Modeling of Thermally Aberrated Optical Cavities for Gravitational Wave Detectors

Lane Scheel*
Valparaiso University, Valparaiso, IN 46383, USA

Advisor: Daniel Brown†
University of Adelaide, Adelaide, SA 5000, Australia

(Dated: July 29, 2024)

As current and next-generation gravitational wave detectors strive for greater sensitivity, higher beam powers result in more thermal absorption within the interferometer cavities and test masses. This causes mirror surface deformation and scattering of the incident beam into higher order modes (HOMs), effects that are not fully understood. While attempts have been made to model these phenomena using the FINESSE software package, its modal basis for constructing optical fields requires an exponentially increasing number of HOMs for high accuracy, leading to computational inefficiency and the impossibility of truly accurate results due to the need for infinite modes. Therefore, we model these thermal effects using the linear canonical transform (LCT) framework. This grid-based numerical approach yields high spatial-frequency results and scales more effectively with increasingly complex optical fields than FINESSE. We present a discussion and analysis of these thermal effects, along with a brief exploration of lost power within LIGO's power recycling cavity.

I. BACKGROUND

A. Gravitational Wave Interferometry

First predicted by Albert Einstein in 1916 [7], gravitational waves are ripples in space-time produced by the acceleration of mass. The energetic release of such an events must be immense in order to produce a signal large enough for us to detect here on Earth. Almost a full century later in 2015, the existence of gravitational waves was confirmed by the Laser Interferometer Gravitational-wave Observatory (LIGO) Scientific Collaboration [9]. A binary black hole system roughly 60 times the mass of the sun inspiraled and merged releasing 3 solar masses of energy in the last fraction of a second of its life. The gravitational waves it shed would propagate through the universe for 1.3 billion years before being detected at an accuracy comparable to measuring the distance between Earth and Alpha Centari to the width of a human hair.

Interferometry is not a new concept. The LIGO interferometers are advanced versions of what Michelson and Morley used in 1887 to try to detect the "luminiferous aether" [10]. A laser beam is split into two, reflected back into the beam splitter, and the interference pattern is measured. Any small change in the arm lengths causes a change in interference, which can be precisely measured.

LIGO's detectors are more precisely known as dual-recycled, Fabry-Perot, Michelson interferometers. Where dual-recycled refers to the power recycling cavity (PRC) and signal recycling cavity (SRC), Farby-Perot refers to the arm cavities being comprised of two parallel reflecting surfaces (input and end test masses), and Michelson refers back to the original design. The dual-recycling cavities and Fabry-Perot arms each greatly enhance the sensitivity of the detector by reducing the noise floor.

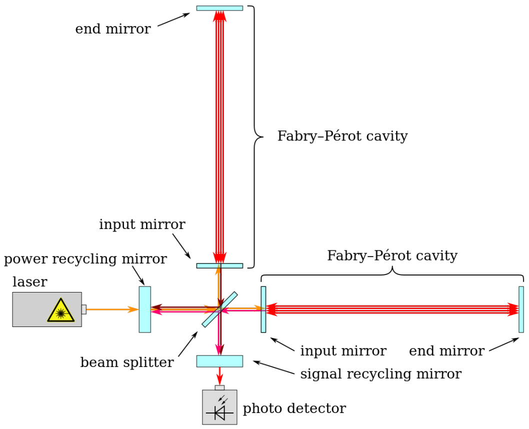

FIG. 1. A simplified diagram of a LIGO interferometer. Pulled from https://simple.wikipedia.org/wiki/LIGO.

Further improvements in the LIGO interferometers include seismic isolation and the use of squeezed light. All of these advanced features come together to allow for an interferometer that can detect changes in distance orders of magnitude smaller than that of an atomic nucleus. This incredible sensitivity enables the detectors to feel the ripples in space-time we call gravitational waves [14].

B. Detecting Gravitational Waves

Data from gravitational wave detectors is collected as *strain*, a measure of the relative change in the length of the arm cavities. When undisturbed, the two laser beams in the interferometer perfectly destructively inter-

---

* lane.scheel@valpo.edu
† daniel.brown@adelaide.ed.au

fere, resulting in zero readout from the sensor. However, when the arms are stretched and compressed to different lengths by a passing gravitational wave, the beams no longer perfectly interfere. The sensor detects this change in interference, which is interpreted as strain. This measurement is taken thousands of times per second during observing runs that last for months.

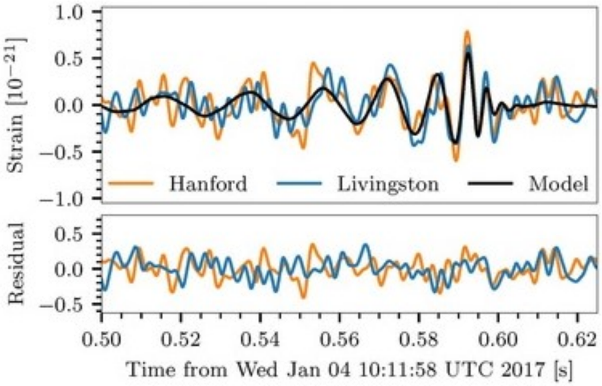

FIG. 2. The gravitational strain detected at both the LIGO Livingston and Hanford sites. Overlayed with a model waveform used for matched filtering. Pulled from https://www.ligo.caltech.edu/LA/image/ligo20170601d

Unfortunately, in such sensitive interferometers as those used by LIGO, noise is a constant challenge. Strain data is incredibly noisy due to seismic activity, quantum vacuum fluctuations, Brownian motion in components, and thermal radiation. To identify signals from passing gravitational waves, the data must be heavily denoised and processed through "matched filtering." This resource-intensive method involves using hundreds of thousands of template waveforms generated by theoretical models of gravitational wave sources, such as binary black holes (BBHs) or binary neutron stars (BNSs). These templates are cross-correlated with the detector strain at each time step to find matches. High correlation flags potential gravitational wave events similar to the theoretical waveform matched to the data. Further analysis with advanced Bayesian inference techniques estimates the source's parameters. After this extensive process, we can determine the likelihood of the data containing an actual gravitational wave signal and identify the type of event that likely produced it [8].

II. MOTIVATION

Despite our advancements in detecting gravitational waves, interferometers at this scale are far from perfect scientific tools. There are still many aspects we do not fully understand, particularly the thermal states of the interferometers. For instance, how do the properties of all optical components change as their temperature increases during operation? Measuring these thermal changes is nearly impossible because it would have to be done while the laser operates at full power. Therefore, we rely on modeling and simulation to gain insights into their behavior.

Of great importance at the moment is understanding the scattering of circulating power in the Fabry-Perot cavities into higher-order-modes (HOMs) due to thermal aberration of the input and end test masses (ITMs & ETMs). Understanding how the optical field scatters into these HOMs is key in determining where the detector is losing power and how we can improve it in the future [12].

Thermal aberration can manifest in various ways. The laser beam heats and deforms mirror surfaces. Dust particles on mirrors become point absorbers that scatter the optical field. Numerous effects need to be understood to optimize the system eventually. These high spatial-frequency features require high resolution and clever modeling techniques to capture their true nature within the detector. Unfortunately, these are computationally expensive, and popular modeling software struggles to handle the job.

III. MODELING

Modeling the propagation of light through an optical cavity is a complex task. First, the optical cavity must be defined using ABCD matrices. These 2x2 matrices are complex and contain information on how the optical field changes as it interacts with each optical component in the cavity. The optical field itself is represented by a complex beam parameter, which carries information on the beam's shape as it propagates through the cavity. This beam parameter is modified by each of the ABCD matrices in the cavity, allowing you to effectively view the beam at any point in the system by operating with the appropriate matrices [11].

The modeling becomes more complicated with the addition of mirror curvatures and other effects which cause scattering of the HG00 fundamental mode into higher order modes. These features are able to be modeled via Hermite-Gauss decomposition. Effectively, the now more complex beam is constructed from a basis of simple Hermite-Gauss modes that have been summed together to produce the desired beam [2].

A. FINESSE 3.0

FINESSE is a popular software in the field of graviational physics that is capable of modeling propagating optical fields via Hermite-Gauss decomposition. Now in its third iteration, FINESSE 3.0, this software can model arbitrary advanced interferometer configurations and provide detailed analysis of the propagating optical field within [2].

The modeling of higher spatial-frequency features in

FINESSE 3.0, such as thermal abberations, point absorbers, and more, is quite difficult and can be unreliable due to the use of a Hermite-Gauss basis. For this method to yield high-accuracy results, the number of HOMs needed grows exponentially. Each of the HOMs must be treated individually within the model, and because of this it very computationally expensive to compute. In fact, to completely and comprehensively describe a more complex optical field, FINESSE would have to use an infinite number of HOMs. This sort of computation is impossible for a computer to perform in a finite amount of time, and therefore we come to the understanding that a optical field composed of higher order modes can almost never be modeled to perfect accuracy.

Despite this, FINESSE has been the go-to software for the last two decades due to its extensive functionality in sensors, components, and analysis, making the small loss in optical field accuracy an acceptable trade-off.

In recent years, understanding the thermal states of LIGO detectors has become a more pressing issue as they are increasingly limited by thermal noise. To further improve sensitivity, we must accurately model high spatial-frequency thermal effects. Therefore, researchers are exploring other methods of modeling optical fields that may yield higher resolution results.

B. The Linear Canonical Transform

The Linear Canonical Transform (LCT) shows promise for modeling high-resolution optical effects. First demonstrated by Collins in 1970 for modeling optical systems [6], the LCT has seen numerous optimizations across various sub-fields. Recently, Ciobanu demonstrated its capability to model circulating optical fields within resonant optical cavities [3].

To use the LCT, the transverse profile of a complex beam amplitude is sampled on an NxM Cartesian grid, defining the model's resolution. Each optical component in the system is represented by an NxM kernel. The power of the LCT lies in convolution, allowing individual kernels to be combined into one large round-trip kernel, simplifying the resulting expression. These individual kernels are derived from the ABCD matrix representations of each optical component, transformed into equivalent NxM matrices (kernels) through linear algebra.

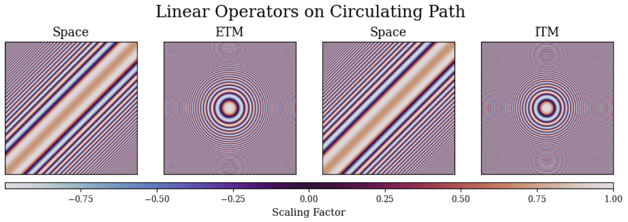

FIG. 3. A visualization of the LCT kernels that represent the LIGO arm optical cavity. Note the the aliasing in the ITM and ETM kernels is a results of low resolution.

The LCT also facilitates easy visualization of the beam shape at any point in the model, but lacks the ability to dissect the beam into individual Hermite-Gauss modes. This feature in FINESSE is useful for diagnosing which higher-order modes are being scattered into.

A deep learning algorithm, such as a Convolutional Neural Network, might be useful for decomposing the LCT beam into its individual Hermite-Gauss constituents. This effort is suggested for future research.

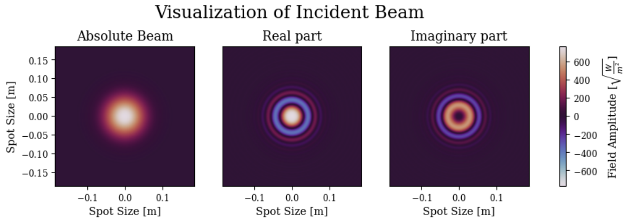

FIG. 4. A visualization of an incident HG00 beam of power 2600W being sampled on the NxM cartesian grid. Also shows individual real and imaginary contributions.

Since the beam is not modeled by individual higher-order modes, it evolves more naturally across the Cartesian grid as it "propagates" through the model, bringing out greater detail in high spatial-frequency effects experienced by interferometers.

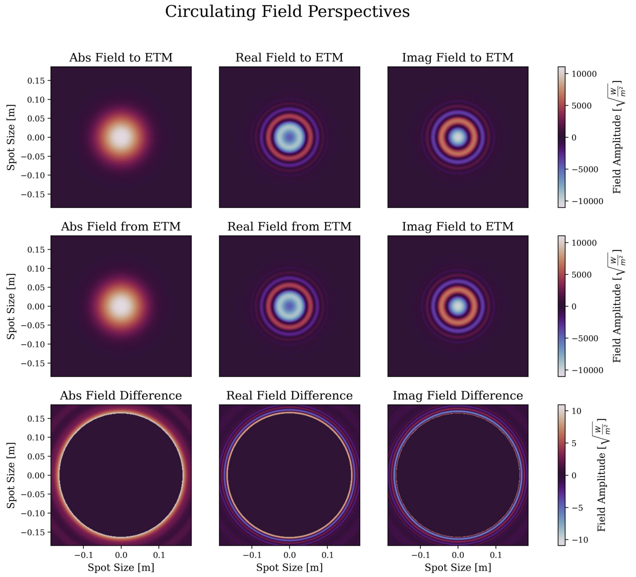

FIG. 5. Optical field visualization from different perspectives within the LCT. We are able to easily visualize the power lost outside of the test mass apertures. Recreated from [4].

IV. METHOD AND RESULTS

It was shown by Ciobanu [4] that the Linear Canonical Transform (LCT) is effective at modeling a circulating optical field in both linear and ring optical cavity configurations. To ensure the correctness of my LCT implementation, I compared its output against FINESSE 3.0 using a variety of input beam modes. After confirming the implementation, I proceeded to compare the numerical precision of the two modeling frameworks.

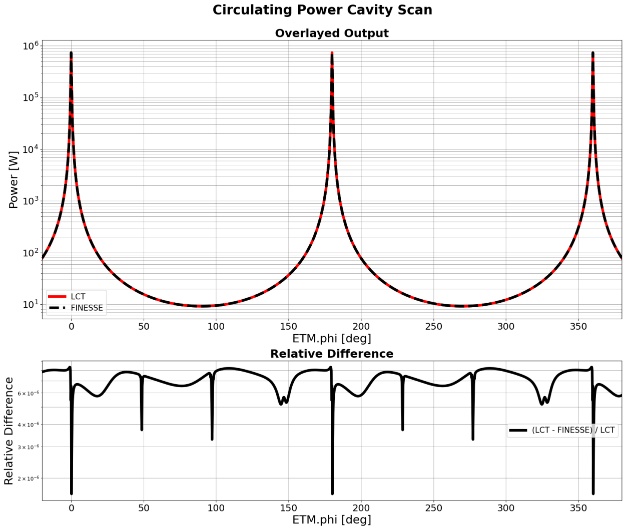

FIG. 6. Overlayed cavity scans of a LIGO arm cavity from the LCT and FINESSE 3.0. Input laser power of 2600W in a pure HG00 mode. Finite apertures are present. Results are imilar to a relative difference of a couple ppm, or 10^{-6}.

As expected, the relative difference between the models increases with the number of high spatial-frequency features within the interferometer, as shown in Table I. In this example, both models apply these features on a 256x256 grid. FINESSE includes higher-order modes only up to order 12, and the relative difference is expected to decrease as more higher-order modes are enabled. This is the simplest beam mode, and the relative difference is expected to increase as the beam becomes more complex in mode makeup.

Features	Maximum Relative Difference
None	5.0e-6
+ Finite Apertures	8.3e-6
+ Thermal Deformation	6.3e-2

TABLE I. Relative difference between LCT and FINESSE 3.0 as more high spatial-frequency features are added. LIGO Arm cavity with 2600W input power in a pure HG00 mode.

The raw numerical precision of the LCT behaves as expected, with loss calculations down to machine precision. While this should also be the case with FINESSE 3.0, results suggest a limit around 10^{-7}. The numerical precision between the two models was compared by calculating the power lost outside the test mass apertures. This was done by calculating the lowest lost eigenvalue at each iteration of the test mass aperture over some range. The eigenvalue is squared and subtracted from one which yields the round-trip loss of the cavity. These results are plotted in Figure 7.

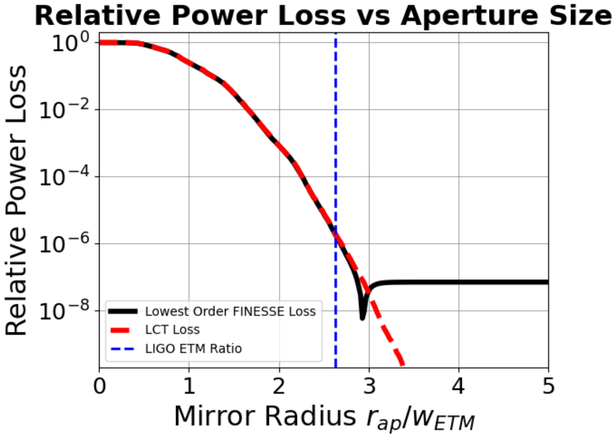

FIG. 7. A comparison of the calculated power lost outside the finite apertures of the test masses. FINESSE 3.0 exhibits an asymptotic behavior that isn't understood. The LCT demonstrates that it should continue down to machine precision.

I suspect FINESSE 3.0 is handling the application of finite apertures quite differently than the LCT. Although both apertures are sampled on a grid of arbitrary resolution, this isn't how FINESSE models the actual beam.

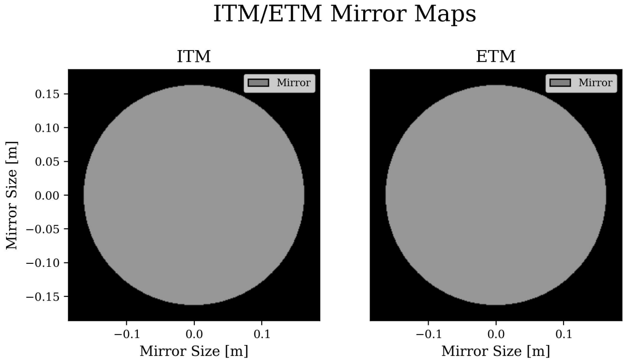

FIG. 8. A visualization of the aperture maps used to simulate the finite apertures of the ITM and ETM. Anything outside the aperture is set to zero.

The computational time required for each framework to perform the same function can vary significantly, especially when FINESSE incorporates higher-order modes. The more modes included, the more computationally expensive the simulation becomes. Depending on the complexity of the optical configuration, there is a point at which FINESSE is faster than the LCT. However, as more accurate solutions are required and more higher-order modes are enabled, the LCT can quickly overtake FINESSE in efficiency.

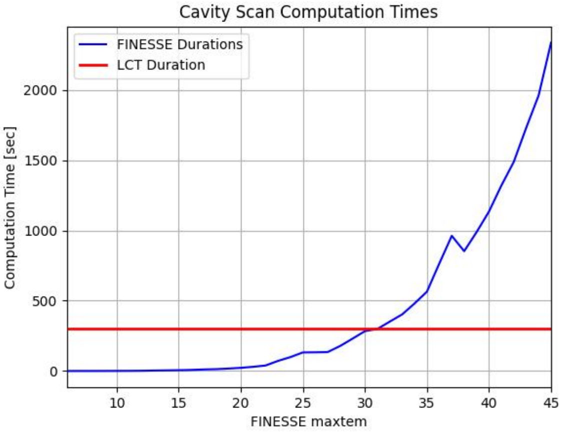

FIG. 9. A comparison of the amount of time necessary for FINESSE 3.0 and the LCT to run a basic cavity scan. All system parameters are identical. Finite apertures are present.

A. Implementation of Thermal Effects

Adding further functionality to the LCT is not very difficult. All aspects of the optical system must simply be represented as an NxM kernel. For mirror heating and deformation, the steps are straightforward: determine the circulating power within the optical cavity, find the fraction of power absorbed by the bulk and coating of the test mass, and understand how the materials of the test masses expand and contract due to heating.

Gathering this information can be challenging. Circulating power is calculated by running the simulation without thermal effects. Absorption rates are usually educated guesses, and material expansion and contraction are approximations that have been developed for over 30 years by researchers Hello and Vinet.

The current LIGO test masses are made of fused silica. The Hello-Vinet framework, implemented in FINESSE 3.0, models deformations using Bessel functions. Instead of implementing this in the LCT, I used FINESSE to derive the deformation solution. This deformation, in meters, must be transformed to a kernel by taking e^{-2kD} over each pixel, where D is the change in surface depth. This transformation accounts for the change in distance and phase that the propagating beam undergoes, storing it as a complex value. We now have an NxM kernel representing the surface deformation of the test mass.

Due to the convolution property of the LCT, the new kernel can be placed appropriately in the propagation order and incorporated into the round trip kernel. There will be kernels for both test masses and any other deformed optical elements.

These thermal deformations introduce scattering into the optical cavity, inducing higher-order modes and losing power from the fundamental HG00 mode. FINESSE models this as well as the LCT, and both agree on the higher-order modes being scattered into.

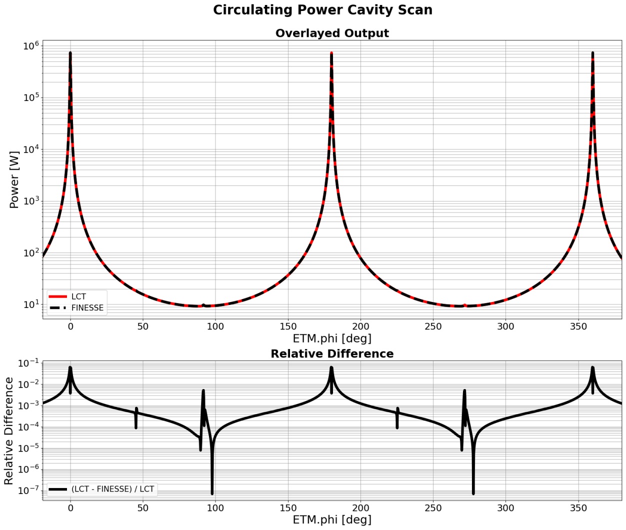

FIG. 10. Overlayed cavity scans from FINESSE 3.0 and the LCT. Both finite apertures and thermal deformation of the test masses are present. FINESSE set to maxtem of 12.

B. Discrepancies

The largest discrepancy between the LCT and FINESSE 3.0 becomes apparent when the beam is more complex than the basic HG00 mode. Comparing the models with various input beams shows that FINESSE tends to underestimate the power scattered into higher-order modes and overestimate the power remaining in the HG00 mode. This effect is more pronounced when FINESSE is limited to fewer higher-order modes but gradually converges to the LCT results as more higher-order modes are included, as shown in Figures 11 & 12.

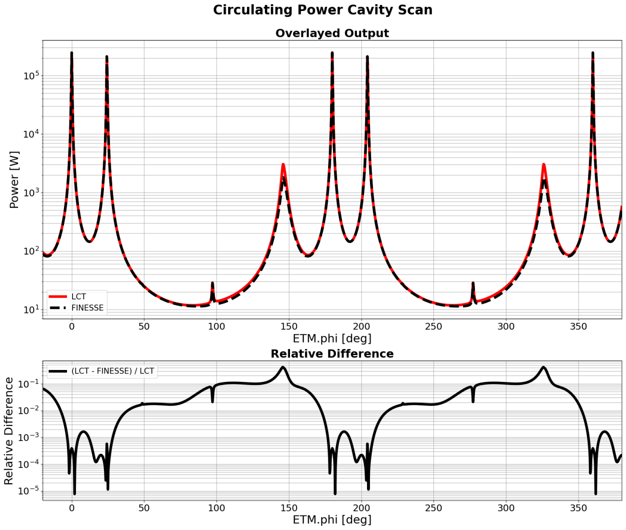

FIG. 11. Overlayed cavity scans from FINESSE 3.0 and the LCT. Finite apertures of the test masses are present. Input beam is comprised of equal parts HG00, HG01, and HG06. FINESSE set to maxtem of 12.

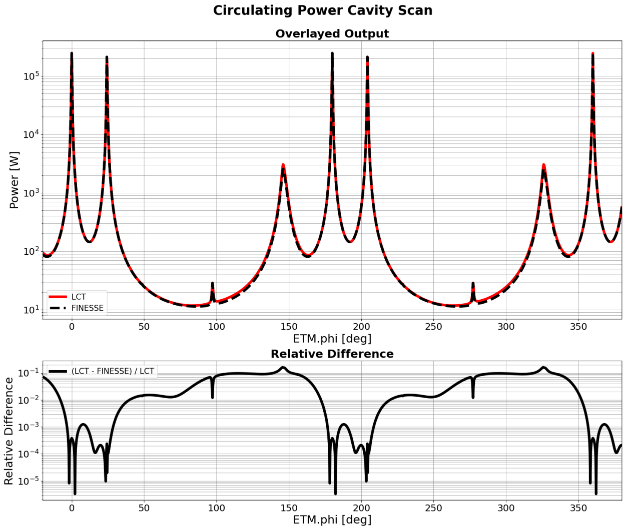

FIG. 12. Overlayed cavity scans from FINESSE 3.0 and the LCT. Finite apertures of the test masses are present. Input beam is comprised of equal parts HG00, HG01, and HG06. FINESSE set to maxtem of 24.

Additionally, after the incorporation of thermal abberation to the test masses, by ramping up the ciruclating power of the cavity and therefore the power absorbed into the optical coatings I could determine if or where FINESSE 3.0 and the LCT began to disagree about the solution. I found that FINESSE 3.0 and the LCT broadly converge to quite similar solutions, but do not converge to precisely the same solution for the amount of power in the resonant HG00 mode. This is characterized in Figures 13 & 14.

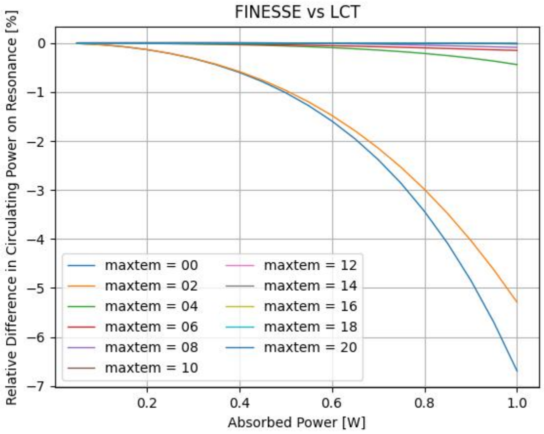

FIG. 13. Relative error on circulating power in resonant HG00 mode between FINESSE 3.0 and the LCT.

Figure 13 implies that the introduction of thermal effects to the mirror surface is primarily a fourth order effect. This is why there is such a large discrepancy between the maxtem=2 and maxtem=4 lines on the plot. Considering that the thermal deformation is a semi-quadratic bump on the mirror surface, this result is easily backed by reason and should be reliable. It should be noted that these thermal deformations are far smaller in magnitude than the curvature of the mirror itself. The thermal deformations are on the scale of microns, while the mirror curvatures are in the thousands of meters.

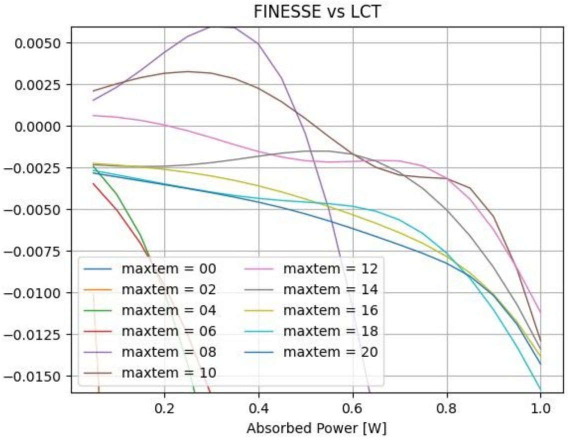

FIG. 14. Relative error on circulating power in resonant HG00 mode between FINESSE 3.0 and the LCT. Zoomed.

We can see from Figure 13 that FINESSE 3.0 and the LCT generally agree, and will converge to what appears to be the same answer with the inclusion of many more higher-order modes in FINESSE. At 1 Watt of absorbed power in each of the optical coatings, as is predicted if the LIGO detectors are to reach there circulating power goal of 1+ mega-Watts, we see that the two modeling frameworks disagree by about 7%. Not awful in the grand scheme of things, but when dealing with a mega-Watt of power, that 7% is equivalent to a discrepancy of 70 kilo-Watts in the two models. Hence why is it import to use as many higher-order modes as possible in FINESSE models dealing with these high spatial-frequency effects.

Taking a closer look, with the inclusion of higher-order modes in FINESSE, we find that the two models agree to within 2/100 of a percent across the absorbed powers. However, we also note that the results do not improve consistently as we include more higher-order modes as one would expect. Certain maxtem's in FINESSE perform more closely to the LCT at different levels of power absorption. It is not currently understood why this happens, but as we know, the LCT evolves the optical field more naturally while FINESSE builds it from discrete layers meaning that a perfectly accurate solution would require infinite higher-order solutions. At this moment, I suspect the discrepancies between the results at higher-order modes due to the numerical error of FINESSE as it tries to deal with both the finite aperture boundaries and each modes that makes up its final solution. However, this is yet to be confirmed due to limited computational resources and requires further inquiry.

V. PRC POWER LOSS

It has been noted that the power circulating within LIGO's power recycling cavity drops to around half of what it should be. A hypothesis by Daniel Sigg, LIGO Hanford's Laad Scientist, is that the 9Mhz sidebands resonant within the PRC are ”leaking” into the main arm cavity if they slip out of anti-resonance once the ITM surface becomes sufficiently thermally distorted.

I constructed a model of the PRC in FINESSE to attempt to solve this mystery. I was able to calculate the amount of power from the 9MHz sidebands being lost into the main arm cavity. These results are shown in Figure 15.

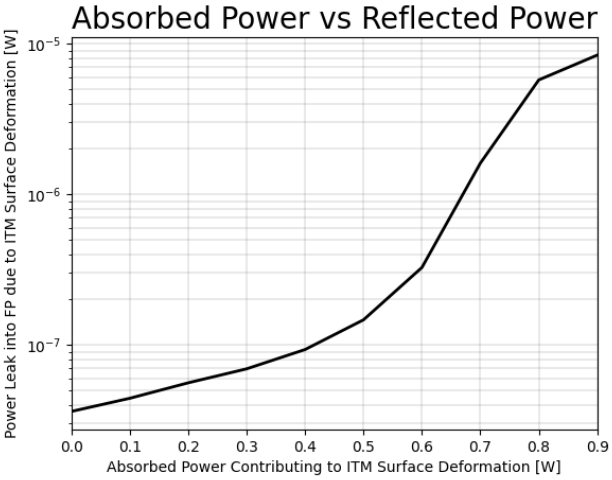

FIG. 15. A plot of the 9 MHz sideband power ”leaking” into the main arm cavity as a function of absorbed power and subsequent thermal deformation of the ITM. Input beam has a power of 1W. So losses are around parts per million.

Unfortunately, time didn't allow for the construction of such a model in the LCT framework, nor did it allow for further exploration of this topic. Based on these preliminary results, it is hard to say if this is truly the reason for the loss of gain in the PRC, but they seem to indicate that this is not the cause of the issue. Further careful modeling of the PRC is warranted.

VI. CONCLUSIONS

From this work, several conclusions can be drawn. Firstly, both FINESSE 3.0 and the LCT are powerful tools capable of accurately describing the optical field within arbitrary interferometer configurations. When many higher-order modes are included, FINESSE converges to the solutions provided by the LCT. However, these calculations in FINESSE are computationally expensive and could take hours to days for large interferometer configurations with high spatial-frequency effects on each optical element.

Secondly, FINESSE 3.0 is a well-developed software that offers greater utility and practicality than the LCT in its current unpackaged and experimental state. With various detectors, configurable optical elements, analysis functions, and more, FINESSE is far more comprehensive and user-friendly.

Therefore, it is advisable to continue using FINESSE 3.0 for interferometer modeling until the LCT is refined and packaged into comprehensive software with similar functionality and ease of use. The results between the models tested in this research are similar enough to be satisfied with FINESSE's performance for scientific research, and the computational cost of including many higher-order modes in FINESSE is justified due to the LCT's unfinished state.

However, there are specific cases where the Linear Canonical Transform would be preferable. For example, research focusing on the thermal states of high-power interferometers or the specific thermal effects of individual components would benefit from an LCT model. In such scenarios, the advantages of an LCT model outweigh the cons of assembling a small, simple model without a well-developed software package.

In the future, with the Linear Canonical Transform in a packaged and refined state, its numerical precision and computational efficiency in modeling high spatial-frequency effects could surpass FINESSE, especially for high-power next-generation gravitational wave detectors. The LCT's reliance on large matrix multiplications allows for significant computational efficiency improvements using high-performance graphics processing units and tensor processing units.

Refining and packaging the Linear Canonical Transform into standardized optical modeling software would make a promising thesis for a future student/researcher and would be a significant addition to this field of work. Tutorials on how to use the Linear Canonical transform for modeling and analysis of resonant optical cavities can be found on the LIGO github. Visit https://git.ligo.org/alexei.ciobanu/lct-tutorials for the original tutorials that enabled this work, and visit https://github.com/LaneScheel/IREU2024-UF for a well documented and easy to follow tutorial that should provide an individual with all the relavent information.

ACKNOWLEDGMENTS

This research was conducted as part of the University of Florida International Research Experience for Undergraduates (IREU) in Gravitational Physics, supported by National Science Foundation grant #1950830. I gratefully acknowledge the guidance and mentorship of Dr. Daniel Brown at the University of Adelaide, the entirety of the OzGrav Collaboration for funding a tour of The Australian National University, as well as those at the University of Florida, Gainesville, for running a culturally enriching and intellectually stimulating program.

---

[1] C. Bond, D. Brown, A. Freise, and K. A. Strain. Interferometric techniques for gravitational-wave detection. *Living Rev. Relativ.*, 13, 2016.

[2] Daniel Brown. FINESSE 3.0 Documentation. https://finesse.ifosim.org/docs/latest/, 2024.

[3] A. Ciobanu, D. Brown, P. Veitch, and D. Ottaway. Modeling circulating cavity fields using the discrete linear canonical transform. *Journal of the Optical Society of America A*, 2021.

[4] Alexei Ciobanu. *Wavefront modelling and sensing for advanced gravitational wave detectors*. Ph.D. thesis, The University of Adelaide, 2021.

[5] Alexei Ciobanu. LCT Tutorials. https://git.ligo.org/alexei.ciobanu/lct-tutorials, 2022.

[6] S. Collins. Lens-system diffraction integral written in terms of matrix optics*. *Journal of the Optical Society of America*, 60.9:1168–1177, 1970.

[7] Albert Einstein. Näherungsweise integration der feldgleichungen der gravitation. *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften*, 1916.

[8] LIGO Lab. LIGOLab Website. https://www.ligo.caltech.edu/, 2024.

[9] LIGO Scientific Collaboration and Virgo Collaboration et al. Observation of gravitational waves from a binary black hole merger. *Physical Review Letters 116.6*, 6, 2016.

[10] A. Michelson and E. Morley. On the relative motion of the earth and the luminiferous ether. *American Journal of Science*, 34:333–345, 1887.

[11] A. Siegman. *Lasers*. Mill Valley, Calif: Univ. Science Books, 1986.

[12] K. A. Strain, K. Danzmann, J. Mizuno, P. G. Nelson, A. Rudiger, R. Schilling, and W. Winkler. Thermal lensing in recycling interferometric gravitational wave detectors. *Physical Letters A*, 194:124–132, 1994.

[13] The LIGO Scientific Collaboration. Overview of advanced ligo adaptive optics. *Applied Optics*, 55:8256–8265, 2016.

[14] The LIGO Scientific Collaboration. LSC Website. https://www.ligo.org/, 2024.