three features of figure 6.5 confirm the analytical predictions of section 5.4.

At $\theta = 0$, essentially all scattered power resides in TEM$_{00}$ (order 0), visible as the bright band along the bottom of the plot. This confirms that the mode-matching procedure of section 5.4.3 achieves $\eta_q \approx 1$: at the zero-crossing the returning beam has the correct beam parameter, and only the tilt-dependent factor $e^{-(kw\theta)^2}$ reduces the TEM$_{00}$ coupling away from perpendicular.

As $|\theta|$ increases, the peak of the distribution shifts to higher mode orders, following the Poisson mean $\mu = (kw\theta)^2$. The rate of this migration is set by the characteristic angular scale $\theta_c = \lambda/(2\pi w) \approx 3.1\,\mu$rad derived in equation (5.39). The distribution is visibly symmetric about $\theta = 0$ because the simulation samples the sinusoidal bender trajectory.

By $|\theta| \approx 8\,\mu$rad, the TEM$_{00}$ fraction is negligible ($f_m \ll 0.01$) and the dominant mode order exceeds $m = 10$. These higher-order modes are far off-resonance in the arm cavity (section 3.2): the frequency splitting $\Delta\nu_{10} \sim 10$ kHz far exceeds the cavity linewidth of ~ 100 Hz, so they do not build up resonantly and do not contribute to the calibration signal at B1. This confirms that the calibration is sensitive only to data taken within the narrow angular window $|\theta| < \theta_c$, as assumed in the fit procedure of section 5.4.1.

As a consistency check, the simulated TEM$_{00}$ fractional power as a function of θ, extracted from the bottom row of figure 6.5, agrees with the Gaussian envelope exp$[-(kw\theta)^2]$ of equation (5.47) across the full range $|\theta| < 10\,\mu$rad. This validates the use of that expression in the calibration without any correction for higher-order mode content.

6.5 Field decay rate from exponential ring-down

Section 5.5.1 derived the double-exponential model for the TEM$_{00}$ power following a transient scattered-light injection (equation (5.56)) where γ is the field decay rate of the arm cavity. This model is now applied to the simulated data to verify the cavity memory implementation and to demonstrate an independent method for measuring γ.

After the bender sweeps through the zero-crossing, the mode-matching efficiency drops to $f_m \approx 0$ and the intracavity perturbation decays freely. The last 25 samples of the P_{00} time series, well after the coupling window has closed, are fitted to equation (5.56) using a Levenberg–Marquardt least-squares solver[1]. The fit is conditioned by shifting the time origin to the start of the tail window and normalising P_{00} to zero mean and unit scale. The physical parameters are recovered by inverting the normalisation after convergence.

[1]The Levenberg–Marquardt algorithm is a standard method for nonlinear least-squares optimisation that interpolates between the Gauss–Newton method and gradient descent. It is particularly well suited to curve-fitting problems where the model depends nonlinearly on a small number of parameters [74, 75].

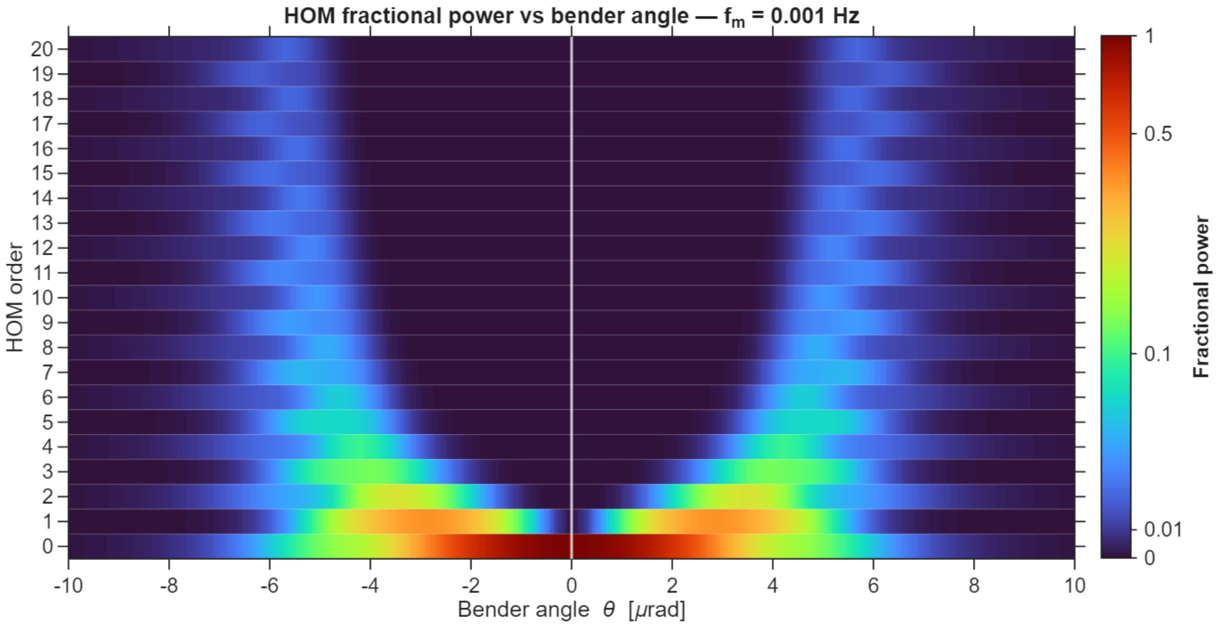

Figure 6.5: Fractional power in each Hermite–Gaussian mode TEM$_{m0}$ as a function of bender tilt angle θ, extracted from the OSCAR simulation at $f_{\text{mod}} = 1$ mHz. The colour scale uses an inverse hyperbolic sine transform to compress the dynamic range. At $\theta = 0$, nearly all power resides in TEM$_{00}$. As $|\theta|$ increases, power migrates to higher mode orders following the Poisson distribution of equation (5.38) in the low mode orders. The white vertical line marks $\theta = 0$.

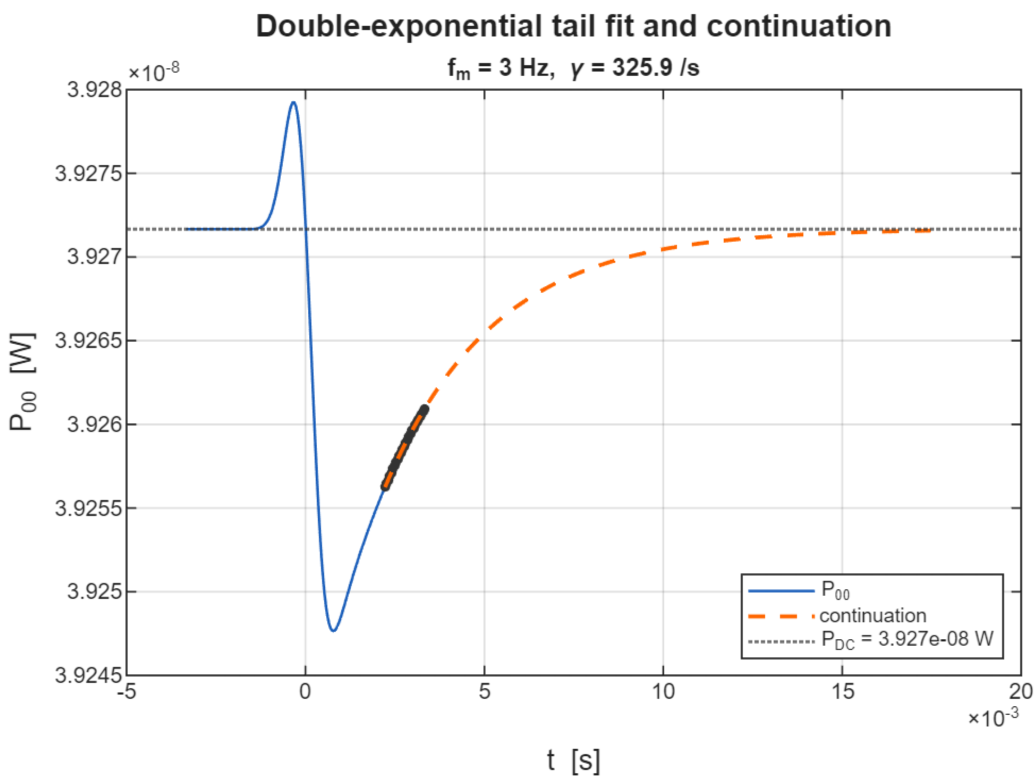

Figure 6.6: TEM$_{00}$ power $P_{00}(t)$ for the $f_{\text{mod}} = 3$ Hz dataset (blue), with the double-exponential ring-down fit (orange dashed) extrapolated beyond the fit window. The grey dotted line marks the unperturbed DC power $P_{\text{DC}} = 3.927 \times 10^{-8}$ W. The fitted decay rate is $\gamma_{\text{fit}} = 325.9\,\text{s}^{-1}$, corresponding to a cavity pole $f_{\text{arm,fit}} \approx 51.9$ Hz.

Figure 6.6 shows the result for the $f_{\text{mod}} = 3$ Hz dataset, which produces the strongest transient because the bender sweeps through the zero-crossing most rapidly. The blue curve is the simulated $P_{00}(t)$, the orange dashed curve is the fitted continuation, and the grey dotted line marks the unperturbed DC power P_{DC}.

The fitted curve converges toward P_{DC} from below without oscillation, consistent with the real-valued coefficients in equation (5.56). The fitted decay rate $\gamma_{\text{fit}} = 325.9\,\text{s}^{-1}$ agrees closely with the value $\gamma_{\text{known}} = 326.1\,\text{s}^{-1}$ computed from the round-trip reflectivity via equation (5.26), corresponding to a cavity pole $f_{\text{arm,fit}} = 51.9$ Hz versus the known $f_{\text{arm}} = 52.1$ Hz.

The small residual discrepancy has three origins. First, the 25-sample tail spans approximately 2 ms at this drive frequency, comparable to the storage time $\tau_{\text{cav}} \simeq 2.9$ ms, so the fit operates on fewer than one full decay time constant. Second, the finite simulation time step Δt introduces a discretisation bias in the extracted rate. Third, after passing the zero-crossing, the injected power does not vanish instantaneously: a small residual coupling persists while f_m decays from its peak value, contaminating the earliest samples in the tail. All three effects diminish at lower drive frequencies, where the tail extends over many storage times and f_m drops to zero more sharply.

This analysis serves two purposes. For the simulation, the agreement between γ_{fit} and γ_{known} validates the effective input-field formulation of equation (5.33): the cavity memory model correctly reproduces the physical decay dynamics. For the future experimental implementation, the ring-down fit provides a measurement of γ that is independent of the finesse measurement [21].

6.6 End-mirror transmission recovery

The preceding sections established all the individual ingredients of the calibration: the sinusoidal fit amplitudes A_{B1} and A_{B8}, the baseline powers, the mode-matching efficiency from B8, and the cavity gain. Combining these via equation (5.71) yields the recovered end-mirror transmission T_{WE}.

Two variants of the gain correction are compared:

1. **Magnitude-only correction.** $G = G_{\text{corr}}$ from equation (6.1), applied to the raw ΔP_{B1} signal.
2. **Full complex correction.** The cavity pole is removed from ΔP_{B1} by the time-domain deconvolution of equation (5.78) before the sinusoidal fit, after which $G = G_{\max}$ is used in equation (5.81).

The recovered $T_{\text{WE}}/T_{\text{WE,ref}}$ using the full complex correction is shown in Figure 6.7, while Figure 6.8 compares both methods across all datasets. The numerical values are listed in Table 6.1.

With the magnitude-only correction, the recovered transmission agrees with the reference to within 0.2 % over the range $f_{\text{mod}} = 1$ mHz to 50 mHz, but increases monotonically at higher drive frequencies, reaching $T_{\text{WE}}/T_{\text{WE,ref}} = 1.15$ at $f_{\text{mod}} = 0.2$ Hz and $T_{\text{WE}}/T_{\text{WE,ref}} = 2.09$ at $f_{\text{mod}} = 3$ Hz. The resulting constraint for the magnitude-only correction is:

$$x_{\text{tip}} \cdot f_{\text{mod}} \lesssim 0.5\ \mu\text{m Hz}, \qquad v_{\text{tip,max}} = 2\pi f_{\text{mod}} x_{\text{tip}} \lesssim \pi\ \mu\text{m s}^{-1}. \tag{6.2}$$

This systematic upward bias tracks the departure of f_{fit} from f_{arm} observed in figure 6.2 and is explained by the neglected phase lag $\phi_s = \arctan(f_{\text{fit}}/f_{\text{arm}})$ introduced by the cavity storage delay (appendix C, equation (C.4)): the magnitude-only correction compensates the amplitude roll-off but does not account for the phase rotation that distorts the signal shape at the zero-crossing and biases the fitted amplitude.

With the full complex correction, the bias is removed: the recovered $T_{\text{WE}}/T_{\text{WE,ref}}$ remains sub-percent accurate across the entire range of drive frequencies tested, from $f_{\text{mod}} = 1$ mHz up to 50 Hz. The deconvolution of equation (5.78) restores the correct signal shape at the zero-crossing by undoing the cavity-induced phase lag in the time domain, so the sinusoidal fit amplitude $A_{B1}^{(\text{deconv})}$ accurately represents the true scattered-light signal. The restriction of equation (6.2) that applied to the magnitude-only correction is therefore no longer necessary.

6.7 Influence of bender length on fitted frequency

As noted in section 6.1, the ratio $f_{\text{fit}}/f_{\text{mod}} \simeq 241$ exceeds the modulation depth $\beta \approx 118$ expected from the longitudinal displacement alone. This discrepancy originates from the finite length of the piezoelectric bender: because the bender is a cantilever, a tip displacement x_{tip} is necessarily accompanied by a tilt $\theta \approx x_{\text{tip}}/\ell$, which adds a transverse phase contribution to the scattered field that modifies the instantaneous frequency at the zero-crossing.

To isolate this effect, the simulation was repeated for several bender lengths while keeping the tip displacement fixed at $x_{\text{tip}} = 10\ \mu$m. The result is seen in figure 6.9. The fitted frequency converges to the known value of the modulation depth.

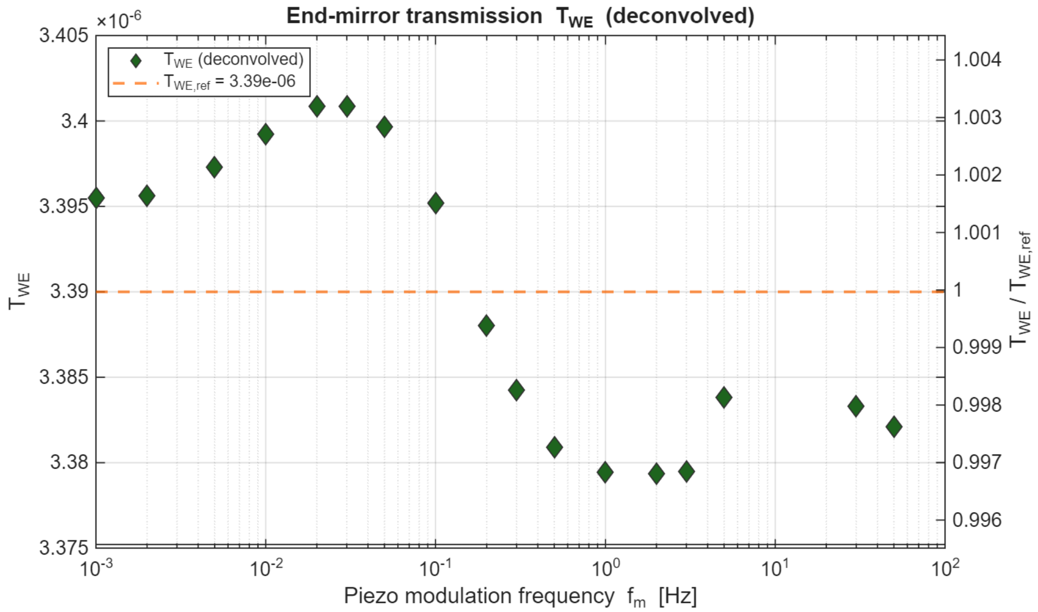

Figure 6.7: Recovered $T_{\text{WE}}/T_{\text{WE,ref}}$ as a function of drive frequency after applying the time-domain deconvolution of equation (5.78). The dashed line marks $T_{\text{WE}}/T_{\text{WE,ref}} = 1$. Sub-percent accuracy is maintained across the full range of drive frequencies tested.

6.8 Summary

The OSCAR simulation validates the scattered-light calibration procedure across the full range of drive conditions tested.

With the magnitude-only pole correction of equation (6.1), the end-mirror transmission is recovered with a relative accuracy of 0.2 % in the regime where $x_{\text{tip}} \cdot f_{\text{mod}} \lesssim 0.5$ μm Hz (equivalently, $f_{\text{fit}} \ll f_{\text{arm}}$). Above the cavity pole, the neglected phase lag introduces a systematic bias that grows to a factor of two at $f_{\text{mod}} = 3$ Hz.

Applying the time-domain deconvolution of equation (5.78) eliminates this bias: the full complex cavity transfer function is accounted for by removing the single-pole filter from the B1 time series before the sinusoidal fit, after which the DC gain G_{max} replaces G_{corr}. With this correction, sub-percent accuracy is maintained across all datasets, from $f_{\text{mod}} = 1$ mHz to 50 Hz. The restriction $x_{\text{tip}} \cdot f_{\text{mod}} \lesssim 0.5$ μm Hz is therefore not a fundamental limitation of the method but rather an artefact of the magnitude-only approximation.

Two additional results confirm ingredients of the calibration that are not directly accessible from the T_{WE} recovery alone. The double-exponential ring-down of P_{00} yields a cavity decay rate in close agreement with the value computed from the mirror reflectivities (section 6.5), validating the cavity memory model of the iterative OSCAR implementation. The modal decomposition of the intracavity field (section 6.4) confirms that TEM$_{00}$ coupling follows the predicted Gaussian envelope exp$[-(kw\theta)^2]$, justifying the restriction of the fit window to the zero-crossing region where mode-matching efficiency is close to unity.

Lastly, the influence of the finite bender length on the fitted signal frequency is characterised in section 6.7.

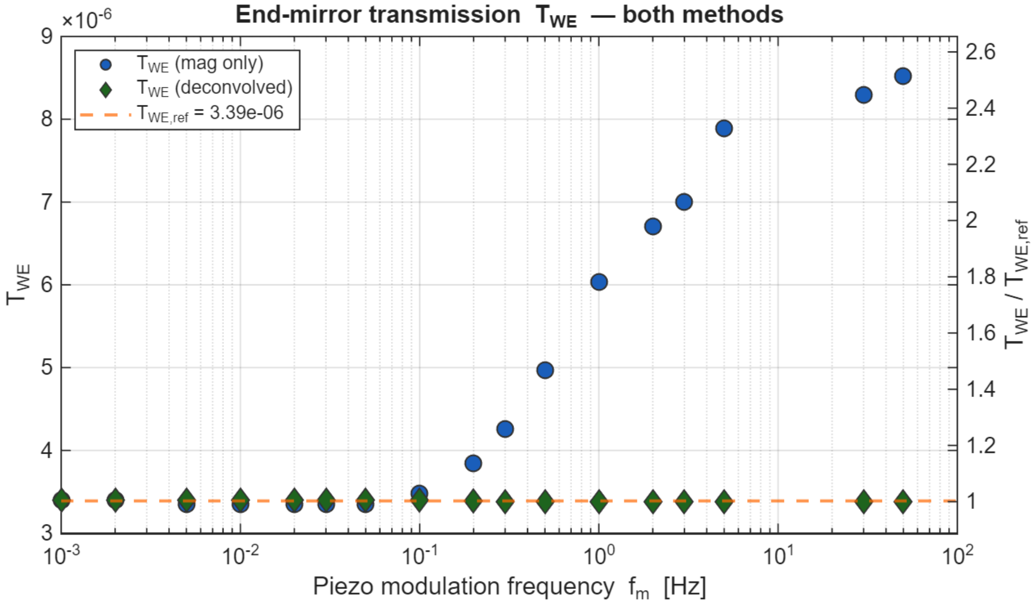

Figure 6.8: Recovered end-mirror transmission as a function of drive frequency. Blue circles: magnitude-only pole correction (equation (6.1)). Green diamonds: full complex correction via time-domain deconvolution (equation (5.78)). The dashed line marks $T_{\text{WE}}/T_{\text{WE,ref}} = 1$. With the deconvolution applied, the bias observed at high f_{mod} in the magnitude-only method is eliminated

f_{mod} [Hz]	f_{fit} [Hz]	$T_{\text{WE}}/T_{\text{WE,ref}}$ (mag.)	$T_{\text{WE}}/T_{\text{WE,ref}}$ (deconv.)
0.001	0.2415	1.0016	1.0016
0.002	0.4829	1.0017	1.0017
0.005	1.1923	0.9869	1.0022
0.01	2.3846	0.9874	1.0027
0.02	4.7692	0.9880	1.0032
0.03	7.1538	0.9883	1.0032
0.05	11.9231	0.9893	1.0029
0.1	23.8462	1.0251	1.0015
0.2	47.6923	1.1340	0.9994
0.3	71.5385	1.2554	0.9983
0.5	119.2308	1.4645	0.9973
1	238.4615	1.7801	0.9969
2	476.9231	1.9785	0.9969
3	715.3846	2.0656	0.9969
5	1207.3579	2.3261	0.9982
30	7244.1472	2.4473	0.9980
50	12073.5786	2.5129	0.9977

Table 6.1: Recovered end-mirror transmission for all datasets, comparing the magnitude-only pole correction with the full complex correction via deconvolution. The reference value is $T_{\text{WE,ref}} = 3.39 \times 10^{-6}$.

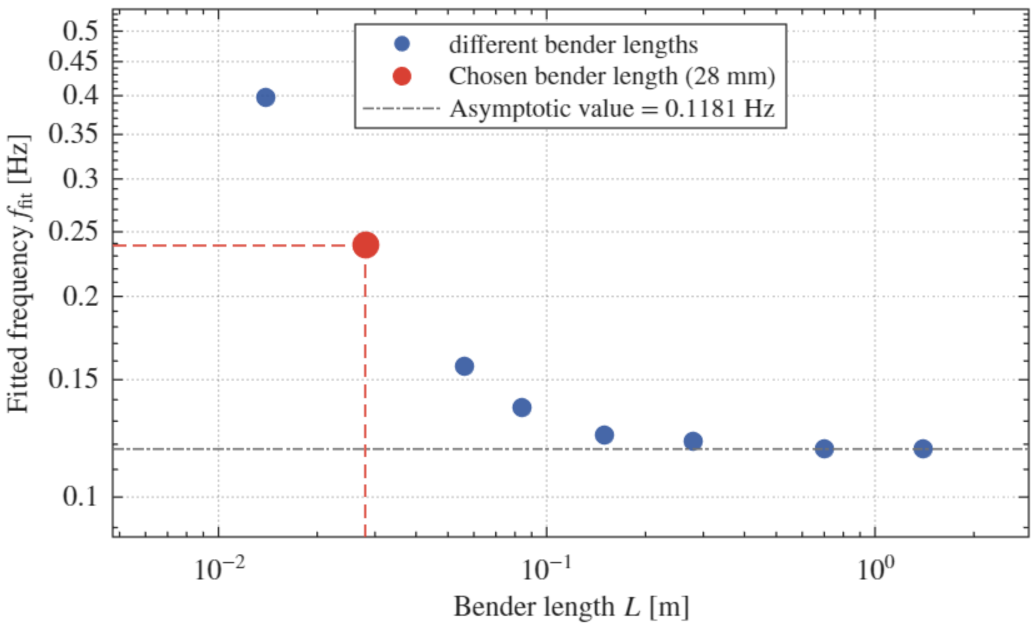

Figure 6.9: Signal frequency as a function of bender length with $f_{mod} = 0.001\ Hz$. The value f_{fit} asymptotically reaches $\beta \cdot f_{\text{mod}}$.

7

Experimental setup for scattered light generation

This chapter describes the experimental hardware that will be used for generating controlled scattered light on the Advanced Virgo end bench. The setup consists of three main components: a piezoelectric bender that provides the oscillating scattering surface (section 7.1), a motorised linear translation stage that positions the bender into or out of the transmitted beam path (section 7.2), and a laser source used for laboratory characterisation of the system before installation in the detector (section 7.3). The chapter also reports on the first laboratory tests, discusses a thermal issue encountered during initial operation, and outlines the remaining steps toward deployment on the Advanced Virgo suspended west end bench.

7.1 Piezoelectric bender

The calibration of the Virgo detector requires precise, controllable, and reproducible actuation to induce a well-characterised optical phase modulation. This modulation is recorded as a strain signal at the photodiode detector and used in the calibration process. For the implementation in Virgo, the Thorlabs PB4NB2S piezoelectric bimorph bender has been chosen. The bender is mounted on the end bench via a linear translation stage (described in section 7.2). This section presents the operating principle of the device, its electromechanical specifications, its electrical driving configuration, and its role in the end bench calibration scheme.

7.1.1 Working principle

A piezoelectric bender converts an applied electrical voltage into mechanical bending by exploiting the piezoelectric effect. In a bimorph configuration, two piezoelectric ceramic layers are bonded together so that when a voltage is applied, one layer expands while the other contracts. This differential strain causes the entire structure to deflect, with a deflection amplitude proportional to the applied voltage, making the device well suited for precision positioning applications.

During the calibration process, the bender is positioned such that the laser beam transmitted through the end mirror of the arm cavity is incident on the bender surface. The bender is driven to oscillate at a specific frequency, inducing a sinusoidal mechanical displacement that modulates the incoming light in phase and in propagation angle. The modulated light is subsequently re-injected into the arm cavity.

7.1.2 Specifications

The key electrical and mechanical specifications of the PB4NB2S piezoelectric bender are listed in table 7.1. Of particular importance for the calibration is the resonant frequency. To operate in a regime where the displacement is predictable and the device response is well approximated by a linear, frequency-independent transfer function, the drive frequency must remain well below resonance. Operating above resonance would produce an unpredictable displacement response unsuitable for the calibration.

7.1.3 Electrical connections and driving configuration

The piezoelectric bender is controlled via three colour-coded wires connected to the electrodes. The bender can be operated in two configurations, as described below. A diagram of the bender with its wiring is shown in figure 7.1.

Single-side voltage control

In this configuration, only one half of the bimorph is driven. Applying a positive voltage between the red and white wires (0 V to 150 V) causes the bender to deflect downward. The silver plus marker on the device defines the upward direction. It is important that the red

Parameter	Value
Drive voltage range	0 – 150 V
Tip displacement at 150 V (no load)	± 450 µm $\pm$ 15 %
Hysteresis	< 15 %
Free (active) length	28 mm
Blocking force at 150 V	1.5 N (0.33 lbs)
Resonant frequency (no load, fixed)	370 Hz
Dissipation factor	< 2.0 %
Capacitance (per bimorph side)	550 nF $\pm$ 15 %
Operating temperature range	-25 to 130 °C
Curie temperature	230 °C
Vacuum compatibility	10^{-7} Torr ($\approx 1.3{\times}10^{-7}$ mbar)
Piezo dimensions ($L \times W \times T$)	$32.0 \times 7.8 \times 0.8$ mm
Holder dimensions ($L \times W \times T$)	$11.0 \times 19.0 \times 3.0$ mm

Table 7.1: Key electromechanical specifications of the Thorlabs PB4NB2S piezoelectric bimorph bender (all values at 25 °C unless stated otherwise) [76].

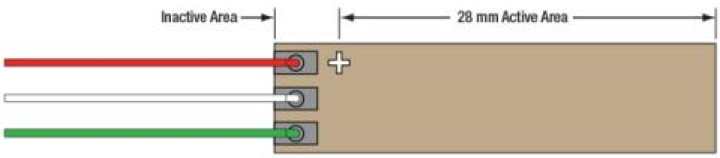

Figure 7.1: Diagram of the PB4NB2S piezoelectric bender showing the three colour-coded electrode wires [76].

and green wires are never simultaneously energised in this configuration, as the resulting differential voltage may exceed the 150 V maximum rating of the ceramic and cause fractures in the piezoelectric layer.

Differential voltage

The differential voltage is the preferred configuration for the calibration application, as it allows continuous and symmetric actuation by driving both ceramic layers simultaneously. Two options are available [76]:

- **Option 1:** the green and red wires are held at 150 V and 0 V (ground), respectively, while the white wire is varied from 0 V to 150 V.
- **Option 2:** the green and red wires are held at $+75$ V and -75 V, respectively, while the white wire is varied from -75 V to $+75$ V.

The importance of differential voltage control lies in the ability to sweep smoothly from full downward deflection through the neutral position to full upward deflection. For the calibration method, the neutral position is of particular importance because it maximises the coupling of the reflected light back into the cavity (section 5.4.2). A symmetric deflection profile provides the most information about this central point.

7.1.4 High-voltage controller

The piezoelectric bender is driven by the Thorlabs MDT694B, a single-channel precision high-voltage piezo controller. It provides an adjustable output voltage in the range 0 V to 150 V DC, with a maximum output current of 60 mA. The full voltage range is available only when the rear-panel voltage limit switch is set to the 150 V position [77].

The dynamic performance of the MDT694B–PB4NB2S system is governed by two limitations:

- **Electrical bandwidth.** The piezo capacitance (550 nF per side) forms a first-order low-pass filter with the controller output impedance ($150\,\Omega$). The resulting -3 dB bandwidth is

$$\text{BW} = \frac{1}{2\pi R_{\text{out}}(C_{\text{piezo}} + C_{\text{out}})} = \frac{1}{2\pi \times 150\,\Omega \times (550\,\text{nF} + 1\,\text{nF})} \approx 1.9\,\text{kHz}. \tag{7.1}$$

  This bandwidth is well above the intended drive frequencies. However, the 10 kΩ protection resistors introduced in section 7.1.5 reduce the effective system bandwidth to approximately 29 Hz, which becomes the practical limiting factor for drive waveform fidelity.

- **Current limit.** The output current is limited to 60 mA. For a sinusoidal drive of amplitude V_{amp} and frequency f, the peak current demand is

$$I_{\text{peak}} = 2\pi f\, V_{\text{amp}}(C_{\text{piezo}} + C_{\text{out}}). \tag{7.2}$$

For a full-range drive of 150 V at 100 Hz, the required peak current is approximately 52 mA. Under the realistic operating conditions for the calibration ($f \ll 100$ Hz, $V_{\text{amp}} \ll 150$ V), the current limit is therefore not a constraint.

7.1.5 Drive circuit and physical connections

The final drive circuit implementing the differential voltage control configuration (Option 1) is shown in figure 7.2. The red electrode is connected to the fixed 150 V bias, the green electrode to ground, and the white electrode receives the time-varying modulation signal from the MDT694B controller.

Following the discharge safety recommendation, two 10 kΩ protection resistors are inserted in series between the 150 V rail and ground, with the signal node tapped at the midpoint. This ensures that the charge stored on the 550 nF piezo capacitance always has a controlled dissipation path and that no direct short-circuit between any pair of electrodes is possible.

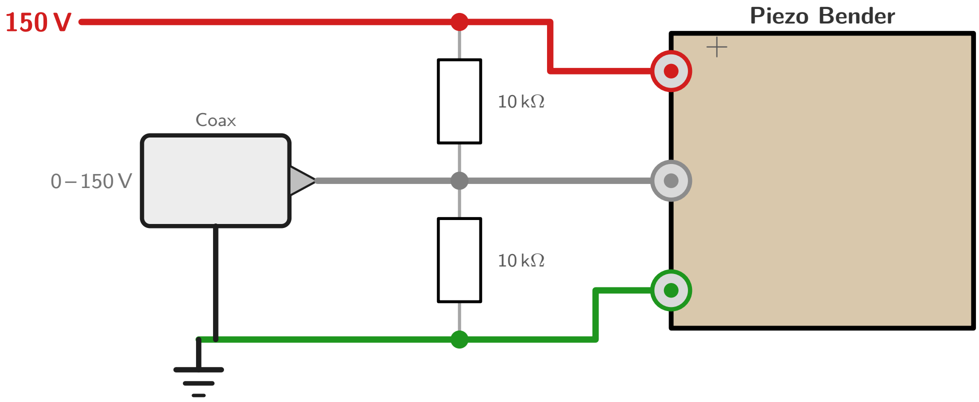

Figure 7.2: Wiring diagram of the PB4NB2S piezoelectric bender in differential voltage control mode (Option 1). The red electrode is held at a fixed 150 V and the white electrode receives the 0–150 V modulation signal from the coaxial cable. The green electrode is connected to ground. Two 10 kΩ protection resistors decouple the signal rail from the supply and ground rails, limiting current in the event of a fault and providing a safe discharge path for the piezo capacitance after driving.

A side effect of the protection resistors is that they form a first-order RC low-pass filter with the piezo capacitance. The equivalent circuit is shown in figure 7.3: the 10 kΩ protection resistor appears in parallel with the piezo capacitance $C_{\text{piezo}} = 550$ nF, which acts as the load driven by the signal source V_{in}.

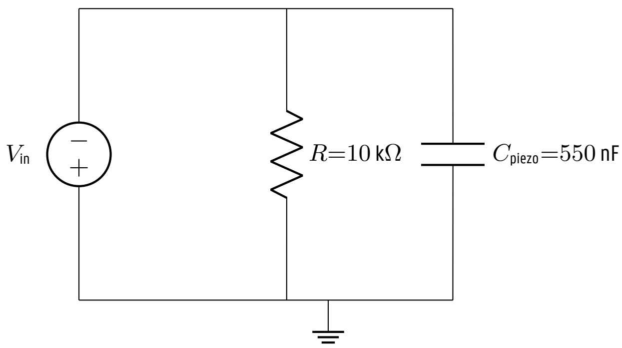

Figure 7.3: Equivalent circuit of the piezoelectric bender drive. The 10 kΩ protection resistor R appears in parallel with the piezo capacitance $C_{\text{piezo}} = 550$ nF.

The impedance of this parallel combination is

$$Z(\omega) = \frac{R \cdot \frac{1}{j\omega C_{\text{piezo}}}}{R + \frac{1}{j\omega C_{\text{piezo}}}} = \frac{R}{1 + j\omega R C_{\text{piezo}}}. \tag{7.3}$$

This is a first-order low-pass response with -3 dB cutoff frequency

$$f_{-3\,\mathrm{dB}} = \frac{1}{2\pi R C_{\mathrm{piezo}}} = \frac{1}{2\pi \times 10\,\mathrm{k}\Omega \times 550\,\mathrm{nF}} \approx 29\,\mathrm{Hz}. \tag{7.4}$$

At frequencies well below $f_{-3\,\mathrm{dB}}$, the capacitor impedance is large compared to R, the parallel combination is dominated by the resistor, and the full drive voltage appears across the piezo. Above $f_{-3\,\mathrm{dB}}$, the capacitor increasingly short-circuits the resistor, the impedance falls as $1/f$, and the voltage across the piezo is attenuated accordingly. A growing phase lag $\phi = -\arctan(f/f_{-3\,\mathrm{dB}})$ also develops, approaching $-90°$ at high frequencies. This cutoff is significantly lower than the 1.9 kHz intrinsic bandwidth of the MDT694B and sets the effective bandwidth of the complete drive system. The drive frequency for the calibration must therefore remain well below 29 Hz to ensure that the bender displacement faithfully follows the drive waveform without significant attenuation or phase distortion.

The modulation signal is delivered from the MDT694B to the circuit via a coaxial cable terminating in an SMA connector. The SMA interface provides a robust, low-noise connection to the BNC output of the MDT694B via a standard SMA-to-BNC adapter [77].

7.2 Linear translation stage

To position the piezoelectric bender within or outside the path of the transmitted beam, a motorised linear translation stage is required. The stage must provide micron-level positioning precision, long-term mechanical stability, and remote computer controllability. The Newport Agilis AG-LS25-27 fulfils these requirements: it is a compact piezo-motor-driven linear stage with a travel range of 27 mm, more than sufficient to move the bender fully into and out of the beam path. The stage is controlled via the AG-UC2 two-axis controller over a USB interface, allowing automated positioning from the measurement PC. A schematic of the stage function is shown in figure 7.4, a photograph in figure 7.5, and a technical drawing in figure 7.6. In figure 7.7, a realistic drawing of the piezobender connected to the stage is shown. This section presents the working principle of the stage, its specifications, and a calibration method for converting commanded steps to physical displacements.

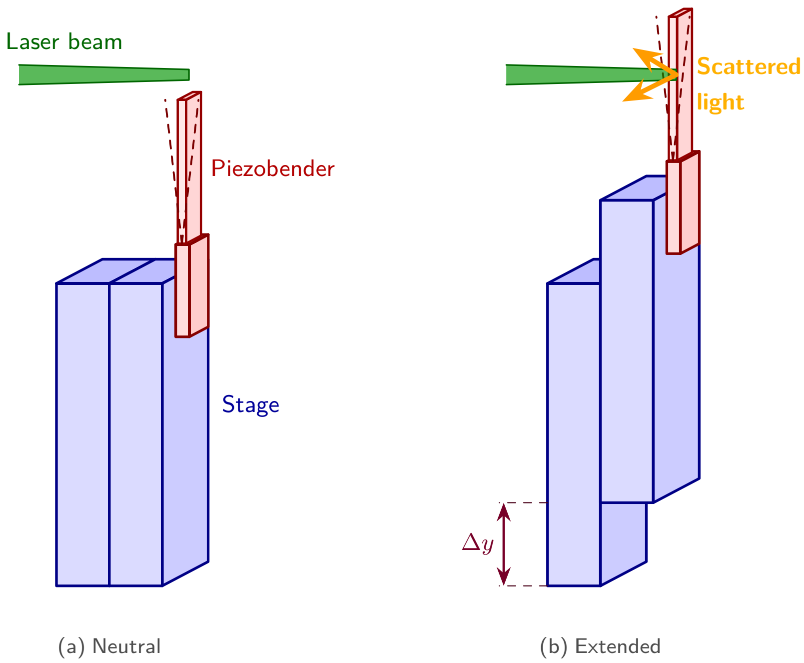

Figure 7.4: 3D schematic of the linear translation stage and piezoelectric bender. (a) Neutral: the movable stage is at rest and the bender tip lies below the laser beam. (b) Extended: the stage is raised by Δy, bringing the bender tip into the beam path where it scatters the incident light back toward the arm cavity.

7.2.1 Working principle: non-resonant stick-slip piezo motor

The AG-LS25-27 is driven by a non-resonant piezo motor operating on the stick-slip principle [79], distinct from both classical screw-and-gear motors [80] and ultrasonic piezoelectric motors.

A piezoelectric actuator element is placed in direct mechanical contact with the moving carriage. Motion is produced through a two-phase voltage cycle:

Figure 7.5: Photograph of the Newport Agilis AG-LS25-27 linear translation stage [78].

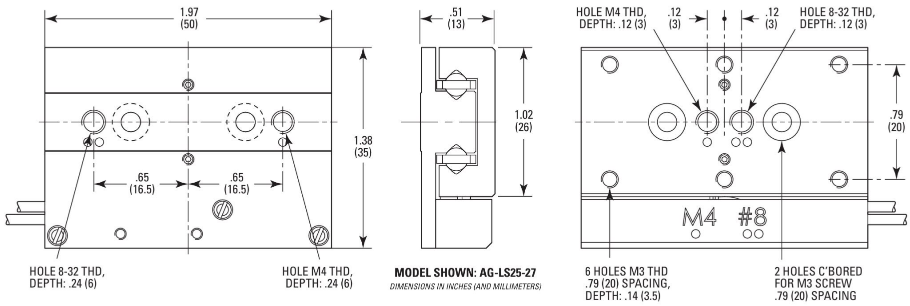

Figure 7.6: Technical drawing of the Newport Agilis AG-LS25-27 showing the mounting hole pattern and overall dimensions [78].

1. *Stick phase:* the piezo element expands slowly. Static friction causes the carriage to stick to the piezo tip and be dragged along with it.
2. *Slip phase:* the piezo element retracts rapidly. Due to the inertia of the carriage, it cannot follow the fast retraction and remains stationary while the piezo slips back underneath it.

Each cycle produces a small net displacement of the carriage in one direction. Repeating this cycle at high frequency produces continuous motion. Reversing the waveform reverses the direction of travel. Examples of drive waveforms used in stick-slip actuators are shown in figure 7.8.

This design offers an important advantage for implementation in Virgo: the absence of intermediate screws or gears eliminates backlash and hysteresis [81, 82].

The relevant technical specifications of the AG-LS25-27, as provided by Newport, are summarised in table 7.2.

Parameter	**Value**
Travel range	27 mm
Minimum incremental motion (MIM)	100 nm
Angular deviation (pitch/yaw)	$< 200\ \mu$rad
Drive mechanism	Non-resonant stick-slip piezo motor
Bearings	Pre-stressed linear ball bearings
Material	Thermally matched stainless steel
Load capacity (normal)	3.5 N
Controller	AG-UC2 (USB)
Interface	USB (921 600 baud)

Table 7.2: Specifications of the Newport Agilis AG-LS25-27 linear translation stage [78].

7.2.2 Controller and communication

The AG-UC2 Agilis controller provides USB computer control for two Agilis axes and features two sets of push buttons per axis: one for step size settings and one for precise low-speed adjustments.

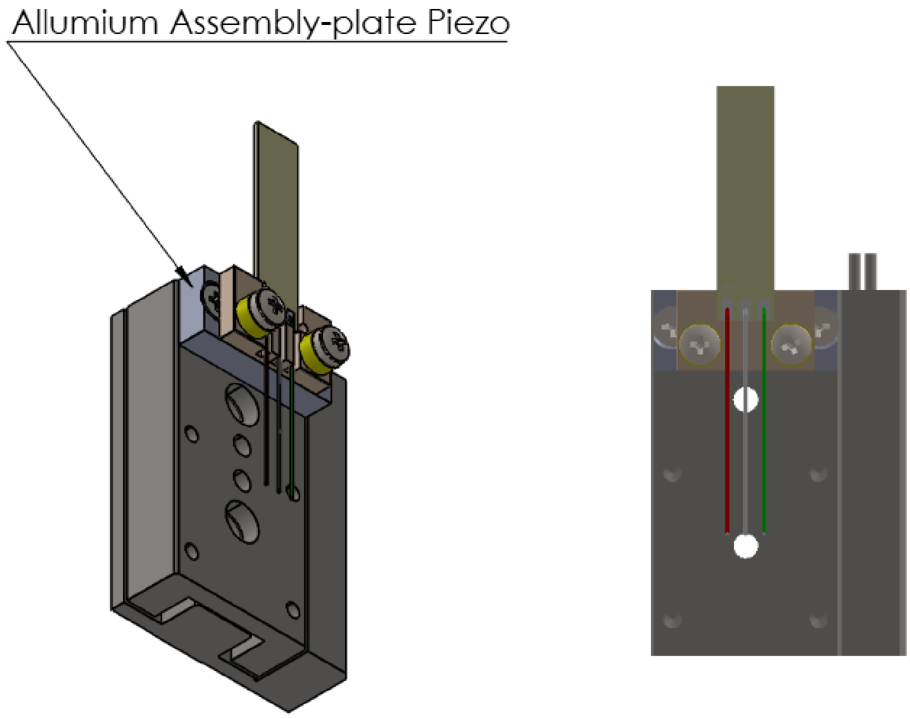

Figure 7.7: 3D figure of the piezobender attached to the stage stage (credits: Yves Israel)

The AG-UC2 accepts ASCII commands over the USB serial interface. Newport also supplies a Windows DLL and LabVIEW VI library, enabling fully remote, automated control of the stage position from a measurement script, essential for the calibration measurements.

In remote mode, all physical push buttons are disabled and the stage is controlled exclusively through software commands. The most important commands are:

- `PR n`: relative move of n steps (positive or negative).
- `PA n`: move to absolute position n (expressed as a percentage of total travel, after homing).
- `MV` ± 4: continuous motion at the highest speed in the positive or negative direction (used for homing).
- `SU n`: set the step amplitude (step size parameter, 1–50).

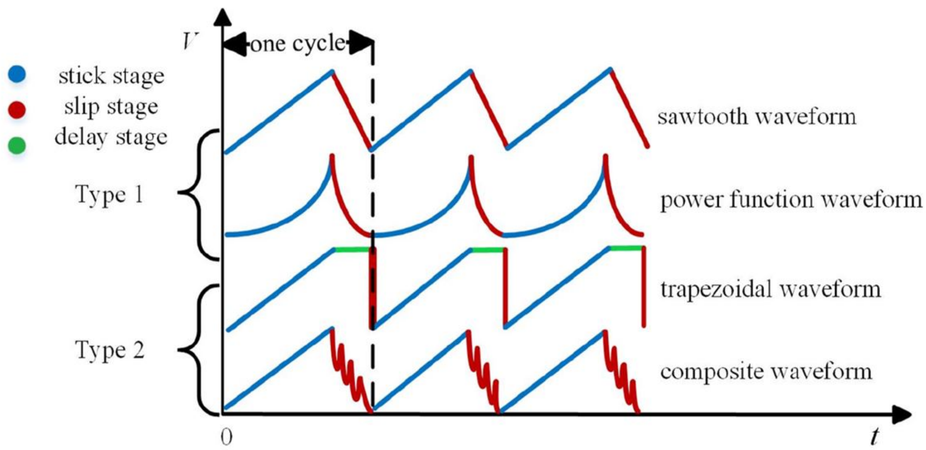

Figure 7.8: Examples of drive waveforms that can be applied to the piezo element of a stick-slip actuator [83].

7.2.3 Step size

The controller does not operate in physical length units. Instead, it commands a number of steps, where each step corresponds to a physical displacement that depends on the *step amplitude* — a dimensionless integer from 1 to 50 that controls the voltage amplitude applied to the piezo element in each stick-slip cycle. A higher step amplitude delivers more energy per cycle, resulting in a larger physical displacement per step.

The step sizes for forward and backward motion are independently adjustable and may differ from each other, reflecting the inherently asymmetric nature of the stick-slip waveform: the slow and fast phases may differ in efficiency depending on direction. The actual physical step size also depends on the load on the stage (the mounted bender will reduce the step size relative to the unloaded value), the temperature (which affects lubricant viscosity and piezo response), and the orientation of the stage (a vertically mounted stage is more strongly influenced by gravity than a horizontally mounted one).

The actual average step size under operating conditions must therefore be determined through a calibration procedure.

7.2.4 Stage calibration

Since the stage operates without a position encoder, an external calibration procedure is necessary to establish a conversion factor between commanded steps and physical displacement.

The AG-LS25-27 features electronic limit switches at both ends of travel. By driving the stage from one limit to the other and counting the total number of steps N_{steps} traversed across the full 27 mm travel distance d_{travel}, the average step size $\bar{s}$ is obtained as

$$\bar{s} = \frac{d_{\text{travel}}}{N_{\text{steps}}}. \tag{7.5}$$

Because the average step size is constant under fixed conditions (load, temperature, orientation), this calibration allows subsequent positioning to an absolute location without encoder feedback, using the PA command.

For the purposes of this experiment, nanometre-level absolute positioning precision is not required: the stage need only move the bender fully into and out of the beam path. This requirement is easily satisfied within the 27 mm travel range.

Minimum displacement of the stage

The translation stage must be able to move the piezoelectric bender fully out of the path of the transmitted beam when the calibration system is not in use. On the Advanced Virgo end bench, the beam transmitted through the end mirror passes through a lens telescope that reduces the beam size before it reaches the beam dump area. At the intended bender position (in front of the beam dump, approximately 10 cm before the dump surface), the beam radius is approximately $w \approx 1$ mm [10].

To ensure that no part of the beam is clipped by the bender or its mounting hardware when retracted, the stage must displace the bender by at least five times the beam diameter:

$$\Delta y_{\text{min}} = 5 \times 2w \approx 5 \times 2\,\text{mm} = 10\,\text{mm}. \tag{7.6}$$

This margin ensures that residual diffraction and alignment tolerances do not cause stray light to scatter off the bender in its parking position, which would introduce uncontrolled scattered-light noise into the detector output during observing runs. The 27 mm travel range of the AG-LS25-27 comfortably exceeds this requirement, leaving margin for alignment adjustments. The minimum incremental motion of 100 nm is far smaller than needed for positioning accuracy, since the bender simply needs to be moved fully into or fully out of the beam path.

In addition, an absorbing element should be placed around the bender in its parking position to capture any residual scattered light and prevent it from re-entering the arm cavity [10].

7.3 Laser source

The scattered light calibration measurements require a stable, low-noise laser source at the wavelength of the Advanced Virgo main beam, $\lambda = 1064$ nm. For the laboratory implementation of the end bench scattered light injection scheme, the Coherent Mephisto has been chosen. This section describes its operating principle, the specifications relevant to the calibration, and its noise properties.

The Mephisto is a continuous-wave diode-pumped solid-state (DPSS) laser based on a monolithic Nd:YAG crystal in a non-planar ring oscillator (NPRO) configuration [84, 85]. In the NPRO geometry, the laser beam follows a non-planar closed path inside a single monolithic crystal. The combination of the non-planar geometry and the Faraday effect due to an applied magnetic field enforces unidirectional lasing, eliminating spatial hole burning and ensuring single-longitudinal-mode operation at all times [84]. Because the entire resonator is contained within a single crystal, the optical path length is mechanically extremely stable, resulting in an ultra-narrow linewidth.

The output frequency can be tuned in two ways. Slow, large-range tuning is achieved thermally: changing the crystal temperature shifts the refractive index and therefore the optical path length, with a tuning coefficient of approximately -3 GHz/K and a total thermal tuning range of 30 GHz [85]. Fast, small-range tuning is achieved via a piezoelectric transducer (PZT), which induces a stress-dependent refractive index change. The PZT provides a tuning coefficient of approximately 1 MHz/V over a range of ± 65 MHz with a response bandwidth of 100 kHz [85].

Intensity noise is actively suppressed by the Coherent Noise Eater, an intracavity electro-optic modulator driven by a feedback loop that measures and cancels power fluctuations. A photograph of the laser source is shown in figure 7.9.

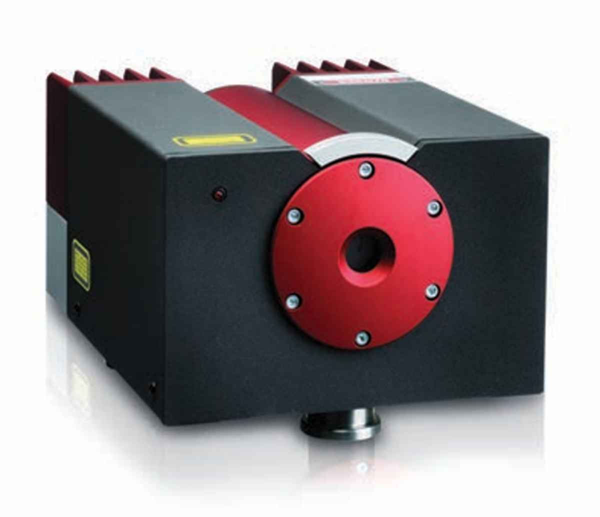

Figure 7.9: Photograph of the Coherent Mephisto NPRO laser source [85].

The key specifications of the Mephisto relevant to the calibration measurements are summarised in table 7.3.

Parameter	**Value**
Wavelength	1064 nm
Output power	up to 2 W
Operational mode	Continuous-wave
Spatial mode	TEM_{00} ($M^2 < 1.1$)
Spectral linewidth	≤ 3 kHz
Coherence length	> 1 km
Relative intensity noise (RIN)	< -140 dB/Hz ($f > 10$ kHz)
Intensity noise (10 Hz–2 MHz)	< 0.03 % rms
Frequency stability	~ 1 MHz/min
Thermal tuning coefficient	-3 GHz/K
Thermal tuning range	30 GHz
PZT tuning coefficient	~ 1 MHz/V
PZT tuning range	± 65 MHz
PZT response bandwidth	100 kHz

Table 7.3: Key specifications of the Coherent Mephisto laser relevant to the scattered light calibration measurements [85].

7.4 Thermal absorption during initial laboratory tests

During the first laboratory tests of the piezoelectric bender, an unexpected thermal problem was encountered. After several seconds of continuous exposure, the bender surface became noticeably hot. A slight smell was detected, likely originating from the adhesive layer bonding the two piezoelectric ceramic layers of the bimorph structure, which had been heated beyond its safe operating temperature. Prolonged exposure at these power densities also risks thermal damage to the piezoelectric ceramic itself, which could degrade the actuator performance. It should be noted that the beam was not optimally aligned during this initial test and may have partially illuminated the adhesive region directly.

The heating arises from optical absorption at the bender surface. The active area of the PB4NB2S is a bare piezoelectric ceramic (PZT), which is neither polished nor optically coated. At $\lambda = 1064$ nm, such ceramic surfaces have a low reflectivity and absorb the majority of the incident optical power. For an incident beam power P_{in} and a surface reflectivity R_{s}, the absorbed power is

$$P_{\text{abs}} = (1 - R_{\text{s}})\, P_{\text{in}}. \tag{7.7}$$

With R_{s} of only a few percent and P_{in} up to 2 W, the absorbed power can approach 2 W. This is a substantial thermal load for a thin ceramic structure with dimensions $32 \times 7.8 \times 0.8\,\text{mm}^3$ and poor thermal conductivity.

This thermal problem is not limited to the laboratory setup. On the Advanced Virgo end bench, the beam transmitted through the end mirror passes through a lens telescope that reduces the beam radius to approximately $w \approx 1$ mm at the intended bender position [10]. The transmitted power reaching the beam dump area is $P_{\text{trans}} = T_{\text{WE}} \times P_{\text{circ}}$, which is expected to reach approximately 1.6 W during O5 [10]. The entire transmitted power is concentrated into a spot of order 1 mm radius on the bender surface, resulting in a power density of order $50\,\text{W}/\text{cm}^2$. Under vacuum, where convective cooling is absent and only radiative and conductive cooling paths are available, this thermal load poses a serious risk of damage to the bender and of outgassing from the heated surface.

The heating poses three risks. First, the adhesive bonding the ceramic layers can degrade or fail if the temperature exceeds its rated limit, potentially destroying the bender [76]. Second, thermally induced changes in the piezoelectric coefficient d_{31} alter the relationship between drive voltage and tip displacement, compromising the reproducibility of the calibration signal [86]. Third, heating under vacuum can cause outgassing of volatile compounds from the adhesive or the ceramic itself, potentially contaminating the optics in the end station [87].

7.5 Outlook

7.5.1 Solution to the heating problem

The thermal problem identified in section 7.4 requires either reducing the optical power absorbed by the bender surface or improving the heat dissipation. Since the full transmitted beam power (up to $\sim$2 W) is focused to a spot of $\sim$1 mm radius on the bender surface both in the laboratory and on the Advanced Virgo end bench, this issue must be addressed for the in-situ application as well. Several approaches are considered.

Reflective coating on the bender surface. The most direct solution is to increase the surface reflectivity R_{s} by applying a thin metallic coating to the active face of the bender. A thin layer of aluminium (of order 100 nm, deposited by evaporation [88]) can achieve reflectivities above 90 % at 1064 nm, reducing the absorbed power by roughly an order of magnitude [89, 88, 90]. The coating must be thin enough that it does not significantly alter the mechanical stiffness or the resonant frequency of the bender, and it must adhere reliably to the ceramic substrate under the cyclic bending motion. A reflective coating has the additional advantage of increasing the parameter f_r (the fraction of light reflected back toward the cavity), which directly increases the scattered-light signal amplitude and improves the signal-to-noise ratio of the calibration measurement. The reflectivity at 1064 nm must be measured after deposition to provide the value of f_r needed for the calibration procedure [91].

Spectralon diffuse reflector. An alternative approach, suggested in the Advanced Virgo Plus technical design report [10], is to attach a Spectralon element to the bender surface. Spectralon is a near-Lambertian diffuse reflector with reflectivity exceeding 99 % at 1064 nm, which would virtually eliminate optical absorption. However, the added mass and stiffness of the Spectralon element would modify the mechanical dynamics of the bender (resonant frequency, tip displacement per volt), and these effects would need to be characterised before deployment.

Thermal management. As a complementary measure, the thermal path from the bender to the translation stage body can be improved by using a thermally conductive mounting bracket (for example, copper or aluminium) between the bender holder and the stage carriage. In vacuum, where convective cooling is absent, this conductive path is the primary available cooling mechanism and is therefore critical. The combination of a reflective coating (to reduce the absorbed power) and a conductive mount (to remove the residual heat) is the baseline strategy.

Verification. The thermal behaviour of the complete assembly under representative conditions (up to 2 W incident power, beam radius $\sim$700 μm, vacuum pressure $\lesssim 10^{-7}$ Torr ($\approx 1.3\times10^{-7}$ mbar)) must be verified before installation on the Advanced Virgo end bench. A dedicated test setup consisting of a vacuum chamber with a viewport, mounted on an optical bench, is available at the BelGrav laboratory for this purpose. The tests will verify that no surface damage occurs, that the outgassing rate remains within acceptable limits during prolonged illumination, and that the bender displacement response is not degraded by the thermal load.

7.5.2 Future tests

Before the system can be deployed on the Advanced Virgo end bench, several characterisation measurements remain to be performed.

The first priority is to calibrate the bender displacement as a function of drive voltage and frequency. This requires an independent displacement sensor (such as a laser vibrometer, a capacitive probe, or a quadrant photodiode (QPD) to monitor the beam position, as is standard practice in gravitational wave detectors) to measure the tip displacement at the intended operating voltages and frequencies, and to verify that the displacement-to-voltage transfer function is linear and frequency-independent in the regime $f \ll f_{\text{res}} = 370$ Hz. The hysteresis of the bender, specified as less than 15 % by the manufacturer, must also be characterised, since the calibration procedure assumes that the tip displacement follows the drive voltage without significant phase distortion.

The second measurement is the optical reflectivity of the bender surface. The fraction of incident power reflected back toward the cavity enters the calibration through the parameter f_r (equation 4.2). While the simulation uses $f_r = 10^{-6}$ as a representative value, the actual reflectivity of the bender surface (which is a bare piezoelectric ceramic, not an optical mirror) must be measured at $\lambda = 1064$ nm. This measurement can be performed using the Mephisto laser source described in section 7.3 together with a calibrated power meter.

Finally, the combined system (bender mounted on the translation stage) should be tested for mechanical stability and reproducibility over multiple insertion-retraction cycles, to verify that the stage positioning does not introduce systematic errors in the bender alignment relative to the beam axis.

7.5.3 Vacuum compatibility

All components of the calibration system must be compatible with the ultra-high vacuum environment of the Advanced Virgo end station, where the residual pressure is maintained at $\lesssim 10^{-9}$ Torr ($\approx 1.3 \times 10^{-9}$ mbar). Outgassing from materials introduced into the vacuum system can degrade the vacuum level, contaminate the optical surfaces of the test masses, and increase the residual gas noise contribution to the detector sensitivity [87].

Both the piezoelectric bender and the translation stage are vacuum-compatible by manufacturer specification (the bender to 10^{-7} mbar, table 7.1 and the stage body is constructed from thermally matched stainless steel). The primary concern for vacuum operation is the outgassing behaviour under prolonged laser illumination, as discussed in section 7.4, rather than the intrinsic vacuum compatibility of the components.

7.5.4 Installation in Advanced Virgo

The calibration system is designed for installation on the Suspended end benches of the Advanced Virgo detector. The installation will proceed in several stages. First, the bender and translation stage assembly will be mounted directly on the suspended end bench. The assembly must be positioned such that, when the stage is extended, the bender tip intercepts the transmitted beam at the intended position on the end bench. The electrical connections (the coaxial cable for the bender drive signal and the USB cable for the stage controller) will be routed through vacuum feedthroughs to the external control electronics.

After installation, a commissioning phase will verify the system performance in situ. This includes confirming that the bender oscillation produces the expected signals at B8, that the translation stage reliably positions the bender into and out of the beam path, and that no excess noise is introduced into the detector output during operation. The scattered-light calibration measurement can then be performed during dedicated commissioning time, following the procedure described in section 5.5. The measurement is expected to be compatible with the detector's observing schedule, since the bender can be retracted from the beam path when not in use, leaving no residual effect on the detector sensitivity.

8 Conclusion

The use of backscattered light from the end benches as a calibration tool for gravitational wave detectors was first proposed and demonstrated by Was et al. [5], who showed during the third observing run (O3) that the scattered-light signal can be used to reconstruct a calibrated strain-equivalent signal and to independently verify the absolute amplitude of the detector response. Building on this foundational work, this thesis has extended the method by developing a detailed simulation framework for a dedicated scattering actuator (a piezoelectric bender) and by characterising the experimental hardware for its deployment on the Advanced Virgo end bench.

The theoretical foundation of the method was established by deriving the scattered-light field model for the Virgo end-bench geometry. The signal amplitudes at the main output photodiode (B1) and the end-bench photodiode (B8) depend on the end-mirror transmission T_{WE} in different ways: the B1 signal depends on the product $G\,T_{\mathrm{WE}}\sqrt{f_{r,m}}$, where G is the frequency-dependent cavity amplitude gain and $\sqrt{f_{r,m}}$ is the fraction of scattered light that is reflected back toward the cavity and mode-matched to the cavity eigenmode, while the B8 signal depends on $\sqrt{f_{r,m}}$ alone. This asymmetry allows T_{WE} to be extracted from the ratio of the two signals without requiring knowledge of the absolute laser power or the intracavity field amplitude.

A time-domain simulation framework was constructed using the OSCAR FFT-based optical code to model the full scattering process. The standard frequency-domain description of the Fabry–Perot cavity response is impractical for this problem because the bender tip displacement of ~ 10 μm at $\lambda = 1064$ nm produces a modulation depth $\beta \approx 118$, distributing power across hundreds of sidebands. The cavity dynamics were therefore reformulated as a discrete-time update equation, and its exact equivalence with the frequency-domain transfer function was established analytically. The resulting effective input-field formulation allows the cavity memory to be incorporated into OSCAR's steady-state solver without modifying the code itself.

The tilt-induced coupling of the scattered field into higher-order Hermite–Gaussian modes was analysed in detail. Because the piezoelectric bender is a cantilever, a tip displacement is necessarily accompanied by a tilt that deflects the reflected beam away from the cavity axis. The power coupled into each transverse mode follows a Poisson distribution whose mean grows quadratically with the tilt angle, with a characteristic angular scale of $\theta_c \approx 3.1$ μrad for the Virgo arm cavity parameters. The calibration signal is therefore concentrated in narrow time windows around the bender zero-crossings, where the tilt vanishes and the mode-matching efficiency peaks at unity.

The simulation results validate the calibration procedure across the full range of drive frequencies. The principal findings are:

1. The end-mirror transmission is recovered with sub-percent accuracy across the full range of drive frequencies tested ($f_{\mathrm{mod}} = 1$ mHz to 50 Hz), using only measured power ratios. The method is self-calibrating: it requires no knowledge of the absolute laser power. At low drive frequencies, where the signal frequency lies well below the cavity pole, a magnitude-only correction of the cavity gain is sufficient. At higher drive frequencies, the cavity storage-time phase lag is removed by deconvolving the single-pole cavity transfer function from the B1 time series in the time domain (equation (5.78)) before the sinusoidal fit. This replaces the frequency-dependent G_{corr} with the DC gain G_{max} and eliminates the systematic bias that the magnitude-only approximation produces above the cavity pole. The restriction $x_{\mathrm{tip}} \cdot f_{\mathrm{mod}} \lesssim 0.5$ μm Hz is therefore not a fundamental limitation of the calibration method.

2. The cavity field decay rate measured from the double-exponential ring-down following a transient scattered-light injection yields $\gamma_{\mathrm{fit}} = 325.9$ s^{-1}, in close agreement with the known value $\gamma_{\mathrm{known}} = 326.1$ s^{-1} computed from the mirror reflectivities. This validates the cavity memory model and provides an independent method for measuring the cavity pole that does not rely on finesse measurements.

3. The Hermite–Gaussian modal decomposition of the intracavity field confirms that the TEM$_{00}$ coupling efficiency follows the predicted Gaussian envelope $\exp[-(kw\theta)^2]$, where $k = 2\pi/\lambda$ is the optical wavenumber, w the beam radius at the end mirror, and θ the instantaneous bender tilt angle, justifying the restriction of the calibration to the zero-crossing window.

4. The influence of the finite bender length on the fitted signal frequency was characterised, showing that the instantaneous frequency at the zero-crossing exceeds the value expected from the longitudinal modulation depth alone due to the transverse phase contribution from the tilt. This geometric effect does not affect the fitted amplitude that enters the calibration.

On the experimental side, the hardware for the calibration system has been partially assembled and characterised. The piezoelectric bender provides up to $\pm 450\,\mu$m tip displacement with a resonant frequency of 370 Hz and is vacuum-compatible to 10^{-7} Torr ($\approx 1.3 \times 10^{-7}$ mbar). The linear translation stage enables automated, remote-controlled positioning of the bender into or out of the transmitted beam path. Initial laboratory tests revealed a thermal issue caused by the absorption of the transmitted laser beam (up to $\sim$2 W focused to a spot of $\sim$1 mm radius) by the bare ceramic surface of the bender. Since the bender has low reflectivity at 1064 nm, the majority of the incident power is absorbed, leading to rapid heating that risks damaging both the piezoelectric ceramic and the adhesive layer, and causing outgassing under vacuum. This problem must be addressed before deployment, either by applying a reflective metallic coating to the bender surface to reduce the absorbed fraction or by improving the thermal conduction from the bender to the stage body.

When combined with an independent laboratory measurement of T_{WE}, the scattered-light method provides a test of the absolute detector calibration that is fully independent of both the photon calibrator and the Newtonian calibrator. The three methods together, each with fundamentally different dominant systematics, form a set of complementary cross-checks that strengthen confidence in the calibration of the global gravitational wave detector network.

Looking ahead, the next step is to complete the laboratory characterisation of the bender–stage assembly, after which experimental validation of the method on the Advanced Virgo detector can proceed. The remaining tasks include resolving the thermal absorption issue, completing the bender characterisation (displacement calibration, surface reflectivity measurement, and vacuum compatibility verification), installing the system on the suspended end benches, and performing the first in-situ scattered-light calibration measurement during a dedicated commissioning period. The deconvolution requires only the time derivative of the B1 signal, which is straightforward to compute from sampled data. For experimental data, where sensor noise is present, a low-pass differentiator (such as a Savitzky–Golay filter [73]) should be applied to suppress noise amplification at high frequencies. If the experimental results confirm the sub-percent accuracy demonstrated by the simulation, the scattered-light calibration could become a valuable addition to the calibration toolkit of current and future gravitational wave observatories.

Bibliography

[1] D. Estevez, P. Lagabbe, A. Masserot, L. Rolland, M. Seglar-Arroyo, and D. Verkindt, "The Advanced Virgo photon calibrators," *Classical and Quantum Gravity*, vol. 38, p. 075007, 2021.

[2] B. P. Abbott *et al.*, "Observation of gravitational waves from a binary black hole merger," *Phys. Rev. Lett.*, vol. 116, p. 061102, Feb. 2016.

[3] F. Acernese *et al.*, "Interferometer Sensing and Control for the Advanced Virgo Experiment in the O3 Scientific Run," *Galaxies*, vol. 8, no. 4, p. 85, 2020.

[4] B. P. Abbott *et al.*, "A gravitational-wave standard siren measurement of the Hubble constant," *Nature*, vol. 551, pp. 85–88, 2017.

[5] M. Was, R. Gouaty, and R. Bonnand, "End benches scattered light modelling and subtraction in advanced virgo," *Classical and Quantum Gravity*, vol. 38, p. 075020, mar 2021.

[6] A. Freise and K. Strain, "Interferometer techniques for gravitational-wave detection," *Living Reviews in Relativity*, vol. 13, Feb. 2010.

[7] D. Baumann, "General relativity." Lecture notes, Institute of Theoretical Physics, University of Amsterdam.

[8] The Virgo Collaboration, "Overview of Advanced Virgo Plus Phase II ," Official Document VIR-1179C-19, Virgo Collaboration, Nov. 2019.

[9] M. Was, "Laser power and SR transmission choice for O4," Advanced Virgo Activity Report VIR-0406A-19, Virgo Collaboration, Apr. 2019. Internal LKV document. https://tds.virgo-gw.eu/ql/?c=14234.

[10] Virgo Collaboration, "Advanced virgo plus for o5 – technical design report," Technical Report VIR-0499A-25, Virgo Collaboration, May 2025. Release A, accessed May 2026.

[11] A. Allocca, D. Bersanetti, J. Casanueva Diaz, C. De Rossi, M. Mantovani, A. Masserot, L. Rolland, P. Ruggi, B. Swinkels, E. N. Tapia San Martin, M. Vardaro, and M. Was, "Interferometer sensing and control for the advanced virgo experiment in the o3 scientific run," *Galaxies*, vol. 8, no. 4, 2020.

[12] C. Mueller, "The Advanced LIGO input mode cleaner," Presentation, Detector Characterization Meeting G1400096, LIGO Scientific Collaboration, Feb. 2014.

[13] Y. Akiyama *et al.*, "Vibration isolation system with a compact damping system for power recycling mirrors of KAGRA," *Class. Quantum Grav.*, 2019.

[14] F. Acernese *et al.*, "Advanced virgo: a second-generation interferometric gravitational wave detector," *Classical and Quantum Gravity*, vol. 32, p. 024001, dec 2014.

[15] F. Bondu, "Dark fringe shot noise sensitivity," Virgo Note VIR-NOT-OCA-1390-243, Virgo Collaboration, Apr. 2003. https://tds.virgo-gw.eu/ql/?c=1451.

[16] B. Swinkels, "Requirements on sensing noise for the Pound-Drever-Hall technique for the lock of the arm-cavities," Virgo Note VIR-0467A-12, Virgo Collaboration, Dec. 2012. Internal LKV document. https://tds.virgo-gw.eu/ql/?c=9290.

[17] ESPCI, "Measurements on the beamsplitter substrate," Virgo Note VIR-NOT-ESP-1390-77, Virgo Collaboration, Jan. 1997. https://tds.virgo-gw.eu/ql/?c=85.

[18] L. Pinard, "Advanced Virgo beam splitter characterization report — BS02 (coatings C14041+C14045)," Specification/Characterisation VIR-0446B-14, Virgo Collaboration, Nov. 2014. Internal Virgo document. https://tds.virgo-gw.eu/ql/?c=10497.

[19] T. T. Fricke, N. D. Smith-Lefebvre, R. Abbott, R. Adhikari, K. L. Dooley, M. Evans, P. Fritschel, V. V. Frolov, K. Kawabe, J. S. Kissel, B. J. J. Slagmolen, and S. J. Waldman, "DC readout experiment in Enhanced LIGO," *Class. Quantum Grav.*, vol. 29, no. 6, p. 065005, 2012.

[20] F. Acernese, M. Agathos, K. Agatsuma, D. Aisa, N. Allemandou, A. Allocca, J. Amarni, P. Astone, G. Balestri, G. Ballardin, F. Barone, J.-P. Baronick, M. Barsuglia, A. Basti, F. Basti, T. Bauer, V. Bavigadda, M. Bejger, M. Beker, and J.-P. Zendri, "Advanced virgo: A second-generation interferometric gravitational wave detector," *Classical and Quantum Gravity*, vol. 32, 01 2015.

[21] C. Jacquet, Optical Tuning Team, and Commissioning Team, "Arm cavity characterization," Virgo Technical Document VIR-0540A-24, Virgo Collaboration, June 2024. Presented at Virgo Week, 5th June 2024.

[22] L. Pinard, "Advanced Virgo input mirror characterization report — IM03 (coatings C118046+C18013) — spare IM," Specification/Characterisation VIR-0326A-23, Virgo Collaboration, Apr. 2023. Internal Virgo document. https://tds.virgo-gw.eu/ql/?c=19196.

[23] L. Pinard, "Advanced Virgo end mirror characterization report — EM03 (coatings C14042/20 + C14035/20)," Specification/Characterisation VIR-0270A-15, Virgo Collaboration, June 2015. Internal Virgo document. https://tds.virgo-gw.eu/ql/?c=10991.

[24] C. Bond, D. Brown, A. Freise, and K. A. Strain, "Interferometer techniques for gravitational-wave detection," *Living Rev. Relativ.*, vol. 19, p. 3, 2016.

[25] J. Mizuno, K. A. Strain, P. G. Nelson, J. M. Chen, R. Schilling, A. Rüdiger, W. Winkler, and K. Danzmann, "Resonant sideband extraction: a new configuration for interferometric gravitational wave detectors," *Phys. Lett. A*, vol. 175, no. 5, pp. 273–276, 1993.

[26] J. Mizuno, *Comparison of optical configurations for laser-interferometric gravitational-wave detectors.* PhD thesis, Universität Hannover, Hanover, Germany, 1995.

[27] E. Polini, "New High Finesse AdV+ Output Mode Cleaner (OMC)," Presentation, LVK Meeting VIR-0261A-21, Virgo Collaboration, Mar. 2021. Virgo Technical Document System.

[28] L. Di Gallo, "Mirror Surface Studies for Advanced Virgo OMC," Scientific & Technical Note VIR-0267A-12, Virgo Collaboration, July 2012. Virgo Technical Document System. Publicly available.

[29] M. Ducrot, R. Bonnand, R. Gouaty, F. Marion, A. Masserot, B. Mours, L. Rolland, and M. Was, "Measurements of Advanced Virgo OMC Cavities Finesse," Virgo Note VIR-0458A-14, Virgo Collaboration, Oct. 2014. Virgo Technical Document System. Publicly available.

[30] A. Di Virgilio, "Seismic and thermal noise upper limits at 10 Hz for the Virgo suspensions," Virgo Note VIR-NOT-PIS-1390-334, Virgo Collaboration, Dec. 2006. https://tds.virgo-gw.eu/ql/?c=1668.

[31] t. Accadia, F. Acernese, M. Alshourbagy, P. Amico, F. Antonucci, S. Aoudia, N. Arnaud, C. Arnault, K. Arun, P. Astone, S. Avino, D. Babusci, G. Ballardin, F. Barone, G. Barrand, L. Barsotti, M. Barsuglia, A. Basti, T. Bauer, and Z. Zhang, "Virgo: A laser interferometer to detect gravitational waves," *Journal of Instrumentation*, vol. 7, p. P03012, 03 2012.

[32] A. Basti, V. Boschi, P. Chessa, V. Dattilo, R. Passaquieti, and P. Ruggi, "The seismic isolation system of advanced virgo plus, phase ii," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 1048, p. 168021, 2023.

[33] P. Amico, L. Bosi, L. Carbone, L. Gammaitoni, F. Marchesoni, M. Punturo, F. Travasso, and H. Vocca, "The Advanced Virgo monolithic fused silica suspension," *Nucl. Instrum. Meth. A*, vol. 824, pp. 644–645, 2016.

[34] G. Losurdo, M. Bernardini, S. Braccini, C. Bradaschia, C. Casciano, V. Dattilo, R. De Salvo, A. Di Virgilio, F. Frasconi, A. Gaddi, A. Gennai, A. Giazotto, H. B. Pan, F. Paoletti, A. Pasqualetti, R. Passaquieti, D. Passuello, R. Taddei, and Z. Zhang, "An inverted pendulum preisolator stage for the Virgo suspension system," *Rev. Sci. Instrum.*, vol. 70, no. 5, pp. 2507–2515, 1999.

[35] G. Losurdo, G. Calamai, E. Cuoco, L. Fabbroni, G. Guidi, M. Mazzoni, R. Stanga, F. Vetrano, L. Holloway, D. Passuello, G. Ballardin, S. Braccini, C. Bradaschia, R. Cavalieri, G. Cella, V. Dattilo, A. Di Virgilio, F. Fidecaro, F. Frasconi, A. Gennai, A. Giazotto, P. La Penna, F. Paoletti, A. Pasqualetti, R. Passaquieti, R. Poggiani, R. Taddei, A. Viceré, and Z. Zhang, "Inertial control of the mirror suspensions of the Virgo interferometer for gravitational wave detection," *Rev. Sci. Instrum.*, vol. 72, no. 9, pp. 3653–3661, 2001.

[36] A. Allocca *et al.*, "Simulations for the Locking and Alignment Strategy of the DRMI Configuration of the Advanced Virgo Plus Detector," *Galaxies*, vol. 10, no. 6, p. 115, 2022.

[37] S. Hild, H. Grote, J. Degallaix, S. Chelkowski, K. Danzmann, A. Freise, M. Hewitson, J. Hough, H. Lück, M. Prijatelj, K. A. Strain, J. R. Smith, and B. Willke, "DC-readout of a signal-recycled gravitational wave detector," *Classical and Quantum Gravity*, vol. 26, no. 5, p. 055012, 2009.

[38] R. L. Ward, R. Adhikari, B. Abbott, R. Abbott, D. Barron, R. Bork, T. Fricke, V. Frolov, J. Heefner, A. Ivanov, O. Miyakawa, K. McKenzie, B. Slagmolen, M. Smith, R. Taylor, S. Vass, S. Waldman, and A. Weinstein, "DC readout experiment at the Caltech 40 m prototype interferometer," *Class. Quantum Grav.*, vol. 25, no. 11, p. 114030, 2008.

[39] K. Somiya, Y. Chen, S. Kawamura, and N. Mio, "Frequency noise and intensity noise of next-generation gravitational-wave detectors with RF/DC readout schemes," *Phys. Rev. D*, vol. 73, p. 122005, 2006.

[40] C. M. Caves, "Quantum-mechanical noise in an interferometer," *Phys. Rev. D*, vol. 23, pp. 1693–1708, Apr 1981.

[41] V. B. Braginsky and F. Y. Khalili, *Quantum Measurement.* Cambridge: Cambridge University Press, 1992.

[42] J. Abadie *et al.*, "A gravitational wave observatory operating beyond the quantum shot-noise limit," *Nature Physics*, vol. 7, p. 962–965, Sept. 2011.

[43] G. M. Harry, H. Armandula, E. Black, D. R. M. Crooks, G. Cagnoli, J. Hough, P. Murray, S. Reid, S. Rowan, P. Sneddon, M. M. Fejer, R. Route, and S. D. Penn, "Thermal noise in interferometric gravitational wave detectors due to dielectric optical coatings," *Class. Quantum Grav.*, vol. 19, pp. 897–917, 2002.

[44] M. Granata, A. Amato, G. Cagnoli, M. Coulon, J. Degallaix, D. Forest, L. Mereni, C. Michel, L. Pinard, B. Sassolas, and J. Teillon, "Progress in the measurement and reduction of thermal noise in optical coatings for gravitational-wave detectors," *Appl. Opt.*, vol. 59, pp. A229–A235, 2020.

[45] S. D. Penn, G. M. Harry, A. M. Gretarsson, S. E. Kittelberger, P. R. Saulson, J. J. Schiller, J. R. Smith, and S. O. Swords, "Mechanical loss in tantala/silica dielectric mirror coatings," *Classical and Quantum Gravity*, vol. 20, no. 13, pp. 2917–2928, 2003.

[46] S. Di Pace, "Status of advanced virgo and upgrades before next observing runs," *Physica Scripta*, vol. 96, 12 2021.

[47] S. Karki *et al.*, "The Advanced LIGO photon calibrators," *Review of Scientific Instruments*, vol. 87, p. 114503, 2016.

[48] F. Acernese *et al.*, "Calibration of Advanced Virgo and reconstruction of the detector strain $h(t)$ during the observing run O3," *Classical and Quantum Gravity*, vol. 39, p. 045006, 2022.

[49] D. Bhattacharjee *et al.*, "Calibrating the global network of gravitational wave observatories via laser power calibration at NIST and PTB," *Metrologia*, vol. 61, p. 054002, 2024.

[50] F. Aubin, E. Dangelser, D. Estevez, A. Masserot, B. Mours, T. Pradier, A. Syx, and P. V. Hove, "The Virgo Newtonian calibration system for the O4 observing run," *Classical and Quantum Gravity*, vol. 42, p. 035007, 2025.

[51] D. Estevez, B. Mours, and T. Pradier, "Newtonian calibrator tests during the Virgo O3 data taking," *Classical and Quantum Gravity*, vol. 38, p. 075012, 2021.

[52] H. Kogelnik and T. Li, "Laser beams and resonators," *Applied Optics*, vol. 5, no. 10, pp. 1550–1567, 1966.

[53] Wikipedia contributors, "Gaussian beam — Wikipedia, the free encyclopedia," 2026. [Online; accessed 11-May-2026].

[54] D. Z. Anderson, "Alignment of resonant optical cavities," *Applied Optics*, vol. 23, no. 17, pp. 2944–2949, 1984.

[55] B. E. A. Saleh and M. C. Teich, *Fundamentals of Photonics*. Wiley, 2nd ed., 2007.

[56] Ansys, "Fabry-perot cavity (interconnect)." https://optics.ansys.com/hc/en-us/articles/360042323754-Fabry-Perot-cavity-INTERCONNECT, 2025. Accessed: 2026-05-11.

[57] A. E. Siegman, *Lasers*. University Science Books, 1986.

[58] C. Jacquet, "Alternative method for the arms Gouy phase measurement," Virgo Technical Document VIR-1074A-24, Virgo Collaboration, Dec. 2024. Presented at the Optical Tuning Meeting, 10th December 2024.

[59] F. W. J. Olver, A. B. Olde Daalhuis, D. W. Lozier, B. I. Schneider, R. F. Boisvert, C. W. Clark, B. R. Miller, B. V. Saunders, H. S. Cohl, and M. A. McClain, "Nist digital library of mathematical functions," 2023. Release 1.1.10 of 2023-06-15.

[60] E. D. Black, "An introduction to Pound–Drever–Hall laser frequency stabilization," *American Journal of Physics*, vol. 69, no. 1, pp. 79–87, 2001.

[61] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward, "Laser phase and frequency stabilization using an optical resonator," *Applied Physics B*, vol. 31, no. 2, pp. 97–105, 1983.

[62] S. Bordo, "Analyse van verstrooid licht voor de kalibratie van de virgo interferometer," master's thesis, Universiteit Antwerpen, Antwerp, Belgium, jun 2026.

[63] J. Degallaix, "Oscar: A matlab based package to simulate realistic optical cavities," *SoftwareX*, vol. 12, p. 100587, 2020.

[64] J. Degallaix, "Oscar a matlab based optical fft code," *Journal of Physics: Conference Series*, vol. 228, p. 012021, may 2010.

[65] Stellafane Amateur Telescope Making, "Measuring sagitta (depth of curve)." https://stellafane.org/, n.d. Accessed: 2026-05-19. Instructions on measuring mirror sagitta using mechanical methods.

[66] M. J. Lawrence, B. Willke, M. E. Husman, E. K. Gustafson, and R. L. Byer, "Dynamic response of a fabry–perot interferometer," *J. Opt. Soc. Am. B*, vol. 16, pp. 523–532, Apr 1999.

[67] M. Rakhmanov, "Doppler-induced dynamics of fields in fabry–perot cavities with suspended mirrors," *Applied Optics*, vol. 40, p. 1942, Apr. 2001.

[68] M. E. Zucker, "Transmission of optical power transients by Fabry–Perot cavities: Laser power stabilization and ringdown measurement applications," ligo technical memorandum, LIGO Project, Oct. 1990. 31 October 1990.

[69] M. Rakhmanov, R. L. Savage Jr., D. H. Reitze, and D. B. Tanner, "Dynamic resonance of light in Fabry–Perot cavities," *Physics Letters A*, vol. 305, pp. 239–244, 2002.

[70] F. Bayer-Helms, "Coupling coefficients of an incident wave and the modes of a spherical optical resonator in the case of mismatching and misalignment," *Applied Optics*, vol. 23, no. 9, pp. 1369–1380, 1984.

[71] D. Estevez, P. Lagabbe, A. Masserot, L. Rolland, M. Seglar-Arroyo, and D. Verkindt, "The Advanced Virgo photon calibrators," Tech. Rep. VIR-0705A-20, Virgo Collaboration, 2020.

[72] A. V. Oppenheim, A. S. Willsky, and S. H. Nawab, *Signals and Systems*. Prentice Hall, 2nd ed., 1997.

[73] A. Savitzky and M. J. E. Golay, "Smoothing and differentiation of data by simplified least squares procedures," *Analytical Chemistry*, vol. 36, no. 8, pp. 1627–1639, 1964.

[74] K. Levenberg, "A method for the solution of certain non-linear problems in least squares," *Quarterly of Applied Mathematics*, vol. 2, no. 2, pp. 164–168, 1944.

[75] D. W. Marquardt, "An algorithm for least-squares estimation of nonlinear parameters," *Journal of the Society for Industrial and Applied Mathematics*, vol. 11, no. 2, pp. 431–441, 1963.

[76] Thorlabs, Inc., *PB4NB2S Piezoelectric Bimorph with Holder, 150 V, $\pm$450 µm Travel – Specification Sheet*. Thorlabs, Inc., Newton, NJ, USA, Feb. 2024. Document CTN010297-S01, Rev. E.

[77] Thorlabs, Inc., *MDT693B and MDT694B Open-Loop Piezo Controllers – User Guide*. Thorlabs, Inc., Newton, NJ, USA, Apr. 2023. Document TTN011225-D02, Rev. F.

[78] Newport Corporation, *AG-LS25-27 Agilis Piezo Linear Stage, User Manual*. Newport Corporation, 2023.

[79] M. Hunstig, "Piezoelectric inertia motors—a critical review of history, concepts, design, applications, and perspectives," *Actuators*, vol. 6, no. 1, p. 7, 2017.

[80] K. Spanner and B. Koc, "Piezoelectric motors, an overview," *Actuators*, vol. 5, no. 1, p. 6, 2016.

[81] M. Hunstig, T. Hemsel, and W. Sextro, "Stick–slip and slip–slip operation of piezoelectric inertia drives—Part I: Ideal excitation," *Sensors and Actuators A: Physical*, vol. 200, pp. 79–89, 2013.

[82] M. Hunstig, T. Hemsel, and W. Sextro, "Stick–slip and slip–slip operation of piezoelectric inertia drives—Part II: Frequency-limited excitation," *Sensors and Actuators A: Physical*, vol. 200, pp. 79–89, 2013.

[83] Z. et al., "A review of recent studies on piezoelectric stick-slip actuators," *Precision Engineering*, 2025.

[84] T. J. Kane and R. L. Byer, "Monolithic, unidirectional single-mode Nd:YAG ring laser," *Opt. Lett.*, vol. 10, no. 2, pp. 65–67, 1985.

[85] Coherent, Inc., *Mephisto/Mephisto S: Ultra-Narrow Linewidth CW DPSS Laser*, 2022. Document MC-013-13-0M0522 Rev. H.

[86] F. Li, Z. Xu, X. Wei, and X. Yao, "Determination of temperature dependence of piezoelectric coefficients matrix of lead zirconate titanate ceramics by quasi-static and resonance method," *Journal of Physics D: Applied Physics*, vol. 42, p. 095417, apr 2009.

[87] M. Boccara, J.-F. Bouchaudy, B. Mours, and F. Richard, "Ultra-high vacuum qualification of the prototype module for the 2×3 km arms of the VIRGO interferometer," *Vacuum*, vol. 57, pp. 283–297, 2000.

[88] R. Lugolole and S. K. Obwoya, "The effect of thickness of aluminium films on optical reflectance," *Journal of Ceramics*, vol. 2015, no. 1, p. 213635, 2015.

[89] M. Yang, C. Zhou, and D. Jiang, "Comparison of different strategies to realize highly reflective thin film coatings at 1064nm," *Infrared Physics & Technology*, vol. 51, no. 6, pp. 572–575, 2008.

[90] A. D. Rakić, A. B. Djurišić, J. M. Elazar, and M. L. Majewski, "Optical properties of metallic films for vertical-cavity optoelectronic devices," *Applied Optics*, vol. 37, pp. 5271–5283, 1998.

[91] D. I. Yakubovsky, A. V. Arsenin, Y. V. Stebunov, D. Y. Fedyanin, and V. S. Volkov, "Optical constants and structural properties of thin gold films," *Optics Express*, vol. 25, no. 21, pp. 25574–25587, 2017.

[92] United Nations, "Transforming our world: The 2030 agenda for sustainable development." https://sdgs.un.org/2030agenda, 2015. Resolution A/RES/70/1, adopted by the General Assembly on 25 September 2015.

[93] F. W. J. Olver, A. B. Olde Daalhuis, D. W. Lozier, B. I. Schneider, R. F. Boisvert, C. W. Clark, B. R. Miller, B. V. Saunders, H. S. Cohl, and M. A. McClain, "NIST digital library of mathematical functions," 2024. Orthogonality: §18.3, Table 18.3.1; Addition theorem: §18.18(ii), Eq. 18.18.14.

[94] R. J. Potton, "Reciprocity in optics," *Reports on Progress in Physics*, vol. 67, no. 5, pp. 717–754, 2004.

Societal reflection

Fundamental research and its place in society

The work in this thesis 'calibrating a gravitational wave detector using scattered light' sits firmly within fundamental physics research. It does not aim to create a product, address a public health problem, or solve an environmental crisis. That makes it reasonable to ask: what is the societal value of this work, and how does it connect to broader goals like the United Nations Sustainable Development Goals (SDGs) [92]?

The most direct link is to SDG 9: Industry, Innovation and Infrastructure, which calls for stronger scientific research and upgraded technological capabilities. Gravitational wave astronomy is one of the most demanding measurement challenges ever attempted, requiring strain sensitivities approaching 10^{-23} Hz$^{-1/2}$ near 100 Hz. Reaching and sustaining this sensitivity has driven continuous advances in precision optics, laser technology, vibration isolation, digital signal processing, and vacuum engineering. Many of these developments have found uses well beyond fundamental physics: laser stabilisation techniques have influenced telecommunications and metrology, seismic isolation concepts have informed earthquake engineering, and advanced mirror coating technology has benefited semiconductor lithography and space optics. The calibration method developed in this thesis contributes to this broader ecosystem of precision measurement.

International collaboration and knowledge sharing

This thesis was carried out within the LIGO–Virgo–KAGRA (LVK) Collaboration, one of the largest international scientific partnerships in physics, involving more than 1500 scientists across dozens of countries. This connects directly to SDG 17: Partnerships for the Goals and to SDG 4: Quality Education. The LVK Collaboration is also a training ground for early-career researchers in experimental physics, data analysis, and engineering. The skills developed during this thesis (spanning optics, electronics, vacuum technology, numerical simulation, and data analysis) are transferable to a wide range of careers in science, technology, and industry.

Sustainability and environmental considerations

Running a gravitational wave detector carries a real environmental cost. Advanced Virgo consumes significant electrical power for its laser systems, vacuum pumps, computing infrastructure, and climate control. Its 3 km beam tubes must be held under ultra-high vacuum continuously. From the perspective of SDG 13: Climate Action, it is worth asking whether this resource use is justified.

Two points are relevant here. First, the total energy consumption of the global detector network is small compared to other large-scale research facilities, and negligible compared to industrial energy use. Second, improving calibration accuracy, which is the goal of this thesis, makes better use of data already being collected. A more accurately calibrated detector extracts more scientific information per unit of observing time, reducing the need for longer observing runs or additional detectors to achieve a given result. In this way, calibration improvements help justify the resources already invested in the detector infrastructure.

Broader perspective

Fundamental research operates on timescales that rarely match immediate societal needs. The scientific and technological returns from gravitational wave astronomy may take decades to fully emerge. History, however, consistently shows that investments in fundamental measurement science yield transformative results. The laser, developed from quantum mechanics with no obvious application at the time, now underpins global telecommunications, manufacturing, and medicine. The World Wide Web was created at CERN to help particle physicists share data. It is impossible to predict which specific technologies will emerge from gravitational wave research, but the pattern of innovation flowing from precision measurement science is well established.

This thesis contributes a small but concrete step to the long-term effort to improve gravitational wave detector performance. By providing an independent method to verify a critical optical parameter of the Virgo detector, it strengthens the foundation on which all subsequent astrophysical measurements rest.

A

Tilt-induced higher-order mode coupling

This appendix derives the coupling coefficients between a tilted Gaussian beam and the Hermite–Gaussian eigenmodes of the arm cavity, leading to the Poisson distribution stated in equation (3.48).

The normalised Hermite–Gaussian modes of the cavity factorise as $u_{mn}(x, y) = \varphi_m(x)\,\varphi_n(y)$, where

$$\varphi_m(x) = \left(\frac{2}{\pi}\right)^{1/4} \frac{1}{\sqrt{2^m\,m!\,w}}\, H_m\!\left(\frac{\sqrt{2}\,x}{w}\right) \exp\!\left(-\frac{x^2}{w^2}\right), \tag{A.1}$$

If the beam is perfectly matched in size and waist position, the incident field differs from the TEM$_{00}$ mode only by the phase ramp:

$$E_{\text{sc}}(x, y) = \varphi_0(x)\,\varphi_0(y)\,e^{ikx\alpha}. \tag{A.2}$$

The coupling coefficient into TEM$_{mn}$ is

$$c_{mn} = \iint_{-\infty}^{\infty} u^*_{mn}(x, y)\; E_{\text{sc}}(x, y)\; dx\,dy = \underbrace{\int_{-\infty}^{\infty} \varphi^*_m(x)\,\varphi_0(x)\,e^{ikx\alpha}\,dx}_{I_m} \cdot \underbrace{\int_{-\infty}^{\infty} \varphi^*_n(y)\,\varphi_0(y)\,dy}_{\delta_{n0}}. \tag{A.3}$$

Since the tilt lies entirely in the x-direction, the y-integral reduces to the orthonormality relation [93], and only modes with $n = 0$ are excited: $c_{mn} = I_m\,\delta_{n0}$.

Evaluation of I_m

Substituting Equation (A.1):

$$I_m = \sqrt{\frac{2}{\pi}}\, \frac{1}{\sqrt{2^m\,m!\;w}} \int_{-\infty}^{\infty} H_m\!\left(\frac{\sqrt{2}\,x}{w}\right) \exp\!\left(-\frac{2x^2}{w^2}\right)\, e^{ikx\alpha}\,dx. \tag{A.4}$$

With the substitution $u = \sqrt{2}\,x/w$ this becomes

$$I_m = \frac{1}{\sqrt{\sqrt{\pi}\,2^m\,m!}} \int_{-\infty}^{\infty} H_m(u)\, \exp\!\left(-u^2 + i\xi u\right) du, \tag{A.5}$$

where the dimensionless tilt parameter is

$$\xi \equiv \frac{kw\alpha}{\sqrt{2}}. \tag{A.6}$$

The integral is evaluated by completing the square in the exponent,

$$-u^2 + i\xi u = -\left(u - \frac{i\xi}{2}\right)^2 - \frac{\xi^2}{4}, \tag{A.7}$$

and shifting the integration variable to $v = u - i\xi/2$:

$$\int_{-\infty}^{\infty} H_m(u)\,e^{-u^2 + i\xi u}\,du = e^{-\xi^2/4} \int_{-\infty}^{\infty} H_m\!\left(v + \tfrac{i\xi}{2}\right)\, e^{-v^2}\,dv. \tag{A.8}$$

The translated Hermite polynomial is expanded using the addition theorem [93]

$$H_m(v+a) = \sum_{j=0}^{m} \binom{m}{j} (2a)^{m-j} H_j(v), \tag{A.9}$$

with $a = i\xi/2$. Integrating term by term and using the orthogonality relation [93] $\int_{-\infty}^{\infty} H_j(v)\, e^{-v^2}\, dv = \sqrt{\pi}\, \delta_{j0}$, only the $j = 0$ term survives:

$$\int_{-\infty}^{\infty} H_m\Big(v + \tfrac{i\xi}{2}\Big)\, e^{-v^2}\, dv = \sqrt{\pi}\, (i\xi)^m. \tag{A.10}$$

Substituting back into Equation (A.5):

$$I_m = \frac{(i\xi)^m}{\sqrt{2^m\, m!}}\, e^{-\xi^2/4}. \tag{A.11}$$

Power coupling and Poisson distribution

The power coupled into TEM$_{m0}$ is $|c_{m0}|^2 = |I_m|^2$. Writing $\mu \equiv \xi^2/2 = (kw\alpha)^2/4$:

$$\boxed{|c_{m0}|^2 = \frac{\mu^m}{m!}\, e^{-\mu}, \qquad \mu = \frac{(kw\alpha)^2}{4} = \left(\frac{\pi w \alpha}{\lambda}\right)^2.} \tag{A.12}$$

The power follows a **Poisson distribution** in the mode index m, with mean $\langle m \rangle = \mu$. One can verify that $\sum_{m=0}^{\infty} |c_{m0}|^2 = 1$, confirming power conservation. For $m = 0$ one recovers the TEM$_{00}$ mode-matching efficiency:

$$\eta(\alpha) = \eta_0\, |c_{00}|^2 = \eta_0\, \exp\left(-\frac{(k\, w(z_{\text{EM}})\, \alpha)^2}{4}\right), \tag{A.13}$$

where η_0 accounts for any residual on-axis mismatch. For small tilts ($\mu \ll 1$), the dominant higher-order contribution is the TEM$_{10}$ mode, with $|c_{10}|^2 \approx \mu$, confirming that beam tilt couples primarily into the first-order mode along the tilt direction.

B

Reciprocity

In the derivation of Section 4.2, the transmission coefficients of the cavity mirrors are assumed to be identical for light propagating in either direction through the optic. This appendix summarises the principle of optical reciprocity that justifies this assumption, and clarifies why the reflectivities of the test masses are nevertheless direction-dependent.

Fundamental principles of optical reciprocity

Optical reciprocity is a symmetry property of Maxwell's equations in linear, time-invariant media with symmetric constitutive tensors. In practical terms, reciprocity states that if a source located at point A produces a field detected at point B, then exchanging source and detector leaves the coupling unchanged [94].

Within linear scattering theory, reciprocity implies that the transmission amplitudes between two ports satisfy

$$t_{1\to 2} = t_{2\to 1}, \tag{B.1}$$

provided the surrounding media are identical on both sides of the device. Reciprocity does not require geometrical symmetry, nor does it require the absence of absorption. The essential conditions are linearity and symmetric constitutive relations, i.e. the absence of non-reciprocal magneto-optic effects.

It is useful to distinguish reciprocity from time-reversal symmetry. Time-reversal symmetry requires a lossless system, whereas reciprocity remains valid even in the presence of absorption or scattering losses. A system may therefore be reciprocal without being time-reversal invariant.

Reciprocity in dielectric multilayer coatings

The cavity mirrors in Virgo are realised as dielectric multilayer stacks consisting of alternating high- and low-refractive-index layers with in the center a substrate. Each layer can be represented by a transfer matrix, and the full optical response follows from the ordered product of these matrices.

For non-magnetic dielectric materials such as silica (SiO_2) and tantala (Ta_2O_5) (used in Advanced Virgo), the permittivity tensor is symmetric. Each layer therefore satisfies the conditions for reciprocity, and so does the full multilayer coating. In particular, the transmission amplitude through the coating is identical in both propagation directions:

$$t_{\text{forward}} = t_{\text{backward}}. \tag{B.2}$$

However, reciprocity does *not* require equal reflectivity for illumination from opposite sides. This distinction is central for the asymmetric AR/HR mirror geometry used in Virgo.

AR/HR mirror configuration in Virgo

A Virgo test mass consists of a fused-silica substrate with different coatings on its two faces:

- a high-reflective (HR) coating on the cavity-facing side,
- an anti-reflective (AR) coating on the rear side.

The HR coating forms the Fabry–Perot arm cavity mirror, while the AR coating suppresses parasitic reflections and ghost beams from the external surface.

Because the two surfaces are physically different optical structures, the boundary-value problem depends on the direction of illumination. Consequently,

$$R_{\text{HR side}} \neq R_{\text{AR side}}. \tag{B.3}$$

This directional asymmetry is expected and does not violate reciprocity. A schematic overview of the two processes is shown in figure B.1.

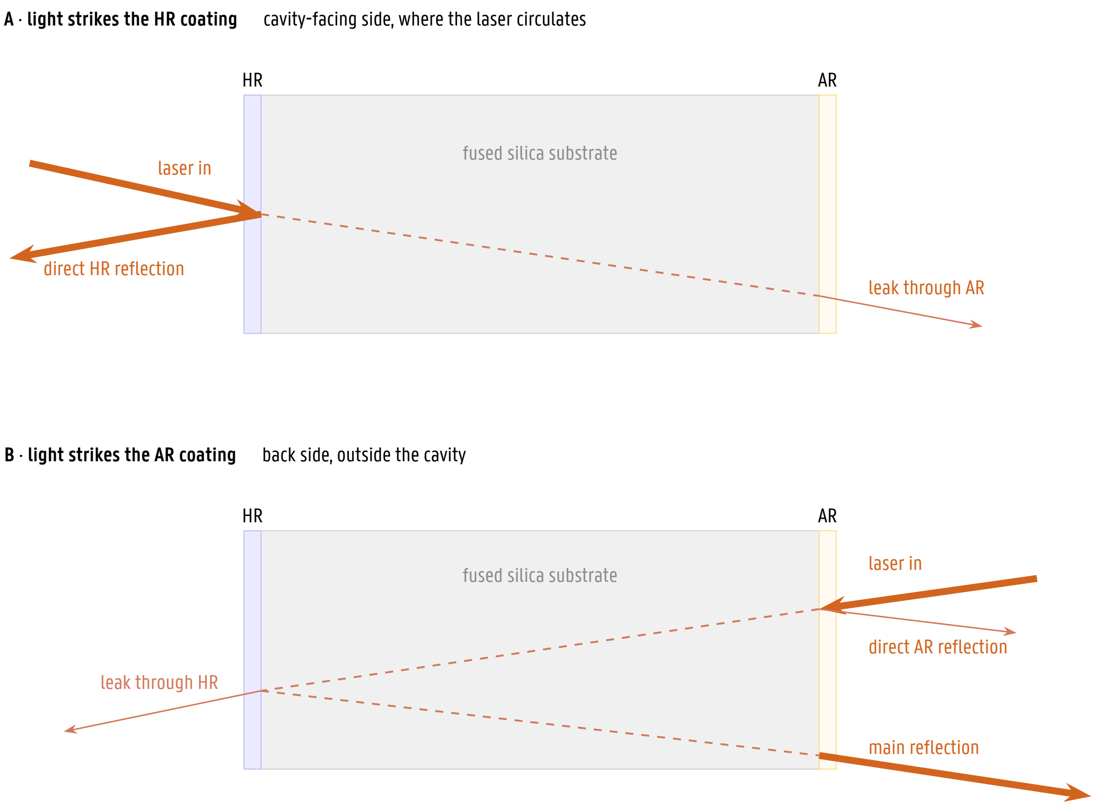

Figure B.1: Directional asymmetry of a Virgo test mass with HR coating on the cavity side and AR coating on the rear side. Reciprocity preserves equal transmission amplitudes in opposite directions.

Energy conservation and direction-dependent reflection

For either direction of illumination, energy conservation requires

$$R + T + \Lambda_{arm} = 1, \tag{B.4}$$

where R is reflected power, T transmitted power, and Λ_{arm} lost power (absorbed or scattered).

For incidence on the HR side, nearly all optical power is reflected:

$$R_{\text{HR}} \approx 1, \qquad T_{\text{HR}} \ll 1. \tag{B.5}$$

For incidence on the AR side, the beam first transmits through the AR coating, propagates through the substrate, reflects from the HR coating, and returns through the substrate and AR coating. The returned power is approximately

$$R_{\text{AR}} \approx T_{\text{AR}}^2 R_{\text{HR}}, \tag{B.6}$$

with additional reduction due to scattering and absorption on the double pass.

The difference in reflectivity therefore arises from propagation through an asymmetric optical structure.

Consistency with reciprocity

Although the reflectivities differ, reciprocity remains fully satisfied. Reciprocity constrains transmission amplitudes between ports:

$$t_{\text{left}\to\text{right}} = t_{\text{right}\to\text{left}}, \tag{B.7}$$

whereas in general

$$R_{\text{left}} \neq R_{\text{right}}. \tag{B.8}$$

The Virgo test masses are therefore fully reciprocal optical elements. Their directional asymmetry in reflectivity is simply a consequence of the deliberately asymmetric HR/AR coating design. This also proves the assumption: $T_{WE}^{l\to r} = T_{WE}^{r\to l}$.

C

Derivation of the DC readout power including scattered-light terms

In DC readout, the interferometer is operated with a small differential arm offset ψ, which leaks carrier light to the antisymmetric port. A passing gravitational wave of strain $h(t)$ induces an additional differential phase shift. This appendix derives the full expression for the power at the B1 photodiode (which is an extension of the work by Was et al. [5]), keeping the cavity gain $G(f)$ complex throughout so that the result remains valid at all signal frequencies f, not only in the low-frequency limit $f \ll f_{\text{arm}}$. This derivation follows the simplified end-bench model presented in [5]. For the signals used in Chapter 5, a slightly modified version of the derivation is applied to be consistent with the end-bench model shown in Figure 5.3.

Complex cavity gain

The single-pole cavity transfer function derived in section 3.2 is

$$G(f) = \frac{1}{1-r} \cdot \frac{1}{1 + i f / f_{\text{arm}}}, \tag{C.1}$$

where $r = \sqrt{1 - T_{\text{WI}} - \Lambda_{\text{arm}}}$ and $f_{\text{arm}} \approx (1-r)\, c/(4\pi L)$ is the cavity pole frequency. For $f \ll f_{\text{arm}}$ the gain is real and positive, recovering the standard scalar amplification. For $f \not\ll f_{\text{arm}}$, the gain becomes complex and can be written in polar form:

$$G(f) = |G(f)|\, e^{-i\varphi(f)}, \tag{C.2}$$

with magnitude and phase

$$|G(f)| = \frac{1}{1-r} \cdot \frac{1}{\sqrt{1 + (f/f_{\text{arm}})^2}}, \tag{C.3}$$

$$\varphi(f) = \arctan\left(\frac{f}{f_{\text{arm}}}\right). \tag{C.4}$$

The phase φ represents the cavity storage lag: at $f = 0$ it vanishes, at $f = f_{\text{arm}}$ it equals $\pi/4$, and as $f \to \infty$ it approaches $\pi/2$. Physically, the storage lag arises because the cavity builds up and releases energy over the storage time $\tau = 1/\gamma$, so signals oscillating faster than f_{arm} are not only attenuated but also phase-shifted.

Intracavity fields

The intracavity fields in the west and north arms are

$$E_{\text{W}}(t) = E_0\, \exp\left[i\left(\frac{\psi}{2} + \delta_h\right)\right] \left(1 + a_s\, e^{i\phi_{\text{sc}}}\right), \tag{C.5}$$

$$E_{\text{N}}(t) = E_0\, \exp\left[-i\left(\frac{\psi}{2} + \delta_h\right)\right], \tag{C.6}$$

where the compact definitions

$$\delta_h \equiv 2\pi G\, \frac{hL}{\lambda}, \tag{C.7}$$

$$a_s \equiv G\, T_{\text{WE}} \sqrt{f_{r,m}} \tag{C.8}$$

collect the gravitational wave phase shift and the scattered light amplitude, respectively. Three physical assumptions are built into these expressions.

First, the DC offset $\psi \ll 1$ is a small differential phase applied to place the interferometer slightly off the dark fringe, providing the local oscillator required for DC readout (section 2.4.2).

Second, the gravitational wave contribution $\delta_h \ll 1$ is treated to first order, valid for all astrophysical signals within the detection band.

Third, the scattered light term $a_s e^{i\phi_{\text{sc}}}$ appears only in the west arm because the scattering element is located on the west end bench. The north arm is assumed to be unperturbed.

Taylor expansion of the phase factors

Both ψ and δ_h are small, so the complex exponentials can be expanded. Keeping terms up to first order in ψ and δ_h, and retaining the quadratic term $\psi^2/4$ because it contributes to the DC power at the output port:

$$e^{\pm i(\psi/2+\delta_h)} = 1 \pm i\left(\frac{\psi}{2}+\delta_h\right) - \frac{1}{2}\left(\frac{\psi}{2}+\delta_h\right)^2 + \cdots \tag{C.9}$$

Expanding the square and dropping the cross term $\psi\,\delta_h$ (second order in two small quantities) and the term δ_h^2 (second order in δ_h):

$$e^{\pm i(\psi/2+\delta_h)} \simeq 1 \pm i\left(\frac{\psi}{2}+\delta_h\right) - \frac{1}{2}\left(\frac{\psi}{2}\right)^2 . \tag{C.10}$$

The retained quadratic term $-\psi^2/8$ is necessary because it contributes to the constant (DC) power at the output. All higher-order terms are negligible for the signal amplitudes of interest.

Field at the antisymmetric port

The field at the antisymmetric port is the difference of the two arm fields after the beam splitter. A contrast defect parameter $\epsilon \ll 1$ accounts for any asymmetry between the two arms (differences in mirror reflectivities, beam overlap, etc.):

$$\Delta E \equiv E_{\text{W}} - (1-\epsilon)\,E_{\text{N}}. \tag{C.11}$$

Substituting equations (C.5) and (C.6) and applying the expansion (C.10):

$$\begin{aligned}\Delta E = E_0\left[1 + i\left(\frac{\psi}{2}+\delta_h\right) - \frac{\psi^2}{8}\right]\left(1 + a_s e^{i\phi_{\text{sc}}}\right)\\ - (1-\epsilon)\,E_0\left[1 - i\left(\frac{\psi}{2}+\delta_h\right) - \frac{\psi^2}{8}\right].\end{aligned} \tag{C.12}$$

Collecting terms and dropping products of two small quantities ($\epsilon\,\delta_h$, $\epsilon\,\psi$, $a_s\,\delta_h$, $a_s\,\psi$):

$$\Delta E \simeq E_0\left(i\psi + \epsilon + 2i\delta_h + a_s e^{i\phi_{\text{sc}}}\right). \tag{C.13}$$

The four terms have clear physical interpretations: $i\psi$ is the carrier leakage from the DC offset (the local oscillator), ϵ is the residual carrier from the contrast defect, $2i\delta_h$ is the gravitational wave signal (the factor of 2 arises from the differential phase in both arms), and $a_s e^{i\phi_{\text{sc}}}$ is the scattered light contribution from the west end bench. For clarity, we define $\Delta E = E_0(A+B)$ with

$$A \equiv i\psi + \epsilon, \qquad B \equiv 2i\delta_h + a_s e^{i\phi_{\text{sc}}}. \tag{C.14}$$

Detected power

The power at B1 is $P_{\text{B1}} = T_{\text{IM}}|\Delta E|^2$. Expanding and dropping terms quadratic in B (second order in the small quantities h and $f_{r,m}$):

$$|\Delta E|^2 = |A|^2 + 2\ \text{Re}(A^*B). \tag{C.15}$$

The DC term gives $|A|^2 = \psi^2 + \epsilon^2$.

The cross term expands as

$$\begin{aligned}A^*B &= (-i\psi + \epsilon)\left(2i\delta_h + a_s e^{i\phi_{\text{sc}}}\right)\\ &= 2\psi\delta_h - i\psi a_s e^{i\phi_{\text{sc}}} + 2i\epsilon\delta_h + \epsilon a_s e^{i\phi_{\text{sc}}}.\end{aligned} \tag{C.16}$$

Taking the real part of each term separately, using $\delta_h = |\delta_h|e^{-i\varphi}$ and $a_s = |a_s|e^{-i\varphi}$:

$$\mathrm{Re}(2\psi\delta_h) = 2\psi|\delta_h|\cos\varphi, \tag{C.17}$$

$$\mathrm{Re}(-i\psi a_s e^{i\phi_{\mathrm{sc}}}) = \psi|a_s|\sin(\phi_{\mathrm{sc}} - \varphi), \tag{C.18}$$

$$\mathrm{Re}(2i\epsilon\delta_h) = 2\epsilon|\delta_h|\sin\varphi, \tag{C.19}$$

$$\mathrm{Re}(\epsilon a_s e^{i\phi_{\mathrm{sc}}}) = \epsilon|a_s|\cos(\phi_{\mathrm{sc}} - \varphi). \tag{C.20}$$

Full expression for P_{B1}

Assembling all terms with $|\delta_h| = 2\pi|G|hL/\lambda$ and $|a_s| = |G|T_{\mathrm{WE}}\sqrt{f_{r,m}}$ gives the following expression for P_{B1}:

$$\begin{aligned} P_{\mathrm{B1}}(t) = \frac{T_{\mathrm{IM}}}{2}|E_0|^2 \Bigg[& \underbrace{\psi^2 + \epsilon^2}_{\text{DC power}} \\ & \underbrace{+ 2\psi \cdot \frac{4\pi|G|L}{\lambda} h(t)\cos\varphi + 2\epsilon \cdot \frac{4\pi|G|L}{\lambda} h(t)\sin\varphi}_{\text{gravitational wave signal}} \\ & \underbrace{+ 2\psi\,|G|T_{\mathrm{WE}}\sqrt{f_{r,m}}\,\sin(\phi_{\mathrm{sc}} - \varphi)}_{\text{scattered light — phase coupling}} \\ & \underbrace{+ 2\epsilon\,|G|T_{\mathrm{WE}}\sqrt{f_{r,m}}\,\cos(\phi_{\mathrm{sc}} - \varphi)}_{\text{scattered light — amplitude coupling}} \Bigg]. \end{aligned} \tag{C.21}$$

This is the central result of this appendix. The standard from the main text is recovered by setting $\varphi = 0$, corresponding to the limit $f \ll f_{\mathrm{arm}}$.

Physical interpretation

DC power. The term $\psi^2 + \epsilon^2$ sets the constant power level at B1. It comes from the deliberately leaked carrier (through ψ) and the residual arm asymmetry (through ϵ). This baseline power determines the shot noise floor but contains no gravitational wave information.

Gravitational wave signal. When the gravitational wave frequency is well below the cavity pole ($f_{\mathrm{GW}} \ll f_{\mathrm{arm}}$), the cavity responds without any phase lag ($\varphi \to 0$), and the signal takes the familiar form $2\psi\,(4\pi|G|L/\lambda)\,h(t)$. As the signal frequency increases toward and beyond the cavity pole, the finite storage time of the cavity introduces a growing phase lag φ. This progressively shifts the signal from the $\cos\varphi$ quadrature, which couples through ψ, into the $\sin\varphi$ quadrature, which couples through the much smaller contrast defect ϵ. Combined with the amplitude roll-off $|G| \propto 1/f_{\mathrm{GW}}$, (section 3.2.2) this is the origin of the detector's sensitivity loss above the cavity pole.

Scattered light. The scattered light behaves in the same way as the gravitational wave signal, but shifted in phase by $-\varphi$. At low scattering frequencies, it appears in the $\sin\phi_{\mathrm{sc}}$ quadrature and couples through ψ. At higher frequencies, the cavity phase lag rotates the signal into the $\cos\phi_{\mathrm{sc}}$ quadrature, where it couples through ϵ instead. Above the cavity pole, the scattered light is increasingly suppressed by $|G| \propto 1/f_{\mathrm{sc}}$, regardless of the operating point ψ.