

Wave Optics in Gravitational Lensing: Part II

Yury Holubeu *

August 12, 2025

This draft is not aimed for distribution.

Contents

VII	Review of Other Theories	4
VIII	Other My Theories	5
IX	Other Special Theories, Models	6
7	Wave effect on PBH microlensing constraints	6
7.1	Introduction	6
7.2	The wave optics and finite source size effects on PBH microlensing	8
7.2.1	Microlensing basics for a point source	8
7.2.2	Effect of wave optics	9
7.2.3	Effect of finite source size	12
7.2.4	Effect of exposure time average	13
7.3	Implications on the PBH constraints from the Subaru HSC microlensing search	17
7.4	Discussion	20
7.5	Conclusion	21
8	Wave optics for rotating stars by Béatrice Bonga, Feldbrugge, Metidieri	22
8.1	Introduction	22
8.2	Rotating stars	23
8.2.1	The time delay	23
8.2.2	Geometric optics	27
8.2.3	Wave optics	29
8.2.4	Validating the lens amplitude	30
8.3	Results	31
8.4	Discussion	35
8.5	Unfolding of the point lens	36
8.6	The small spin approximation	37
9	Wave asymptotics and their application to astrophysical plasma lensing by Grillo, Cordes	42
9.1	Introduction	43
9.2	Zeroth and first order geometric optics	45
9.2.1	Geometrical picture	45
9.2.2	Zeroth order gain	47
9.2.3	The 2D Kirchoff diffraction integral	47
9.2.4	Accuracy and regions of applicability	50
9.3	Second order geometric optics	51
9.3.1	Complex rays	51
9.3.2	Caustic location and extent of the caustic zone	52
9.3.3	Gain inside the caustic zone: catastrophe theory and uniform asymptotics	54

*<https://yuriholubeu.github.io/>, yuri.holubev@gmail.com

9.3.4	Uniform asymptotics in plasma lensing	57
9.3.5	Advantages of second order geometric optics	58
9.4	Astrophysical applications	58
9.4.1	TOA perturbations	58
9.4.2	Dispersion measure perturbations	61
9.5	Summary and conclusions	62
9.6	Solving the KDI using the FFT	63
9.7	Estimation of the value of the gain at the caustic	65
9.8	Numerics	65
9.9	More numerical examples and lens colormaps	67

X Other Special Theories, Models 73

XI Method of Path Integral for WO in GL 74

10	Proper time path integrals for gravitational waves: an improved wave optics framework by Braga, Ricciardone, et al.	74
10.1	Introduction	74
10.2	The diffraction integral	77
10.2.1	Field vs particle description	79
10.3	Proper time path integral for a scalar wave	80
10.3.1	Adding a new time parameter	81
10.3.2	Hamiltonian derivation	83
10.4	The high-frequency limit	83
10.4.1	Geometric Optics: the ray equation	83
10.4.2	WKB form of the propagator	85
10.4.3	Recovering the diffraction integral under the paraxial approximation	87
10.5	Perturbation theory	90
10.5.1	Setting up the the expansion	90
10.5.2	The free relativistic particle propagator	90
10.5.3	First order	91
10.5.4	Higher orders	91
10.5.5	Perturbation theory with the Generating Functional	92
10.6	Applications	94
10.6.1	Coulomb potential: an explicit example	94
10.6.2	Massive scalar field	97
10.7	Lensing by a Black Hole: adding the polarization	97
10.7.1	Preliminary about Black Hole perturbation theory	98
10.7.2	Proper time solution of Teukolsky equation	99
10.7.3	Weak field limit and spin effects	100
10.8	Conclusions	101
10.9	Free propagator computations	102
10.10	Coulomb potential computations	103

XII Preparation to Wave Optics for Gravitational Lensing 105

11	Oscillatory path integrals for radio astronomy by Feldbrugge, Pen, Turok	105
11.1	Introduction	105
11.2	From Feynman to Fermat to Fresnel-Kirchhoff	109
11.3	Evaluating the Fresnel-Kirchhoff integral	111
11.3.1	The geometric optics limit	113

11.3.2	Catastrophe theory	116
11.4	Example: a one-dimensional lens	119
11.4.1	Geometric optics approximation	119
11.4.2	Finding the thimbles	122
11.4.3	Integrating along the thimbles	127
11.5	The elementary catastrophes	129
11.5.1	The fold A_2	129
11.5.2	The cusp A_3	131
11.5.3	The swallowtail A_4	134
11.5.4	The elliptic umbilic D_4^-	140
11.5.5	The hyperbolic umbilic D_4^+	143
11.6	Two dimensional localized lenses	146
11.6.1	A generic peak	146
11.6.2	A degenerate peak	150
11.6.3	The swallowtail caustic	151
11.6.4	The elliptic umbilic caustic	154
11.7	Signatures of caustics in fast radio bursts	157
11.8	Conclusions	158
11.9	Defining oscillatory integrals	158
11.9.1	One-dimensional integral	159
11.9.2	Multi-dimensional integrals	160
11.10	Young's double-slit experiment	161
XIII	Other Experiments	166
XIV	Appendix	167
.1	Appendix	167
.1.1	Literature	167
	Bibliography	168

Preface

Parts of the note are:

Main theory of other effects

Other experiments

Other Special Theories, Models

Articles about Already Known Topics

Other Experiments

Key sections are:

(write links to the table of contents)

I used the following literature.

Amazing facts

(I'll reveal it later)

Puzzles for motivation

(I'll reveal it later)

Part VII

Review of Other Theories

(not important articles will be summarized here)

Part VIII

Other My Theories

Part IX

Other Special Theories, Models

7 On the wave optics effect on primordial black hole constraints from optical microlensing search by Sugiyama, Sugiyama, Toshiki, Masahiro

Abstract

Micro lensing of stars, e.g. in the Galactic bulge and Andromeda galaxy (M31), is among the most robust, powerful method to constrain primordial black holes (PBHs) that are a viable candidate of dark matter. If PBHs are in the mass range $M_{\text{PBH}} \lesssim 10^{-10} M_{\odot}$, its Schwarzschild radius (r_{Sch}) becomes comparable with or shorter than optical wavelength (λ) used in a microlensing search, and in this regime the wave optics effect on microlensing needs to be taken into account. For a lensing PBH with mass satisfying $r_{\text{Sch}} \sim \lambda$, it causes a characteristic oscillatory feature in the microlensing light curve, and it will give a smoking gun evidence of PBH if detected, because any astrophysical object cannot have such a tiny Schwarzschild radius. Even in a statistical study, e.g. constraining the abundance of PBHs from a systematic search of microlensing events for a sample of many source stars, the wave effect needs to be taken into account. We examine the impact of wave effect on the PBH constraints obtained from the r -band (6210Å) monitoring observation of M31 stars in Niikura et al. (2019), and find that a finite source size effect is dominant over the wave effect for PBHs in the mass range $M_{\text{PBH}} \simeq [10^{-11}, 10^{-10}] M_{\odot}$. We also discuss that, if a denser-cadence (10 sec), g -band monitoring observation for a sample of white dwarfs over a year timescale is available, it would allow one to explore the wave optics effect on microlensing light curve, if it occurs, or improve the PBH constraints in $M_{\text{PBH}} \lesssim 10^{-11} M_{\odot}$ even from a null detection.

7.1 Introduction

The nature of dark matter (DM) is one of the most tantalizing problems in cosmology and physics. Unknown stable elementary particle(s) beyond the Standard Model of Particle Physics, the so-called Weakly Interacting Massive Particle(s) (WIMP), has been thought of as a viable candidate of DM, but has yet to be detected either in direct experiments, collider experiments, or indirect searches [e.g. 20, 18, 4]. Primordial black holes (PBHs) [17, 8] are alternative, viable candidate of DM [11, 9, 34]. Recently the PBH DM scenario has got attention again, partly because of recent claims that PBHs of $10 M_{\odot}$ mass scales can be progenitors of binary black holes whose gravitational wave have been detected by the LIGO/Virgo experiment [e.g. 33, 7].

Given these growing interests, there are many observational attempts to search for or constrain PBHs of various mass scales [10, 12, 19, 35]. Gravitational microlensing is the most powerful, robust method of constraining PBHs [32, 15], because it is a gravitational effect and can probe mass of a lensing compact object, if detected, regardless of whether or not the lensing object is visible [6, 3, for future prospects]. In such a microlensing search we should keep in mind a discovery potential: if we have an even single, secure candidate of microlensing event indicating a mass scale of $M \lesssim$ a few M_{\odot} and if the counter object is confirmed as a black hole (or extremely invisible object) based on any follow-up, deep observation in various wavelengths, it can be a smoking gun evidence of PBH because any supernova explosion or other astrophysical process cannot make such a light-mass BH of \lesssim a few M_{\odot} . The pioneer work was done by the MACHO and EROS experiments that used monitoring observations of stars in the Large Magellanic Cloud to search for microlensing events and then obtained upper bounds

on the abundance of compact objects over a wide range of mass scales $M \simeq [10^{-7}, 10]M_{\odot}$ [1, 2, 39]. This constraint was recently updated in Niikura et al. [31] that used the public OGLE microlensing events [25].

Niikura et al. [30], where authors of this paper are co-authors, used the new wide-field prime-focus camera at the 8.2m Subaru telescope, Hyper Suprime-Cam (HSC), to carry out very dense cadence observation (2 min cadence) of the Andromeda galaxy (hereafter M31). Thanks to the wide field-of-view and large aperture of HSC/Subaru, they were able to monitor many stars in M31 (about 770 kpc in distance or 24.4 mag for the distance modulus) and to search for microlensing events of much shorter timescales than previously done. They found one possible PBH microlensing event compared to the theoretically-expected number of events up to 1000 events if PBHs make up all DM in the Milky Way and M31 halo regions. The results were then translated into most stringent upper bounds on the abundance of PBHs over the range of mass scales, $M_{\text{PBH}} \simeq [10^{-11}, 10^{-7}]M_{\odot}$. As discussed in Niikura et al. [30], there is a fundamental limitation to constrain PBHs in $M \lesssim 10^{-11}M_{\odot}$, from an optical microlensing search, due to the finite source size effect and the wave optics effect [see 21, 5, for similar discussion].

There are several earlier works discussing the wave optics effect on gravitational lensing [36, 14, 27, 28, 38, 24, 26]. These considered lensing of gamma-ray burst or gravitational waves, which are in much shorter or longer wavelengths than optical light. If PBHs are lighter than $10^{-10}M_{\odot}$, the Schwarzschild radius becomes comparable with or even shorter than optical wavelengths, then we cannot ignore the wave optics effect on microlensing due to interference and diffraction effects. In an extreme case even a PBH, e.g. if lighter than $10^{-11}M_{\odot}$, cannot bend the path of optical light from a star. The HSC/Subaru data of M31 was the first kind of data to realize that the wave optics effect can be important for optical microlensing observation.

Hence the purpose of this paper is to study the effect of wave optics on the optical microlensing search, with particular focus on the Subaru HSC M31 data in Niikura et al. [30]. For comprehensiveness, we also study the effect of finite source size on microlensing [40, 13]¹. The wave effect itself for optical microlensing is quite interesting because it gives a direct evidence of PBH, if the effect is measured from the microlensing light curve even for a single event, because any astrophysical object cannot have such a tiny Schwarzschild radius (their physical size is bigger than the light wavelength). After carefully studying the two effects on microlensing light curve, we will discuss how these affect the microlensing constraints on the abundance of PBHs in such a light mass range. Our study will give a more quantitative study of the results in Niikura et al. [30]. We will also discuss how these constraints can be improved by using a bluer optical data, such as g -band data, than used in Niikura et al. [30], because the wave effect would be smaller in shorter wavelengths.

The structure of this paper is as follows. In Section 7.2 we review the effects of wave optics and finite source size on microlensing in optical wavelengths. In Section 7.3, we discuss the implications of wave effect and finite source size effect on the PBH constraints obtained from the Subaru HSC microlensing search of stars in M31. In Section 7.4 we will give a discussion of how the optical microlensing constraints on PBHs can be improved if a microlensing search based on bluer-filter data (g -band filter) is used. We then give conclusion in Section 11.8. Throughout this paper we adopt the natural units, $c = 1$ (c is the speed of light).

¹Matsunaga & Yamamoto [24] also studied both effects of wave optics and finite source size on gravitational lensing phenomena, but their study is more general, and did not discuss the consequences for optical microlensing search

7.2 The wave optics and finite source size effects on PBH microlensing

In this section, we review the effects of wave optics effect and finite source size on PBH microlensing, following the papers [27, 38, 24]

7.2.1 Microlensing basics for a point source

The characteristic angular scale of microlensing for a star in M31 due to a PBH is the Einstein radius on the sky [32] [also see 36, 23]:

$$\begin{aligned}\theta_E &\equiv \frac{R_E}{d_L} \\ &\simeq 10^{-3} \mu\text{as} \left(\frac{M_{\text{PBH}}}{10^{-10} M_\odot} \right)^{1/2} \left(\frac{d_S}{d_{\text{M31}}} \right)^{-1/2} \left(\frac{1-x}{x} \right)^{1/2},\end{aligned}\quad (7.1)$$

where R_E is the Einstein radius defined as

$$R_E \equiv \sqrt{4GM_{\text{PBH}}d_L \left(1 - \frac{d_L}{d_S} \right)},\quad (7.2)$$

d_L and d_S are distances to a lensing PBH and a source star, respectively, M_{PBH} is the PBH mass, and $x \equiv d_L/d_S$. We assume that a source star is at the distance to M31 for which we assume $d_{\text{M31}} = 770$ kpc, and plugged in $M_{\text{PBH}} = 10^{-10} M_\odot$ into the above equation as a working example. If a source star is closer than θ_E in separation from a lensing PBH on the sky, the source star is multiply imaged by its lensing. However, as implied by the above equation, the separation between multiple images is too small to be resolved by an optical telescope – this phenomena is “microlensing”. What is observed instead is a magnification of the total flux of two images relative to that of the original single image. The microlensing magnification for a point source, under the geometrical optics approximation (i.e. when ignoring the wave optics effect), is given as

$$A_{\text{nw}}^{\text{p}}(u) = \frac{u^2 + 2}{u\sqrt{u^2 + 4}},\quad (7.3)$$

where u is the dimension-less impact parameter between lens and source in units of the Einstein radius, the subscript “nw” denote *no wave effect*, and the superscript “p” denotes *point source*. If $u \leq 1$, a source star is multiply imaged, and the total magnification $A \geq 1.34$ (i.e. the observed flux is brightened by a factor of 1.34 compared to the original brightness), which we will often consider a clear event of microlensing phenomena to be detected in the following.

Because a PBH and a source star have a relative motion with respect to an observer, it causes the observed flux of a source star to vary with observation epoch (time), leading a characteristic light curve of the observed star flux. In this way a microlensing event can be identified from the observation and is distinguishable from other variable stars, as done in many experiments [e.g. 2, 39, 37, 25]. A typical timescale of the microlensing light curve can be estimated from a crossing time of the Einstein radius for a lensing PBH with respect to a source star:

$$\begin{aligned}t_E &\equiv \frac{R_E}{v} \\ &\simeq 9.8 \text{ min} \left(\frac{M_{\text{PBH}}}{10^{-10} M_\odot} \right)^{1/2} \left(\frac{v}{200 \text{ km/s}} \right)^{-1} \\ &\quad \times \left(\frac{d_S}{d_{\text{M31}}} \right)^{1/2} (x(1-x))^{1/2},\end{aligned}\quad (7.4)$$

where v is the relative velocity for a observer-lens-source system [30]. Here we assumed $v = 200$ km/s for a typical velocity of PBH in the halo regions of Milky Way and M31 as implied by their rotation curves (any object following the gravity in the Milky Way and M31 halo regions should have a similar velocity structure). More exactly speaking, the velocity relevant for a microlensing light curve is the velocity component on the two-dimensional plane perpendicular to the line-of-sight direction, and one needs to take into account variations in the velocity component in the plane, as well as the dependence of lens distance. For this reason, the timescale of light curve even for a fixed mass PBH has a wide distribution [30]. Nevertheless, the above equation implies that the light curve timescale is sensitive to a mass of lensing object; a PBH of $10^{-10}M_{\odot}$ mass scales would give a minute timescale, while a lensing object of solar mass scales would give a few months timescale for its light curve, e.g. as shown in MACHO experiment [2]. In other words, if we can identify a *secure* microlensing event that has a very short timescale such as minute timescale, it would be a smoking gun detection of light-mass PBH such as $10^{-10}M_{\odot}$, because any astrophysical processes cannot produce a compact object of such small mass scales.

7.2.2 Effect of wave optics

If the Schwarzschild radius of a lensing object is sufficiently larger than the wavelength of light used in a microlensing observation, the microlensing event is well described by “geometrical optics approximation” [36]. On the other hand, if the Schwarzschild radius becomes comparable with or shorter than the light wavelength, we need to take into account the wave optics effect (interference and diffraction effects) on the lensing magnification. In the following, we briefly review the wave optics effect on microlensing for a point source, following Takahashi & Nakamura [38].

The wave effect is characterized by the parameter “ w ”, defined as

$$\begin{aligned} w &\equiv \frac{8\pi GM_{\text{PBH}}}{\lambda} \\ &= 4\pi \frac{r_{\text{Sch}}}{\lambda} = 5.98 \left(\frac{M_{\text{PBH}}}{10^{-10}M_{\odot}} \right) \left(\frac{\lambda}{6210\text{\AA}} \right)^{-1}, \end{aligned} \quad (7.5)$$

where r_{Sch} is the Schwarzschild radius of lensing PBH, and λ is the characteristic wavelength of light in an observation (here we assumed 6210\AA corresponding to a central wavelength of r -band in the Subaru telescope as our default choice). Once we fix the light wavelength, the wave parameter depends only on the mass of PBH, so the wave effect is independent of the distance of PBH. This is because wave effect is not a geometrical effect which depends on the path of light, but rather a local effect around the lensing object.

The magnification including the wave optics effect, $A_{\text{w}}^{\text{p}}(w, u)$, is given in Schneider et al. [36] [also see 27, 28, 38] as

$$A_{\text{w}}^{\text{p}}(w, u) = \frac{\pi w}{1 - e^{-\pi w}} \left| {}_1F_1 \left(\frac{i}{2}w, 1; \frac{i}{2}wu^2 \right) \right|^2, \quad (7.6)$$

where ${}_1F_1$ is the confluent hypergeometric function. The maximum magnification is realized by setting the impact parameter $u = 0$:

$$A_{\text{w}}^{\text{p}}|_{\text{max}} = A_{\text{w}}^{\text{p}}(w, 0) = \frac{\pi w}{1 - e^{-\pi w}}, \quad (7.7)$$

where the subscript “w” denotes “wave effect”. From this equation we can find that, even if microlensing occurs and if $w \ll 1$, the maximum magnification can be significantly reduced

as $A_w^p|_{\max} \approx 1 + \pi w/2$ ($A_w^p = 1$ means no lensing magnification) compared to $A \rightarrow \infty$ in the geometrical optics approximation case. This is because the gravitational potential induced by the lightest PBH is too weak to bent the path of optical light.

When the wavelength becomes very short compared to the Schwarzschild radius of lensing object, i.e. $w \gg 1$, the magnification under the geometrical optics approximation is realized: $A_w^p(w, u) \rightarrow A_{\text{geo}}^p(w, u)$:

$$A_{\text{geo}}^p(w, u) = \frac{u^2 + 2}{u\sqrt{u^2 + 4}} + \frac{2}{u\sqrt{u^2 + 4}} \sin \left[w \left\{ \frac{1}{2}u\sqrt{u^2 + 4} + \log \left| \frac{\sqrt{u^2 + 4} + u}{\sqrt{u^2 + 4} - u} \right| \right\} \right]. \quad (7.8)$$

The second term of the above equation is a rapidly oscillatory function of u for a fixed w ($w \gg 1$). In practice the second term averages out, i.e. is irrelevant in an actual observation, if we consider a finite source size (see below) or if we consider a finite range of light wavelengths in an optical filter and/or a finite exposure time of an observation. For this reason, the second term is often ignored under the geometrical optics approximation, which becomes equivalent to Eq. (7.3).

Fig. 7.1 compares lensing magnification taking into account wave optics with that of geometrical optics approximation for a point source in M31. When $wu \gtrsim \mathcal{O}(1)$, the two curves agree with each other, which implies that the geometrical optics approximation holds valid. On the other hand, when $wu < \mathcal{O}(1)$, the wave effect becomes significant, leading to less magnification compared to the that of geometrical optics. This means that in the regime of $wu < \mathcal{O}(1)$ the geometrical approximation is not valid. For $wu \ll 1$ as an extreme case, $A_w^p \rightarrow 1$, meaning no magnification due to the wave optics effect.

Fig. 7.2 shows lensing magnification as a function of impact parameter u for PBHs of different mass scales, assuming the optical wavelength $\lambda = 6000 \text{ \AA}$ as in Fig. 7.1. For a microlensing system we are interested in, the impact parameter is time-varying as given by $u = u_{\min} + vt$, where v is the relative velocity of lens-source-observer system (v is considered constant for a microlensing phenomenon of interest), u_{\min} is the minimum separation between lens and source, and t is time from the minimum separation. Therefore the curves in this figure are equivalent to a light curve of microlensing as a function of observation time. As expected from Fig. 7.1, the heavier the lens object is, more rapidly the lensing magnification oscillates with u . This feature plays an important role when considering the effect of finite source size. We also note the maximum lensing magnification has an asymptotic behavior at $u \rightarrow 0$: $A_w^p \rightarrow \text{constant}$, while the geometrical optics approximation (Eq. 7.3) gives $A_{\text{nw}}^p \rightarrow u^{-1}$ at $u \rightarrow 0$ as explicitly shown by the dashed curve.

As obvious from Figs. 7.1 and 7.2, characteristic features of the wave optics effect on PBH microlensing are oscillatory features in the light curve. If we can find any microlensing event, even a single event, that shows such an oscillatory feature in its light curve from an optical observation, it would be a smoking gun evidence of PBH having light mass in the range $M_{\text{PBH}} \simeq [10^{-11}, 10^{-10}]M_{\odot}$, because any astrophysical, compact object cannot have such a tiny Schwarzschild radius comparable with optical light wavelength (the physical size of any compact object should be much greater than the wavelength). However, to detect such an oscillatory feature in the light curve, we need a sufficiently dense cadence observation to well sample the light curve. This requires a careful design of the observation strategy. Here the lower cut $10^{-11}M_{\odot}$ comes from the fact that there is no lensing effect for PBHs with masses below the mass limit, because of too strong wave effect. Hence the wave effect in the optical microlensing is worth to explore from an actual observation.

We have so far assumed a static spectrum of light, and implicitly assumed that light of different paths due to lensing has an interference when taking into account the wave effect.

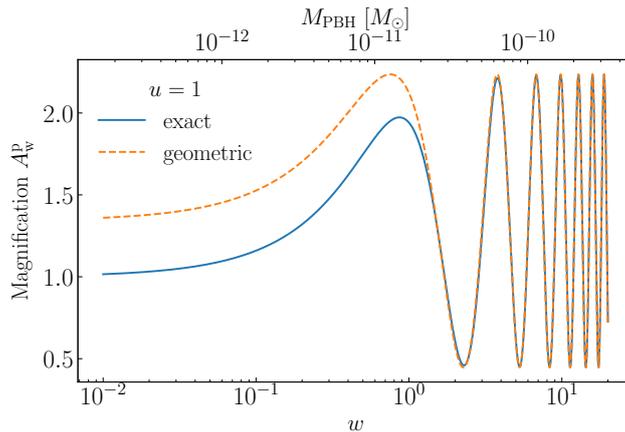


Figure 7.1: Lensing magnification of a point source, A_w^P , as a function of the “wave parameter”, $w \equiv 8\pi M_{\text{PBH}}G/\lambda = 4\pi r_{\text{Sch}}(M_{\text{PBH}})/\lambda$, where M_{PBH} is a PBH mass, r_{Sch} is its Schwarzschild radius, and λ is the wavelength of light used in a microlensing observation. For an optical-wavelength observation, the w parameter in the x -axis can be read as PBH mass scale; the values in the upper x -axis correspond to PBH masses for the r -band filter wavelength ($\lambda = 6000\text{\AA}$). For this plot, we fixed $u = 1$ for the impact parameter in units of the microlensing Einstein radius, $u = b/R_E = 1$; it leads to $A = 1.34$ for the geometrical optics limit denoted by the dashed curve. Although the curve has a rapid oscillation at $w \gg 1$, what is actually observed is the averaged magnification, e.g. over a finite exposure time, which indeed gives $A = 1.34$. The solid curve shows the result taking into account the wave effect. When $w \lesssim 1$, where the wavelength becomes longer than the PBH Schwarzschild radius, the dashed and solid curves start to deviate from each other, meaning that the wave effect becomes significant. It has an asymptotic limit, $A \rightarrow 1$ at $w \ll 1$, i.e. no lensing magnification due to the wave effect.

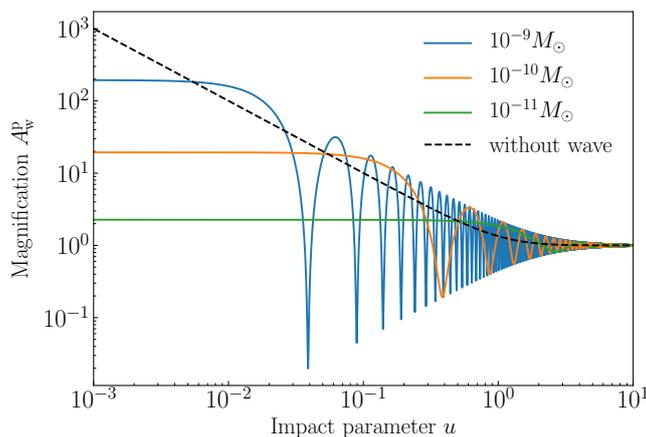


Figure 7.2: Lensing magnification of a point source as a function of the impact parameter u for a lensing PBH of different mass scales. As in Fig. 7.1, we adopted $\lambda = 6000\text{\AA}$ for the light wavelength. Each solid curve is equivalent to a light curve of microlensing, but in an actual observation we need to take into account the average of magnification over a finite exposure time and/or a finite range of wavelengths in a filter, which can be compared to the geometrical optics limit (dashed curve). Each curve becomes flattened at small impact parameters due to the wave effect. The maximum magnification at small impact parameters is lower for lighter-mass PBHs due to more significant wave effect.

Exactly speaking, a source might have a temporal variation in its flux, and this might prevent such an interference. Here we comment on the temporal coherence in the lensing phenomena.

Since light from a star is a black-body radiation that is characterized by its surface temperature (T_S), we can naively consider that a star is a static object, i.e. has no temporal variation in the flux. Nevertheless we consider a possible shortest temporal variation or a possible shift in its light frequency; such a shift might be from light in an absorption line of stellar spectrum, which is caused by a motion of gas in the stellar atmosphere. Compared to this, a continuum spectrum of a star can be safely considered static (i.e. much longer time variation). The shortest timescale of such an absorption line can be estimated by the Doppler effect due to a thermal motion of gas elements in the stellar surface [16], which causes a shift in the frequency of light by $\Delta\nu/\nu \simeq v_{\text{th}}$:

$$\begin{aligned} \tau_c &\sim \frac{1}{\Delta\nu} \\ &\sim \frac{\lambda}{v_{\text{th}}} \simeq 5 \times 10^{-11} \text{ sec} \left(\frac{m}{m_p} \right)^{1/2} \left(\frac{T_S}{6000\text{K}} \right)^{-1/2} \left(\frac{\lambda}{6210\text{\AA}} \right), \end{aligned} \quad (7.9)$$

where v_{th} is the velocity of gas elements in thermal equilibrium for which we assume $mv_{\text{th}}^2/2 = 3T_S/2$, we considered the hydrogen mass (m_p) for gas element for simplicity, we assumed $T_S = 6000$ K for the surface temperature as in the Sun, and λ is a typical wavelength in the r -band. On the other hand, for microlensing, we want to consider an interference or superposition of light rays that arrive at an observe after traveling along the two light paths that are bent by a lensing PBH. The relevant timescale is a ‘‘time delay’’ that refers to as a difference in their arrival time of the two light paths, given in Schneider et al. [36] as

$$\Delta t_{\text{lens}} \simeq 2.0 \times 10^{-15} \text{ sec} \left(\frac{M}{10^{-10} M_\odot} \right) K(u), \quad (7.10)$$

where

$$K(u) \equiv \frac{1}{2} u \sqrt{u^2 + 4} + \log \left| \frac{\sqrt{u^2 + 4} + u}{\sqrt{u^2 + 4} - u} \right|. \quad (7.11)$$

Note $K(u) \sim \mathcal{O}(1)$ for $u < 1$ in the multiple image regime. From Eqs. (7.9) and (7.11), we can find $\Delta t_{\text{lens}} \ll \tau_c$, meaning that the lensing time delay is much shorter than a possible temporal variation of light from a source star even if we consider a light around an absorption line. Thus we can safely consider a temporal coherence of a star light, and in other words we can safely consider an interference of light in the microlensing case (the use of Eq. (7.6) is valid). We again note that the Subaru r -band photometry mainly measures a portion of a continuum spectrum in a black-body radiation of a star, which has much longer temporal variations, even if it has, than the lensing time delay. Thus we can safely consider that Eq. (7.6) is valid for our study.

7.2.3 Effect of finite source size

The effect of finite source size on microlensing is characterized by the ratio of source size to the Einstein radius on the sky [40]. We use the parameter ρ_* to denote the ratio, defined as

$$\begin{aligned} \rho_* &\equiv \frac{\theta_S}{\theta_E} = \frac{R_S/d_S}{R_E/d_L} \\ &\simeq 5.9 \left(\frac{R_S}{R_\odot} \right) \left(\frac{d_S}{d_{\text{M31}}} \right)^{-1/2} \left(\frac{M_{\text{PBH}}}{10^{-10} M_\odot} \right)^{-1/2} \left(\frac{x}{1-x} \right)^{1/2}, \end{aligned} \quad (7.12)$$

where R_S is the finite source size for which we assumed the solar radius. If $\rho_* \gtrsim 1$, the source size effect becomes significant for the microlensing magnification. Even if $\rho_* \sim O(0.1)$, the effect is not negligible as we will show below. Throughout this paper we mainly consider the solar radius, R_\odot , for the source star size following Niikura et al. [30] that use a monitoring observation of (mainly main-sequence) stars in M31 with Subaru/HSC. Even if the physical source size is fixed, the the source size effect varies with the relative position of lensing PBH to an observer, as given by $\rho_* \propto [x/(1-x)]^{1/2}$; for a smaller x (a PBH closer to an observer), ρ_* becomes smaller. This means that the source size effect is stronger when a PBH is lighter or if a PBH is nearer to the source.

To model the finite source size effect, we assume that a source star is a circular disk with constant surface brightness. This is a simplified assumption, but is sufficient for our purpose. Assuming that the source center is separated from a lensing PBH by u (i.e. the impact parameter between the source center and lens), we can compute the lensing magnification for such a disk-shaped star from an average of the lensing magnification (Eq. 7.3) over the source disk:

$$A^f(u, \rho_*) \equiv \frac{1}{\pi\rho_*^2} \int_{|\mathbf{y}| \leq \rho_*} d^2\mathbf{y} A_{\text{nw}}^p(|\mathbf{u} - \mathbf{y}|), \quad (7.13)$$

where the integral variable \mathbf{y} moves within the circular source disk of radius ρ_* , we set the origin ($\mathbf{y} \equiv \mathbf{0}$) to the source center, the superscript ‘‘f’’ in A^f denotes the ‘‘finite source size effect’’, and we set a lensing PBH to be located at $\mathbf{u} = (u, 0)$ due to symmetry of the circular disk star. Note $\int_{|\mathbf{y}| \in \rho_*} d^2\mathbf{y} = \pi\rho_*^2$. When further taking into account the wave optics effect, we replace A_{nw}^p in the above equation with A_w^p (Eq. 7.6) and then perform the numerical integration to obtain the magnification for a given set of parameters (u, ρ_*, w) .

Fig. 7.3 shows lensing magnification as a function of the impact parameter for different input values of ρ_* corresponding to different physical source sizes for a fixed PBH mass, $M_{\text{PBH}} = 10^{-10}M_\odot$. When taking into account the finite source size, the lensing magnification becomes flattened at separations, $u \lesssim \rho_*$. For $\rho_* = 0.1$, the finite source size effect is not significant, but not negligible. On the other hand, the finite source size effect smears out oscillatory features in the light curve, because different points in a source (disk-like shape source) have different phases in the wave optics effects of lensing magnification, and the oscillatory features are averaged out when integrating the lensing magnification over the source region. Thus the finite source size makes it difficult to extract an oscillatory feature in the microlensing light curve. Nevertheless we should note that the finite source size effect depends on a lens distance; the effect becomes relatively less important for a lensing PBH closer to an observer, while the wave effect is independent of lens distance. This difference might help distinguish the wave effect from an observed microlensing light curve.

7.2.4 Effect of exposure time average

For an actual observation of the PBH microlensing search, we need to further take into account the effect of finite exposure time. Here an ‘‘exposure’’ time means a time duration during which photons from a source star are collected and then the total number of photons (more exactly electrons converted from the photons) in one exposure is stored into the hard drive after closing a shutter of the camera. In the following we denote the exposure time as t_{exp} .

We throughout this paper adopt $t_{\text{exp}} = 90$ sec as our fiducial parameter following the observation in Niikura et al. [30]. The impact parameter of a given microlensing event would

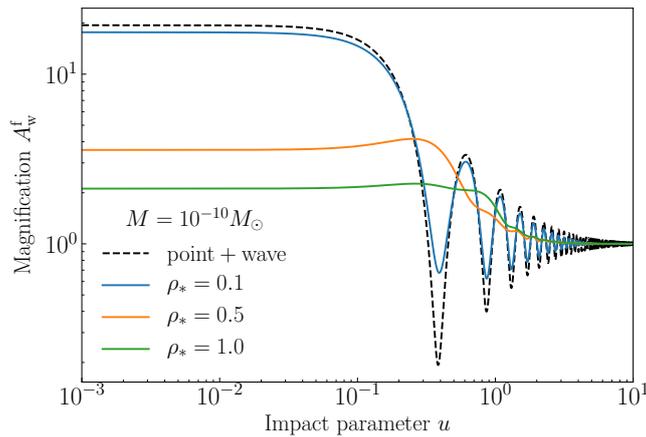


Figure 7.3: As in Fig. 7.2, but the solid curves show the lensing magnification when further taking into account the finite source size effect in addition to the wave effect. Note that the plotting range of y -axis is different in this plot from that in Fig. 7.2. We consider several cases for the source size as given in the legend. The source size is defined in units of the Einstein radius of a lensing PBH: e.g., $\rho_* = 1$ corresponds to $\theta_S = \theta_E$ on the sky.

be changed during the exposure time by an amount of

$$\begin{aligned} \Delta u &\simeq \frac{vt_{\text{exp}}}{R_E} \\ &= 0.15 \left(\frac{v}{200 \text{ km/s}} \right) \left(\frac{t_{\text{exp}}}{90 \text{ sec}} \right) \\ &\quad \times \left(\frac{M_{\text{PBH}}}{10^{-10} M_\odot} \right)^{-1/2} \left(\frac{d_S}{d_{\text{M31}}} \right)^{-1/2} (x(1-x))^{-1/2}, \end{aligned} \quad (7.14)$$

Thus what we can observe is the microlensing magnification averaged over the exposure time or equivalently the interval Δu around the fiducial point of u . Because of $\Delta u \propto v/R_E$ for a fixed exposure time, the averaging effect is more significant when v is greater or R_E is smaller. Hence the lensing magnification including the effect of a finite exposure time can be estimated as

$$\bar{A}(u) = \frac{1}{\Delta u} \int_{u-\Delta u/2}^{u+\Delta u/2} du' A(u'). \quad (7.15)$$

More exactly speaking, however, the averaged magnification needs to be computed from a given trajectory around the fiducial point u . In the above equation, we simply assume that the trajectory is along the path towards the center of source star (i.e. the path with the minimum impact parameter $u_{\text{min}} = 0$). This is a simplified treatment, but is enough for our purpose. We will use this equation to estimate the impact of finite exposure time on our results.

Now we come back to a question of whether the wave effect on the microlensing light curve can be measured by a dense cadence of an observation such as that done in Niikura et al. [30], where an imaging data of M31 was taken every 2 min (90 sec for exposure plus about 30 sec readout of the data). Fig. 7.4 shows a simulated light curve for a typical case of microlensing for such an observation. As can be found, an oscillatory feature due to the wave effect is difficult to be captured because of the insufficient cadence.

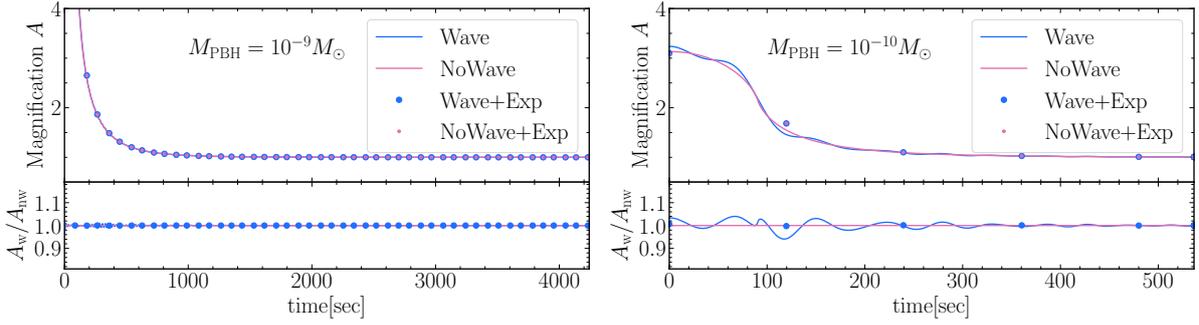


Figure 7.4: A simulated light curve for a microlensing due to a PBH with $M_{\text{PBH}} = 10^{-9}M_{\odot}$ (left panel) and $M_{\text{PBH}} = 10^{-10}M_{\odot}$ (right panel), for a star in M31. Here we consider a case that a lensing PBH is at distance $d_L = 10$ kpc, has a perpendicular velocity of 100 km/s, and assume an observed wavelength of $\lambda = 6210\text{\AA}$ in the Subaru r -band and a solar radius for a source star size, $R_S = R_{\odot}$, as a working example. In this case the wave effect parameter $w = 59$ (5.9) and the finite source size relative to the Einstein radius $\rho_* = 0.21$ (0.66) for the case of $M_{\text{PBH}} = 10^{-9}M_{\odot}$ ($10^{-10}M_{\odot}$). Furthermore, we consider the discrete sampling effect in a measurement of the light curve; we here assume a 2 min cadence meaning that the light curve is sampled every 2 min as denoted by circle points. We here compare the results with and without the wave effect. The lower panels show the ratio. An oscillatory feature in the light curve due to the wave effect is difficult to be captured because of the finite source size effect for $M_{\text{PBH}} = 10^{-9}M_{\odot}$ and exposure time average effect for $10^{-10}M_{\odot}$, respectively.

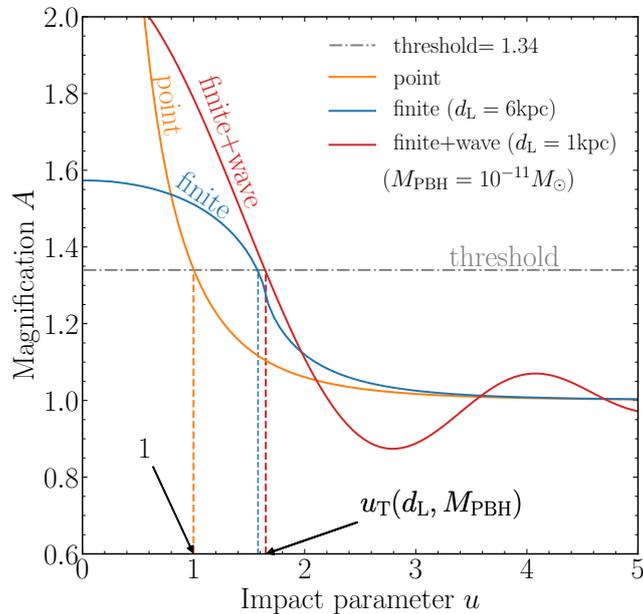


Figure 7.5: Comparison of the microlensing light curves for a point source or when taking into account the finite source size effect and/or the wave optics effect. Here we consider $M_{\text{PBH}} = 10^{-11}M_{\odot}$ for the PBH mass. The vertical line denotes $A = 1.34$, which corresponds to a nominal threshold of lensing magnification leading to a detection of the event in an observation. For this, the vertical lines gives the impact parameter threshold, u_T , corresponding to the magnification threshold for each case. The lensing magnification $A \geq 1.34$ at $u \leq u_T$.

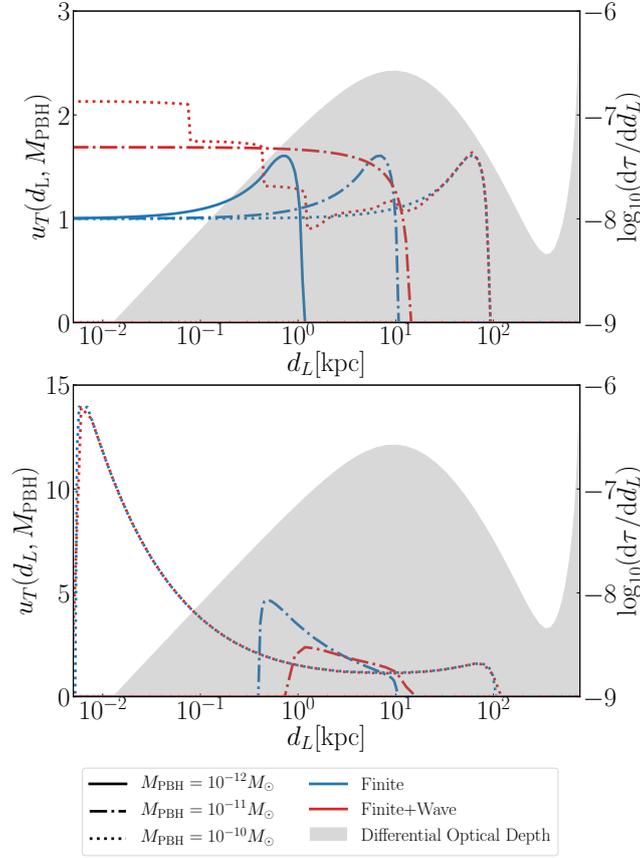


Figure 7.6: The impact parameter threshold u_T as a function of PBH distance, where the lensing magnification for a source star in M31 becomes greater than the detection threshold, $A > 1.34$ when $u < u_T$ as discussed in the previous figure. Note $u_T = 1$ for a point source independently of the lens distance. The different curves, as indicated by the legend at bottom, show the results when taking into account the finite source size effect and/or the wave optics effect, for PBHs of different mass scales as indicated. When PBHs are too light, no PBH can cause a detectable microlensing due to the source size and wave effects. This is the case for $M_{\text{PBH}} = 10^{-12} M_{\odot}$ ($u_T = 0$ for this case). The shaded region shows the differential contribution to the microlensing optical depth showing how PBHs at each distance contribute the total optical depth (the integrand of Eq. 7.16). The difference between the upper and lower plots is that, for the lower panel, we include the effect of finite exposure time, 90 sec here; when computing a more realistic light curve of microlensing, we compute the lensing magnification averaged over the exposure time assuming a typical velocity of lensing PBH with respect to a source star at each distance (for which we assume a velocity expected from the velocity dispersion of DM halo at each distance). When PBHs become too light, the light curve has a rapid oscillatory feature as a function of separation between lens PBH and source star as indicated in the previous figure, so it causes a more significant average of the lensing magnification.

7.3 Implications on the PBH constraints from the Subaru HSC microlensing search

In this section we study the impact of wave optics effect and finite source size effect on a search of microlensing for source stars in M31 due to intervening PBHs that exist between the Earth and M31, if PBHs make up DM in the Milky Way and M31 halo regions, by a certain mass fraction. To do this we employ the parameters in Niikura et al. [30] to model the spatial and velocity distributions of DM including PBHs in the Milky Way and M31 regions.

The optical depth of microlensing for a single source star in M31 due to PBHs is given as

$$\tau_{\text{PBH}} = \frac{\Omega_{\text{PBH}}}{\Omega_{\text{DM}}} \int_0^{d_s} dd_L \frac{\rho_{\text{DM}}(d_L)}{M_{\text{PBH}}} \pi R_E^2, \quad (7.16)$$

where a factor $\Omega_{\text{PBH}}/\Omega_{\text{DM}}$ gives the mass fraction of PBHs to DM that exists along the line-of-sight direction in the Milky Way and M31 regions, and $\rho_{\text{DM}}(d_L)$ is the mass density distribution of dark matter along the line-of-sight direction up to the source star. Here we assume a single source plane; i.e. we assume that all source stars in M31 are at the same distance, which is a good approximation because the spatial extent of stellar distribution in M31 (~ 10 kpc) is very small compared to the distance of M31 (770 kpc). To model the spatial distribution of DM, we employ Navarro-Frenk-White (NFW) models [29, 22] to reproduce the flat rotation curve for each of the Milky Way or M31, respectively; we compute the total DM distribution $\rho_{\text{DM}}(r)$ by a sum of the contributions of two NFW profiles at the distance of a lens from the Milky Way center and the M31 center, respectively [see 30, for details]. Here, we simply included ‘‘microlensing’’ events in consideration if lensing PBH and source star are separated by less than the Einstein radius on the sky. Since $R_E^2 \propto M_{\text{PBH}}$, the above optical depth does not depend on PBH mass. In the following we will more carefully consider how the detection threshold for a microlensing observation could be changed if taking into account the finite source size effect and/or wave optics effect.

The differential event rate of PBH microlensing for a single source star in M31 is defined [see Eq. 18 in 30] as

$$\begin{aligned} \frac{d\Gamma_{\text{PBH}}}{d\hat{t}} = & 2 \frac{\Omega_{\text{PBH}}}{\Omega_{\text{DM}}} \int_0^{d_s} dd_L \int_0^{u_T(d_L, M_{\text{PBH}})} du_{\text{min}} \\ & \times \frac{1}{\sqrt{u_T^2(d_L, M_{\text{PBH}}) - u_{\text{min}}^2}} \frac{\rho_{\text{DM}}(d_L)}{M_{\text{PBH}} v_c^2(d_L)} v^4 \exp \left[-\frac{v^2}{v_c^2(d_L)} \right], \end{aligned} \quad (7.17)$$

where $v = 2R_E \sqrt{u_T^2 - u_{\text{min}}^2}/\hat{t}$. The units of $d\Gamma_{\text{PBH}}/d\hat{t}$ is [events/hours/hours] giving the event rate for a single star per unit observation time [hours] per unit light curve timescale (\hat{t}) [hours]. For the velocity distribution of DM, therefore PBHs, we assume the isotropic, random virial motion according to the NFW halo at a given radius of lensing PBH from the Milky Way center or M31 center.

The integration with respect to the minimum impact parameter u_{min} in Eq. (7.17) is in the range $u_{\text{min}} = [0, u_T]$, where u_T is the threshold impact parameter defined as follows. We observationally detect a microlensing event of PBH if the magnification of a source star is large enough and then the light curve of a source star is detected under observational conditions. More exactly speaking, to estimate the detection sensitivity of microlensing events, it requires detailed simulations of microlensing light curve taking into account observation conditions as well as combinations of model parameters, as done in Niikura et al. [30]. This is beyond the scope of this paper. Here we simply assume that a microlensing event can be detected if the maximum magnification is greater than a threshold value $A_T = 1.34$, which corresponds to the magnification at the impact parameter $u = 1$ for a point source under geometrical optics

approximation (Eq. 7.3). Even if we consider the wave optics effect and/or the finite source size effect, we can determine the threshold impact parameter, u_T , corresponding to $A_T = 1.34$, once the model parameters (distance to PBH, PBH mass, source size, and optical wavelength) are fixed. If $0 \leq u_{\min} \leq u_T$, the microlensing event has a magnification with $A \geq 1.34$. The integration over u_{\min} in Eq. (7.17) includes microlensing events with $A \geq 1.34$.

Fig. 7.5 illustrates how the threshold impact parameter u_T is changed when taking into account the finite source size effect and the wave optics effect, assuming typical values of the model parameters. For this particular set of the model parameters, the finite source size effect and the wave effect both increase the threshold value of the impact parameter u_T compared to the point source and geometrical optics case ($u_T = 1$). Hence the effects increase a cross section of the microlensing. The oscillatory feature in the light curve appears at $u > u_T$ or $A < A_T$ for this case. For other parameters, such oscillatory feature can appear at $u < u_T$. In this way we can estimate the numeric value of u_T as a function of the model parameters.

In Fig. 7.6 we study how the threshold impact parameter varies as a function of the distance to PBH for different PBH mass scales as well as the dependence on the finite source size effect and/or the wave optics effect. For a point source, $u_T = 1$ independently of the lens distance. First of all, when taking into account the finite source size effect and/or the wave optics effect, only PBHs at particular distances can contribute the microlensing events. If PBH is too light such as $M_{\text{PBH}} \lesssim 10^{-11} M_\odot$ as we will show below, PBHs can not cause a detectable microlensing event. Such an undetectability of lightest PBH leads to a cutoff in the PBH abundance constraint at mass scales below the critical mass $M_{\text{cut}} \lesssim 3.3 \times 10^{-12} M_\odot$ ($2.5 \times 10^{-12} M_\odot$) when we use a microlensing observation of r -band(g -band). The step-like feature in the result for $M_{\text{PBH}} = 10^{-10} M_\odot$ arises from the oscillatory feature in the light curve when taking into account the wave effect, if u_T is determined according to the method in Fig. 7.5. However, as shown in the lower panel, a finite exposure time has a significant impact on the threshold calculation, where we used Eq. (7.15) to estimate the effect; the finite exposure time (90 sec) averages out the oscillatory feature in the light curve, leading to no detectable microlensing event, i.e. $u_T \rightarrow 0$, if a lensing PBH is closer to an observer². Even if u_T appears to have a greater value for PBHs at $d_L \lesssim 10^{-1}$ kpc for $M_{\text{PBH}} = 10^{-10} M_\odot$, the contribution of PBHs at such small distances to the microlensing event rate is very small as indicated by the gray-shaded region, and therefore the contribution to the event rate is negligible.

Fig. 7.7 is the main result of this paper. Here we use the microlensing search results from one-night Subaru HSC r -filter data of M31 in Niikura et al. [30], where they found only one possible candidate of PBH microlensing compared to many expected events if all DM is made up of PBHs. The results were translated into an upper bound on the abundance of PBHs at each mass scale assuming a monochromatic mass spectrum for PBHs. Fig. 7.7 compares the results for three cases: i) if we assume a point source and ignore the wave optics effect (most optimistic case), ii) if we take into account the finite source size effect, and iii) if we take into account both the finite source size effect and the wave optics effect. Here we assume the solar radius for source size for all stars in M31, and include the “selection” function of microlensing events that give a probability to recover microlensing events of a given light curve timescale under the Subaru HSC observation conditions (taken from Figure 18 in Niikura et al.). The figure shows almost no difference between the results of cases ii) and iii), meaning that the finite source size effect is a dominant effect for low-mass PBHs in $M_{\text{PBH}} = [10^{-11}, 10^{-7}] M_\odot$. In other words, the wave effect is very difficult to distinguish in this kind of statistical study of microlensing search based on the r -band data.

²In this case, the Einstein crossing time (Eq. 7.4) becomes shorter for a typical velocity of PBH, so the exposure time average smears out the light curve, leading to a smaller net magnification.

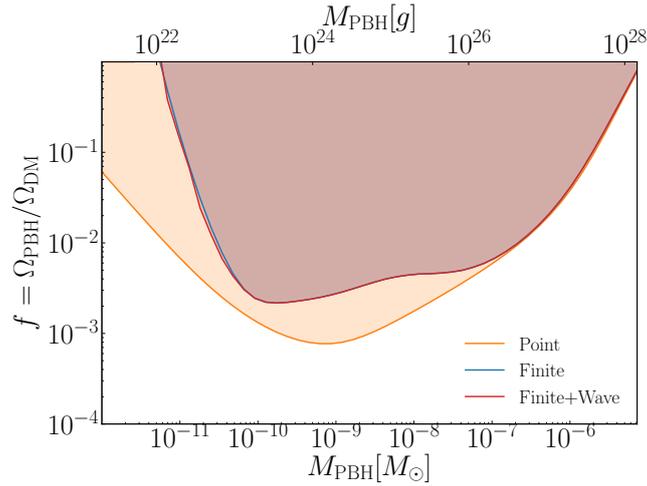


Figure 7.7: The impact of the finite source size effect and the wave effect on the 95% C.L. upper bound on the PBH mass fraction to DM in the halo regions of the Milky Way and M31 in Niikura et al. [30], which are based on the r -band filter data of M31 taken with the Subaru HSC. Note that the r -band filter has 6210\AA for the central wavelength, and they used 8.7×10^7 stars for the constraints. Compared to the point source result, the source size and wave effects cause a sharp cut at $M_{\text{PBH}} \lesssim 3.3 \times 10^{-12} M_{\odot}$, i.e. no constraint for such light-mass PBHs, as denoted by the vertical line. For the finite source size effect we assume the solar radius for all source stars.

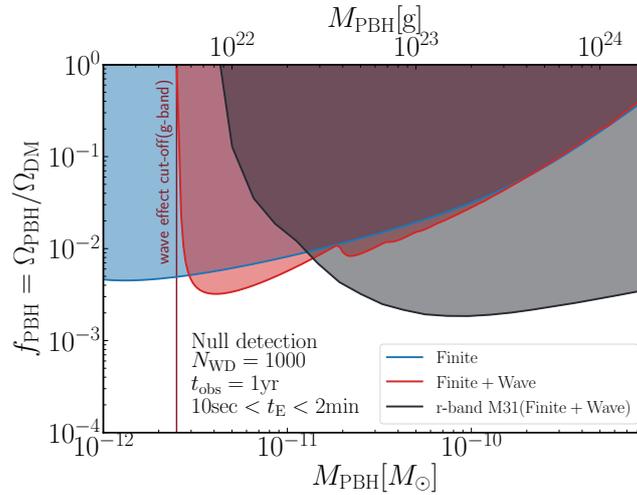


Figure 7.8: An expected constraint on PBH abundance for a hypothetical dense-cadence monitoring observation of white dwarfs in a g -band filter (centered at $\lambda = 4730\text{\AA}$). Here we assume that we can have a monitoring observation for 1000 white dwarfs in LMC, that all white dwarfs have the same radius of $0.01R_{\odot}$, that we can perfectly recover the microlensing light curve in the timescale range of $10 \text{ sec} \leq t_E \leq 2 \text{ min}$, if occurs, and $t_{\text{obs}} = 1 \text{ year}$ for the net, total observation time (see text for details). The red curve shows the expected 95% C.L. upper limit on the PBH abundance, if no lensing event is detected (i.e. null detection), when taking into account both the wave optics and finite source size effects. The blue curve shows the result when including the finite source size effect alone. Hence the cutoff mass scale at $M_{\text{PBH}} \simeq 2.5 \times 10^{-12} M_{\odot}$, as denoted by the vertical line, is due to the wave optics effect for the g -band observation. For comparison, the gray curve shows the red curve in Fig. 7.7.

7.4 Discussion

Fig. 7.7 clearly shows a fundamental limitation of the PBH constraint from the Subaru r -band data; no constraint on PBHs at $M \lesssim 10^{-11} M_{\odot}$. How can we improve the microlensing constraints on PBH abundance from an optical observation? One possibility is to use a shorter-wavelength filter for the microlensing search because the ratio of PBH's Schwarzschild radius to the light wavelength becomes greater, which reduces the impact of wave optics effect. For example, if we can use a g -band filter whose central wavelength ($\lambda = 4730\text{\AA}$) is shorter than that of r -band filter ($\lambda = 6210\text{\AA}$) we have considered, we could still monitor many stars. In particular, white dwarfs are bluer than main sequence stars, and have a smaller radius which is typically smaller than the solar radius by a factor of 100. This reduces the impact of finite source size effect. However, the abundance of white dwarfs is smaller than that of main sequence stars, by a factor of 10 as implied from the MOA or OGLE microlensing events for stars in the Large Magellanic Cloud (LMC) or the Galactic bulge region [37, 25, 31]. In addition, white dwarfs are much fainter than main sequence stars (by more than 5 magnitudes than main sequence stars³). So a large-aperture telescope is needed, or we need to monitor white dwarfs at closer distances than in M31, e.g. white dwarfs in the Galactic bulge or LMC so that white dwarfs are bright enough to be detected by a reasonably large-aperture telescope such as the 8.2m Subaru telescope.

To explore microlensing event(s) due to PBHs in the mass range $M_{\text{PBH}} \lesssim 10^{-11} M_{\odot}$, we need a much denser-cadence observation of white dwarfs than done in the Subaru r -band 2 min-cadence observation of Niikura et al. [30] (90 sec exposure plus 30 sec exposure), in order to well sample a microlensing light curve of much shorter timescale than 2 min. The upcoming Large Synoptic Survey Telescope (LSST)⁴ will allow for 2 sec readout time that is much shorter than 30 sec of the Subaru HSC. Furthermore, if a large-format CMOS image sensor is available for a large aperture telescope, it will allow a much faster readout imaging of source stars including white dwarfs. Here we assume that we can have a monitoring observation for a sample of 1000 white dwarfs, compared to 8.7×10^7 stars in M31 in Niikura et al. [30], using a sufficiently dense cadence observation with a g -band filter at a large aperture telescope. We assume that the 1000 white dwarfs are at 50 kpc in distance (e.g., LMC) so that those are bright enough, and that all the white dwarfs have the same radius, $R_S = 0.01 R_{\odot}$ for simplicity. Then we assume that we can well sample a light curve of microlensing, if occurs, in the range of $10 \text{ sec} \leq t_E \leq 2 \text{ min}$ for the typical light curve timescale. More exactly we assume that we can perfectly recover a microlensing event of the timescale range, if occurs. Finally we assume we have a 1-year amount of data for the net observation time. Such a telescope/detector would be interesting to explore, and seems feasible within next 10 years or so. Here we simply use the same models in Niikura et al. [30] to study forecasts for such a microlensing search (so we assume the white dwarf sample is in the direction to M31). In addition, a microlensing of white dwarfs would be more suitable to search for the wave optics effect due to the relatively weaker impact of the finite source size effect.

Fig. 7.8 shows the expected upper bound on the PBH abundance from the g -band microlensing search we described above, assuming no secure microlensing event (i.e. null detection). Compared to Fig. 7.7, such a g -band dense-cadence monitoring observation can constrain PBHs at lighter mass scales, down to a few times $10^{-12} M_{\odot}$. The cut off at the low PBH mass is from the wave optics effect; PBHs in the lighter mass scales cannot cause microlensing. This result can be understood as follows. The expected number of microlensing events is roughly given as $N_{\text{exp}} \sim N_S \times \tau \times (t_{\text{obs}}/t_E)$, where N_S is the number of source stars (here white dwarfs), τ is the optical depth of PBH microlensing for a single source star, and t_{obs} is the net observation

³E.g., see <http://sci.esa.int/gaia/60209-white-dwarfs-in-gaia-s-hertzsprung-russell-diagram/>

⁴<https://www.lsst.org>

time. The optical depth $\tau \simeq 10^{-6}$ for the Milky Way halo if PBHs make up all DM in the Milky Way halo region. Hence, if we search for a microlensing event of 10 sec timescale, the expected number is $N_{\text{exp}} \sim 10^3 \times 10^{-6} \times (3 \times 10^7 \text{ sec}) \times (10 \text{ sec})^{-1} \sim 3000$. If we cannot find any event, i.e. null detection, it gives an upper limit that such PBHs are not allowed by more than $1/3000 = 0.003$ for the mass fraction of PBHs to DM, which roughly reproduces the result in Fig. 7.8. More quantitatively, in the figure we took into account the distribution of light curve timescales due to the velocity distribution of PBHs in the Milky Way halo region.

7.5 Conclusion

In this paper, we studied the effects of wave optics and finite source size on the optical microlensing search for stars in M31 due to PBHs that would exist in the Milky Way and M31 halos if DM is made up of PBHs. If PBH is in mass scales of $M_{\text{PBH}} \lesssim 10^{-10} M_{\odot}$, its Schwarzschild radius (r_{Sch}) becomes comparable with or shorter than the optical filter wavelength (e.g. the r -band filter, centered at 6210\AA), and the wave effect on microlensing needs to be considered for such PBHs. For PBHs with $r_{\text{Sch}} \ll \lambda$, even PBHs cannot bend the path of optical light, so causes no microlensing magnification, which gives the fundamental limit for an optical microlensing search ($M_{\text{PBH}} \lesssim 10^{-11} M_{\odot}$ for the optical r -band filter). Nevertheless, if we can find a secure microlensing event for PBH with $\sim 10^{-10} M_{\odot}$ in an optical filter observation, it is a smoking-gun evidence of PBHs because any other astrophysical object cannot have such a tiny Schwarzschild radius (their physical size is much larger than the optical wavelength). This is an interesting possibility to explore from an actual observation. However, we showed that the finite source size effect is equally important or even more significant, if the microlensing observation targets main-sequence stars, and likely erases characteristic signatures of the wave effect in the microlensing light curve. As the main result of this paper, we studied the impact of wave effect and finite source size effect on the PBH constraints obtained from the Subaru HSC microlensing search for stars in M31 in Niikura et al. [30]. We showed that the effects are significant, and the finite source size effect is a dominant effect compared to the wave effect, if source stars have a size of the solar radius as expected for main sequence stars (see Fig. 7.7). Nevertheless, upcoming wide-area surveys such as LSST, Euclid and WFIRST would be very powerful to search for microlensing events due to PBHs over a wide range of mass scales, so the results shown in this paper are relevant for microlensing search from the upcoming surveys.

In order to “detect” the wave optics effect of microlensing in an optical observation, we need to use a source star which has a smaller size. Such a source is white dwarf, which has a smaller size (typical a few % the solar radius), and can be observed in optical wavelengths. However, the number of white dwarfs are smaller than that of main sequence stars, by a factor of 10, and white dwarfs are much fainter than main sequence stars, by more than 5 magnitudes (a factor of 100 in the flux). Hence the microlensing search for white dwarf sources is challenging, but it is not impossible. For example, a CMOS sensor detector would be promising, compared to CCD camera, because it allows a much denser-cadence monitoring observation of source stars and therefore enables to search for ultrashort timescale microlensing events (the readout time of Subaru HSC is 30 sec). We discussed that such a monitoring observation in g -band filter can improve the PBH abundance in the mass scales, $M \lesssim 10^{-11} M_{\odot}$, and therefore such a CMOS camera with large field-of-view at a large-aperture telescope would be powerful to further explore PBH signatures.

8 Wave optics for rotating stars by Béatrice Bonga, Feldbrugge, Metidieri

Abstract

Gravitational lensing in wave optics is a rich field combining caustic singularities, general relativity and interference phenomena. We present a detailed evaluation of wave optics effects resulting from the frame-dragging of a rotating star modeled by a Lense-Thirring spacetime. We demonstrate that, contrary to what was previously stated in the literature (see e.g. Ref. [1]), the spin of the star leaves an intricate imprint on the interference fringes and the caustics of the lensed source. This interference pattern can in principle be used to directly measure the spin of the lens.

8.1 Introduction

Einstein's theory of gravity has reshaped our understanding of the Universe, both at large scales and about the objects within it, ranging from black holes and gravitational waves to gravitational lensing. Each phenomenon is of great interest and can potentially reveal new physics. In this paper, we reconsider gravitational lensing in the wave optics regime by a rotating star.

Over the last decades, tremendous progress has been made in the theory and observation of gravitational lensing in astronomy using the geometric optics approximation, treating light as rays following null geodesics in curved spacetime. The geometric approximation suffices for most astrophysical settings. However, wave optics are important when you have coherent (nearly) monochromatic waves. In such cases, as infinitely many rays with nearly equal path lengths converge at the observer, the resulting phenomena of diffraction and interference invalidate geometrical optics. With the advent of numerous Fast Radio Bursts and pulsar observations [2], and the first gravitational waves detections [3], all emitting coherent long wavelength radiation, it may become possible to observe such interference effects due to gravitational lensing for the first time. This is particularly important in the vicinity of caustics, where the intensity is amplified and the geometric approximation breaks down.

The wave nature of radiation in gravitational lensing was first formulated by [2, 1]. Recently, much progress has been made on the numerical evaluation of the interference patterns resulting from lensing in wave optics using both Fourier [25], and more general complex analysis methods [6, 7, 8, 9, 10]. For a path integral analysis of gravitational lensing in wave optics see [11].

The effect of rotation of the lens is clearly also important. This is already evident in the geometrics optics limit. As the star rotates, radiation traveling in the direction of rotation of the object will move past the massive object faster than radiation moving against the rotation, as seen by a distant observer. This is known as frame-dragging or the Lense-Thirring effect [12, 13]. Relativistic frame-dragging was recently observed in relativistic jets [14] and in a pulsar white dwarf binary system [15]. Moreover, in the coming years, the Lense-Thirring effect on the orbits of the S2 star around the supermassive black hole in the center of our Milky Way may be detected with the GRAVITY instrument of the Very Large Telescope [16].

The simplest situation in which wave optics are important is the case in which the lens is modelled as a point mass and the source, lens and observer are all far away from each other, so that the thin lens approximation applies. This case has been extensively studied in the weak field regime in which (post-)Newtonian theory applies. Remarkably, the intensity pattern in this case can be evaluated completely analytically [2]. Generalizations of this scenario include a singular isothermal sphere lens [17, 18, 19] and a replacement of a single point mass by a binary [6].

In this article, we investigate the effect of a rotating lens. We account for the spin by using the Lens-Thirring metric, *i.e.*, the slow-spin approximation of the Kerr metric. This was previously studied in [1], in which the authors concluded that the interference pattern of a rotating star has the same *shape* as for the non-rotating lens, but that the wave pattern is shifted *translationally* in the direction perpendicular to the angular momentum vector projected onto the lens plane. Other papers using similar arguments also concluded that the interference pattern of the rotating lens is degenerate with that of the non-rotating lens up to this translational shift [20, 21, 22]. As a result, it was argued that a rotating lens is not distinguishable from a non-rotating one (unless one has an independent way of knowing the precise location of the source). We find that this is in fact incorrect: rotation of the lensing object does change the interference pattern beyond simply shifting it. Earlier results were based on a seemingly innocent mathematical transformation that is in fact highly singular. By not using this transformation, we find that the phenomenology of the rotating lens is much richer than merely shifting the interference pattern. Hence, in principle, the rotating lens is not degenerate with the non-rotating one and one can determine how fast the lens is rotating based on its interference pattern.

The outline of this paper is as follows. In Sec. 8.2, we introduce the setup and describe our methods for evaluating the interference pattern. The results are discussed in Sec. 8.3, where we show explicit interference patterns and contrast these with the shifted interference pattern as argued for in [1, 20, 21, 22]. In App. 8.5, we demonstrate explicitly how sensitive critical and caustic curves are to small perturbations. App. 8.6 discusses the small spin approximation and some of its limitations.

8.2 Rotating stars

In this section, we provide a lightning review of the key quantity relevant for wave optics: the Kirchhoff-Fresnel integral with its integrand determined by the time-delay function. Next, we sketch the derivation of the time delay function for a rotating star, evaluate the caustics, and propose a method to efficiently evaluate the resulting Kirchhoff-Fresnel integral.

In the presence of a lens, the wave amplitude of a point source assumes the form of a Kirchhoff-Fresnel integral

$$\Psi(\mathbf{y}) = \left(\frac{w}{2\pi i}\right)^{d/2} \int_L e^{iwT(\mathbf{x},\mathbf{y})} d\mathbf{x}, \quad (8.1)$$

with the angular frequency of the radiation w , the point \mathbf{x} on the d -dimensional lens plane, the relative position of the observer and source \mathbf{y} in the image plane and the time delay function

$$T(\mathbf{x}, \mathbf{y}) = \frac{(\mathbf{x} - \mathbf{y})^2}{2} - \varphi(\mathbf{x}), \quad (8.2)$$

governed by a geometric term $(\mathbf{x} - \mathbf{y})^2/2$ and the phase variation φ modeling the effect of the lens. The integral in Eq. (8.1) can be interpreted as the superposition of all possible rays propagating from the source to the observer and intersecting the lens plane at \mathbf{x} weighted by the phase e^{iwT} [23]. Note that the interference integral is given in dimensionless units, w , \mathbf{x} , and \mathbf{y} and that the amplitude is normalized with respect to the unlensed case, *i.e.*, when $\varphi(\mathbf{x}) = 0$, the amplitude $\Psi(\mathbf{y}) = 1$. For more details, see [24, 2].

8.2.1 The time delay

The time delay function for a rotating star (see Eq. 8.4) has appeared in the literature before [25, 26, 27, 1], here we review its derivation. A busy reader may skip this subsection.

Gravitational lensing is governed by the time as measured by the observer that light rays take to travel from the source to the observer while following geodesics in the curved spacetime. The source is modelled by a point source that emits spherical, monochromatic waves with frequency ω . The emitted light then propagates past a rotating gravitational lens before reaching a distant observer. The space-time metric of a rotating star in the weak field limit is

$$ds^2 = - (1 + 2U(\mathbf{r})) c^2 dt^2 + (1 - 2U(\mathbf{r})) d\mathbf{r}^2 - \frac{4G\varepsilon_{jkm} J^k x^m dx^j dt}{c^2 \|\mathbf{r} - \mathbf{r}^*\|^3}, \quad (8.3)$$

which can be interpreted as the Minkowski metric, plus a small disturbance due to the gravitational potential $U(\mathbf{r}) = -GM/(c^2 \|\mathbf{r} - \mathbf{r}^*\|)$ of the star of mass M located at \mathbf{r}^* , and a frame-dragging term resulting from the spin \mathbf{J} of the lens.

When the particle's deflection occurs over a relatively small region, relative to the distance from the source to the observer, one can use the thin lens approximation. In this approximation, the time delay function consists of three components modulo some constant terms that do not affect the wave amplitude. First, the distance of the lensed path exceeds that of a straight line from the source to the observer by $\frac{1}{2} \frac{D_{OL} D_{OS}}{D_{OS} - D_{OL}} \|\boldsymbol{\theta} - \boldsymbol{\theta}_S\|^2$ with the angular diameter distances from the observer to the lens D_{OL} and the observer to the source D_{OS} respectively. $\boldsymbol{\theta}_S$ stands for the angular position of the source as measured from the line of sight, and $\boldsymbol{\theta}$ is the angular position of the ray crossing the lens plane $\boldsymbol{\theta}$ with respect to the position of the lens on the sky (see Fig. 8.1 for a comprehensive schematic of the setup). Second, the line integral of the potential $\int U(\mathbf{r}) dr$ for small angles assumes the form $4GM/c^2 \log \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|$ up to a constant. Third, the frame-dragging term yields a speed-up/delay for passing along/opposite the spin direction of the star by $\frac{4G}{c^3 D_{OL}} \frac{(\mathbf{J} \times \mathbf{n}) \cdot \boldsymbol{\theta}}{\|\boldsymbol{\theta}\|^2}$. In terms of dimensionless units, we obtain the time delay function

$$T(\mathbf{x}, \mathbf{y}) = \frac{(\mathbf{x} - \mathbf{y})^2}{2} - \log x + \frac{\boldsymbol{\alpha} \cdot \mathbf{x}}{x^2}, \quad (8.4)$$

with the norm $x = \|\mathbf{x}\|$, the spin vector

$$\boldsymbol{\alpha} = \frac{\mathbf{J} \times \mathbf{n}}{cMr_E}, \quad (8.5)$$

where \mathbf{n} is a unit vector pointing along the line of sight towards the source, and the Einstein radius

$$r_E = \sqrt{\frac{4GM}{c^2} \frac{D_{OL}(D_{OS} - D_{OL})}{D_{OS}}}. \quad (8.6)$$

In this formula, the angles

$$\mathbf{x} = \frac{\boldsymbol{\theta}}{\theta_E}, \quad \mathbf{y} = \frac{\boldsymbol{\theta}_S}{\theta_E}, \quad (8.7)$$

are normalized with respect to the Einstein angle

$$\theta_E = r_E/D_{OL} = \sqrt{\frac{4GM}{c^2} \frac{D_{OS} - D_{OL}}{D_{OL} D_{OS}}}. \quad (8.8)$$

In addition, we work with the dimensionless angular frequency $w = 4GM\omega/c^2$, which is defined as the frequency of the radiation ω normalized using the gravitational radius of the lens. The

frequency also enters in the definition of the wavelength of the radiation, defined as $\lambda = 2\pi c/\omega$ (see [1]). In astrophysics, the dimensionless angular frequency w is generally a large number.

It is often convenient to express the spin vector in terms of the Kerr parameter $\mathbf{a} = \mathbf{J}/(cM)$:

$$\boldsymbol{\alpha} = \frac{\mathbf{a} \times \mathbf{n}}{r_E}. \quad (8.9)$$

For black holes, the Kerr parameter takes values between 0 and M with M corresponding to an extremal black hole. As we are modeling rotating stars instead of black holes and work with the small spin limit, the Kerr parameter should be considered small, *i.e.* $a/M < 1$. This implies that for realistic settings $\alpha \ll 1$, given that in the thin lens approximation both D_{OL} and D_{OS} are large so that r_E is large as well ($r_E \sim \sqrt{2GM D_{OL}/c^2}$). While the angular momentum \mathbf{J} is a three-vector, we will from hereon interpret the spin vector $\boldsymbol{\alpha}$ as a two-vector in the lens plane, normal to the projection of the spin vector \mathbf{J} onto the lens plane $\mathbf{J} - (\mathbf{n} \cdot \mathbf{J})\mathbf{n}$. As a result, frame dragging only alters the time delay function when the angular momentum vector \mathbf{J} is misaligned with the line of sight \mathbf{n} and for a given value of J , the effect is maximal when the angular momentum vector lies in the lens plane. Using spherical coordinates

$$\mathbf{J} = J(\sin \theta_0 \cos \phi_0, \sin \theta_0 \sin \phi_0, \cos \theta_0), \quad (8.10)$$

with the line of sight $\mathbf{n} = (0, 0, 1)$, where θ_0 is the angle between the angular momentum and the line of sight, the spin vector assumes the form

$$\boldsymbol{\alpha} = \frac{a \sin \theta_0}{r_E} (-\sin \phi_0, \cos \phi_0), \quad (8.11)$$

with the norm $a = \|\mathbf{a}\|$. The angle θ_0 determines the amplitude of the spin vector, while the angle ϕ_0 determines its orientation.

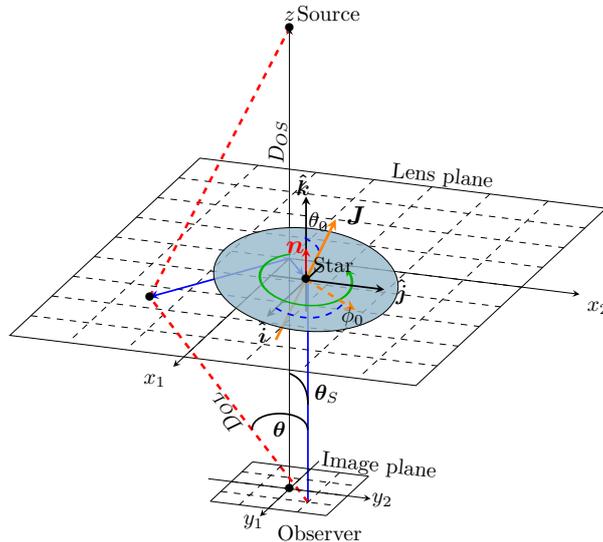


Figure 8.1: Schematic representation of the lensed null rays that arrive at the observer. In the thin lens approximation, we consider that the rays travel along straight lines (as they would in Minkowski space), and the trajectory's deviation due to the rotating star's presence only occurs in the lens plane, represented by the dashed grid. The rotation of the star drags the neighboring spacetime, which has been represented in blue. The line connecting the source to the observer is the line of sight. We consider the star to lie at an angular distance θ_S as measured from the line of sight. A ray intersects the lens plane at an angular distance θ as measured from the position of the star. We define a Cartesian coordinate system centered at the star with the z axis along the line of sight and pointing towards the source. Similarly, the spherical coordinate system (θ_0, ϕ_0) is defined in the usual way. The distance separating the observer from the lens is denoted D_{OL} while the distance separating the observer from the source is D_{OS} . Both distances are much larger than the characteristic lengths in the lens plane. In orange, we represent the angular momentum vector of the star \mathbf{J} and in red a unit vector \mathbf{n} parallel to the line of sight.

In the literature, the time delay function of a rotating star is sometimes expressed as [1, 20, 21]

$$T(\mathbf{x}, \mathbf{y}) = \frac{(\mathbf{x} - \mathbf{y})^2}{2} - \log \|\mathbf{x} - \boldsymbol{\alpha}\| - \sum_{n=2}^{\infty} \frac{\alpha^n \cos n\bar{\theta}}{nx^n}, \quad (8.12)$$

whereby the frame dragging term $\boldsymbol{\alpha} \cdot \mathbf{x}/x^2$ is absorbed into the logarithm and the higher order corrections, with $\bar{\theta}$ the angle between $\boldsymbol{\alpha}$ and \mathbf{x} . These higher-order corrections, often written as $\mathcal{O}(\alpha^2)$, are then neglected in those references. To first order in the spin parameter, this time delay function is equivalent to the time delay function of the non-rotating point source with $\mathbf{x} \mapsto \mathbf{x} - \boldsymbol{\alpha}$. Based on this observation, it is then argued that the interference pattern of a rotating star is the same as the interference pattern of the non-rotating star shifted by the spin vector $\boldsymbol{\alpha}$ (see for example [1, 20, 21, 22]).

While this may appear completely reasonable at first, this seemingly innocent trick is in fact mathematically rather violent: The residue $\sum_{n=2}^{\infty} \frac{\alpha^n \cos n\bar{\theta}}{nx^n}$ diverges logarithmically at both the origin $\mathbf{0}$ and the point $\boldsymbol{\alpha}$ (where the sum reduces to the harmonic series $\sum_{n=2}^{\infty} n^{-1}$). While the first singularity can be understood as the position of the lens and also appears in the time delay function in Eq. (8.4), the latter singularity is not physical. It is merely an artifact of rewriting Eq. (8.4) as Eq. (8.12). Since the Kirchhoff-Fresnel integral ranges over all points \mathbf{x} in the lens plane, this way of rewriting the system is not a small perturbation of the original system, even for small α . Moreover, it is important to realize that the rotational symmetry of the phase variation of the non-rotating lens $\varphi(\mathbf{x}) = \log x$ makes the interference pattern degenerate (in the sense of catastrophe theory). A small fluctuation in the phase variation φ — thereby breaking the rotational symmetry — will dramatically change the caustics and interference pattern (see appendix 8.5 for some concrete examples). The frame-dragging contribution certainly breaks the rotation symmetry as light rays moving along the spin direction of the star are sped up while light rays propagating opposite of the spin direction are delayed by frame dragging. By approximating the phase variation of the rotating star by

$$\varphi(\mathbf{x}) \sim -\log \|\mathbf{x} - \boldsymbol{\alpha}\|, \quad (8.13)$$

the rotational symmetry and consequently the degeneracy of the non-rotating point lens is preserved. Hence, this is not a good approximation as we will also explicitly see in Sec. 8.3.

8.2.2 Geometric optics

Following Fermat's principle, the classical rays correspond to the stationary point of the time delay function,

$$\nabla_{\mathbf{x}} T(\mathbf{x}, \mathbf{y}) = \mathbf{x} - \mathbf{y} - \frac{\mathbf{x} - \boldsymbol{\alpha}}{x^2} - \frac{2(\boldsymbol{\alpha} \cdot \mathbf{x})\mathbf{x}}{x^4} = 0. \quad (8.14)$$

Upon solving for \mathbf{y} , we obtain the geometric optics map,

$$\boldsymbol{\xi}(\mathbf{x}) = \mathbf{x} - \frac{\mathbf{x} - \boldsymbol{\alpha}}{x^2} - \frac{2(\boldsymbol{\alpha} \cdot \mathbf{x})\mathbf{x}}{x^4}, \quad (8.15)$$

sending points in the lens plane to the observer. In the *geometric optics approximation*, the intensity assumes the form

$$I_{\text{geometric}}(\mathbf{y}) = \sum_{\mathbf{x} \in \boldsymbol{\xi}^{-1}(\mathbf{y})} \frac{1}{|\det \nabla \boldsymbol{\xi}(\mathbf{x})|} \quad (8.16)$$

including a contribution for each classical ray propagating from the source to the observer passing the lens at $\boldsymbol{\xi}^{-1}(\mathbf{y}) = \{\mathbf{x} \mid \boldsymbol{\xi}(\mathbf{x}) = \mathbf{y}\}$ [24]. The intensity (8.16) spikes when the deformation tensor $\nabla \boldsymbol{\xi}$ is singular. In the lens plane, this corresponds to the critical curve

$$\mathcal{C} = \{\mathbf{x} \mid \det \nabla \boldsymbol{\xi}(\mathbf{x}) = 0\}, \quad (8.17)$$

$$= \{\mathbf{x} \mid \mathbf{x}^2(\mathbf{x}^4 - 1) - 4\boldsymbol{\alpha} \cdot (\boldsymbol{\alpha} + \mathbf{x}) = 0\}. \quad (8.18)$$

When mapping the critical curve to the image plane, we obtain the caustic curve,

$$\boldsymbol{\xi}(\mathcal{C}) = \{\boldsymbol{\xi}(\mathbf{x}) \mid \mathbf{x} \in \mathcal{C}\}, \quad (8.19)$$

consisting of the geometric pattern at which the intensity of the geometric optics approximation spikes. As we will see below, the caustics are the places where the geometric optics approximation starts to fail. Moreover, the caustics mark the locations in the image plane at which the intensity pattern in wave optics qualitatively changes. As we approach a fold/cusp caustic in the image plane the two/three rays $\boldsymbol{\xi}^{-1}(\mathbf{y})$ coalesce at a point on the critical curve in the lens plane.

For the non-rotating point lens, the lens map $\boldsymbol{\xi}(\mathbf{x}) = \mathbf{y}$ yields two real classical rays [2]

$$\mathbf{x}_{\pm} = \frac{\mathbf{y}}{2y} \left(y \pm \sqrt{4 + y^2} \right). \quad (8.20)$$

The critical curve \mathcal{C} , consisting of the unit circle $\{\|\mathbf{x}\| = 1\}$ in the lens plane, is known as the Einstein ring. The caustic curve consist of the point $\boldsymbol{\xi}(\mathcal{C}) = \{\mathbf{0}\}$ in the image plane. The observation that the one-dimensional unit circle is mapped to a zero-dimensional point signals that this set-up is degenerate resulting from the radial symmetry of the phase-variation $-\log x$. A small perturbation, breaking the radial symmetry of the phase variation in the lens plane dramatically changes the caustic curve into an astroid, consisting of a fold curve with four cusp points (see appendix 8.5). A similar phenomenon can be observed in the unfolding of the caustic resulting from a symmetric liquid dropped lens [28, 29].

Including the frame-dragging effect breaks the radial symmetry of the lens and removes the degeneracy of the non-rotating point lens. The lens now has up to 5 classical rays corresponding to stationary points of the time delay function. The critical and caustic curve gradually change as we increase the spin parameter $\alpha = \|\boldsymbol{\alpha}\|$ (see Fig. 8.2):

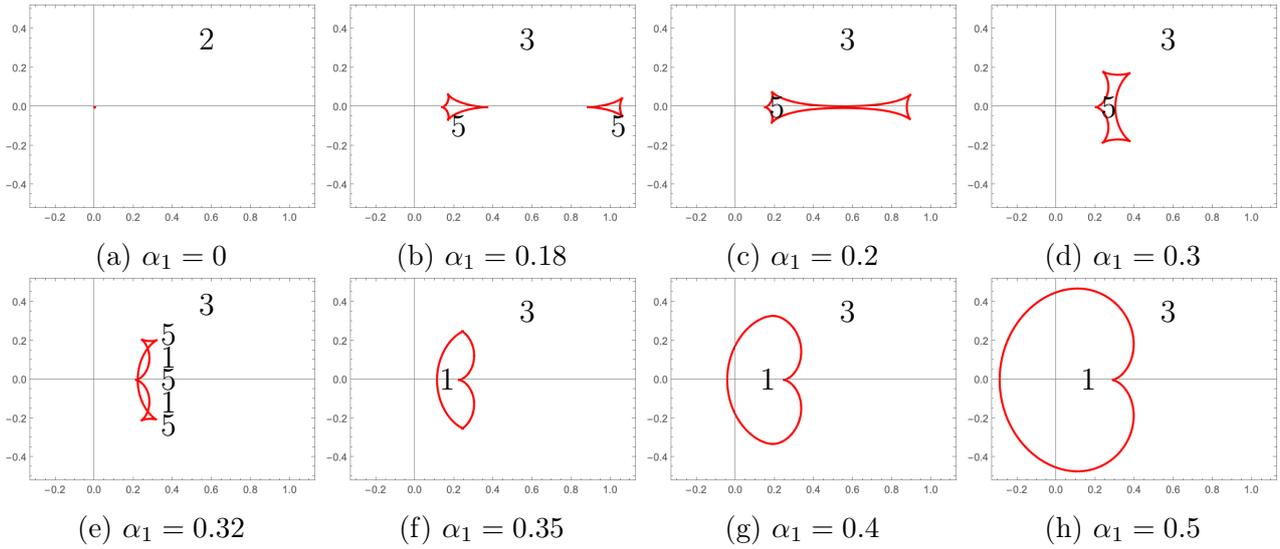


Figure 8.2: The caustic curves for $\alpha = (\alpha_1, 0)$ for $\alpha_1 = 0, 0.18, 0.2, 0.3, 0.32, 0.35, 0.4, 0.5$. The numbers indicate the number of images in each region bounded by the caustic curve.

- For $0 < \alpha < \frac{1}{3\sqrt{3}} \approx 0.192$, the critical curve consists of two loops. The corresponding caustics form a quadrangle and a triangle enclosing two 5-image regions. Outside these caustic curves, there exist 3 real classical rays.
- At $\alpha = \frac{1}{3\sqrt{3}}$ the two 5-image regions merge to form a larger 5-image regions.
- At $\alpha \approx 0.31$, the right horizontal fold line starts to overtake the left horizontal caustics. When the right fold passes over the left fold, we observe the formation of two 1-image regions.
- At $\alpha = \frac{1}{7}\sqrt{\frac{13+16\sqrt{2}}{7}} \approx 0.322$ the two 1-image regions merge.
- At $\alpha \approx 0.35$, the two remaining 5-image regions collapse to a point in a swallowtail caustic. The 5-image regions vanish for larger α .
- For larger α , we obtain a 1-image region surrounded by a 3-image region. The caustic curve consists of a fold curve and a single cusp point. As α increases, the caustic inflates. These findings for large α are in agreement with [30].

For the rotating star, the outside region always consists of three classical rays. For each point \mathbf{y} outside the caustics, the time delay function has three real and two complex critical points (forming a complex conjugate pair) in \mathbf{x} when solving Eq. (8.14). When crossing a caustic while moving \mathbf{y} , two things can happen: the two complex critical points merge on the real plane at the critical curve and become real classical rays forming a five-image region, or two real critical points merge on the critical curve and form a complex conjugate pair of complex saddle points. The complex critical points (sometimes known as complex rays) do not play a role in the geometric optics analysis of the rotating lens but do influence the interference pattern in wave optics (see for example [8]).

8.2.3 Wave optics

For a non-rotating point lens, the radial symmetry allows us to write the lensing amplitude as a radial integral,

$$\Psi(\mathbf{y}) = \frac{w}{2\pi i} \int e^{iw(\frac{1}{2}(\mathbf{x}-\mathbf{y})^2 - \log x)} d\mathbf{x} \quad (8.21)$$

$$= -ie^{iwy^2/2} \int_0^\infty J_0(wry) e^{iw(\frac{1}{2}r^2 - \log r)} r dr, \quad (8.22)$$

using polar $\mathbf{x} = r(\cos \theta, \sin \theta)$ and Cartesian coordinates $\mathbf{y} = (y_1, y_2)$, the norm $y = \|\mathbf{y}\|$, and the integral representation of the Bessel function

$$2\pi J_0\left(\sqrt{z_1^2 + z_2^2}\right) = \int_0^{2\pi} e^{i(z_1 \cos \theta + z_2 \sin \theta)} d\theta. \quad (8.23)$$

Remarkably, the radial integral can be evaluated

$$\Psi(\mathbf{y}) = 2^{-1-\frac{iw}{2}} (-iw)^{1+\frac{iw}{2}} \Gamma\left(-\frac{iw}{2}\right) L_{-\frac{iw}{2}}\left(\frac{iwy^2}{2}\right) \quad (8.24)$$

in terms of the Laguerre function L_n and the gamma function Γ . The intensity of the radiation assumes the form

$$|\Psi(\mathbf{y})|^2 = \frac{\pi w}{1 - e^{-\pi w}} \left| {}_1F_1\left(\frac{iw}{2}, 1; \frac{iwy^2}{2}\right) \right|^2, \quad (8.25)$$

with the Kummer confluent hypergeometric function

$${}_1F_1(a, b; z) = \sum_{n=0}^{\infty} \frac{a^{(n)} z^n}{b^{(n)} n!}, \quad (8.26)$$

where we use the rising factorial $a^{(n)} = a(a+1)\dots(a+n-1)$ (see [2] for details).

For a rotating lens, the frame-dragging term in polar coordinates is

$$\frac{\boldsymbol{\alpha} \cdot \mathbf{x}}{x^2} = \frac{\alpha_1 \cos \theta + \alpha_2 \sin \theta}{r}, \quad (8.27)$$

with $\boldsymbol{\alpha} = (\alpha_1, \alpha_2)$. This makes the lensing amplitude for the non-rotating lens take on the following form

$$\begin{aligned} \Psi(\mathbf{y}) &= -iwe^{iwy^2/2} \\ &\times \int_0^\infty J_0(wr\|\mathbf{y} - \boldsymbol{\alpha}/r^2\|) e^{iw(r^2/2 - \log r)} r dr. \end{aligned} \quad (8.28)$$

Note that this is only a slight variation on the amplitude evaluated by [1] using the approximation in Eq. (8.13):

$$\begin{aligned} \Psi(\mathbf{y}) &\sim -iwe^{iwy^2/2} \\ &\times \int_0^\infty J_0(wr\|\mathbf{y} - \boldsymbol{\alpha}\|) e^{iw(r^2/2 - \log r)} r dr, \end{aligned} \quad (8.29)$$

which can be evaluated analytically as this is simply the Kirchoff-Fresnel integral for the non-rotating lens but shifted

$$\Psi(\mathbf{y}) \stackrel{[1]}{=} 2^{-1-\frac{iw}{2}} (-iw)^{1+\frac{iw}{2}} \Gamma\left(-\frac{iw}{2}\right) L_{-\frac{iw}{2}}\left(\frac{iw\|\mathbf{y} - \boldsymbol{\alpha}\|^2}{2}\right). \quad (8.30)$$

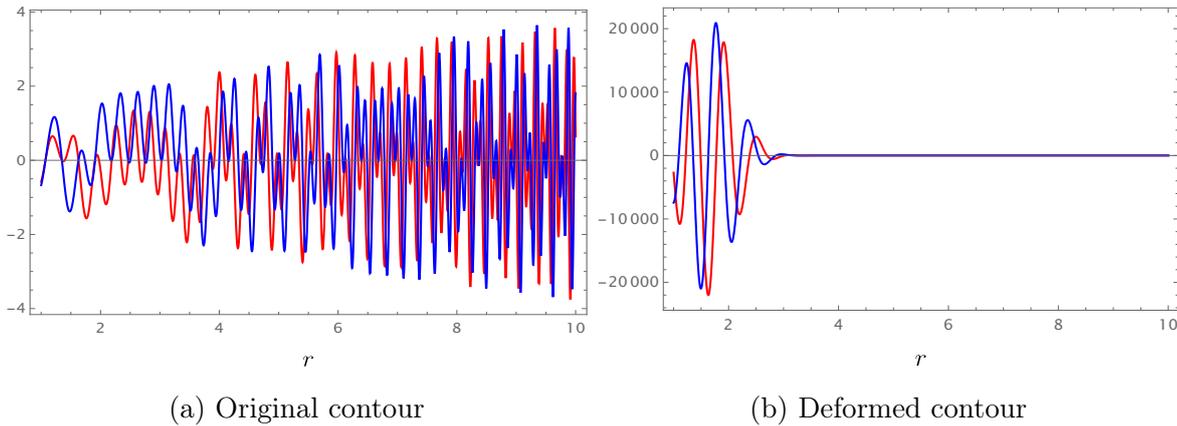


Figure 8.3: The integrand of the radial integral in Eq. (8.28) along the real line $[1, \infty)$ (left), and deformed contour $[1, e^{i\pi/4}\infty)$. The blue and red lines indicate the real and imaginary parts, respectively.

Unfortunately, the radial integral in Eq. (8.28) cannot be evaluated using special functions and is highly oscillatory for large r . However, as the integrand is analytic, and the analytic continuation is dominated by the Gaussian term $e^{i\omega r^2/2}$ for large $|r|$ in the complex plane, we can safely deform the half line $[0, \infty)$ to an integration contour starting at the origin 0 tangential to the real axis and ending at $e^{i\pi/4}\infty$. This deformation suppresses the integrand for large $|r|$ and enables the efficient evaluation of the amplitude $\Psi(\mathbf{y})$ using conventional integration techniques. As the analytic continuation diverges around $r = 0$ in the complex plane, we deform the real half-line $(0, \infty)$ into the two line segments $(0, 1] \cup [1, e^{i\pi/4}\infty)$ (see Fig. 8.3). This is a rudimentary application of Picard-Lefschetz theory [8]. More intricate deformations of the original integration domain can further improve the convergence of the radial integral. However, this simple deformation suffices for the purpose of this paper.

8.2.4 Validating the lens amplitude

To validate our numerical method for the evaluation of the wave amplitude in the previous section, we approximate the Kirchhoff-Fresnel integral in three ways.

First, the Eikonal approximation of $\Psi(\mathbf{y})$ is a good approximation in the semi-classical regime, bridging the geometric optics approximation (8.16) and the full Kirchhoff-Fresnel integral,

$$\Psi(\mathbf{y}) \approx \sum_{\mathbf{x} \in \xi^{-1}(\mathbf{y})} \frac{e^{i\omega T(\mathbf{x}, \mathbf{y}) - in(\mathbf{x})\pi/2}}{\sqrt{|\det \nabla \xi(\mathbf{x})|}}, \quad (8.31)$$

with the Morse index n of the critical point \mathbf{x} (0 for minima, 1 for saddle points, and 2 for maxima of the time delay T). For a detailed exposition of the Eikonal approximation see [24]. In the present discussion, we will only include the real rays, though the Eikonal approximation can be extended by including relevant complex rays [8]. In Fig. 8.4, we compare the geometric optics approximation, the Eikonal approximation and the Kirchhoff-Fresnel integral for a rotating lens. The geometric optics approximation captures the main behavior of the lensing pattern but misses the interference of the three real rays in the triple-image regions. Additionally, the geometric optics approximation diverges at the caustics. The Eikonal approximation is a significant improvement as it captures the interference in the triple-image regions. However, it fails to capture the oscillations on the left side of the single-image region. These oscillations result from the interference of the real ray and a complex ray (associated with the left fold

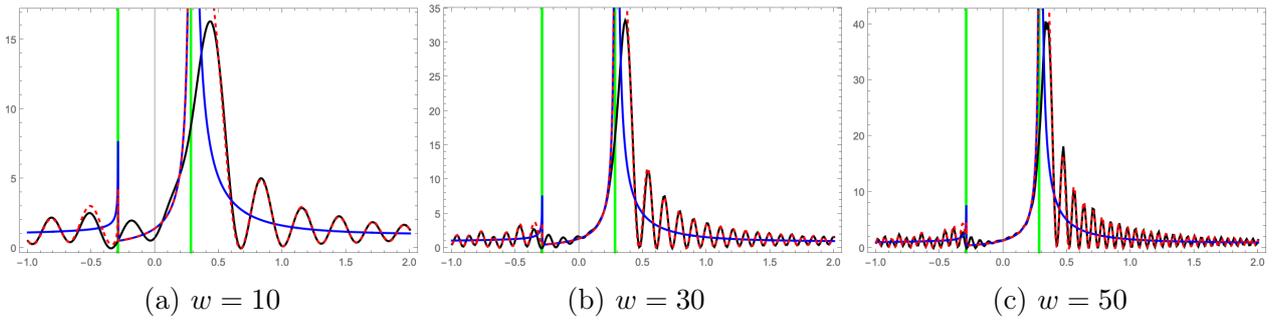


Figure 8.4: The intensity $|\Psi(\mathbf{y})|^2$ for y along the equatorial plane/for $y = (y_1, 0)$ with $-1 \leq y_1 \leq 2$ for $\boldsymbol{\alpha} = (1/2, 0)$ in geometric optics (blue), the Eikonal approximation (red), and the numerical evaluation of the Kirchhoff-Fresnel integral (black). The vertical green lines mark the caustics separating the one-image region (between the two green lines) from the triple-image regions (the regions on the outsides of the green lines).

caustic). Both the geometric optics and Eikonal approximations overestimate the intensity near the left cusp caustic. Given that the oscillations of the Eikonal approximation and the numerical evaluation of the Kirchhoff-Fresnel integral line up, and the deviations can be understood, we are confident about the accuracy of the numerical evaluation.

The second method to validate our numerical results is by expanding the radial integral in the Kirchhoff-Fresnel integral (8.28) in powers of the rotation parameter, for which we can obtain analytic results. As the expansion of the Bessel function,

$$J_0(wr\|\mathbf{y} - \boldsymbol{\alpha}/r^2\|) \quad (8.32)$$

$$= J_0(wry) + w\boldsymbol{\alpha} \cdot \mathbf{y} \frac{J_1(wry)}{ry} - w^2(\boldsymbol{\alpha} \cdot \mathbf{y})^2 \frac{J_0(wry)}{2r^2y^2} + w(y_1^2(\alpha_1^2 - \alpha_2^2) + 4\alpha_1\alpha_2y_1y_2 - y_2^2(\alpha_1^2 - \alpha_2^2)) \frac{J_1(wry)}{2r^3y^3} + \mathcal{O}(\boldsymbol{\alpha}^3),$$

only features the Bessel functions $J_0(wry)/r^m$ and $J_1(wry)/r^m$ for positive integers m , the closed form integral

$$\int_0^\infty J_\nu(wyr)e^{iw(\frac{1}{2}r^2 - \log r)} r^n dr$$

$$= 2^{-\frac{1+\nu-n+iw}{2}} (-iw)^{-\frac{1+\nu+n-iw}{2}} (wy)^\nu \Gamma\left(\frac{1+\nu+n-iw}{2}\right) {}_1\tilde{F}_1\left(\frac{1+\nu+n-iw}{2}; 1+\nu; -\frac{iy^2w}{2}\right) \quad (8.33)$$

allows us to evaluate the radial integral and approximate the integral to any order in $\boldsymbol{\alpha}$. Fig. 8.5 shows a comparison of this perturbative expansion and the numerically evaluated Kirchhoff-Fresnel integral. The expansion matches the numerical evaluation well for small $\|\boldsymbol{\alpha}\|$ but requires many terms to capture the oscillations of the interference pattern.

Finally, the numerical evaluation matches qualitatively the shifted interference pattern of the non-rotating star for small α (see the next section for details).

8.3 Results

When evaluating the Kirchhoff-Fresnel pattern as a function of \mathbf{y} we obtain an interference pattern matching the caustics obtained from the geometric optics approximation (see Fig. 8.6). For small $\|\boldsymbol{\alpha}\|$, the caustic curve is small compared to the typical length scale of the interference pattern and the interference pattern is close to the shifted pattern predicted by [1] (although see below for an explicit comparison). For larger $\|\boldsymbol{\alpha}\|$, the interference pattern becomes more intricate neatly following the caustics. Note that the oscillations in the interference pattern in

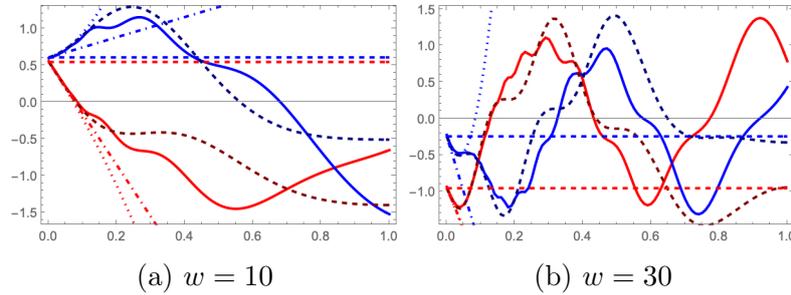


Figure 8.5: A comparison of the exact evaluation (solid curves) and approximations of the lens amplitude as a function of $\boldsymbol{\alpha} = (\alpha, 0)$ for $\omega = 10$ (left panel) and $\omega = 30$ (right panel) and $\mathbf{y} = (1, 1)$. We compare the numerical evaluation with the zeroth-order (dashed curves), first-order (dashed-dotted curves), second-order approximation (dotted curves), and the shifted non-rotating star (dark dashed curves). The real and imaginary parts of Ψ are plotted in red and blue, respectively.

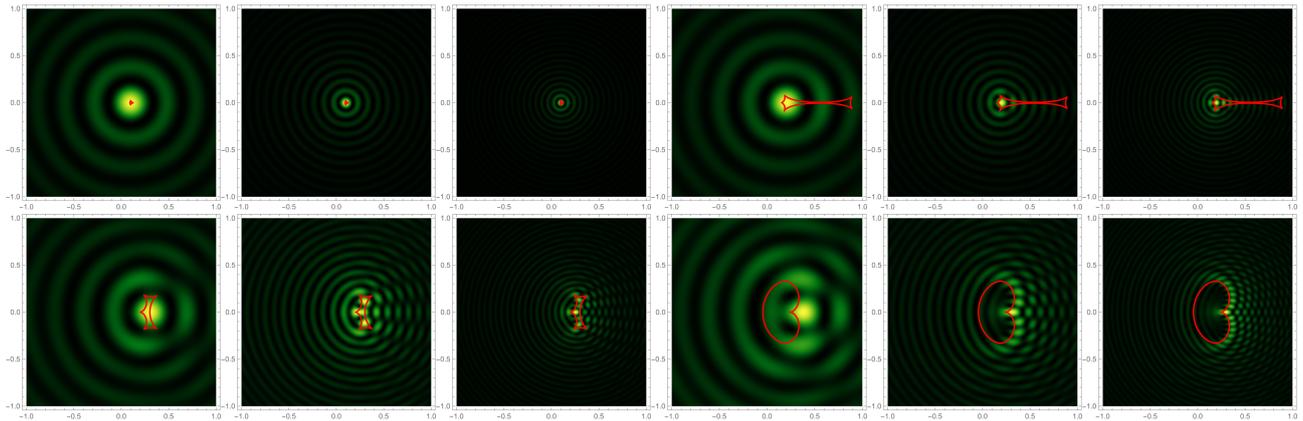


Figure 8.6: Interference patterns of the rotating lens for $\boldsymbol{\alpha} = (\alpha_1, 0)$ for $\alpha_1 = 0.1, 0.2, 0.3, 0.4$ (top to bottom) and the angular frequencies $w = 10, 30, 50$ (left to right). The red curves are the caustics: in the geometric optics approximation, the intensity becomes infinite there.

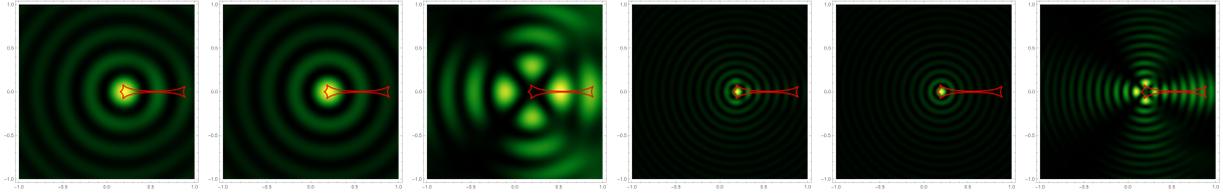


Figure 8.7: A comparison of the exact Kirchhoff-Fresnel integral (left) and the shifted non-rotating interference pattern (center) and their difference (right) for $\boldsymbol{\alpha} = (0.2, 0)$, for $w = 10$ (the upper panels) and $w = 30$ (the lower panels). The difference between the two patterns is reminiscent of the light emitted by a lighthouse.

the single image region for $\|\boldsymbol{\alpha}\| \geq 0.322$ are the result of a relevant complex ray corresponding to a complex critical point of the time delay function.

We compare the interference pattern with the shifted non-rotating lens proposed by [1] in Fig. 8.7. Though the two interference patterns appear similar, they show a systematic difference even for small α : this difference is highlighted in the right column. This difference can in principle allow one to directly infer the spin of the rotating star through frame dragging from the observed fringes.

We want to highlight that using the shifted non-rotating time delay to compute the interference pattern generated by the rotating star is *incorrect* (see also the discussion around Eq. (8.12) in Sec. (8.2.1)). The underlying reason why the approximation

$$T(\mathbf{x}, \mathbf{y}) = \frac{(\mathbf{x} - \mathbf{y})^2}{2} - \log\|\mathbf{x} - \boldsymbol{\alpha}\| - \sum_{n=2}^{\infty} \frac{\alpha^n \cos n\bar{\theta}}{nx^n} \quad (8.34a)$$

$$\approx \frac{(\mathbf{x} - \mathbf{y})^2}{2} - \log\|\mathbf{x} - \boldsymbol{\alpha}\|, \quad (8.34b)$$

fails despite its seeming correctness (and its extensive usage in the literature [1, 20, 21, 22]) is that the residue $\sum_{n=2}^{\infty} \frac{\alpha^n \cos n\bar{\theta}}{nx^n}$ diverges for $\mathbf{x} = \mathbf{0}$, and $\mathbf{x} = \boldsymbol{\alpha}$ (following from the fact that we approximate a function $\log\|\mathbf{x}\|$ which diverges at $\mathbf{x} = \mathbf{0}$ by the function $\log\|\mathbf{x} - \boldsymbol{\alpha}\|$ which diverges at $\mathbf{x} = \boldsymbol{\alpha}$). The time delays (8.34a) and (8.34b) give rise to radically different interference patterns. This can be seen by estimating the interference pattern generated by the correction term $\sum_{n=2}^{\infty} \frac{\alpha^n \cos n\bar{\theta}}{nx^n}$. The Kirchhoff-Fresnel integral for the time delay (8.34a) is

$$\Psi(\mathbf{y}) = \frac{w}{2\pi i} \int e^{iw \left[\frac{(\mathbf{x} - \mathbf{y})^2}{2} - \log\|\mathbf{x} - \boldsymbol{\alpha}\| - \sum_{n=2}^{\infty} \frac{\alpha^n \cos n\bar{\theta}}{nx^n} \right]} d\mathbf{x}. \quad (8.35)$$

Using the identity $e^x = \sum_{k=0}^{\infty} \frac{x^k}{k!}$, we obtain

$$\begin{aligned} \Delta\Psi(\mathbf{y}) &= \Psi(\mathbf{y}) - \Psi_{\text{shifted}}(\mathbf{y}) \\ &= -\frac{w^2}{2\pi} \int e^{iw \left[\frac{(\mathbf{x} - \mathbf{y})^2}{2} - \log\|\mathbf{x} - \boldsymbol{\alpha}\| \right]} \frac{\alpha^2 \cos 2\bar{\theta}}{2x^2} d\mathbf{x} + O(\alpha^3) \end{aligned} \quad (8.36)$$

where $\Psi(\mathbf{y})$ is given by Eq. (8.35) and Ψ_{shifted} is

$$\Psi_{\text{shifted}}(\mathbf{y}) = \frac{w}{2\pi i} \int e^{iw \left[\frac{(\mathbf{x} - \mathbf{y})^2}{2} - \log\|\mathbf{x} - \boldsymbol{\alpha}\| \right]} d\mathbf{x}. \quad (8.37)$$

For the purpose of the argument, it is enough to consider the correction of order α^2 to the shifted non-rotating interference pattern. If considering the approximation to the time delay given by Eq. (8.34b) was correct, then the next order correction in α to the time delay in

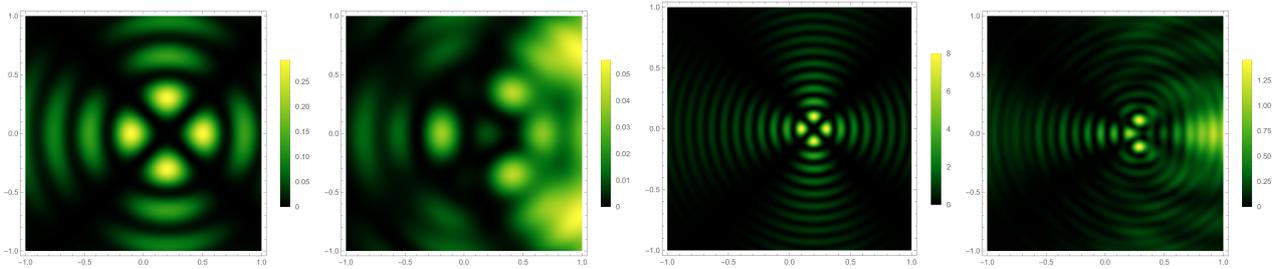


Figure 8.8: The second order correction in α to the shifted nonrotating interference pattern (left) and its difference with respect to the exact Kirchhoff-Fresnel integral (right), with the latter capturing the higher-order contributions given by $O(\alpha^3)$. We used $\boldsymbol{\alpha} = (0.2, 0)$ and $w = 10$ in the first row and $w = 30$ in the second. The brightness of the image is also shown for comparison.

Eq. (8.36) either 1) does not significantly impact the interference pattern, or 2) if it does, the intensity of the correction is smaller than roughly α^2 . As we will see in the following, neither of these two statements is true when evaluating Eq. (8.36). Recall that $\bar{\theta}$ is a function of \boldsymbol{x} and $\boldsymbol{\alpha}$, in particular

$$\frac{\alpha^2 \cos 2\bar{\theta}}{2x^2} = \frac{2(\boldsymbol{x} \cdot \boldsymbol{\alpha}) - \alpha^2 x^2}{2x^4}. \quad (8.38)$$

Using polar coordinates centered at $\boldsymbol{x} = \boldsymbol{\alpha}$

$$\boldsymbol{x} - \boldsymbol{\alpha} = r(\cos \theta, \sin \theta), \quad (8.39)$$

we can rewrite expression (8.36) as

$$\begin{aligned} \Delta\Psi = & -\frac{w^2 e^{iw\frac{(\boldsymbol{\alpha}-\boldsymbol{y})^2}}{2}}}{4\pi} \int d\theta [(\alpha_1^2 - \alpha_2^2) \cos 2\theta + 2\alpha_1\alpha_2 \sin 2\theta] \times \\ & \int dr e^{iw\left[\frac{r^2}{2} + \frac{(\boldsymbol{\alpha}-\boldsymbol{y})^2}{2} + r(\alpha_1 - y_1) \cos \theta + r(\alpha_2 - y_2) \sin \theta\right]} r^{-1-iw} + O(\alpha^3) \end{aligned} \quad (8.40)$$

where we only kept terms up to second order in α^2 , *i.e.*,

$$\frac{\alpha^2 \cos 2\bar{\theta}}{2x^2} = \frac{(\alpha_1^2 - \alpha_2^2) \cos 2\theta + 2\alpha_1\alpha_2 \sin 2\theta}{2r^2} + O(\alpha^3). \quad (8.41)$$

In practice, this means that for this term we can use the transformation $\boldsymbol{x} = r(\cos \theta, \sin \theta)$ rather than Eq. (8.39) since the factor of $\boldsymbol{\alpha}$ in the transformation enters as a higher order correction.

The radial integral

$$I_1 = \int_0^\infty dr e^{iw\left[\frac{r^2}{2} + \frac{(\boldsymbol{\alpha}-\boldsymbol{y})^2}{2} + rb\right]} r^{-1-iw}, \quad (8.42)$$

regularized using analyticity [31], can be evaluated in terms of the Kummer's confluent hypergeometric function

$$I_1 = 2^{-1-iw/2} e^{\pi w/4} w^{iw/2} \left\{ (i-1)b\sqrt{w}\Gamma\left[\frac{1}{2} - \frac{iw}{2}\right] {}_1F_1\left[\frac{1}{2} - \frac{iw}{2}, \frac{3}{2}, -\frac{iw}{2}b^2\right] + \Gamma\left[-\frac{iw}{2}\right] {}_1F_1\left[-\frac{iw}{2}, \frac{1}{2}, -\frac{iw}{2}b^2\right] \right\}. \quad (8.43)$$

We did not find a closed-form expression for the remaining integral over θ in $\Delta\Psi$. However, as this integral is well-behaved, and runs over a finite domain, we evaluate it numerically. In Fig. 8.8 we show the interference pattern created by the second order correction term in α

in Eq. (8.36) and its difference with the right column of Fig. 8.7 for $w = 10$ (first row) and $w = 30$ (second row). The right column of Fig. 8.8 represents the higher-order corrections not captured by Eq. (8.36). The second-order correction (left column in Fig. 8.8) already encodes the “lighthouse” effect that we observed in the right column of Fig. 8.7. This shows that using the shifted non-rotating time delay leads to a loss of structure in the interference pattern, and the first statement above is incorrect.

Notice that the intensity of the second-order correction (left column in Fig. 8.8) is comparable (for $w = 10$) or even higher (for $w = 30$) to the magnitude of $\alpha \approx 0.2$. Therefore, this difference is not negligible, which also invalidates the second statement above. Also note that the intensity of the higher order correction depicted in the right column of Fig. 8.8 is one order of magnitude smaller than the second order effect in α (left column). This is expected and hints at the possibility of approximating $\sum_{n=2}^N \frac{\alpha^n \cos n\bar{\theta}}{nx^n}$ by a finite number of terms $N > 2$. As seen by looking at the intensity bars in Fig. 8.8, for higher frequencies, N needs to be larger to capture the structure of the interference pattern up to order α . However, we do not see any advantage in using Eq. (8.34a) truncated up to a certain N rather than simply using Eq. (8.4) since we have explicitly shown that in both cases, the $N = 2$ correction *must* be included, thereby already breaking the degeneracy of the caustics.

8.4 Discussion

We have calculated the interference pattern of a rotating lens in the small spin approximation using the Lense-Thirring metric. Despite previous claims in the literature, the resulting interference pattern is not merely a translated version of that produced by a non-rotating lens of equivalent mass; rather, it exhibits distinctive characteristics. A comparative analysis of the interference patterns generated by the rotating lens and the displaced non-rotating lens reveals a prominent feature that bears resemblance to the light emitted from a lighthouse (refer to Fig. 8.7). If these features can be observed, this would be a new and independent methodology to measure the spin of rotating objects, a task that is generally considered to be challenging.

In strong gravitational lensing, the lens bends the paths of the radiation and multiple images, arcs, or Einstein rings appear. In weak gravitational lensing, the deflection due to the presence of a lens is impossible to detect from a single background source. The presence of the foreground lens mass has to be statistically identified through the systematic alignment of multiple background sources around the lensing mass. In the present paper, we consider systems for which radiation experiences a strong deflection but for which the different images are not resolved. Rather, we observe the interference of the different images. This is known as gravitational microlensing. Not unlike Young’s double slit experiment, the complementarity principle indicates that when we can identify the path of the rays, the interference pattern disappears. As a result, gravitational microlensing is often washed out. However, it is relevant for coherent radiation with long wavelengths, like the recently observed fast radio bursts and gravitational waves.

The interference patterns presented in this paper are not directly observed in the sky given that we do not have access to the entire image in the image plane. Rather, the effect of a gravitational microlens is a transient astronomical event. As the relative alignment of the source, lens and observer changes, the source’s apparent brightness changes. The lens system traces a line in the presented interference pattern observed as fluctuations in intensity. The presented interference patterns can thus be interpreted as a collection of waveforms, where the caustics mark the most striking features. A frame-dragging detection by a microlens would thus require one to fit the presented model of the interference fringes to the light curve.

Note that our results indicate that frame-dragging could in principle also be observed in strong gravitational lensing, through the relative positions of the images in the sky (as one

should consider the time delay (8.4) rather than the approximation (8.13)).

The discussion in this paper is tailored to the language of wave optics as observable in electromagnetic radiation. However, wave optics may also be observable in gravitational wave observations and the relevant mathematics is analogous. The observable in this context is the waveform. The lensed waveform is computed from the unlensed one by convolving it with the Kirchoff-Fresnel integral in the frequency domain. The prospect of observing wave optics with gravitational waves at low frequencies with the Laser Interferometer Space Antenna (LISA) is particularly promising. In the point mass approximation for the lens, lens masses as “small” as $\sim 10^3 M_\odot$ can be inferred from the observed gravitational wave from a massive black hole binary with a total mass of $\sim 10^6 M_\odot$ [32]. It would be interesting to investigate the observability of a spinning lens using gravitational waves.

The foundations laid down in this article provide a stepping stone for a full derivation of a rotating black hole as lens. This would require a derivation of the time delay function using the Kerr metric instead of its small spin limit and take strong field effects into account (such as the possibility of rays disappearing into the black hole and rays going around the black hole multiple times before reemerging).

8.5 Unfolding of the point lens

The time delay of the non-rotating point lens $T(\mathbf{x}, \mathbf{y}) = (\mathbf{x} - \mathbf{y})^2/2 - \log x$ has a one-dimensional critical curve $\mathcal{C} = \{\|\mathbf{x}\| = 1\}$ and a degenerate zero-dimensional caustic curve $\xi(\mathcal{C}) = \{\mathbf{0}\}$. To study the unfolding of the caustic, let us perturb the time delay function,

$$T(\mathbf{x}, \mathbf{y}) = \frac{(\mathbf{x} - \mathbf{y})^2}{2} - \log x + \delta T(\mathbf{x}) \quad (8.44)$$

with a quadratic fluctuation $\delta T(\mathbf{x}) = c + \mathbf{x} \cdot \boldsymbol{\eta} + \frac{1}{2} \mathbf{x}^T \Sigma \mathbf{x}$ for a constant scalar c , 2-vector $\boldsymbol{\eta}$ and symmetric 2×2 matrix Σ . The constant term is irrelevant to the caustics. The linear term can be absorbed into the geometric contribution to the time delay function. For the quadratic term, we use the rotation symmetry of the unperturbed lens to rotate to the eigenframe of Σ ,

$$\Sigma = \begin{pmatrix} \gamma_1 + \gamma_2 & 0 \\ 0 & \gamma_1 - \gamma_2 \end{pmatrix}. \quad (8.45)$$

The generalized time delay thus assumes the form

$$T(\mathbf{x}, \mathbf{y}) \sim \frac{(\mathbf{x} - (\mathbf{y} + \boldsymbol{\eta}))^2}{2} - \log x + \frac{\gamma_1}{2}(x_1^2 + x_2^2) + \frac{\gamma_2}{2}(x_1^2 - x_2^2), \quad (8.46)$$

up to constant contributions. The corresponding lens map

$$\xi(\mathbf{x}) = \mathbf{x} - \frac{\mathbf{x}}{x^2} + \mathbf{x} \cdot \begin{pmatrix} \gamma_1 + \gamma_2 \\ \gamma_1 - \gamma_2 \end{pmatrix} - \boldsymbol{\eta} \quad (8.47)$$

shows that the linear term in the fluctuation δT shifts the caustic by $\mathbf{y} \mapsto \mathbf{y} - \boldsymbol{\eta}$. The term $\gamma_1 x^2$ preserves the rotational symmetry of the lens and the degeneracy of the caustic curve. The critical curve is a circle with radius $1/\sqrt{1 + \gamma_1}$ and the caustic curve is the point $\{-\boldsymbol{\eta}\}$. The shear term $\frac{1}{2}\gamma_2(x_1^2 - x_2^2)$ breaks the rotation symmetry of the point lens. For small γ_2 , the caustic is a one-dimensional astroid, consisting of four fold lines joined by four cusp points (see Fig. 8.9). This pattern is stable with respect to small perturbations.

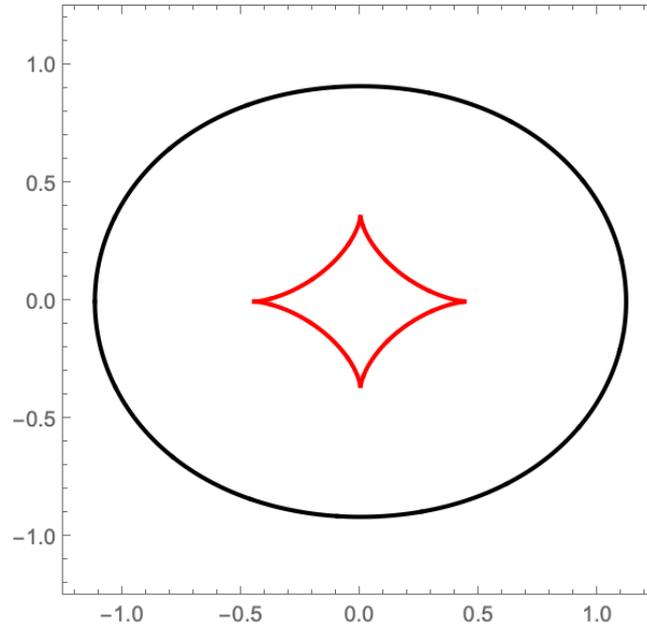


Figure 8.9: The critical curve (black) in the lens plane and caustic curve (red) in the image plane of a non-rotating point lens with a shear term superimposed on top of each other ($\gamma_1 = 0$ and $\gamma_2 = 0.2$).

8.6 The small spin approximation

In the main text, we use the small spin approximation, in which we assume that the modulus of the spin parameter α is small. However, we have not yet discussed the meaning of “small” in this context. Effectively, when deriving the time delay function, we used the Lense-Thirring metric, the weak field approximation of the Kerr metric in ACMC (asymptotically Cartesian and mass centered) coordinates [33, 34], *i.e.*, the Kerr metric expanded to first order in r_s/r and the Kerr parameter a . The limitations of the weak field approximation are known: the Lense-Thirring metric is suitable to describe the spacetime surrounding stars, but does not apply to neutron stars or black holes [27]. The extension of the lens to a Kerr black hole will be discussed elsewhere. In this appendix, we will be concerned with the regularity of the next-order term in the small spin approximation. As it has extensively been discussed in the main text, one must be careful when taking the small spin limit since the terms accompanying the small spin parameter might diverge. Here, we compute the next order correction in the Kerr parameter $\sim a^2$ to the time delay and check that this term does not include singularities (apart from the physical singularity at $\boldsymbol{x} = 0$ discussed in Sec. 8.2.1).

The derivation presented here is for a simplified setup, which differs from the one detailed in the main text, but makes the derivation more transparent. The setup is represented in Fig. 8.10, where the coordinate system is located on the black hole in the lens plane with the spin vector aligned with the z -axis. The black hole is located along the line of sight (represented by a dashed line) such that the y -axis lies along the line of sight and points towards the source. The lens is located at a distance D_{OL} from the observer, while the source lies at a distance D_{OS} . These distances are considered to be much larger than the characteristic Schwarzschild radius $r_s = 2M$ (in geometrized units $G = c = 1$). Given the set-up of the problem, in which the characteristic distances from the observer to the lens D_{OL} and from the lens to the source $D_{LS} = D_{OS} - D_{OL}$ are much larger than the size of the black hole r_s , the signal from the source will propagate in a spacetime that is very similar to the Minkowski spacetime everywhere but close to the lens.

The Kerr metric in ACMC coordinates to quadratic order in the Kerr parameter a and

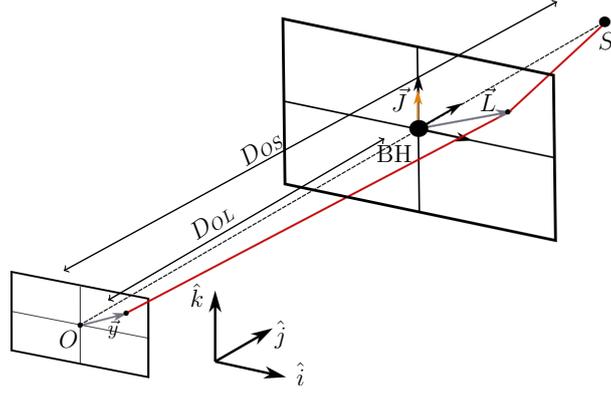


Figure 8.10: Schematics of a null ray being lensed by a slowly rotating Kerr black hole. The distance from the observer to the lens is D_{OL} , and the total distance from the source to the observer is D_{OS} . The coordinate system is placed on the black hole, with the spin vector pointing in the vertical direction. The ray intersects the lens plane at \vec{L} .

linear order in r_s/r reads [33]

$$g_{tt} = -1 + \frac{r_s}{r} - \frac{3r_s a^2 \cos^2 \theta}{2r^3} \quad (8.48a)$$

$$g_{t\phi} = -\frac{ar_s \sin^2 \theta}{r} \quad (8.48b)$$

$$g_{rr} = 1 + \frac{r_s}{r} - \frac{a^2}{r^2} - a^2 \frac{r_s(4 + \cos^2 \theta)}{2r^3} \quad (8.48c)$$

$$g_{\theta\theta} = r^2 \left(1 + \frac{a^2}{r^2}\right) \quad (8.48d)$$

$$g_{\phi\phi} = r^2 \sin^2 \theta \left(1 + \frac{a^2}{r^2} + \frac{r_s a^2 \sin^2 \theta}{r^3}\right) \quad (8.48e)$$

$$g_{r\theta} = -\frac{r_s a^2 \cos \theta \sin \theta}{r^2}. \quad (8.48f)$$

From this metric, it is straightforward to compute the null geodesic equation

$$\dot{t} = E \left(1 + \frac{r_s}{r}\right) - \frac{aJ}{r^2} \frac{r_s}{r} - 3E \frac{a^2}{2r^2} \frac{r_s}{r} \cos^2 \theta \quad (8.49a)$$

$$\dot{\phi} = \frac{J}{r^2 \sin^2 \theta} + \frac{aE}{r^2} \frac{r_s}{r} - \frac{a^2 J}{r^4} \left(\frac{1}{\sin^2 \theta} + \frac{r_s}{r}\right) \quad (8.49b)$$

$$0 = (r^2 + a^2)\ddot{\theta} + 2r\dot{r}\dot{\theta} - \frac{(r^2 - a^2)J^2}{r^4 \sin^2 \theta} \cot \theta \quad (8.49c)$$

$$- \frac{a^2 r_s}{4r^3} \sin 2\theta (3E^2 - 3\dot{r}^2 + 2r\ddot{r})$$

$$\dot{r}^2 = E^2 \left(1 + \frac{a^2(r + r_s \sin^2 \theta)}{r^3}\right) - \frac{2aJr_s E}{r^3} \quad (8.49d)$$

$$- \frac{r - r_s}{r^3} (r^4 \dot{\theta}^2 + J^2 \csc^2 \theta)$$

$$+ \frac{a^2}{4r^5} \left(6J^2 r_s + r^4(3r_s - 8r)\dot{\theta}^2 - \frac{6r_s J^2}{\sin^2 \theta}\right)$$

$$- r_s r^3 \dot{\theta} (r\dot{\theta} \cos 2\theta - 4\dot{r} \sin 2\theta),$$

where E and J are the energy and angular momentum of a test particle moving on a null geodesic. The dot indicates a derivative with respect to the proper time along the geodesic trajectory. Notice that up to order we are working in, that is, $\mathcal{O}(a^3, r_s a^2, r_s^2)$ the geodesic equation coincides with the usual geodesic equation of the Kerr metric in Boyer-Lindquist coordinates.

The time delay is the proper time as measured by an observer at rest between two events: the emission of light from a faraway source and the reception of the ray. The relationship between the proper time of the observer and the coordinate time to first order in r_s/D_{OL} is simply [34]

$$\Delta\tau = \sqrt{-g_{tt}}\Delta t = \left(1 - \frac{r_s}{2D_{OL}} + \frac{3a^2}{4D_{OL}^2} \frac{r_s}{D_{OL}} \cos^2\theta\right) \Delta t \quad (8.50)$$

where Δt is the coordinate time it takes the ray to travel from the source to the observer after being deviated by the lens. Taking into account that we work in the thin lens approximation so that the distance D_{OL} is large compared to r_s , we can approximate

$$\Delta\tau \approx \Delta t. \quad (8.51)$$

Hence, we define the time delay as

$$T = \Delta t \Big|_{\text{source}}^{\text{observer}} \quad (8.52)$$

up to order $\mathcal{O}(a^3, r_s a^2, r_s^2, r_s/D_{OL})$, which using the geodesic equation Eq. (8.49) reads

$$T = \int_0^1 \dot{t} d\tau \quad (8.53a)$$

$$= \int_0^1 \left(E \left(1 + \frac{r_s}{r}\right) - \frac{aJ r_s}{r^2 r} - 3E \frac{a^2 r_s}{2r^2 r} \cos^2\theta \right) d\tau \quad (8.53b)$$

where the integral over the proper time τ is taken along the geometric path (described in Fig. 8.10) that the ray would follow in the absence of the lens. Namely,

$$\dot{t} = E \quad (8.54a)$$

$$\dot{\phi} = \frac{J}{r^2 \sin^2\theta} \left(1 - \frac{a^2}{r^2}\right) \quad (8.54b)$$

$$0 = (r^2 + a^2)\ddot{\theta} + 2r\dot{r}\dot{\theta} - \frac{(r^2 - a^2)J^2}{r^4 \sin^2\theta} \cot\theta \quad (8.54c)$$

$$\dot{r}^2 = E^2 \left(1 + \frac{a^2}{r^2}\right) - (2a^2 + r^2)\dot{\theta}^2 - \frac{J^2 \csc^2\theta}{r^2}. \quad (8.54d)$$

This system of differential equations can be solved by realizing that the Kerr parameter only enters as a quadratic correction to the trajectory. Hence, we need to expand all of the relevant quantities as

$$t(\tau) = t_0(\tau) + a^2 t_2(\tau) \quad (8.55a)$$

$$r(\tau) = r_0(\tau) + a^2 r_2(\tau) \quad (8.55b)$$

$$\theta(\tau) = \theta_0(\tau) + a^2 \theta_2(\tau) \quad (8.55c)$$

$$\phi(\tau) = \phi_0(\tau) + a^2 \phi_2(\tau) \quad (8.55d)$$

$$E = E_0 + a^2 E_2 \quad (8.55e)$$

$$J = J_0 + a^2 J_2 \quad (8.55f)$$

The solution to the geodesic equation to zeroth order in the Kerr parameter is

$$r_0(\tau) = \sqrt{D_{LS}^2(\tau - 1)^2 + L^2\tau^2} \quad (8.56a)$$

$$\theta_0(\tau) = \arctan\left(\frac{\sqrt{D_{LS}^2(\tau - 1)^2 + L_1^2\tau^2}}{L_2\tau}\right) \quad (8.56b)$$

$$\phi_0 = \arctan\left(\frac{D_{LS}(1 - \tau)}{L_1\tau}\right) \quad (8.56c)$$

$$J_0 = -D_{LS}L_1 \quad (8.56d)$$

$$E_0 = \sqrt{D_{LS}^2 + L^2} \quad (8.56e)$$

where we have used that the position of the source and the lens in Cartesian coordinates are $\mathbf{x}(\tau = 0) = (0, D_{LS}, 0)$ and $\mathbf{x}(\tau = 1) = (L_1, 0, L_2)$, and where we have defined $L = \|\mathbf{L}\| = \sqrt{L_1^2 + L_2^2}$. This solution corresponds to a straight line in Minkowski space. For the correction of order a^2 we obtain

$$r_2(\tau) = \frac{\tau(\tau - 1)}{\sqrt{D_{LS}^2(\tau - 1)^2 + L^2\tau^2}} \quad (8.57a)$$

$$\theta_2(\tau) = \frac{L_2(L^2 - D_{LS}^2)\tau(1 - \tau)^2}{2L^2\sqrt{D_{LS}^2(\tau - 1)^2 + L_1^2\tau^2}(D_{LS}^2(1 - \tau)^2 + L^2\tau^2)} \quad (8.57b)$$

$$\phi_2(\tau) = \frac{L_1(L^2 - D_{LS}^2)\tau(1 - \tau)}{2(L^2)D_{LS}(D_{LS}^2(1 - \tau)^2 + L_1^2\tau^2)} \quad (8.57c)$$

$$J_2 = -\frac{L_1(L^2 + D_{LS}^2)}{2D_{LS}L^2} \quad (8.57d)$$

$$E_2 = \frac{1}{\sqrt{D_{LS}^2 + L^2}}, \quad (8.57e)$$

where we have fixed the integration constant by fixing the extremum of the path to its initial and final values according to Fig. 8.10.

The time delay from the source to the lens plane is computed using Eq. (8.53) together with the path (8.55)-(8.57). Similar to the path and the constants of motion E and J , we can expand the time delay function in terms of a

$$T = T_0 + aT_1 + a^2T_2 + \mathcal{O}(a^3, r_s^2), \quad (8.58)$$

where the time delay to zeroth order in the spin is

$$T_0 = \int_0^1 E_0 \left(1 + \frac{r_s}{r_0}\right) d\tau, \quad (8.59)$$

to first-order

$$T_1 = -J_0 r_s \int_0^1 \frac{1}{r_0^3} d\tau, \quad (8.60)$$

and to second-order

$$T_2 = \int_0^1 \left[E_2 \left(1 + \frac{r_s}{r_0}\right) - E_0 \frac{r_s}{r_0} \left(\frac{r_2}{r_0} + \frac{3}{2r_0^2} \cos^2 \theta_0\right) \right] d\tau. \quad (8.61)$$

The last term codifies the time delay due to the rotating background to the second order in the spin. Evaluating these expressions is cumbersome but otherwise straightforward using Eqs. (8.55)-(8.57). The first term yields

$$T_0 = \sqrt{D_{LS}^2 + L^2} + r_s \log\left(\frac{L^2 + L\sqrt{D_{LS}^2 + L^2}}{D_{LS}\sqrt{D_{LS}^2 + L^2} - D_{LS}^2}\right), \quad (8.62)$$

the linear term in the spin is

$$T_1 = \frac{L_1 r_s (D_{LS} + L)}{D_{LS} L^2}, \quad (8.63)$$

while the quadratic term in the spin is

$$T_2 = \frac{1}{\sqrt{D_{LS}^2 + L^2}} + r_s \frac{-2D_{LS}^3 L_2^2 + D_{LS}^2 L(L_1^2 + L^2) + 2D_{LS} L_1^2 L^2 - L_2^2 L^3}{2D_{LS}^2 L^4 \sqrt{D_{LS}^2 + L^2}}. \quad (8.64)$$

Next, we need to evaluate the time delay along the second half of the path $(r(\tau), \theta(\tau), \phi(\tau))$, which can be obtained by replacing $\tau \rightarrow 1 - \tau$ and $D_{LS} \rightarrow -D_{OL}$ in Eqs. (8.55)-(8.57) for an observer located at the line of sight. Solving the geodesic equation with these boundary conditions and evaluating Eqs. (8.58)-(8.61) yields

$$T_0 = \sqrt{D_{OL}^2 + L^2} + r_s \log \left(\frac{D_{OL}^2 + D_{OL} \sqrt{D_{OL}^2 + L^2}}{L \sqrt{D_{OL}^2 + L^2} - L^2} \right). \quad (8.65)$$

for the zeroth order contribution to the time delay. The first and second-order contributions for the observer-lens path can be obtained by replacing $D_{LS} \rightarrow -D_{OL}$ in Eqs. (8.63) and (8.64).

Combining the time delays for the paths between the source and lens and the lens and image plane, and expanding to first order in $1/D_{LS}$ and $1/D_{OL}$, and zeroth order in r_s/D_{LS} and r_s/D_{OL} yields

$$T = D_{OS} + \frac{D_{OS} L^2}{2D_{LS} D_{OL}} - 2r_s \log \left(\frac{L}{2\sqrt{D_{LS} D_{OL}}} \right) + \frac{2ar_s L_1}{L^2} - a^2 \frac{2L_2^2 r_s}{L^4} \quad (8.66)$$

Using an analogous normalization to the one discussed in the main text (8.7), but taking into account that \mathbf{L} is a *distance* in the lens plane rather than an angle,

$$\boldsymbol{\alpha} = \frac{\mathbf{a} \times \mathbf{n}}{r_E} = -\frac{a}{r_E} (1, 0), \quad \mathbf{x} = \frac{\mathbf{L}}{r_E}, \quad (8.67)$$

with the Einstein radius

$$r_E = \sqrt{2r_s \frac{D_{LS} D_{OL}}{D_{OS}}}, \quad (8.68)$$

and $n = (0, 1, 0)$, we obtain

$$T(x) = 2r_s \left(\phi_0 + \frac{\|\mathbf{x} - \mathbf{y}\|^2}{2} - \log x - \frac{\alpha x_1}{x^2} - \frac{\alpha^2 x_2^2}{x^4} \right) \quad (8.69)$$

where $x = \|\mathbf{x}\|$ is the norm of the vector $\mathbf{x} = (x_1, x_2)$ and we have taken into account the shift in the image plane given by \mathbf{y} (see Fig. 8.10). The phase ϕ_0 is the collection of the constant terms (independent of \mathbf{L}) in Eq. (8.66). In covariant form Eq. (8.69) reads

$$T(x) = 2r_s \left(\phi_0 + \frac{\|\mathbf{x} - \mathbf{y}\|^2}{2} - \log x + \frac{\boldsymbol{\alpha} \cdot \mathbf{x}}{x^2} - \frac{(\boldsymbol{\beta} \cdot \mathbf{x})^2}{x^4} \right), \quad (8.70)$$

where $\boldsymbol{\beta} = \mathbf{a}/r_E$. From the last term in Eq. (8.70), it follows that the second-order correction in the Kerr parameter to the time delay is well-behaved, given that no unphysical singularities are introduced. This discussion shows that the slow spin expansion is well-defined to this order.

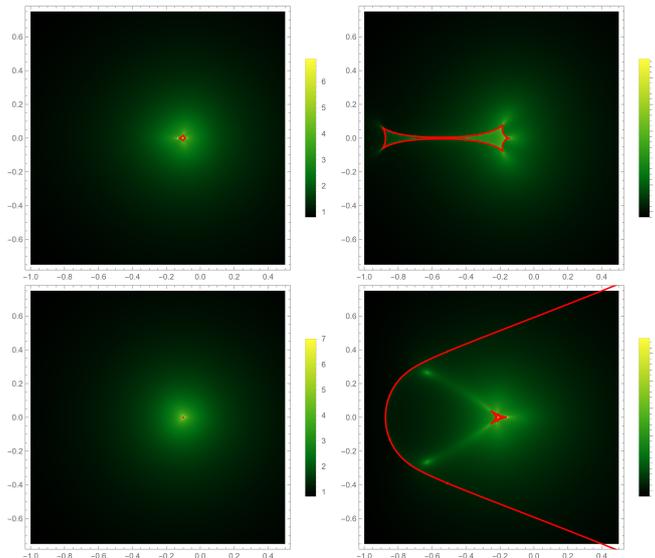


Figure 8.11: Intensity profile within the geometric optics approximation (Eq. (8.16)) for $\alpha = 0.1$ (first column) and $\alpha = 0.2$ (second column). We used the time delay to first order in the Kerr parameter in Eq. (8.4) (first row) and to second order in a as in Eq. (8.70) (second row). The caustics are depicted in red.

To estimate the validity of the slow spin approximation, we compare the intensity (computed using the geometric optics approximation Eq. (8.16)) obtained with and without the second-order correction in the Kerr parameter to the time delay. The comparison for $\alpha = 0.1$ and $\alpha = 0.2$ are depicted in the first and second columns of Fig. 8.11. The first row in Fig. 8.11 has been computed using Eq. (8.4), while the second includes the second order correction in a to the time delay in Eq. (8.70). The images in the first column (for the normalized spin $\alpha = 0.1$) display the same features, and in fact, their intensity profile differs only slightly. For $\alpha = 0.2$ we can already see that the intensity profile computed using the second-order correction in α displays different features to our first-order approximation. Hence, the small limit approximation we have described in this paper is valid for $\alpha \lesssim 0.1$. Notice that this is not a very stringent restriction for stars as rotating lenses given that astrophysical values of α will be tiny (since r_E is typically very large).

Finally, a piece of warning is in order: this estimate on the validity of the small spin approximation should be regarded with caution, as it is intended as a rough order of magnitude calculation. We have used the small field approximation of the Kerr metric to compute the next order correction in a to the time delay. Given that the Kerr uniqueness theorem [35] is only applicable to perfectly axisymmetric bodies, it is not guaranteed that the exterior of our rotating star is well approximated by the metric in Eq. (8.48) (that is to say, Birkhoff's theorem does not generalize to rotating stars).

9 Wave asymptotics and their application to astrophysical plasma lensing by Grillo, Cordes

Abstract

Plasma lensing events can have significant observational consequences, including flux density modulations and perturbations in pulse arrival times. In this paper we develop and apply a formalism that extends geometrical optics to describe the effects of two dimensional plasma lenses of arbitrary shape. We apply insights from catastrophe theory and the study of uniform asymptotic expansions of integrals to describe the lensing amplification close

to fold caustics and in shadow regions, and explore the effects of image appearance and disappearance at caustics in the time of arrival (TOA) perturbations due to lensing. The enhanced geometric optics approach successfully reproduces the predictions from wave optics and can be efficiently used to simulate multifrequency TOA residuals during lensing events. Lensing will introduce perturbations both in the way the residuals change as a function of frequency and also in the magnitude and sign of the residuals averaged over a frequency band. The deviations from the expected dispersive ν^{-2} scaling will be most significant when including observations at low frequencies. We examine the consequences of lensing in the context of precision pulsar timing and touch on its potential relevance to the study of FRBs.

9.1 Introduction

The phenomenon of astrophysical plasma lensing has attracted considerable attention ever since the first detections of so-called “extreme scattering events” (ESEs) in the late 1980s (Fiedler et al. 28) and early 1990s (Cognard et al. 12), during which the measured flux density of the observed objects (a millisecond pulsar in the latter case, and a quasar in the former) underwent large fluctuations with a frequency dependent structure over a period of time of the order of months. Subsequent works describing observations of ESEs, such as those by Fiedler et al. [27] and Clegg et al. [10] mentioned the idea, introduced in Cognard et al. [12], that these events were the result of plasma overdensities in the interstellar medium that act as lenses as they cross the line of sight between the Earth and the source of radiation, refracting the incoming radio waves and creating observable regions of focusing and defocusing.

Clegg et al. [11] gave a detailed exposition of the geometric optics of one dimensional Gaussian lenses and performed numerical simulations to find appropriate lens parameters that could match the observed flux fluctuations of specific ESEs, and some subsequent works have also aimed to derive the characteristics of specific lenses deemed to be responsible for particular ESE observations (Pushkarev et al. 50, Bannister et al. 1, Tuntsov et al. 61, Vedantham et al. 63, Kerr et al. 36, Main et al. 43).

More recently, plasma lensing has also been suggested as a possible mechanism to explain certain properties of FRBs (Cordes et al. 23, Dai & Lu 25), and other works have examined different kinds of lens models, as well as their possible observational signatures (Pen & King 48, Er & Rogers 26), although most of the analysis so far has been done in only one dimension and for a few specific lens shapes.

Plasma lensing events do not only have observable effects in the source’s light curve, they also introduce perturbations in the times of arrival (TOAs) of the radiation, via a combination of geometric and dispersive effects. Thus, plasma lensing events can have potentially important consequences for pulsar timing, as the possible detection of low frequency gravitational waves via this method is dependent on our ability to detect $\lesssim 100$ ns deviations in pulse arrival times. In fact, some plasma lensing events have been inferred by their effects on observed pulsar TOAs (Lam et al. 41) and dispersion measures (DMs) (Coles et al. 14), instead of their effects on measured flux density, since in some cases the presence of strong scintillations can effectively mask whatever effects the lensing events have on the source’s light curve.

In contrast to the random fluctuations in the electron column density that are responsible for scintillation, plasma lensing events are produced by larger scale inhomogeneities in the ISM, motivating the use of geometrical optics. Nevertheless, it has been useful for some authors modelling scintillation phenomena to study the effect of nonturbulent phase screens, particularly in the transition regime from weak to strong scintillations (Watson & Melrose 64, Melrose & Watson 44). Furthermore, the underlying optics based on the Kirchhoff diffraction integral (KDI) is the same for both scintillations and plasma lensing, meaning that a considerable amount of the formalism used in the study of scintillations can be applied in the latter context.

A potentially important effect of plasma lensing is the appearance and disappearance of multiple images of the source as the lens crosses the line of sight. Such multiple imaging has been directly observed in cases in which the angular separation of some of the images has been large enough (Gupta et al. 30, Pushkarev et al. 50), and can be inferred from the existence of fringes in the dynamic spectra of pulsars during certain epochs of observation (Cordes & Wolszcan 24, Gupta et al. 31, Cordes et al. 19). The coalescence of images is associated with regions in which a straightforward calculation of the flux using geometric optics diverges; these regions are known as caustics, and the ability to describe these regions is of importance both in the context of plasma lensing and scintillation (Goodman et al. 29, Melrose & Watson 44, Cordes et al. 23). The geometric optics framework, however, is useful because it provides information about the different images, including their amplitudes, phases, and locations, and at the same time provides a relatively simple way of calculating the total flux without the need of finding a full solution to the KDI. Thus, it is desirable to describe the amplification in the caustic regions without having to abandon the geometric optics point of view. Different authors in the astrophysical context have employed a variety of methods to handle the geometrical optics infinities, but so far the problem has not been solved using wave asymptotic methods derived from the geometrical theory of diffraction (Borovikov & Kinber 6) and catastrophe optics (Poston & Stewart 49, Berry & Upstill 3, Stannes 57, Kravtsov & Orlov 39, Katsaounis et al. 34, Kryukovskii et al. 40), in order to predict the potential observational signatures of two dimensional plasma lenses of arbitrary shape.

Our primary goal in this paper is therefore to use wave asymptotic methods to characterize the effects of astrophysical plasma lensing, develop the resulting formalism that describes the observational effects of two dimensional plasma lenses that cross our line of sight, and present some numerical results based on the application of this formalism. We restrict ourselves to cases in which the source of radiation can be accurately regarded as a point source, and focus on the effects of plasma lensing on pulsar timing. The paper is divided as follows. In §9.2, we present what we call the zeroth and first order geometrical optics of two dimensional lenses, which formally yields infinite flux amplitudes at caustic regions. In §9.3 we use wave asymptotic methods to construct a second order geometric optics description. In §9.4 we use the concepts developed in §9.2 and §9.3 to examine the TOA and DM perturbations due to a specific plasma lens realization, and we summarize conclusions in §9.5. We expect to apply the methodology presented here to specific events in subsequent work.

9.2 Zeroth and first order geometric optics

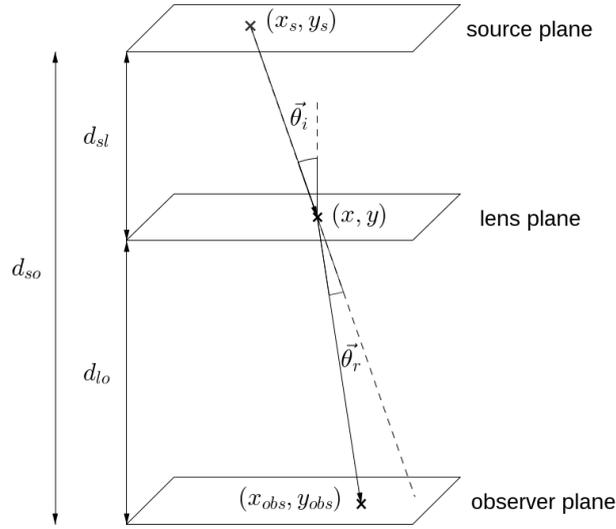


Figure 9.1: Lensing geometry.

9.2.1 Geometrical picture

We follow the basics of the treatment given in Clegg et al. 11 and Cordes et al. 23 but extend their results to two dimensions⁵. We start by defining planes for the source, the lens, and the observer with coordinates \mathbf{x}_s , \mathbf{x} , and \mathbf{x}_{obs} , respectively, with a source-lens distance d_{sl} , a lens-observer distance d_{lo} , and a source-observer distance $d_{so} = d_{sl} + d_{lo}$, as depicted in Figure 9.1. The geometric optics approximation treats the radiation emitted from the source as a cone of rays, and the effects of lensing can be described by the way the lens affects the mapping of the rays from the source plane to the observer plane. From the geometry in the figure, we see that the 2D angle of incidence of a ray into the lens plane θ_i and its deviation angle are given (in the paraxial approximation) by

$$\theta_i = \frac{\mathbf{x}_s - \mathbf{x}}{d_{sl}} \quad (9.1)$$

$$\theta_r = \frac{\mathbf{x}_{obs} - \mathbf{x}}{d_{lo}} - \theta_i. \quad (9.2)$$

Combining into a single equation in terms of θ_r gives the lens equation,

$$\mathbf{x}_s \left(\frac{d_{lo}}{d_{so}} \right) + \mathbf{x}_{obs} \left(\frac{d_{sl}}{d_{so}} \right) = \mathbf{x} + \theta_r \left(\frac{d_{sl}d_{lo}}{d_{so}} \right). \quad (9.3)$$

We now define a new set of coordinates \mathbf{x}' as a combination of the source and observer coordinates scaled by the distances, namely

$$\mathbf{x}' \equiv \mathbf{x}_s \left(\frac{d_{lo}}{d_{so}} \right) + \mathbf{x}_{obs} \left(\frac{d_{sl}}{d_{so}} \right). \quad (9.4)$$

and write the lens equation in the simpler form

⁵Tuntsov et al. [61] also gives a two dimensional account.

$$\mathbf{x}' = \mathbf{x} + \boldsymbol{\theta}_r \left(\frac{d_{sl}d_{lo}}{d_{so}} \right), \quad (9.5)$$

This expression is perfectly general and not only applies to plasma lensing, but to gravitational lensing as well (Schneider et al. 54). The nature of the lensing is what determines the formula for the deviation angle $\boldsymbol{\theta}_r$. A general expression for this angle can be obtained with the additional assumptions that the lens's surface slope is small, and that the lens's medium is uniform. The result of the ray propagating through the lens is that the lens advances or retards the ray's phase, depending on whether the value of the refractive index n_r is greater or smaller than unity, because the phase velocity v_p will be smaller or greater than c . More precisely, we can write this phase difference $\delta\phi_{\text{lens}}$ as

$$\delta\phi_{\text{lens}} = \omega\tau = kc\tau, \quad (9.6)$$

where τ is the propagation time difference between a lensed ray and an unlensed ray, $k = 2\pi/\lambda$ is the wavenumber and ω is the radiation's angular frequency. By this definition, $\tau < 0$ implies that $\delta\phi_{\text{lens}} < 0$ and therefore $v_p > c$. For a lens of length l parallel to the direction of propagation, this is

$$\tau = \frac{l}{c} (n_r - 1). \quad (9.7)$$

For a cold, unmagnetized plasma, the frequency dependent index of refraction is given by

$$n_r = \sqrt{1 - \left(\frac{\omega_e}{\omega}\right)^2} \approx 1 - \frac{\lambda^2 r_e n_e}{\pi}, \quad (9.8)$$

where $\omega_e^2 = 4\pi n_e e^2/m_e$ corresponds to the square of the electron plasma frequency, e is the elementary charge, m_e is the mass of the electron, r_e is the electron's classical radius, and n_e is the electron number density, and the approximate equality comes from the fact that $\omega_e \ll \omega$ for ω within the radio spectrum. According to geometrical optics, rays propagate in the direction normal to the surfaces of constant phase (Born & Wolf 5, Ch. 3), so the refractive angle $\boldsymbol{\theta}_r$ is given by

$$\boldsymbol{\theta}_r = \frac{1}{k} \nabla \delta\phi_{\text{lens}}. \quad (9.9)$$

When the electron column density or dispersion measure perturbation $\text{DM} = n_e l$ at the lens plane varies as a function of transverse position, $\text{DM} \rightarrow \text{DM}(\mathbf{x})$, $\boldsymbol{\theta}_r \neq 0$, and lensing occurs. Putting everything together, the phase perturbation becomes

$$\delta\phi_{\text{lens}}(\mathbf{x}) = -\lambda r_e \text{DM}(\mathbf{x}), \quad (9.10)$$

which implies that the refractive angle is

$$\boldsymbol{\theta}_r = -\frac{\lambda^2 r_e}{2\pi} \nabla \text{DM}(\mathbf{x}) = -\frac{c^2 r_e}{2\pi\nu^2} \nabla \text{DM}(\mathbf{x}). \quad (9.11)$$

For convenience, we write $\text{DM}(\mathbf{x})$ as the product of a maximum perturbation DM_ℓ and a function with unit maximum $\psi(\mathbf{x})$, and take the origin of the lens plane's coordinate system to coincide with the lens's center. Thus Eq. 9.11 takes the form

$$\boldsymbol{\theta}_r = -\frac{c^2 r_e \text{DM}_\ell}{2\pi\nu^2} \nabla \psi(\mathbf{x}). \quad (9.12)$$

We now define the Fresnel scale as $r_F = \sqrt{cd_{sl}d_{lo}/2\pi d_{so}\nu}$, the lens phase as $\phi_0 = -cr_e \text{DM}_\ell/\nu$, and a new parameter $A = r_F^2 \phi_0$, and substitute Eq. 9.12 in terms of these

new quantities into the lens equation, which yields a more compact form that is specific to plasma lensing,

$$\mathbf{x}' = \mathbf{x} + A\nabla\psi(\mathbf{x}). \quad (9.13)$$

Finally, we define dimensionless coordinates using the characteristic lens scales a_x and a_y , such that $u'_x = x/a_x$ and $u'_y = y/a_y$, and explicitly write Eq. 9.13 in its adimensionalized component form. Using the notation $\psi_{ij} = \frac{\partial^{i+j}\psi}{\partial u_x^i \partial u_y^j}$, and defining $\alpha_{x,y} = A/a_{x,y}^2$,

$$\begin{aligned} \begin{bmatrix} u'_x \\ u'_y \end{bmatrix} &= \begin{bmatrix} u_x + \frac{A}{a_x^2}\psi_{10}(u_x, u_y) \\ u_y + \frac{A}{a_y^2}\psi_{01}(u_x, u_y) \end{bmatrix} \\ &= \begin{bmatrix} u_x + \alpha_x\psi_{10} \\ u_y + \alpha_y\psi_{01} \end{bmatrix}. \end{aligned} \quad (9.14)$$

In general, Eq. 9.14 must be solved numerically using a root finding algorithm. More details on the numerical techniques used to produce the examples presented throughout the paper can be found in Appendix 9.8. The vector $\mathbf{u}'(t)$ changes as a function of time as the Earth, the lens, and the source move with different velocities, and the nature of this change will partly determine the observational signature of a specific lens realization. The number of solutions of the equation corresponds to the number of images of the source as seen by the observer, and in general vary as a function of $\mathbf{u}'(t)$ and the parameters $\alpha_{x,y}$.

9.2.2 Zeroth order gain

A large majority of the existing literature on plasma lensing (Clegg et al. 11, Pen & King 48, Tuntsov et al. 61, Cordes et al. 23, Er & Rogers 26, Vedantham et al. 63) derives the gain (or magnification) for an individual image G_j directly from some version of Eq. 9.14, and the total gain is found by adding together the gains of all n images. More specifically, the image magnification is said to correspond to the absolute value of the inverse of the Jacobian of the mapping between the u and u' planes, evaluated at a solution to the lens equation $\mathbf{u} = \mathbf{u}_j^0$,

$$\begin{aligned} G_j &= |\mathcal{J}|^{-1} \\ &= \left| (1 + \alpha_x\psi_{20})(1 + \alpha_y\psi_{02}) - \alpha_x\alpha_y\psi_{11}^2 \right|^{-1} \end{aligned} \quad (9.15)$$

and the total gain is

$$G = \sum_{j=1}^n G_j. \quad (9.16)$$

We refer to this expression as the “zeroth order” geometrical optics gain. It corresponds to a sum of intensities, and as such it fails to take into account the interference between the images that arises from the phase differences of the corresponding fields. An accurate description of the interference pattern can be obtained by solving the Kirchhoff diffraction integral (KDI), which we introduce below.

9.2.3 The 2D Kirchhoff diffraction integral

Once we adopt a wave description of the radiation, the scalar wavefield as a function of position with respect to the source is given by the time independent Helmholtz equation. The general form of the KDI is a formal solution to the Helmholtz equation (Born & Wolf 5, Ch. 8; Thorne & Blandford 60, Ch. 8). In the paraxial approximation and for the near field, as is

the case for AU sized lenses and astronomical distances, the integral can be written in terms of dimensionless coordinates (Goodman et al. 29, Melrose & Watson 44, Cordes et al. 23)

$$\varepsilon(\mathbf{u}', \nu) = \frac{a_x a_y}{2\pi r_F^2} \iint d^2 u \exp(i\Phi), \quad (9.17)$$

where the phase Φ is the sum of a geometric term and the phase perturbation due to the lens, $\delta\phi_{\text{lens}} = \phi_0 \psi(\mathbf{u})$,

$$\Phi(\mathbf{u}', \mathbf{u}, \nu) = \frac{1}{2r_F^2} [a_x^2(u_x - u'_x)^2 + a_y^2(u_y - u'_y)^2] + \phi_0 \psi(\mathbf{u}). \quad (9.18)$$

The integral is normalized such that in the absence of a lens (ie. $\delta\phi_{\text{lens}} = 0$), $\varepsilon(\mathbf{u}', \nu) = 1$ for all \mathbf{u}' and ν . Analytic solutions to the integral are only available for a few specific forms of ψ (Watson & Melrose 64). As detailed in Appendix 9.6, this representation of Φ allows us to write the integral as a convolution of two functions, which can then be solved numerically by employing the convolution theorem and the Fast Fourier Transform (FFT). However, this method is only adequate for lenses that have sizes that are a small fraction of an AU and in cases where $|\phi_0|$ is relatively small, because the required grid size for proper sampling grows prohibitively large as the oscillations of $\exp(i\Phi)$ become more pronounced.

An approximate solution that grows more accurate as the strength of the lens increases follows by applying the method of stationary phase. For a rapidly oscillating two dimensional integral of the form $I(\mathbf{x}) = \iint d^2 \mathbf{x} g(\mathbf{x}) \exp[if(\mathbf{x})]$, the stationary phase lemma (Bleistein & Handelsman 4) indicates that the principal contributions to the integral's value come from the points in which the phase is stationary, that is, where the derivatives of the phase vanish, $f_{10} = f_{01} = 0$. In the general case where these points $\mathbf{x} = \mathbf{x}_j^0$ are complex, each provides a contribution to the integral I_j given by (Connor 15)

$$I_j = \frac{2\pi i g_j \exp(i f_j)}{\Delta_j^{\frac{1}{2}}}, \quad (9.19)$$

where $f_j = f(\mathbf{x}_j^0)$, $g_j = g(\mathbf{x}_j^0)$, $\Delta_j = f_{20}f_{02} - f_{11}^2$ evaluated at \mathbf{x}_j^0 , and the square root in the denominator is taken to be positive or negative depending on the context. When the stationary point is purely real, the contribution reduces to (Bleistein & Handelsman 4, Cooke 17)

$$I_j = \frac{2\pi g_j \exp[i f_j + \frac{i\pi}{4}(\delta_j + 1)\sigma_j]}{|\Delta_j|^{\frac{1}{2}}}, \quad (9.20)$$

where $\sigma_j = \text{sgn}(\Delta_j)$, $\delta_j = \text{sgn}(f_{02})$, and the square root in the denominator is now taken to be positive. In the case of the KDI as given in the form Eq. 9.17, the points of stationary phase correspond to the points that satisfy the two dimensional equation

$$\begin{bmatrix} \Phi_{10} \\ \Phi_{01} \end{bmatrix} = \begin{bmatrix} \frac{a_x^2(u_x - u'_x)}{r_F^2} + \phi_0 \psi_{10} \\ \frac{a_y^2(u_y - u'_y)}{r_F^2} + \phi_0 \psi_{01} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}. \quad (9.21)$$

A quick examination reveals that this is precisely equivalent to the lens equation Eq. 9.14, given our definitions of the parameters $\alpha_{x,y}$, which therefore implies that solving the KDI by the method of stationary phase leads to geometric optics⁶. Performing the appropriate substitutions

⁶It is also possible to derive the geometric optics quantities by directly solving the Helmholtz equation via WKB methods (see, e.g. Born & Wolf 5, Ch. 3; Katsaounis et al. 34; Poston & Stewart 49, Ch. 12).

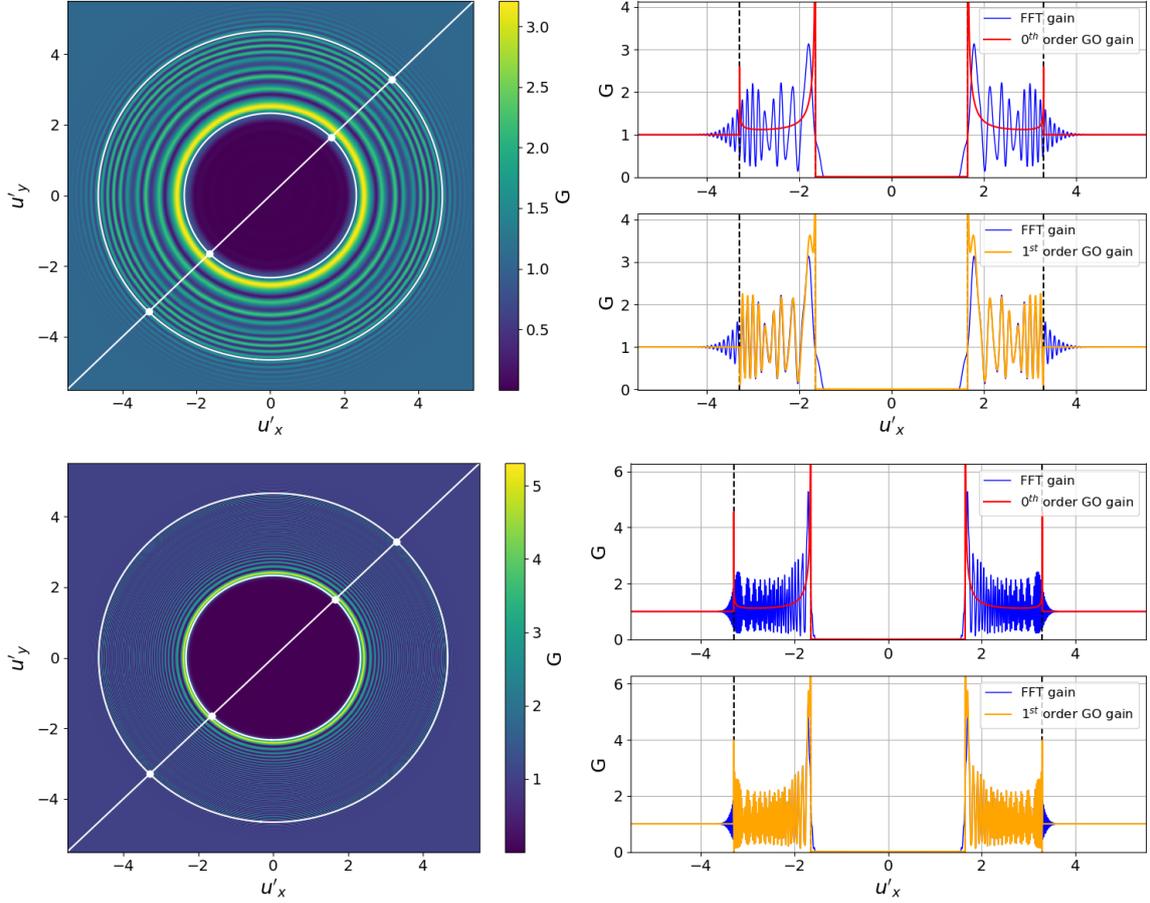


Figure 9.2: Comparison of the gains obtained from a full numerical solution of the KDI, zeroth order geometrical optics, and first order geometrical optics. The top panel corresponds to a lens with $\phi_0 = -50$ rad and the bottom panel corresponds to one with $\phi_0 = -250$ rad (thus $\text{DM}_\ell > 0$ in both cases, and the lenses are diverging). The frequency of observation is $\nu = 0.8$ GHz, $d_{so} = 1$ kpc, $d_{sl} = 0.5$ kpc for both the top and bottom panels. For the top panel, $a_x = a_y = 1.5 \times 10^{-2}$ AU, and for the bottom panel, $a_x = a_y = 1.5\sqrt{5} \times 10^{-2}$ AU. The lens shape is described by a two dimensional Gaussian, $\psi(\mathbf{u}) = \exp(-u_x^2 - u_y^2)$. The left column shows color maps of the gain obtained by solving the KDI via the FFT. The white circles correspond to caustic curves, and the straight white line shows the path of the observer through the u' plane. The right column shows the gain along this path as calculated via the FFT method, zeroth order geometrical optics, and first order geometrical optics. The points of intersection between the caustics and the observer path are marked by white points in the left column and by dashed black lines on the plots in the right column. The geometric optics gain at the caustics is formally infinite, so the GO gains were evaluated up to a short distance away from the caustic.

in Eq. 9.21, we have that the scalar field ε_j^r due to a real stationary point is

$$\begin{aligned}\varepsilon_j^r(\mathbf{u}', \nu) &= \frac{a_x a_y}{r_F^2 |\Delta_j|^{\frac{1}{2}}} \exp \left[i\Phi_j + \frac{i\pi}{4}(\delta_j + 1)\sigma_j \right] \\ &= \frac{\exp \left[i\Phi_j + \frac{i\pi}{4}(\delta_j + 1)\sigma_j \right]}{\left| (1 + \alpha_x \psi_{20})(1 + \alpha_y \psi_{02}) - \alpha_x \alpha_y \psi_{11}^2 \right|^{\frac{1}{2}}},\end{aligned}\quad (9.22)$$

where now we have $\Delta_j = \Phi_{20}\Phi_{02} - \Phi_{11}^2$, $\sigma_j = \text{sgn}(\Delta_j)$, $\delta_j = \text{sgn}(\Phi_{02})$, and all quantities are evaluated at the stationary points, $\mathbf{u} = \mathbf{u}_j^0$. This gives the normalized scalar field due to one image of the source, with a maximum amplitude

$$A_j = |\mathcal{J}|^{-\frac{1}{2}} = \frac{a_x a_y}{r_F^2 |\Delta_j|^{\frac{1}{2}}}\quad (9.23)$$

and an oscillating component with phase

$$\beta_j^r = \Phi_j + \frac{\pi}{4}(\delta_j + 1)\sigma_j.\quad (9.24)$$

The total scalar field due to real solutions of the lens equation is simply the sum of the contributions from the n real stationary points,

$$\varepsilon^r(\mathbf{u}', \nu) = \sum_{j=1}^n \varepsilon_j^r = \sum_{j=1}^n A_j e^{i\beta_j^r}.\quad (9.25)$$

The gain can then be obtained by taking the squared modulus of this last expression, $G = |\varepsilon^r(\mathbf{u}', \nu)|^2$. This is the ‘‘first order’’ geometrical optics gain. The presence of the oscillatory component in each of the images results in interference. As noted above, this is not correctly captured by Eq. 9.16.

9.2.4 Accuracy and regions of applicability

A curious feature of the phasors that emerge from the stationary phase solutions is that they include not only the geometric phase Φ but also a potential phase shift related to the signs of the second derivatives at the stationary points. This phase shift is physically associated with the passage of a ray through a caustic. A caustic corresponds to a surface in parameter space that yields a null Jacobian, $\mathcal{J} = 0$ (Berry & Upstill 3, Kravtsov & Orlov 39). As we approach a caustic, $A_j \rightarrow \infty$, and the zeroth and first order geometric optics approximation fail. The reason for this failure is that the approximations do not take into account diffractive effects that occur due to the finite frequency of the waves. Caustics also correspond to boundaries that separate regions in parameter space that contain different numbers of real solutions to the lens equation.

Figure 9.2 illustrates the difference between the gains obtained from the zeroth and first order approaches in the case of an overdense ($\text{DM}_\ell > 0$), two dimensional Gaussian lens described by $\psi(\mathbf{u}) = \exp(-u_x^2 - u_y^2)$ with equal lens scales, $a_x = a_y$. For a fixed frequency of observation, the wave optics amplification as a function of \mathbf{u}' can be calculated by solving the KDI using the FFT. The left column shows this amplification as a function of \mathbf{u}' for the range $-5.5 \leq u'_x \leq 5.5$ and two different values of ϕ_o , -50 rad (top) and -250 rad (bottom). The white circles correspond to the caustics, and the straight white line denotes the observer’s path along the plane. The right column shows the zeroth (red) and first (orange) order gains along the path superposed with the wave optics gain (blue) for both cases.

From the figure, we can see that unlike the zeroth order approximation, the first order approach is able to reproduce the wave optics oscillations accurately in bright regions that contain more than one real solution to the lens equation. However, wave optics also predicts that in regions with only one real image of the source, the observer should still see an interference pattern that decays (grows) exponentially as she crosses from the caustic's bright (dark) side to the dark (bright) side.

For instance, the Gaussian lens from the figure shows two sets of circular caustics in the u' plane. An observer crossing the u' plane through its center will pass through three regions in which the form of the gain is qualitatively different. Far away from the two caustic zones, $G = 1$, and the intensity shows no modulations due to lensing. This corresponds to the dark side of the outer caustic. As the observer approaches the outer caustic singularity, the intensity starts showing oscillations whose amplitude grows exponentially, even though there is still only one real solution to the lens equation. Crossing into the region between the two caustics, the oscillations's amplitude reach a peak shortly after the boundary, and the observer sees three images corresponding to three real solutions to the lens equation. After that, the amplitude decays and then recovers, peaking right next to the boundary that separates the bright region from the dark side of the inner caustic. This dark region contains a single, highly demagnified image of the source, but there is still a hint of an exponentially decaying interference pattern that disappears a short distance away from the boundary. After crossing the center, an equivalent pattern is observed in reverse as the observer moves from the inner dark side to the bright region and then to the outer dark side. Clegg et al. [11], Melrose & Watson [44], and Cordes et al. [23] studied an analogous lens shape in one dimension. The short paper by Stinebring et al. [58] presents similar two dimensional plots without the geometrical optics curves.

In summary, then, although the inclusion of ray interference dramatically improves the accuracy of the geometric optics approximation, this approach is still unable to reproduce the correct form of the gain close to the caustic singularity (where it blows up), and on the dark side of caustic boundaries (where it fails to account for oscillations).

9.3 Second order geometric optics

9.3.1 Complex rays

So far, we have limited our analysis to the case in which coordinates in the u plane, and the solutions to the lens equation are purely real. In order to reproduce the oscillations that occur in the caustic shadows, however, it is necessary to extend the analysis to the complex plane. When two or more real roots of the lens equation merge at a caustic, they reemerge at the dark side as a complex conjugate pair of solutions to the lens equation that yield a complex phase $\Phi_{\pm} = \Phi_r \pm i\Phi_i$. $\Phi_i > 0$ grows as we move farther into the shadow side in parameter space. From Eq. 9.19, we can write the field ε_{\pm}^c due to this complex conjugate pair as

$$\varepsilon_{\pm}^c(\mathbf{u}', \nu) = A e^{\mp\Phi_i} e^{i\beta_{\pm}^c}, \quad (9.26)$$

where $A = a_x a_y |\Delta_{\pm}|^{-\frac{1}{2}} / r_F^2$ is the same as in the case of a purely real stationary point (Eq. 9.24), and $\beta_{\pm}^c = \Phi_r + \pi/2 - \text{Arg}(\Delta_{\pm})/2$. This expression implies that ε_+^c decreases exponentially as a function of Φ_i , whereas ε_-^c increases exponentially. The exponentially increasing solution can be disregarded as unphysical (Kravtsov et al. 38), but the exponentially decaying contribution can be included as part of the asymptotic approximation to the KDI. Doing so effectively reproduces the shadow side oscillatory pattern predicted by wave optics, as long as we remain far enough away from the caustic. At the caustic, the complex conjugate pair of solutions merge, and $A \rightarrow \infty$.

The idea of looking for complex solutions to the lens equation has surfaced in a variety of contexts. Schramm & Kayser [56] apply the concept to gravitational lensing, and Budden & Terry [7] apply it in the context of radio ray tracing in the atmosphere. There is also a direct connection between complex stationary points, the method of steepest descent, and hyperasymptotics of oscillatory integrals (Kaminski 33, Howls 32).

9.3.2 Caustic location and extent of the caustic zone

In the language of geometric optics, caustics correspond to envelopes of families of rays, and are formed at the surfaces on which rays cross each other. Determining the parameter values for ray crossings to occur is, in general, a non trivial problem in more than one dimension and for an arbitrary lens shape. For a fixed frequency of observation, the necessary condition is that

$$(1 + \alpha_x \psi_{20})(1 + \alpha_y \psi_{02}) - \alpha_x \alpha_y \psi_{11}^2 = 0 \quad (9.27)$$

for at least some value of \mathbf{u} . If this is the case, caustic curves will show up in the u plane, and their form in the u' plane can be determined by mapping these curves via the lens equation. The caustic curves plotted over the colormaps in the left column of Figure 9.2 were constructed using this method.

On the other hand, for a fixed u' coordinate, the locations of caustics in the frequency line need to be determined by solving the set of equations

$$\begin{aligned} \frac{\psi_{10}\psi_{01}}{\Delta u_x \Delta u_y} + \frac{\psi_{20}\psi_{01}}{\Delta u_y} + \frac{\psi_{02}\psi_{10}}{\Delta u_x} + \psi_{20}\psi_{02} - \psi_{11}^2 &= 0 \\ \left(\frac{a_x}{a_y}\right)^2 \frac{\Delta u_x}{\psi_{10}} - \frac{\Delta u_y}{\psi_{01}} &= 0 \end{aligned} \quad (9.28)$$

for \mathbf{u} , where $\Delta u_{x,y} = u'_{x,y} - u_{x,y}$. The caustics will be located at frequencies ν_{caus} , given by

$$\begin{aligned} \nu_{\text{caus}} &= \frac{c}{a_x} \left[\frac{d_{sl} d_{lo} r_e \text{DM}_\ell}{2\pi d_{so} \Delta u_x} \psi_{10} \right]^{\frac{1}{2}} \\ &= \frac{c}{a_y} \left[\frac{d_{sl} d_{lo} r_e \text{DM}_\ell}{2\pi d_{so} \Delta u_y} \psi_{01} \right]^{\frac{1}{2}}, \end{aligned} \quad (9.29)$$

evaluated at the solutions of Eq. 9.28 for which the argument under the square root is positive. Numerical results indicate that the formation of caustics at fixed frequencies, for lenses with Gaussian-like shapes (with a maximum electron column density at the center that falls off relatively quickly) occurs when $\alpha_{x,y} \lesssim -1.2$ (for the positive DM_ℓ case) and $\alpha_{x,y} \approx 0.5$ (for the negative DM_ℓ case). If both α_x and α_y satisfy this condition, two sets of caustics form; if only one does, just one set appears.

A consequence of this requirement is that larger lenses require larger magnitudes of DM_ℓ in order to form caustics in the u' plane at a fixed frequency of observation. Thus, small values of DM_ℓ will only lead to caustic formation in cases involving small lenses or highly elongated lenses. For example, keeping the relevant distances fixed at $d_{so} = 1$ kpc and $d_{sl} = 0.5$ kpc, a value of $\text{DM}_\ell = \pm 10^{-6}$ pc cm⁻³, which corresponds to a lens phase of $\phi_0 \approx \mp 33$ rad at 0.8 GHz, yields a maximum value of $a_{x,y} \approx 2.4 \times 10^{-2}$ AU for the overdense case and $a_{x,y} \approx 3.6 \times 10^{-2}$ AU for the underdense case. Ray crossings for lenses with $a_{x,y} \approx 1$ AU would require a minimum value of $|\text{DM}_\ell| \approx 2 \times 10^{-3}$ pc cm⁻³ for the diverging lens and $|\text{DM}_\ell| \approx 7 \times 10^{-4}$ pc cm⁻³ for the converging lens, which correspond to lens phases at 0.8 GHz of $\phi_o \approx -5.8 \times 10^4$ rad and $\phi_o \approx 2.4 \times 10^4$ rad, respectively. Changing the lens-observer distance d_{so} and source-lens distance d_{sl} also leads to changes in $\alpha_{x,y}$, although not in a very simple way because the value of $d_{lo} = d_{so} - d_{sl}$ also

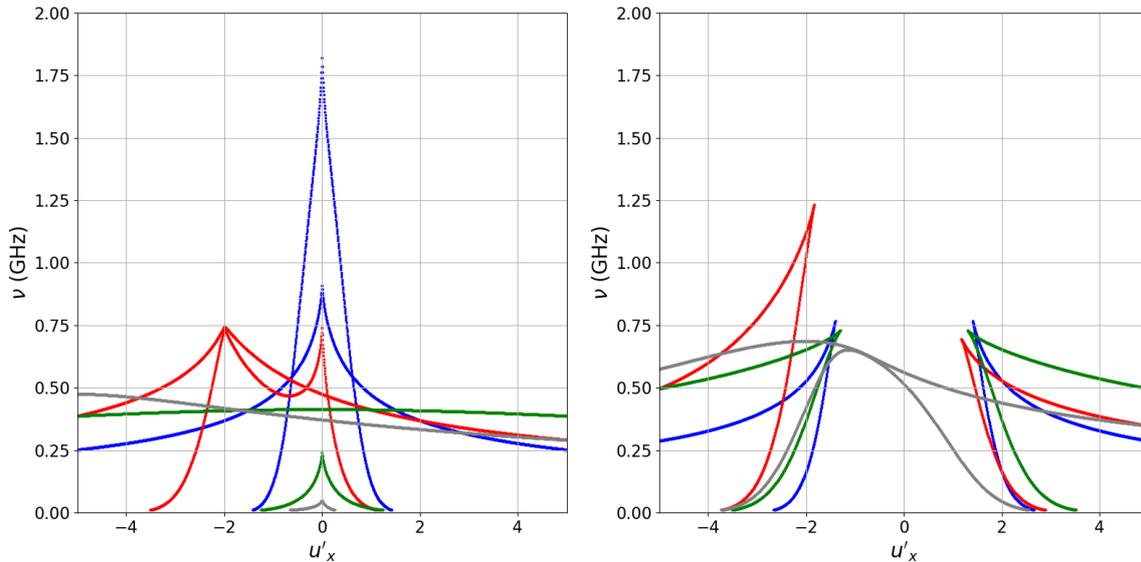


Figure 9.3: Caustic curves in the dynamic spectra of underdense (left) and overdense (right) two dimensional Gaussian lenses for different paths across the u' plane. The blue caustics derive from a path with slope $m = 1$ and y-intercept $n = 0$, the red caustics have $m = 0.5$ and $n = 1$, the green caustics correspond to $m = 0$, $n = 1.5$, and the grey caustics are produced by $m = 0.3$ and $n = 2$. We use a value of $DM_\ell = \pm 10^{-3} \text{ pc cm}^{-3}$, which corresponds to a lens phase $\phi_0 \approx \mp 3 \times 10^4 \text{ rad}$. The source-observer distance $d_{so} = 1 \text{ kpc}$, and the source-lens distance $d_{sl} = 0.5 \text{ kpc}$ in both cases. Both lenses have scales $a_x = 0.5 \text{ AU}$ and $a_y = 1 \text{ AU}$.

factors into the expression. In general, however, if we keep d_{sl} fixed at $d_{so}/2$, increasing d_{so} also increases $\alpha_{x,y}$ and makes caustic formation more likely. The radius of the caustic curves tends to increase linearly with $|\alpha_{x,y}|$ ⁷.

Figure 9.3 shows the caustics in the dynamic spectra of a lensing event for underdense (left) and overdense (right) Gaussian lenses for multiple paths along the u' plane, constructed by repeated application of Eq. 9.28 and Eq. 9.29 over a range of u' coordinates. Although the lens parameters are identical in both cases, it can be seen that flipping the sign of DM_ℓ generates a completely different set of caustic curves, and that the path of the observer through the u' plane can also significantly alter the caustic shapes.

Caustics will show up as a function of ν at a fixed value of u' if we search within a range of frequencies that contains a value of ν that leads to at least one of the $\alpha_{x,y}$ parameters having a magnitude larger than the required minimum. Since $|\alpha_{x,y}| \propto \nu^{-2}$, caustic curves in dynamic spectra, such as the ones depicted in Figure 9.3, will show up only at low frequencies.

In practical terms, it is useful to be able to locate caustics as functions of both \mathbf{u}' and ν . Telescope observations made during an observing epoch correspond roughly to observations made at a fixed \mathbf{u}' . Observations with a large enough frequency range would in principle allow us to see the effects of caustics (under the right circumstances) in a single epoch of observation if a lensing event is taking place. At the same time, since the coordinates in u' change as a function of time, we also expect to see caustic effects in observations made within a narrow frequency band over a range of epochs.

At a caustic boundary, two or more images of the source appear or merge, depending on whether the caustic is crossed from one side or the other. In other words, the number of real roots of the lens equation changes by at least two. The first order geometric approximation breaks down in the vicinity of the caustic when two or more images of the source become

⁷An important exception is the underdense ($DM_\ell < 0$) Gaussian lens with $a_x = a_y$, which presents an infinitely small caustic at the center corresponding to a focus, and a single circular caustic surrounding it.

indistinguishable from each other. As noted by Kravtsov & Orlov [39], a useful operational definition for the width of the caustic zone is the boundary at which the absolute value of the geometrical phase difference $|\Delta\Phi_{ij}|$ between two or more roots is less than π ,

$$|\Delta\Phi_{ij}| \lesssim \pi, \quad (9.30)$$

where i, j are the labels of each of the roots. The number of coalescing images determines the type of caustic, as it describes the kind of singularity, or catastrophe, that occurs within the caustic zone.

A number of previous works (Chako 8, Bleistein & Handelsman 4, Cooke 17, Wong 65, Cordes et al. 23) have dealt with the problem of obtaining the maximum gain within this region by employing an extension of the stationary phase method to approximate the gain at the singularity. Although the derived formulae (some of which are presented in Appendix 9.7) are relatively simple to apply and can be useful for some types of analyses, it is not in general correct because the geometric optics approximation breaks down some distance away from the caustic, close to the boundary defined by Eq. 9.30.

9.3.3 Gain inside the caustic zone: catastrophe theory and uniform asymptotics

Catastrophe theory, first developed by the mathematician René Thom (Thom 59) and subsequently applied to optics by Sir Michael Berry and others (Berry 2, Nye 45, Berry & Upstill 3), provides a useful way of categorizing geometric optics singularities. The basic idea is that close to a caustic, the phase function can be locally mapped into a standard form that is determined by the number of merging images. This standard form is expressed in terms of a fixed number of state and control variables, which are related by the mapping to the physical variables. Solving the KDI for the particular case of this standard form yields a transitional approximation that describes the gain within the caustic region.

In general, it is very difficult to rigorously construct a mapping that takes the global form of the phase to the standard form. Instead, the mapping is performed by expanding the phase in a Taylor series at the point that satisfies both the lens equation and Eq. 9.27, in addition to rotating and scaling the coordinate system such that it is possible to match the coefficients present in this form of the phase to the standard form of the catastrophe. This procedure is described in Kravtsov & Orlov [39], and performed specifically for the case of two dimensional scattering screens in the context of scintillation by Goodman et al. [29].

Watson & Melrose [64] rely on an analogous procedure to derive the one dimensional transitional approximation for the case of two merging images, which corresponds to a fold caustic. The fold catastrophe is the first of the seven elementary catastrophes described by Thom in his original work, and it is the simplest to model and describe. In the vicinity of the fold, the phase can be locally mapped to a cubic, and the KDI can be mapped into the canonical integral (Berry & Upstill 3)

$$\begin{aligned} I_{\text{fold}}(\xi) &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} dt \exp \left[i \left(\frac{t^3}{3} + \xi t \right) \right] \\ &= \sqrt{2\pi} Ai(\xi), \end{aligned} \quad (9.31)$$

where ξ denotes a control variable and t denotes a state variable, and $Ai(\xi)$ is the Airy function. The observer sees no real images on the dark side of the caustic, and two images on the bright side, but the intensity at the dark side does not drop to zero instantly as predicted by first order geometric optics. For practical purposes, it is possible to adopt this transitional form within

the caustic region, and revert back to the regular geometrical optics description far away from the caustic, as in Watson & Melrose [64].

A better, more general solution is to employ the method of uniform asymptotics, as initially developed by Chester et al. [9], Ursell [62], and Ludwig [42] for oscillatory integrals, and later explicitly applied to optics and related to catastrophe theory by Kravtsov [37] and Kravtsov & Orlov [39]. This solution enables us to describe the gain in regions both close and far away from the caustics by the application of a single, global expression that employs the integral of the standard form associated with the type of catastrophe involved, the derivatives of this integral, and some combination of the parameters derived from geometrical optics. Close to the caustic, the expression behaves like the transitional approximation, and far away from it, it matches the field given by the regular geometrical optics approximation.

Uniform asymptotic expressions for the fold caustic have been derived by multiple authors starting with Chester et al. [9]⁸, and in general there are slight variations between each of the presented expressions. We derive it here in an intuitive manner.

The general scheme consists in starting with an ansatz with the same number of terms as there are rays involved in the formation of the caustic, one term involving the function corresponding to the canonical caustic integral, and the rest involving its derivatives. Each of these terms is multiplied by an unknown coefficient, and their sum is multiplied by a phasor. For the fold caustic, it is possible to construct the uniform asymptotic simply by starting with the ansatz and matching the relevant parameters to the geometrical optics coefficients far away from the caustic, by employing the asymptotic forms of the Airy function and its derivative for large negative and positive arguments. Thus, for the bright side, we start with an ansatz of the form,

$$\begin{aligned}\varepsilon_{\text{bright}}(\mathbf{u}', \nu) &= e^{i\chi} [g_1 I_{\text{fold}}(\xi) + g_2 I'_{\text{fold}}(\xi)] \\ &= \sqrt{2\pi} e^{i\chi} [g_1 Ai(\xi) + g_2 Ai'(\xi)],\end{aligned}\quad (9.32)$$

where g_j , χ , and ξ are all potentially functions of \mathbf{u}' . From Eq. 9.25, we have that the first order geometrical optics solution in the case of two real rays can be written as

$$\varepsilon^r(\mathbf{u}', \nu) = A_1 e^{i\beta_1^r} + A_2 e^{i\beta_2^r}.\quad (9.33)$$

The asymptotic forms of the Airy function and its derivative for large negative argument are the well known formulas,

$$Ai(\xi) \approx \frac{1}{\sqrt{\pi}} (-\xi)^{-\frac{1}{4}} \cos \left[\frac{2}{3} (-\xi)^{\frac{3}{2}} - \frac{\pi}{4} \right]\quad (9.34)$$

$$Ai'(\xi) \approx \frac{1}{\sqrt{\pi}} (-\xi)^{\frac{1}{4}} \sin \left[\frac{2}{3} (-\xi)^{\frac{3}{2}} - \frac{\pi}{4} \right],\quad (9.35)$$

which are obtained by applying the one dimensional stationary phase method to the integral in Eq. 9.31. Defining $\gamma = 2(-\xi)^{\frac{3}{2}}/3 - \pi/4$, using Euler's identity, and substituting into Eq. 9.32, we get

$$\begin{aligned}\varepsilon_{\text{bright}}(\mathbf{u}', \nu) &= \frac{e^{i\chi}}{\sqrt{2}} \left\{ e^{i\gamma} \left[g_1 (-\xi)^{-\frac{1}{4}} + i g_2 (-\xi)^{\frac{1}{4}} \right] \right. \\ &= \left. + e^{-i\gamma} \left[g_1 (-\xi)^{-\frac{1}{4}} - i g_2 (-\xi)^{\frac{1}{4}} \right] \right\}.\end{aligned}\quad (9.36)$$

⁸Also see Ludwig [42], Connor [15], Stannnes [57], Borovikov & Kinber [6], Kravtsov & Orlov [39], Qiu & Wong [51], Katsaounis et al. [34]

Matching this to Eq. 9.33, we obtain two sets of equations that can be used to determine g_1 , g_2 , χ , and ξ in terms of the geometrical optics amplitudes A_j , and the phases β_j^r . The first set is

$$\begin{aligned} A_1 &= \frac{1}{\sqrt{2}} \left[g_1(-\xi)^{-\frac{1}{4}} + ig_2(-\xi)^{\frac{1}{4}} \right] \\ A_2 &= \frac{1}{\sqrt{2}} \left[g_1(-\xi)^{-\frac{1}{4}} - ig_2(-\xi)^{\frac{1}{4}} \right]. \end{aligned} \quad (9.37)$$

Solving for g_1 and g_2 gives

$$\begin{aligned} g_1 &= \frac{1}{\sqrt{2}}(A_1 + A_2)(-\xi)^{\frac{1}{4}} \\ g_2 &= \frac{i}{\sqrt{2}}(A_1 - A_2)(-\xi)^{-\frac{1}{4}}. \end{aligned} \quad (9.38)$$

The second set of equations is

$$\begin{aligned} \chi + \gamma &= \beta_1^r \\ \chi - \gamma &= \beta_2^r, \end{aligned} \quad (9.39)$$

which leads to

$$\begin{aligned} \chi &= \frac{1}{2}(\beta_1^r + \beta_2^r) \\ \xi &= - \left[\frac{3}{4} \left(\beta_1^r - \beta_2^r + \frac{\pi}{2} \right) \right]^{\frac{2}{3}}. \end{aligned} \quad (9.40)$$

Putting everything together, we obtain the uniform asymptotic for the fold caustic's bright side,

$$\begin{aligned} \varepsilon_{\text{bright}}(\mathbf{u}', \nu) &= \sqrt{\pi} e^{i\chi} \left[(A_1 + A_2)(-\xi)^{\frac{1}{4}} Ai(\xi) \right. \\ &\quad \left. + i(A_1 - A_2)(-\xi)^{-\frac{1}{4}} Ai'(\xi) \right]. \end{aligned} \quad (9.41)$$

The ambiguity in the labeling is resolved by the condition $\beta_1 - \beta_2 + \pi/2 > 0$. The merging rays will in general have opposite parities, with $\beta_j^r - \Phi_j = 0$ for one ray and $\beta_j^r - \Phi_j = \pm\pi/2$ for the other, so this condition is equivalent to $\Phi_1 - \Phi_2 > 0$. Note that even though A_1 and A_2 diverge as they approach the singularity, the quantity $(A_1 + A_2)(-\xi)^{\frac{1}{4}}$ goes to a finite limit, because $\xi \rightarrow 0$ at the caustic. By the same token, although $(-\xi)^{-\frac{1}{4}}$ goes to infinity at the singularity, the quantity $(A_1 - A_2)(-\xi)^{-\frac{1}{4}}$ does not, because $A_1 - A_2 \rightarrow 0$.

At the caustic's dark side, we know from §9.3.1 that far from the singularity, the geometrical optics field reduces to that of a single complex ray, $\varepsilon_+^c = Ae^{-\Phi_i} e^{i\beta_+^c}$. Therefore, our ansatz no longer contains $Ai'(\xi)$. Instead, we have that

$$\varepsilon_{\text{dark}}(\mathbf{u}', \nu) = \sqrt{2\pi} e^{i\chi} g_0 Ai(\xi). \quad (9.42)$$

The asymptotic of $Ai(\xi)$ for large positive argument is

$$Ai(\xi) \approx \frac{\exp\left(-\frac{2}{3}\xi^{\frac{3}{2}}\right)}{2\xi^{\frac{1}{4}}\sqrt{\pi}}. \quad (9.43)$$

Matching coefficients as before, we obtain $\chi = \beta_+^c$, $\xi = \left[\frac{3}{2}\Phi_i\right]^{\frac{2}{3}}$, and $g_0 = A\xi^{\frac{1}{4}}\sqrt{2}$. Thus,

$$\varepsilon_{\text{dark}}(\mathbf{u}', \nu) = 2\sqrt{\pi} e^{i\beta_+^c} A\xi^{-\frac{1}{4}} Ai(\xi). \quad (9.44)$$

Again, even though at the caustic $A \rightarrow \infty$, the expression does not diverge because $\xi \rightarrow 0$ at the same point.

9.3.4 Uniform asymptotics in plasma lensing

For the present case of plasma lenses and astrophysical distances, the idealized situation presented above involving two real images on the caustic's bright side and no real images on its dark side does not actually occur, as the lens is not opaque and the immense distances allow the initial cone of emitted radiation to grow to a size much larger than that of the lens by the time the two encounter each other. Thus, Eqs. 9.41 and 9.44 cannot be applied directly as given: there will always be at least one real ray involved in the description of the field, and the total number of rays will always be odd. The more general way of dealing with such a situation would be to implement the uniform asymptotic for the next catastrophe in the series, the cusp. The canonical integral in that case is

$$I_{\text{cusp}}(\xi_1, \xi_2) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} dt \exp \left[i \left(\xi_1 t + \frac{\xi_2 t^2}{2} - \frac{t^4}{4} \right) \right], \quad (9.45)$$

which is related to the Pearcey integral $P(\xi_1, \xi_2)$ (Pearcey 47) by the relationship $I_{\text{cusp}}(\xi_1, \xi_2) = P^*(-\sqrt{2}\xi_1, -\xi_2)/\sqrt{\pi}$. The observer sees three images in the bright side and one image in the dark side, which is exactly what happens for the Gaussian lens analyzed in Figure 9.2. The corresponding ansatz for the bright side of the caustic would then be

$$\varepsilon(\mathbf{u}', \nu) = e^{i\chi} \left[g_1 I_{\text{cusp}}(\xi_1, \xi_2) + g_2 \frac{\partial I_{\text{cusp}}}{\partial \xi_1} + g_3 \frac{\partial I_{\text{cusp}}}{\partial \xi_2} \right]. \quad (9.46)$$

Finding the unknown parameters g_j , ξ_j , and χ , however, is not possible via implementation of the same matching procedure we used above for the fold caustic. One reason for this is that the asymptotic forms of the Pearcey integral are much more complicated than those of the Airy integral (Paris 46). Instead, the correct strategy involves obtaining systems of equations for the relevant quantities by exploiting the correspondence between the phase function of the canonical integral and the phase function of the KDI at the stationary points, as described in detail by Connor [16] and Katsaounis et al. [34]. Unfortunately, in the case of cusps and higher order catastrophes, it is not possible to express all the unknown parameters as a function of the geometrical optics quantities in a simple form.

For practical purposes, however, this is rarely necessary. Cusps correspond to points in which three solutions of the lens equation merge. These points are connected to each other by curves which correspond to fold catastrophes, where only two images merge. Far from these cusp points, Eq. 9.46 can be written as the sum of the uniform asymptotic for the fold caustic and the regular geometric optics contribution from each of the n images not involved in the formation of the fold lines. This also holds for higher order catastrophes. Thus, as long as we are not too close to catastrophes of higher order, the total field can be written as

$$\varepsilon(\mathbf{u}', \nu) = \varepsilon_{\text{fold}} + \sum_{j=1}^n A_j e^{i\beta_j^r}, \quad (9.47)$$

where $\varepsilon_{\text{fold}}$ is given by Eq. 9.41 or Eq. 9.44 depending on whether we are at the caustic's dark side or bright side.

Figure 9.4 shows a comparison between the gain obtained from the FFT and that obtained using the uniform asymptotic formulas for a slice across the u' plane, for two different lens shapes ψ and lens phases ϕ_0 . Unlike the case of the circular Gaussian lens with positive DM_ℓ depicted in Figure 9.2, the lenses in these figures show cusps as well as folds. In Figure 9.4, both the elliptical Gaussian with $DM_\ell < 0$ (top panels) and the ring-like lens with $DM_\ell > 0$ (bottom panels) show fold lines interrupted by cusp points at which three roots merge and the curvature of the fold lines is reversed. The number of images that can be seen varies depending on the position in the u' plane and the type of lens. For the negative DM_ℓ elliptical Gaussian, the

observer sees one image in the dark side of the outer caustic zone, three images after crossing the outer caustic boundary, and five images in the central caustic. For the ring-like lens in the bottom panels, the number of images is equal to one outside the caustic zones, three inside the mirrored crescent shaped caustics and in between the two central caustic curves, and five at the center. Other lens shapes can show larger numbers of images and catastrophes of higher order. Some examples are shown in Appendix 9.9.

9.3.5 Advantages of second order geometric optics

As long as $|\phi_o| \gg 1$, second order geometric optics is able to produce remarkably accurate results. Unlike the FFT method, it can be implemented for essentially arbitrary values of a_x , a_y , and ϕ_0 without difficulty. We have applied the second order approach only to the case of slices across the u' plane at a fixed frequency of observation, but the equations for the field hold identically if we were to vary any of the parameters present in the phase function Eq. 9.18. Thus, we can use second order geometric optics to produce accurate plots of the gain as a function of ν at a fixed position in the u' plane. Even for small values of the lens scales, constructing such a plot using the FFT would be extremely computationally expensive, as it would require performing two dimensional FFTs at each frequency of observation.

Using the concepts developed so far, we can construct sections of the dynamic spectrum of a lens event, at least for the case in which these show no cusps. Plots of the gain as a function of position along a line in the u' plane at a single frequency will then correspond to horizontal slices of the dynamic spectrum, whereas plots of the gain as a function of ν at fixed u' coordinates will correspond to vertical slices. This is illustrated further in Figure 9.5. From the figure, it is also apparent that larger magnitudes of the maximum column density $|\text{DM}_\ell|$ induce faster oscillations in the gain, and the contributions from complex rays in the shadow sides of caustics become less important.

9.4 Astrophysical applications

9.4.1 TOA perturbations

One of the important potential effects of plasma lensing, in particular with regards to its consequences to pulsar timing, is the issue of perturbations in pulse arrival times. The importance of these potential perturbations has been clear for a long time (see, e.g. Cordes & Wolszcan 24, Cordes et al. 18) and has resurfaced more recently given the potential of PTAs to detect low frequency gravitational waves (Cordes & Shannon 20, Cordes et al. 21) and in the context of FRBs (Cordes et al. 23, Dai & Lu 25). Our analysis will rely on examples that use parameters that are more likely to be relevant for pulsar timing, where the resulting perturbations are in the order of microseconds, and the distances place the source and lens inside the Milky Way galaxy. Nevertheless, the same concepts can be applied to the FRB case by increasing the distances, the lens sizes, and the magnitude of the maximum dispersion measure perturbations.

Geometry and dispersion

Refraction due to plasma lensing invariably introduces a geometric delay into the time of arrival of radiation, independently of whether the lensing effect is produced by an underdensity or an overdensity in the interstellar medium. By Fermat's principle, an unlensed ray will travel in a straight line from the source to the observer, and lensing introduces a deviation from this straight path. Referring to the geometry of Figure 9.1, we can write the magnitude of the

geometric delay Δt_{geo} as

$$\Delta t_{\text{geo}} = t_{g_x}(u_x - u'_x)^2 + t_{g_y}(u_y - u'_y)^2, \quad (9.48)$$

where the $t_{g_{x,y}} = a_{x,y}^2 d_{so} / 2cd_{sl}d_{lo}$ are the geometrical delay coefficients along the $u_{x,y}$ axes. The location of images in the u plane is determined by the coordinates in the u' plane and the lens equation, so for a given image located at $\mathbf{u} = \mathbf{u}_j^0$, we can express the geometric delay as

$$\Delta t_{\text{geo}} = t_{g_x} \alpha_x^2 \psi_{10}^2 + t_{g_y} \alpha_y^2 \psi_{01}^2. \quad (9.49)$$

Independently of the geometric delay, the lens will also introduce a dispersive perturbation in pulse arrival time due to the plasma's effect on the radiation's group velocity v_g . For a cold plasma, $v_g v_p = c^2$, which means that the dispersive perturbation in the TOA is given by

$$\begin{aligned} \Delta t_{\text{DM}} &= \frac{c r_e \text{DM}_\ell}{2\pi\nu^2} \psi(\mathbf{u}) \\ &= 4.149 \text{ ms} \times \frac{\text{DM}_\ell}{\nu^2} \psi(\mathbf{u}), \end{aligned} \quad (9.50)$$

where the second equality applies for a DM_ℓ in units of pc cm^{-3} and ν in GHz. If $\text{DM}_\ell > 0$, the dispersive perturbation will introduce a TOA delay, as the column density of electrons along the line of sight will increase⁹. On the other hand, if $\text{DM}_\ell < 0$, the lens will constitute a ‘‘pinhole’’ in the interstellar medium, and radiation passing through the lens will experience less of a dispersive delay than radiation traveling outside of it.

The total TOA perturbation for each image Δt_j is simply the sum of the geometric and dispersive perturbations,

$$\Delta t_j = \Delta t_{\text{geo}}^j + \Delta t_{\text{DM}}^j. \quad (9.51)$$

When $\text{DM}_\ell > 0$, both perturbations are positive, and the total TOA perturbation will be positive for any combination of parameters, frequency of observation, and position in the u' plane. On the other hand, when $\text{DM}_\ell < 0$, Δt_j can be either positive or negative depending on the relative magnitudes of Δt_{geo}^j and Δt_{DM}^j . For an observer close to the origin of the u' plane and a lens with a maximum dispersion measure perturbation at the center of the lens plane, the maximum TOA advance will occur for solutions to the lens equation that are within the u plane's central region, since at these points the geometric delay will be minimum and the dispersive advance will be maximum. For a fixed position in the u' plane, the magnitude of the geometric perturbation for an individual image will decrease as ν^{-4} (Rickett 52), whereas the dispersive delay will decrease as ν^{-2} , which means that dispersive delays will dominate geometric perturbations at large frequencies. Geometric delays will grow as a function of the lens size, but larger lenses do not necessarily increase the maximum dispersion measure perturbation, so geometric delays acquire more significance as the lens size grows and DM_ℓ stays constant. In general, the magnitude of the total TOA perturbation per image decreases as a function of frequency.

Figure 9.6 shows a sequence of plots of Δt along a path through the u' plane for frequencies of observation 0.8, 1.0, and 1.2 GHz for overdense and underdense Lorentzian lenses with $\text{DM}_\ell = \pm 5 \times 10^{-4} \text{ pc cm}^{-3}$, which gives a lens phase ϕ_o for each of the frequencies of $\sim \mp 1.6 \times 10^4$ rad, $\mp 1.3 \times 10^4$ rad, and $\mp 1.1 \times 10^4$ rad, respectively. For the overdense lens sequence in the top panels, both the geometric and dispersive perturbations are positive. At $\nu = 0.8$ GHz, the geometric contribution dominates over the dispersive contribution, as is apparent by the facts

⁹Of course, it is possible to have a lens function ψ that is both positive and negative depending on \mathbf{u} , such as $\psi(\mathbf{u}) = \sin(u_x) + \sin(u_y)$. Thus, this statement is correct only for lens realizations that have $\psi > 0$ for all \mathbf{u} , which is the case for all the examples shown throughout this work.

that one, the maximum TOA delay occurs far from the origin of the u' coordinate system, where the geometric perturbation is larger than the dispersive perturbation, and two, the minimum delay in the caustic zone occurs at the origin, where the dispersive delay is maximum and the geometric delay is minimum. Outside of the caustic region, the delay is negligible. As we increase the frequency, it can be seen that the difference between the minimum delay at the center and the maximum delay at the edges of the caustic zone becomes less noticeable, as the magnitude of the geometric delay decreases faster than that of the dispersive delay. The maximum number of images produced in the case of the overdense lens is three, and the caustic pattern is very similar to that of an overdense two dimensional Gaussian like the one depicted in Figure 9.2.

The bottom panels, corresponding to the lens with $DM_\ell < 0$, show a different sequence. This time the maximum number of images (five) is seen along the section of the observer's path through u' that is closer to the center of the caustic region, and the caustic curves form cusps as well as folds. The dispersive perturbation is now negative, and is able to overpower the geometric delay only in regions close to the origin, where the geometric delay is at a minimum. Nevertheless, only one of the five images actually shows a TOA advance.

In both the overdense and the underdense case, we see that the total magnitudes of the perturbations decrease as a function of frequency, and almost no lensing effects are apparent at 1.2 GHz, although this is more dramatic for the underdense lens than for the overdense one. Both the size of the caustic zone and the distance between each of the caustic curves decrease as a function of frequency, because of the weakening of the lens's refractive power.

Telescope observations of TOA perturbations during a lensing event

The examples from the previous section apply only to the unrealistic case of measurements performed at an infinitely narrow frequency band, and ignore the fact that in general a telescope will be unable to resolve individual images. In reality, the incident electric field $E(t)$ is sampled as a function of time by the telescope's receiver, and individual pulse shapes are constructed by taking the Fourier transform of $E(t)$, $\tilde{E}(\nu) = \int dt E(t)e^{-2\pi i\nu t}$ and transforming back after filtering $\tilde{E}(\nu)$ with a bandpass of bandwidth $\Delta\nu_r$ centered on frequency ν_0 . Then, the electric field measured by the telescope across a single band E_{band} can be written as (Cordes & Wasserman 22)

$$E_{\text{band}}(t, \nu_0; \Delta\nu_r) = \int_{\nu_0 - \Delta\nu_r/2}^{\nu_0 + \Delta\nu_r/2} d\nu \tilde{E}(\nu) e^{2\pi i\nu t}. \quad (9.52)$$

The pulse profile for the band can then be constructed by taking the square modulus of Eq. 9.52. The effects of lensing can be quantified as follows. Let an unlensed pulse be described by a normalized electric field $V_0(t)$ and Fourier transform $\tilde{V}_0(\nu)$. Then, the lensed pulse over a band V_{band} will be given by

$$V_{\text{band}}(\mathbf{u}', t, \nu_0; \Delta\nu_r) = \int_{\nu_0 - \Delta\nu_r/2}^{\nu_0 + \Delta\nu_r/2} d\nu \tilde{V}_0(\nu) \varepsilon(\mathbf{u}', \nu) e^{2\pi i\nu t}, \quad (9.53)$$

where $\varepsilon(\mathbf{u}', \nu)$ is the normalized scalar field from the monochromatic KDI, Eq. 9.17. As we showed in §9.3, we can accurately and efficiently solve the KDI using second order geometric optics, by expressing $\varepsilon(\mathbf{u}', \nu)$ as a sum of terms of the same form as Eq. 9.47. In practice, $V_0(t)$ looks like modulated white noise, and the processing of the data captured by the telescope will involve heterodyning to baseband, coherently dedispersing, and the folding of multiple pulses to obtain a better signal to noise ratio. Nevertheless, in the context of numerical simulations, we can get an idea about how lensing events will show up in our data by regarding the unlensed pulse as a unit impulse at $t = 0$, $V_0(t) = \delta(t)$. Then, $\tilde{V}_0(\nu) = 1$, and the deviation of V_{band} from V_0 will be exclusively due to the characteristics of the lens and the observing position \mathbf{u}' .

We can mimic how the perturbation will look in real data by convolving $I = |V_{\text{band}}|^2$ with a suitable pulse template and adding white noise. The TOA perturbation for each band can be calculated afterwards using PyPulse¹⁰.

Figure 9.7 shows numerically simulated pulse dynamic spectra, individual image gains and TOA perturbations, and combined TOA perturbations as a function of frequency for a single epoch of observation (fixed \mathbf{u}'). In the top panel, lensing occurs as a result of an underdensity with Gaussian shape, $\psi(\mathbf{u}) = \exp(-u_x^2 - u_y^2)$, whereas in the bottom panel, the lens is overdense with shape $\psi(\mathbf{u}) = \exp\left[-(u_x^2 + u_y^2)^2\right]$. The parameters used, listed in the figure's caption, lead to caustic formation in both cases. The effects of the lenses in pulse TOAs vary dramatically as a function of frequency, especially close to the caustics, where the signal to noise ratio can be observed to increase as the image magnification becomes large, and sharp discontinuities arise as images appear or disappear.

The case of the underdense lens is especially complex due to the fact that different images can have either positive or negative TOA perturbations, meaning that the overall pulse TOA can be delayed or advanced depending on the frequency and epochs of observation, the lens parameters, and the lens shape. It is a generic feature of underdense lenses that the number of images tends to decrease as a function of frequency, and thus we expect to be able to observe multiple imaging events more often at low frequencies than at large frequencies. The size of the multiple images regions, as well as the distance between caustics, will in general change as we vary the coordinates in the u' plane, the value of DM_ℓ , and the size of the lens.

For the overdense lens in the figure, the region of multiple imaging follows a region of very large demagnification of a single image, a behavior that can also be observed in the case of the stochastic Gaussian lens shown in Figure 9.5, and appears to be generic for the case of Gaussian-like overdense lenses, although more complicated lens shapes can lead to other types of behavior. The signal to noise ratio in the first region is therefore extremely low, and the perturbations are dominated by white noise. The region of multiple imaging shows a gradual increase in the TOA delay as a function of frequency, with the signal to noise ratio increasing as we move closer to the caustic, after which the lensing effects are minimal. Again, we can see sharp discontinuities in the behavior of the perturbations at both caustic points.

9.4.2 Dispersion measure perturbations

Modern pulsar timing models and pulsar timing packages like TEMPO and TEMPO2 operate on the assumption that the frequency dependent delay of incoming radiation is purely dispersive, with the total delay being given by

$$\Delta t = 4.149 \text{ ms} \times \frac{DM}{\nu^2}, \quad (9.54)$$

where DM is in standard units of pc cm^{-3} and ν is in GHz. Physically, DM corresponds to the total integrated column density of electrons along the line of sight between the Earth and the pulsar. As discussed in the previous section, a lens changes the dispersive contribution depending on its characteristic shape and the parameter DM_ℓ , but also introduces a geometric perturbation in the TOAs due to refraction. These perturbations will be different for each image of the source for the cases in which the lensing is strong enough for ray crossings to occur. Thus, during a strong lensing event like the ones we have analyzed in this work, the expected ν^{-2} relationship for the group delay will not in general hold. Furthermore, we would expect that attempts at finding the best value of DM according to Eq. 9.54 will yield different best fit values and different deviations from the expected ν^{-2} scaling depending on the frequency band. This

¹⁰Lam, M. T., 2017, PyPulse, Astrophysics Source Code Library, record ascl:1706.011

follows from the fact that the nature of the frequency dependence of the perturbations due to the lens can change drastically as a caustic is crossed, as illustrated in Figure 9.7. This also means that a lensing event will not necessarily show up in the data as an increase in the ν^{-4} dependence of the residuals, except in cases in which the frequency band across which the data is being analyzed contains only a single image. A more sophisticated analysis, taking into account the details involved in the operational determination of DM and the way it changes in time, as described in Keith et al. [35], is outside the scope of this work.

9.5 Summary and conclusions

We have built on previous works that have studied the phenomenon of astrophysical plasma lensing in the context of ESEs, scintillations, and FRBs by developing a more general formalism that applies to two dimensional plasma lenses formed by both underdensities and overdensities in the ISM, and that can be used to study and predict the many possible ways in which lensing can affect observational quantities such as pulse intensities and TOAs. We showed that the geometrical optics method commonly employed in previous works to construct lensed light curves is unable to properly describe the fluctuations in the gain due to the interference between multiple source images, and is also unable to properly describe the gain within caustic zones.

By incorporating elements of catastrophe theory and the study of uniform asymptotic approximations of highly oscillatory integrals, we have developed an enhanced version of geometric optics that is able to account for such oscillatory features, and that does not break down at caustic curves in which two geometric optic images merge. We showed how this type of geometric optics can be successfully leveraged to construct the flux perturbations due to a variety of lens shapes and sizes, overcoming some of the limitations of other numerical approaches. We also apply some elements of this approach to characterize the possible form of TOA perturbations due to lensing events.

Our results indicate that there are many ways in which lensing effects can present themselves to an observer, depending on the lens shape, the magnitude of the electron density's departure from the surrounding ISM, whether this departure acquires the form of an overdensity or underdensity, and a series of other parameters such as the lens size, distances, and the frequencies of observation. The two dimensional model also adds an important degree of freedom in the form of the observer's path through the u' plane, something that cannot be correctly accounted for by one dimensional models. This extra degree of freedom also leads to the appearance of higher order diffraction catastrophes in parameter space that our approach is presently unable to accurately model. We expect to solve this problem in future work, as the successful implementation of uniform asymptotic methods for catastrophes like the cusp can greatly expand the the volume of parameter space that can be explored accurately in simulations.

Consistent with the results of previous works [29, 44, 64, 58], we find that lensing effects tend to be stronger at lower frequencies since the refractive power of plasma is more pronounced at large wavelengths. We also find our results for the overdense Gaussian lens to be consistent with results presented in previous works [11, 58, 23, 26]. Unlike these studies, however, we also analyze underdense Gaussian lenses, and find that their observational consequences are dramatically different from the overdense case. We also apply the uniform asymptotics approach to other types of lens shapes that have not been explored in the past.

The increasing accuracy of pulsar timing methods and procedures, as well as the growing population of pulsars under observation, imply that relatively rare phenomena like lensing events will be observed more often, and that their impact on the timing residuals will be more noticeable. Thus, being able to model such events will become increasingly more important. We expect to apply the methodology outlined in this work to establish whether chromatic aberrations such as the ones reported recently by Coles et al. [14] and Lam et al. [41] are

indeed the results of lensing phenomena and, if so, develop a model of the lensing structures responsible for such occurrences. The concepts developed here also have direct application to the modelling of ESEs for sources other than pulsars, and it is possible that lensing could be part of the explanation for some of the mysteries surrounding FRBs, which makes future work on this topic all the more important.

9.6 Solving the KDI using the FFT

The two dimensional Kirchoff diffraction integral (KDI) introduced in §9.2.3, gives the normalized wave optics field ε as a function of the observer coordinates \mathbf{u}' by integrating over an angular spectrum of plane waves,

$$\varepsilon(\mathbf{u}', \nu) = \frac{a_x a_y}{2\pi r_F^2} \iint d^2u \exp(i\Phi), \quad (9.55)$$

where Φ is the geometric phase,

$$\Phi(\mathbf{u}', \mathbf{u}, \nu) = \frac{1}{2r_F^2} [a_x^2(u_x - u'_x)^2 + a_y^2(u_y - u'_y)^2] + \phi_0 \psi(\mathbf{u}), \quad (9.56)$$

with r_F the Fresnel scale, a_x and a_y the lens scales, ϕ_0 the lens strength parameter, and ψ the lens shape. Given the form of the phase function, the KDI can be written as a two dimensional convolution integral,

$$\varepsilon(\mathbf{u}', \nu) = \iint d^2u G(\mathbf{u} - \mathbf{u}', \nu) H(\mathbf{u}, \nu), \quad (9.57)$$

where

$$G(\mathbf{u}, \nu) = \frac{a_x a_y}{2\pi r_F^2} \exp \left[\frac{i}{2r_F^2} (a_x^2 u_x^2 + a_y^2 u_y^2) \right], \quad (9.58)$$

$$H(\mathbf{u}, \nu) = \exp [i\phi_0 \psi(\mathbf{u})]. \quad (9.59)$$

From the discrete version of the convolution theorem (Schmidt 53), we have that

$$\varepsilon(\mathbf{u}', \nu) = \mathcal{F}^{-1} \{ \mathcal{F} [G(\mathbf{u}, \nu)] \cdot \mathcal{F} [H(\mathbf{u}, \nu)] \}, \quad (9.60)$$

where \mathcal{F} and \mathcal{F}^{-1} correspond to the discrete Fourier transform and its inverse, respectively, and \cdot denotes element by element multiplication. Thus, it is in principle possible to solve the KDI numerically for arbitrary lens shapes using the Fast Fourier Transform (FFT). The technique is applied for plasma lenses in one dimension by Watson & Melrose [64] and Melrose & Watson [44], and in two dimensions by Stinebring et al. [58] using code developed by Coles et al. [13], and we use it in the main text to show that it is possible to use an enhanced version of geometric optics to reproduce the intensity fluctuations predicted by wave optics.

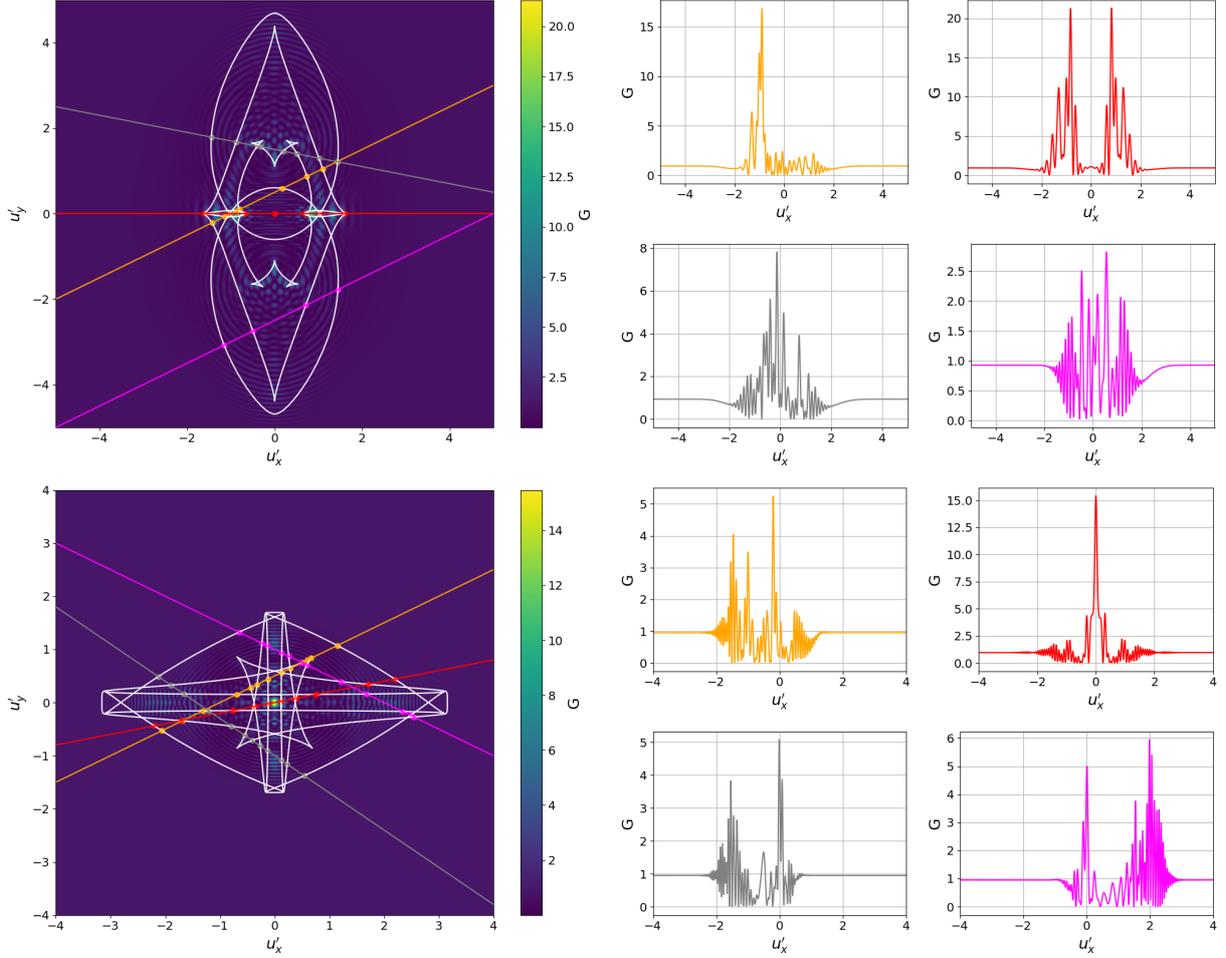


Figure 9.8: *Left*: Colormap of the gain in the u' plane overlaid with the caustic curves in white and slices along the plane in different colors for two different lens shapes. The points of intersection between the slices and the caustic lines are marked by points. The top panel corresponds to a lens with shape $\psi(\mathbf{u}) = 0.74 (u_x^2 + u_y^6) \exp(-u_x^2 - u_y^2)$, and parameters $a_x = a_y = 0.02$ AU, $DM_\ell = -1.5 \times 10^{-6}$ pc cm $^{-3}$, $\nu = 0.8$ GHz, $d_{so} = 1$ kpc, and $d_{sl} = 0.5$ kpc. The bottom lens has $\psi(\mathbf{u}) = \exp(-u_x^4 - u_y^4)$, $a_x = 0.04$ AU, $a_y = 0.05$ AU, $DM_\ell = -2 \times 10^{-6}$ pc cm $^{-3}$, $\nu = 1.0$ GHz, $d_{so} = 5$ kpc, $d_{sl} = 2.5$ kpc. In both cases, $\phi_0 \approx 50$ rad. *Right*: Plots of the gain along the paths shown in the left panel for each lens. Both kinds of lens show folds, cusps, and higher order catastrophes. The top lens can generate up to nine images of the source, whereas the bottom lens can produce up to seventeen.

Although useful, this approach suffers from serious limitations. First, it does not give information about the number of images of the source that can potentially be seen by the observer or the respective amplifications, phases, and TOAs of each of these images. Second, in practice the method can only be applied for a restricted range of lens scales $a_{x,y}$ and relatively small values of ϕ_o . The issue is the grid size necessary to properly sample the oscillations of the functions $G(\mathbf{u}, \nu)$ and $H(\mathbf{u}, \nu)$. We illustrate this for the former case. Consider a lens with characteristic scales $a_x = a_y = a$. By Nyquist's sampling theorem, the maximum array index n_{max} that can be sampled along a given axis is given by (Schmidt 53)

$$n_{max} = \frac{\pi r_F^2}{(\Delta x)^2}, \quad (9.61)$$

where Δx is the grid spacing in physical units. Now, let u'_{max} be the half-width of the u' plane along either of the axes, and N be the size of the array along that axis. Then, the sampling

interval can be written as

$$\Delta x = \frac{2au'_{max}}{N}. \quad (9.62)$$

Setting $N = n_{max}$ and rearranging, we have that the size of the grid along one axis required to ensure proper sampling is

$$N = \frac{4a^2(u'_{max})^2}{\pi r_F^2}. \quad (9.63)$$

This means that if we want to properly calculate the field for a lens with size $a = 1$ AU up to $u'_{max} = 5$ and with distances $d_{sl} = 0.5$ kpc, $d_{so} = 1$ kpc, and frequency of observation $\nu = 0.8$ GHz, we need $N \approx 1.5 \times 10^6$. This might be acceptable for the one dimensional case, but a two dimensional grid with side of size N is too big for even a modern desktop computer to handle. A more detailed analysis of sampling constraints and the numerical simulation of wave propagation using Fourier optics can be found in Schmidt [53].

Perhaps the primary advantage of this numerical strategy is that it does not have any problem calculating the field at caustic regions for any kind of catastrophe, even the higher order ones. Figure 9.8 shows the gain obtained using this method for different paths through the u' plane for a lens that shows higher order catastrophes than the ones in the main text.

9.7 Estimation of the value of the gain at the caustic

For very large values of ϕ_0 , it might be desirable in some cases to find the gain due to a lens using zeroth order geometric optics (Eq. 9.16), since the oscillations due to multiple imaging will give a value of the flux consistent with the prediction from that equation once we take into account the frequency resolution of the observations. Close to the caustics, however, the gain diverges. When ϕ_0 is large and the lens has strong refractive power, the gain can diverge in such a way that the maximum value occurs extremely close to the caustic, and this value can be estimated by the an extension of the method of stationary phase.

This estimate has been derived in more than one dimension by just a handful of authors in the context of asymptotic expansions of integrals, and their results do not necessarily agree with each other. Here we give two of the published formulas, specifically applied to the KDI, although we do not derive them. According to Chako [8] and Wong [65], the gain at the singularity is

$$G_{max} = \frac{a_x^2 a_y^2}{12\pi r_F^4} \frac{\Gamma^2(\frac{1}{3})}{|\Phi_{20}| |\Phi_{03}|^{\frac{2}{3}}}. \quad (9.64)$$

Bleistein & Handelsman [4] and Cooke [17] give a more complicated expression,

$$G_{max} = \frac{a_x^2 a_y^2}{4\pi^2 r_F^4} |\Phi_{20}| \Gamma^2(\frac{1}{3}) \left(\frac{32\pi^2}{3|B|^2} \right)^{\frac{1}{3}}, \quad (9.65)$$

where $B = \Phi_{20}^3 \Phi_{03} - 3\Phi_{20}^2 \Phi_{11} \Phi_{12} + 3\Phi_{20} \Phi_{11}^2 \Phi_{21} - \Phi_{11}^3 \Phi_{30}$. All derivatives of the phase in both equations are evaluated at the degenerate stationary phase point for which $\Phi_{10} = \Phi_{01} = \Phi_{20} \Phi_{02} - \Phi_{11}^2 = 0$. Some numerical experimentation has determined that both formulas give similar but not the same results.

9.8 Numerics

The key behind successful application of geometric optics as presented in the main text is the ability to numerically solve the lens equation, Eq. 9.14. This is essentially a two dimensional nonlinear root finding problem, with the added difficulties that the number of roots can be more than one, and that roots can appear or disappear as a function of the input parameters.

A general method for two dimensional root finding consists in rewriting the system of equations in the form

$$\begin{bmatrix} f(x, y) \\ g(x, y) \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \quad (9.66)$$

where $f(x, y)$ and $g(x, y)$ are the two equations that must be solved simultaneously. Once this is done, we can produce contour plots of both equations in order to find the sets of curves that satisfy $f(x, y) = 0$ and $g(x, y) = 0$. The roots of the two dimensional system will then correspond to the points of intersection of these sets of curves. When implemented properly, this method allows one to find all the roots of a two dimensional system within a range of values for x and y . A similar idea was pursued by Schramm & Kayser [55] to solve the lens equation for gravitational lensing. The disadvantage of this scheme is that it requires the evaluation of both $f(x, y)$ and $g(x, y)$ in a two dimensional grid that spans the area in which we are looking for solutions, which can be very computationally expensive if done repeatedly.

Since we are interested in solving the equation at many different points in parameter space, it is desirable to find a way to solve the lens equation that does not require us to apply the above algorithm at every single point of the independent variable. We can do this by combining it with other, more efficient numerical techniques that have been developed for numerical root finding in an arbitrary number of dimensions. These have existed for a long time, and are available for a variety of programming languages. In Python, some of these routines are available via the SciPy¹¹ library's optimization package. Although more efficient, these algorithms have the limitation that they rely on the user to input a guess solution that must be close enough to the actual solution. Furthermore, if there are multiple roots, they will only find the one closest to the input guess. This means that there is no way to find out exactly how many roots there are for a particular set of parameters.

Our strategy consists in combining the contour plotting method with the optimization algorithms in SciPy. First, we find the caustic locations for the range of parameters that we want to find the solutions of the lens equation for. If we are looking for solutions as a function of \mathbf{u}' , we apply the contour plotting algorithm to find the intersections between the curves in the u' plane that satisfy Eq. 9.25 with the line $u'_y = mu'_x + n$, where m and n parameterize the path through the u' plane. If we are looking for the solutions as a function of ν , we apply the contour plotting algorithm to simultaneously solve the system of equations given in Eq. 9.26.

This step allows us to separate the regions in parameter space that contain different numbers of solutions to the lens equation. Now, we can apply the contour plotting method again to find the number of roots at the center of each region. This results in the method being more reliable, because close to region boundaries, at least two roots will be very close to each other, whereas they will be maximally separated at the region's center. After having found the roots at the center of each of these regions, we find the other roots by iterating forward and backward in parameter space, using the root finding algorithm from SciPy with the previously found roots as the input guess solutions. As long as the distance between neighboring values of the independent variable is small enough, this strategy tends to work. It has the advantage of being much more efficient than applying the contour plotting method repeatedly, and also allows us to find the roots up to a very close distance to the caustic. This method has been tried for a wide variety of lens shapes and parameters, and has been found to be very reliable, especially for finding the real roots of the lens equation.

In order to find the complex rays, we need to apply a modified version of the above procedure that does not rely on contour plotting. The reason is that extending the search of solutions to the complex plane transforms the two dimensional lens equation into a four dimensional equation, and evaluating four different equations in four dimensional space is not practically

¹¹Jones E, Oliphant E, Peterson P, et al. SciPy: Open Source Scientific Tools for Python, 2001-, <http://www.scipy.org/>

feasible. Instead, we exploit the fact that, as discussed in the main text, very close to the shadow side of a caustic, the only important set of complex conjugate solutions to the lens equation is the one that has the smallest magnitude of its imaginary part. The real part of this complex conjugate set will be almost the same as that of the solution to the lens equation that intersects with the singularity. Thus, we use SciPy’s root finding algorithm with a value of the independent variable that falls in the caustic’s shadow side but at the same time is very close to the singularity, and input the value of u at the caustic as the guess solution. From there, we can recursively look for complex solutions that are farther away from the caustic in the same manner as we did for the case of the real solutions.

9.9 More numerical examples and lens colormaps

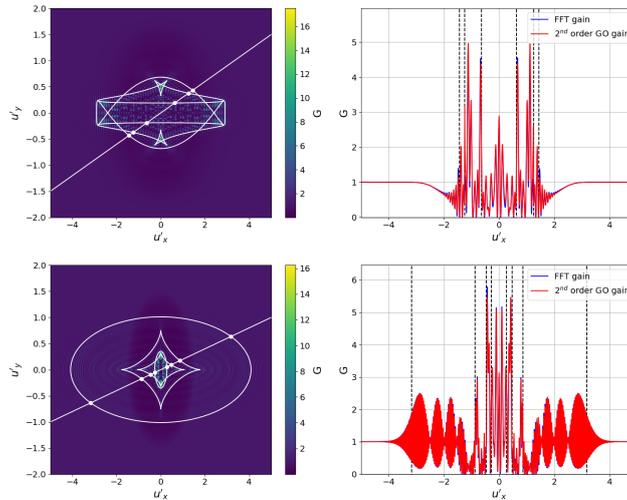


Figure 9.9: Comparison of the gains obtained from a full numerical solution of the KDI and second order geometric optics. The left column shows color maps of the gain obtained by solving the KDI via the FFT. The white circles correspond to caustic curves, and the straight white line shows the path of the observer through the u' plane. The right column shows the gain along this path as calculated via the FFT method and second order geometric optics. The points of intersection between the caustics and the observer path are marked by white points in the left column and by dashed vertical black lines on the plots in the right column. The top panel shows an underdense rectangular Gaussian lens with $\psi(\mathbf{u}) = \exp(-u_x^2 - u_y^4)$, lens phase $\phi_0 = 80$ rad, and lens scales $a_x = 1.5 \times 10^{-2}$ AU and $a_y = 3 \times 10^{-2}$ AU. The bottom panel corresponds to an underdense super-Gaussian lens with $\psi(\mathbf{u}) = \exp[-(u_x^2 + u_y^2)^3]$, lens phase $\phi_0 = 120$ rad, and lens scales $a_x = 2.5 \times 10^{-2}$ AU and $a_y = 4 \times 10^{-2}$ AU. The frequency of observation is $\nu = 1.4$ GHz, $d_{so} = 1$ kpc, $d_{sl} = 0.5$ kpc for both the top and bottom panels.

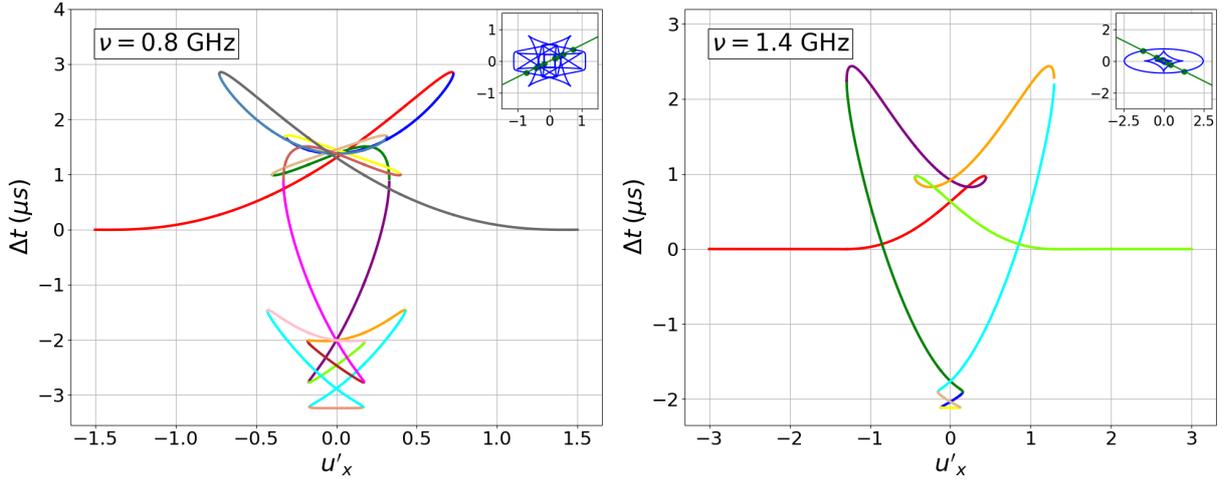


Figure 9.10: Individual image TOAs for two different lenses and different paths through the u' plane. The left panel corresponds to a square Gaussian lens with $\psi(\mathbf{u}) = \exp(-u_x^4 - u_y^4)$, $DM_\ell = -5 \times 10^{-4} \text{ pc cm}^{-3}$, $d_{so} = 1 \text{ kpc}$, $d_{sl} = 0.5 \text{ kpc}$, $a_x = 0.5 \text{ AU}$ and $a_y = 0.6 \text{ AU}$. The frequency of observation is $\nu = 0.8 \text{ GHz}$, which gives a lens phase of $\phi_o \approx 1.63 \times 10^4 \text{ rad}$. The right panel corresponds to a super-Gaussian lens with $\psi(\mathbf{u}) = \exp[-(u_x^2 + u_y^2)^2]$, $DM_\ell = -1 \times 10^{-3} \text{ pc cm}^{-3}$, $d_{so} = 5 \text{ kpc}$, $d_{sl} = 2.5 \text{ kpc}$, $a_x = 0.7 \text{ AU}$ and $a_y = 1 \text{ AU}$, with $\nu = 1.4 \text{ GHz}$, and thus $\phi_o \approx 1.86 \times 10^4 \text{ rad}$. Different colors denote different images, and the top right subplots show the path of the observer through the u' plane and the caustic curves. The maximum number of images in each plot is seventeen and nine, respectively.

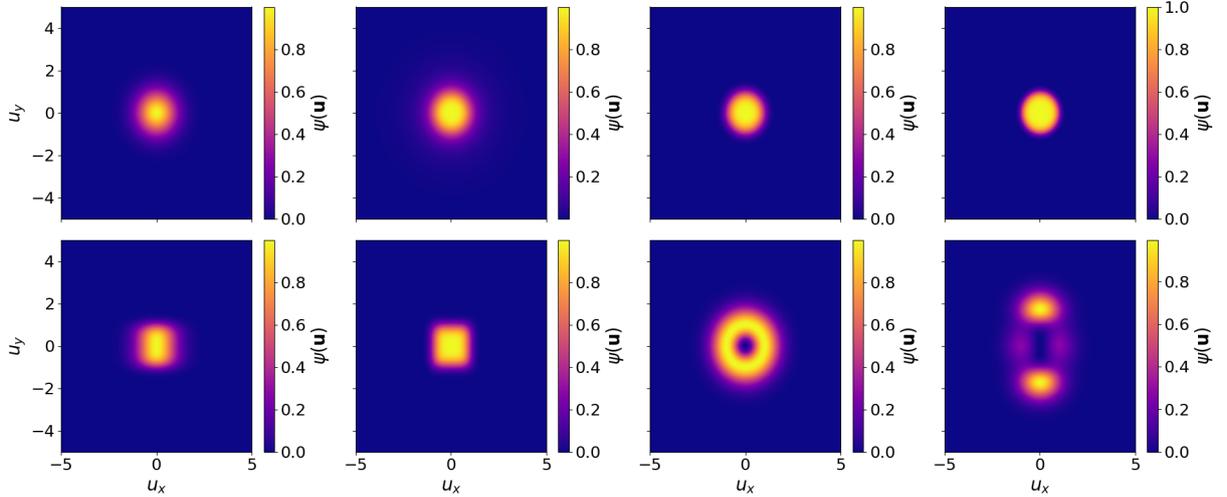


Figure 9.11: Colormaps of the different types of lensing structures used in the text and appendices. *Top row, from left to right:* Gaussian lens, $\psi(\mathbf{u}) = \exp(-u_x^2 - u_y^2)$, Lorentzian lens, $\psi(\mathbf{u}) = \frac{1}{[(u_x^2 + u_y^2)^2 + 1]}$, and super-Gaussian lenses $\psi(\mathbf{u}) = \exp[-(u_x^2 + u_y^2)^2]$ and $\psi(\mathbf{u}) = \exp[-(u_x^2 + u_y^2)^3]$. *Bottom row, from left to right:* Rectangular Gaussian lens, $\psi(\mathbf{u}) = \exp(-u_x^2 - u_y^4)$, square Gaussian lens, $\psi(\mathbf{u}) = \exp(-u_x^4 - u_y^4)$, ring-like lens $\psi(\mathbf{u}) = 2.72(u_x^2 + u_y^2) \exp(-u_x^2 - u_y^2)$, and double lens $\psi(\mathbf{u}) = 0.74(u_x^2 + u_y^6) \exp(-u_x^2 - u_y^2)$.

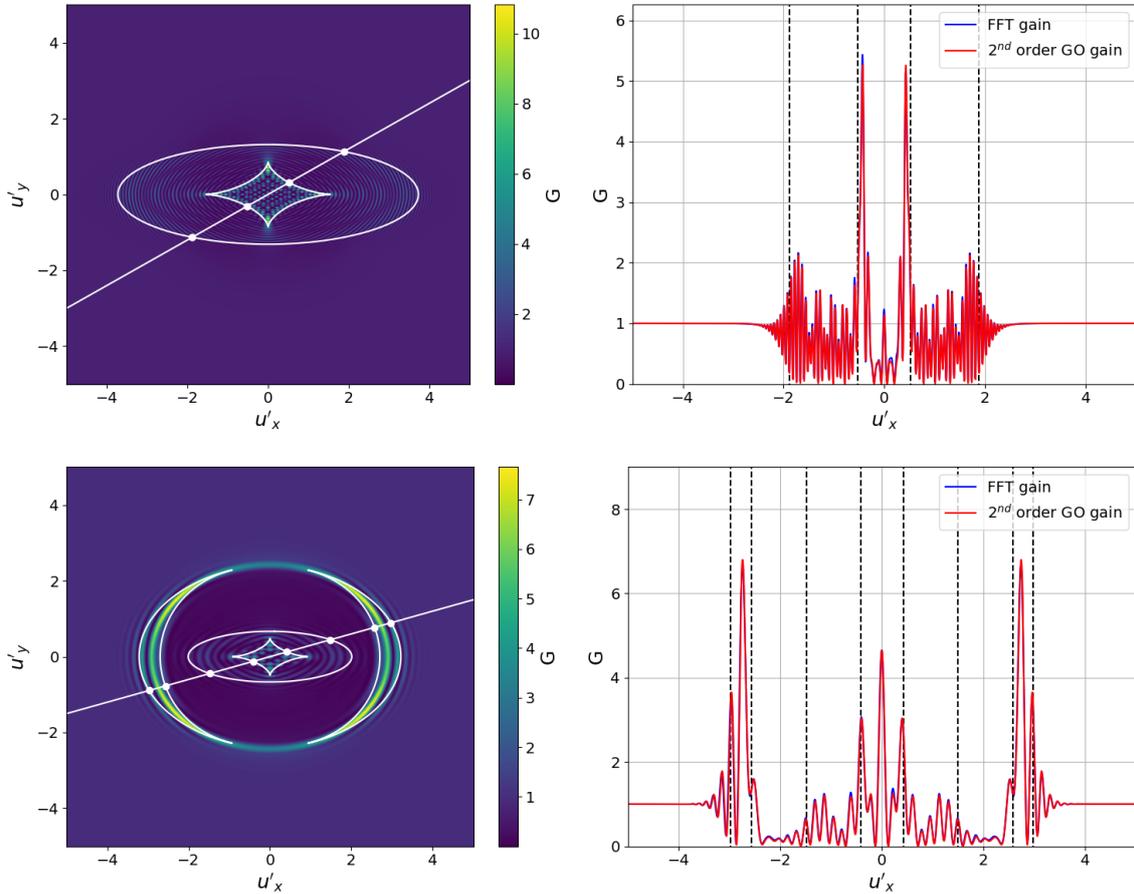


Figure 9.4: Comparison of the gains obtained from a full numerical solution of the KDI and second order geometric optics. The left column shows color maps of the gain obtained by solving the KDI via the FFT. The white circles correspond to caustic curves, and the straight white line shows the path of the observer through the u' plane. The right column shows the gain along this path as calculated via the FFT method (blue) and via second order geometric optics (red). The points of intersection between the caustics and the observer path are marked by white points in the left column and by dashed vertical black lines on the plots in the right column. The top panels shows results for an underdense elliptical Gaussian lens with $\psi(\mathbf{u}) = \exp(-u_x^2 - u_y^2)$, lens phase $\phi_0 = 100$ rad, and lens scales $a_x = 2 \times 10^{-2}$ AU and $a_y = 3 \times 10^{-2}$ AU. The bottom panel corresponds to an overdense ring-like lens with $\psi(\mathbf{u}) = 2.72(u_x^2 + u_y^2) \exp(-u_x^2 - u_y^2)$, lens phase $\phi_0 = -30$ rad, and lens scales $a_x = 2 \times 10^{-2}$ AU and $a_y = 3 \times 10^{-2}$ AU. The frequency of observation is $\nu = 0.8$ GHz, $d_{so} = 1$ kpc, $d_{sl} = 0.5$ kpc for both the top and bottom panels. The central caustic at the center of both u' planes in the left column occur because $a_x \neq a_y$, and is known as a structurally stable caustic of primary aberration (Berry & Upstill 3).

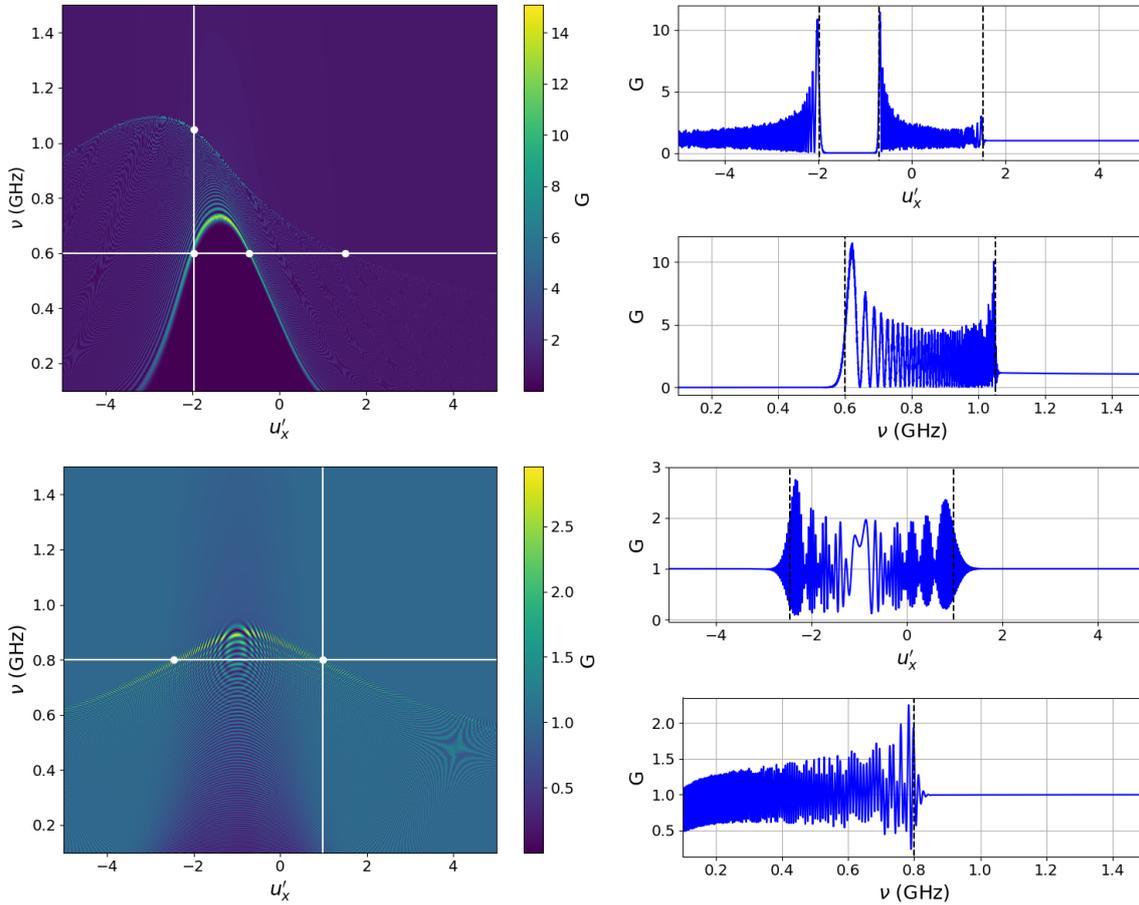


Figure 9.5: Sections of dynamic spectra and slices across them for overdense and underdense perturbed Gaussian lenses with $\psi(\mathbf{u}) = \exp(-u_x^2 - u_y^2) \{1 - A [\sin(Bu_x) + \sin(Bu_y)]\}$ and different DM_ℓ magnitudes. The left column shows the two dimensional spectrum for both lenses, with the top row corresponding to the overdense lens and the bottom row to the underdense lens. The vertical and horizontal lines correspond to the slices across the spectra plotted in the right column. Caustic intersections are marked by white dots in the left column plots and by dashed black lines in the right column plots. The overdense lens has a maximum column density of $DM_\ell = 10^{-4}$ pc cm $^{-3}$ and lens scales of $a_x = 0.1$ AU and $a_y = 0.2$ AU, whereas the underdense lens has $DM_\ell = -10^{-5}$ pc cm $^{-3}$, and $a_x = a_y = 0.04$ AU. Both lenses have perturbation parameters $A = 1.5 \times 10^{-2}$ and $B = 5$, source-observer distance $d_{so} = 1$ kpc, and source-lens distance $d_{sl} = 0.5$ kpc. The path through the u' plane in both cases is a straight line with slope $m = 0.5$ and y-intercept $n = 2.5$.

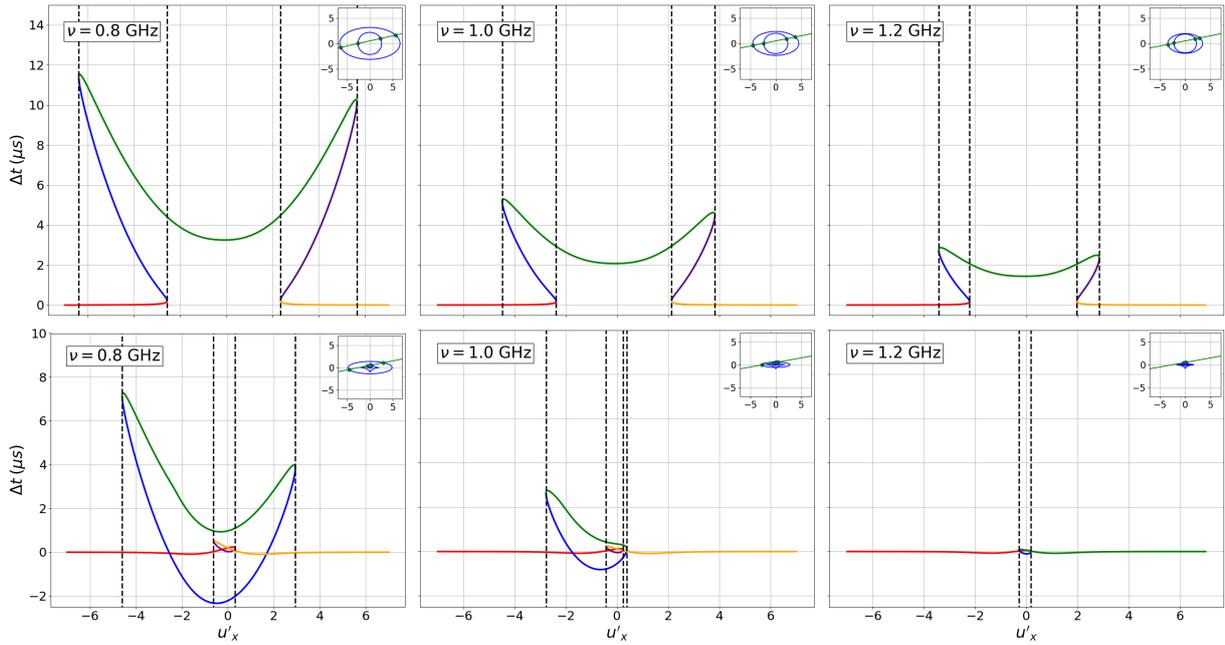


Figure 9.6: Timing perturbations of for pulses corresponding to different images as a function of observer position for overdense (top row) and underdense (bottom row) lenses with $DM_\ell = \pm 5 \times 10^{-4} \text{ pc cm}^{-3}$. Different frames in each row correspond to different frequencies of observation. Both overdense and underdense lenses have a Lorentzian shape with $\psi(\mathbf{u}) = \frac{1}{[(u_x^2 + u_y^2)^2 + 1]}$ and lens scales $a_x = 0.25 \text{ AU}$, $a_y = 0.4 \text{ AU}$. The distances used were $d_{so} = 1 \text{ kpc}$ and $d_{sl} = 0.5 \text{ kpc}$, and the path through the u' plane has slope $m = 0.2$ and y -intercept $n = 0.5$. The subplot in the top corner of each subpanel shows the (blue) caustic curves in the u' plane for the corresponding frequency of observation, together with the (green) path of the observer through the u' plane. The different colors in the Δt vs u'_x plots trace the timing perturbation for each individual image.

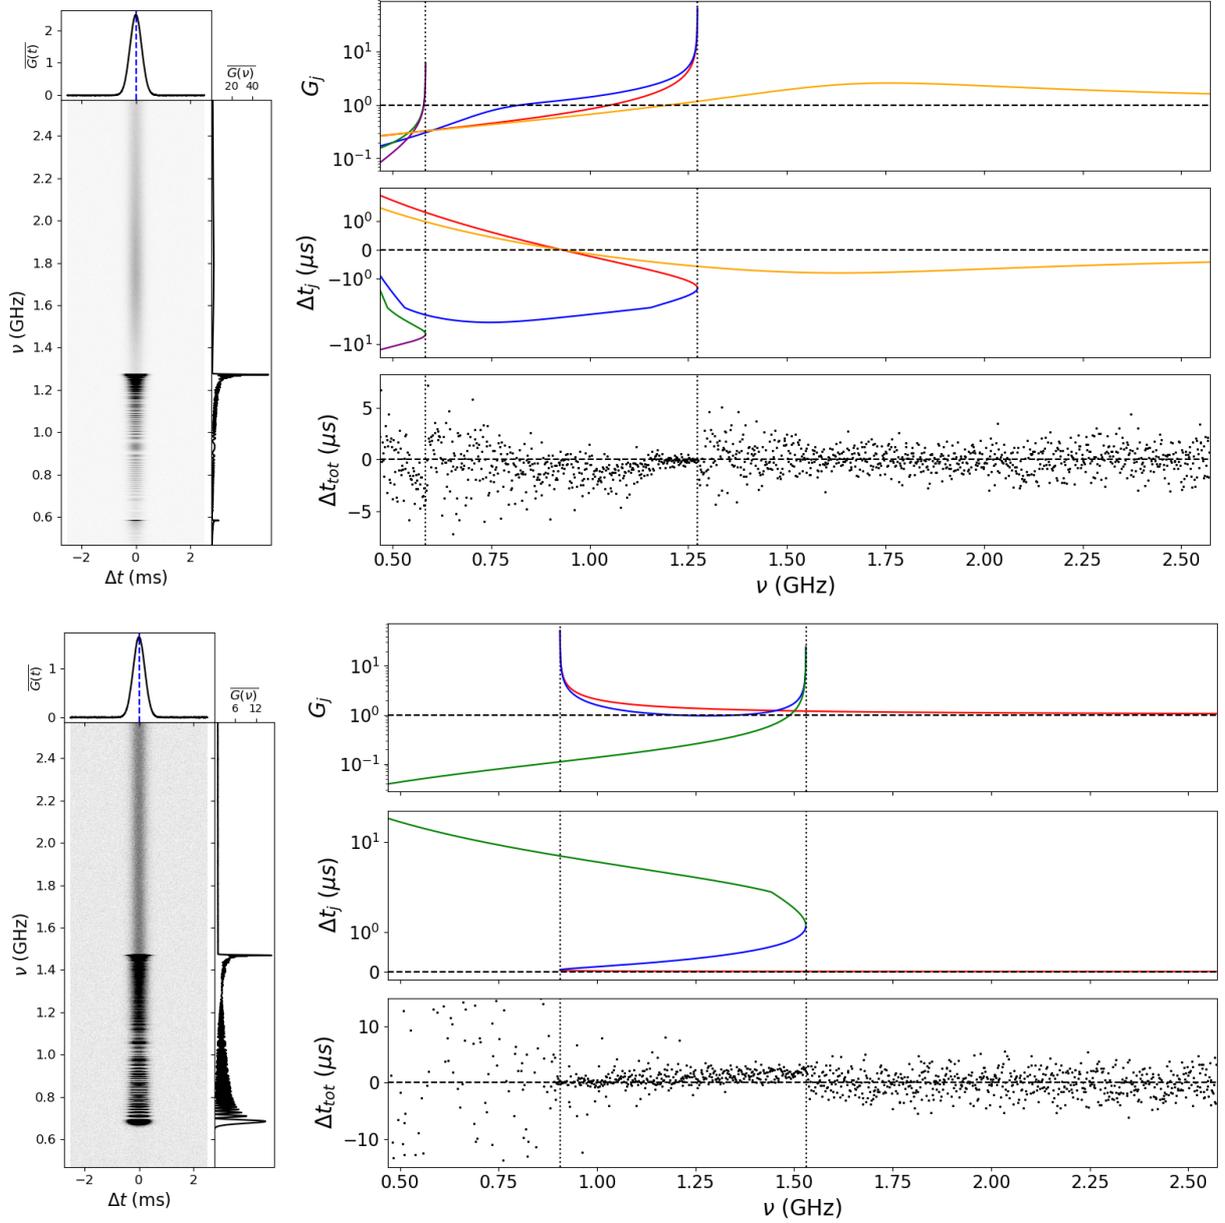


Figure 9.7: Pulse dynamic spectra, individual image gains and TOA perturbations, and combined TOA perturbation as a function of frequency for a single epoch of observation for an underdense lens (top panel) and an overdense lens (bottom panel). The lens in the top panel has Gaussian shape $\psi(\mathbf{u}) = \exp(-u_x^2 - u_y^2)$, $\text{DM}_\ell = -7 \times 10^{-4} \text{ pc cm}^{-3}$, $d_{so} = 1 \text{ kpc}$, $d_{sl} = 0.5 \text{ kpc}$, $a_x = 0.5 \text{ AU}$ and $a_y = 1.1 \text{ AU}$, and the epoch corresponds to a position in the u' plane with coordinates $\mathbf{u}' = (0.1, 0.1)$. The lens in the bottom panel has shape $\psi(\mathbf{u}) = \exp[-(u_x^2 + u_y^2)^2]$, $\text{DM}_\ell = 10^{-3} \text{ pc cm}^{-3}$, $d_{so} = 2 \text{ kpc}$, $d_{sl} = 1.5 \text{ kpc}$, $a_x = 0.8 \text{ AU}$, $a_y = 1.1 \text{ AU}$, and $\mathbf{u}' = (-1.5, -0.55)$. The pulse profile contains 2048 bins, and the pulse repetition period is $T = 5 \text{ ms}$, giving an integration time $\Delta t_{\text{int}} \approx 2.44 \mu\text{s}$. The channel bandwidth is $\Delta\nu_r = 1.5 \text{ MHz}$. We use a Gaussian template to model the pulse shape.

Part X

Other Special Theories, Models

Part XI

Method of Path Integral for WO in GL

10 Oscillatory path integrals for radio astronomy by Feldbrugge, Pen, Turok

Abstract

We introduce a new method for evaluating the oscillatory integrals which describe natural interference patterns. As an illustrative example of contemporary interest, we consider astrophysical plasma lensing of coherent sources like pulsars and fast radio bursts in radioastronomy. Plasma lenses are known to occur near the source, in the interstellar medium, as well as in the solar wind and the earth's ionosphere. Such lensing is strongest at long wavelengths hence it is generally important to go beyond geometric optics and into the full wave optics regime. Our computational method is a spinoff of new techniques two of us, and our collaborators, have developed for defining and performing Lorentzian path integrals. Cauchy's theorem allows one to transform a computationally fragile and expensive, highly oscillatory integral into an exactly equivalent sum of absolutely and rapidly convergent integrals which can be evaluated in polynomial time. We require only that it is possible to analytically continue the lensing phase, expressed in the integrated coordinates, into the complex domain. We give a first-principles derivation of the Fresnel-Kirchhoff integral, starting from Feynman's path integral for a massless particle in a refractive medium. We then demonstrate the effectiveness of our method by computing the interference patterns of Thom's caustic catastrophes, both in their "normal forms" and within a variety of more realistic, local lens models, over all wavelengths. Our numerical method, implemented in a freely downloadable code, provides a fast, accurate tool for modeling interference patterns in radioastronomy and other fields of physics.

10.1 Introduction

Interference is one of the most universal phenomena in nature. In classical physics, the linear superposition of sound waves, surface waves, radio waves, light or gravitational waves all exhibit the same characteristic patterns of constructive and destructive interference. Interference is also fundamental to quantum physics. The basic quantum amplitudes describing particles or fields are most elegantly formulated as path integrals – sums over trajectories weighted by the phase factor $e^{i\mathcal{S}/\hbar}$, with \mathcal{S} the action and \hbar Planck's constant. As ubiquitous as interference and interference patterns are, they are generally hard to compute. The oscillatory integrals involved are only conditionally and not absolutely convergent, meaning they converge slowly and artefacts such as dependence on unphysical cutoffs may be hard to avoid. Likewise, if the integrals are performed iteratively, as is often the only practicable method, conditional convergence is in general insufficient to guarantee uniqueness, since the order in which partial integrals are taken can affect the result.

In quantum mechanics, these difficulties run deep. In fact, so far they have thwarted all efforts to rigorously define nontrivial real-time Feynman path integrals, even in non-relativistic quantum mechanics [1]. The only available existence proofs involve a Wick rotation from real, Lorentzian time to imaginary, Euclidean time, which maps the phase factor to a real Boltzmann weight (for a recent review see, *e.g.*, [2]). Unfortunately, securing mathematical rigour this way comes at a high price: the system's dynamics can only be described in imaginary time instead of real time where experiments and observations actually take place. Analytic continuation back to real time is often only possible for certain quantities, such as perturbative S-matrix elements

and, even then, is often hard. Furthermore, for some theories, including general relativity and quantum condensed matter models with a “sign problem,” *e.g.* the Hubbard model, the Wick rotation trick does not work.

This paper represents a step towards a new, broadly applicable method for defining and computing Lorentzian path integrals. Here, we study the interference of relativistic waves, emitted from coherent sources and propagating through a region in which the refractive index varies in space, *i.e.*, a lens. As we shall show, the quantum mechanical path integral amplitude reduces, in this case, to an ordinary, finite dimensional integral.

The study of optical interference patterns dates back over two centuries, long predating Maxwell’s equations, but remains of enduring interest. Starting in the 1970’s, Berry, Nye and collaborators studied examples of “diffraction catastrophes” – the characteristic patterns created by diffraction about each of Thom’s stable caustic catastrophes, and compared intricate mathematical calculations with beautiful experiments [3, 4, 5, 6, 7]. Recently, the need to accurately and efficiently compute similar patterns has arisen in radioastronomy where bright, coherent sources of radio waves like pulsars and fast radio bursts are being detected in rapidly growing numbers [8, 9, 10]. These objects are beacons lighting up the universe. They will potentially provide a vast new source of information for astrophysics and cosmology. Typically, they are lensed by diffuse astrophysical plasmas intervening along the line of sight. Since plasma lensing is strongest at long wavelengths, this lensing must be modeled in the full, wave optics regime [11, 12, 13]. Although challenging, such modeling will likely be vital to our ability to draw precise inferences from these sources [14].

Motivated by this contemporary need, we shall use astrophysical plasma lensing as our main example. However, as should be clear to the reader, the principles involved are far more broadly relevant. The interference patterns created by astrophysical plasma lenses and observed over astronomical or even cosmological distances are governed by exactly the same physics at play in Young’s double slit experiment or X-ray crystallography. This is both a striking example of universality in physics and a reminder of how the universe increasingly provides us with a powerful laboratory for studying fundamental physics.

Spatial variations in the refractive index of astrophysical plasmas can arise due to turbulence in the interstellar medium or other sources of heating [15, 16]. Pulsar observations have provided examples where plasma lensing amplifies the brightness of a coherent radio source by factors approaching a hundred [17, 13, 18]. It has been pointed out that plasma lensing is likely to play an important role in the phenomenology of Fast Radio Bursts (FRBs) [12, 13]. So-called Extreme Scattering Events (ESEs), where the brightness of radio sources is seen to change by factors of a few, are also thought likely to be due to as yet unexplained plasma lensing [19, 20]. Recently, there has been growing interest in the idea that coherent gravitational wave pulses and trains, emitted from black hole or neutron star mergers, could be gravitationally lensed and thereby magnified. In this situation it will again be important to go beyond geometric optics and include wave diffraction [21]. In all these examples, when the line of sight between source and observer passes through a caustic of a lens, at a given frequency, the observed intensity may be enhanced leading to a pulse in frequency, time, or both. These situations have mainly been studied for one-dimensional lenses near fold and cusp caustics [14, 22]. Here, we shall explore more complex, two-dimensional examples including the swallowtail, elliptic and hyperbolic umbilic catastrophes which we describe below.

There is already an extensive astrophysical literature on the computation of interference patterns in wave optics [23, 24], but published methods tend to converge slowly [25]. They are expensive to implement and the results are sometimes inconclusive. In this paper, we shall present faster and more reliable methods. Our approach builds on Picard-Lefschetz theory, a general, exact approach to multidimensional oscillatory integrals based upon saddle point and steepest descent techniques (for an introduction [26]; for applications to quantum cosmology,

see [27, 28, 29, 30, 31] and to relativistic quantum mechanics, see [32]). As we shall show, our methods allow for the fast and reliable computation of even very intricate “diffraction catastrophe” patterns. The calculations of these patterns by Berry, Nye and collaborators were an analytical *tour de force*, but relied heavily on the particularities of Thom’s canonical “normal forms” of catastrophes, and the mathematical properties of the related special functions, with each case treated separately. Unfortunately, while the normal forms represent the correct universal forms locally, they diverge at large distances. Hence, they are unrealistic as models for natural lenses. Realistic modeling requires a more versatile method which can be efficiently and straightforwardly implemented numerically. We present just such a method here.

Our method applies uniformly, with modest restrictions, to generic lens models. It is simple to implement numerically and computes interference patterns in polynomial time. As far as we have been able to check, our results agree perfectly with those aforementioned. The only requirement of our method is that it should be possible to analytically extend the interference phase into the complexified space of the spatial coordinates over which the integral is taken. Such functions embrace a very large class of lens models including, for example, any rational function, and should be more than sufficient for most modeling purposes. For simplicity, we shall not consider phases which possess branch cuts in the space of complexified coordinates. However, there are physical cases of interest where such phases do occur and an extension of our approach to this more general setting is an interesting problem for the future.

Instead of using specific properties of special functions and symmetries, our method exploits Cauchy’s theorem to exactly transform an integral of an oscillatory phase factor into a sum of absolutely convergent integrals taken over “Lefschetz thimbles” in the space of complexified coordinates. These “thimble” integrals are fast to compute numerically, requiring only polynomial time. They are insensitive to numerical cutoffs and may be performed iteratively in any order with no change to the result. In this paper, we demonstrate the efficacy of our method by computing the interference patterns for one- and two-dimensional thin lenses. We study the most observationally accessible catastrophes, both in their “normal forms” and in a set of more realistic, localized lens models where these catastrophes appear. Our one-dimensional numerical code, capable of handling generic one-dimensional lenses, is now publicly available [online](#)¹².

A simple example of the type of integral we are interested in is

$$\Psi(\mu, \alpha, \nu) = \left(\frac{\nu}{\pi}\right)^{1/2} \int_{-\infty}^{\infty} dx e^{i\phi(x)\nu}, \quad \text{where} \quad \phi(x) = (x - \mu)^2 + \frac{\alpha}{1 + x^2}. \quad (10.1)$$

Here, $\Psi(\mu, \alpha, \nu)$ is the amplitude whose square $|\Psi(\mu, \alpha, \nu)|^2$ gives the intensity of light observed at a position, frequency and lens strength controlled by the parameters μ , ν and α . The control parameter μ is determined by the transverse positions of the observer and the source relative to the lens (See Fig. 11.1 and Eq. (11.13) below). The frequency of the waves is proportional to ν so the spacing of interference fringes shrinks as ν is increased. The eikonal limit is $\nu \rightarrow \infty$. Finally, α controls the strength of the lens which, in this example, is taken to have a Lorentzian profile. The integral (11.1) is analytically intractable. However, it is simple to compute numerically, for reasonable values of ν , μ and α , using the methods we shall describe below.

In the eikonal limit of large ν , only real saddle point solutions – real stationary points of the phase $\phi(x)$ – contribute significantly to the amplitude. Each one corresponds to a particular ray. For $\alpha < 1$ the lens is “weak” and there is only one real solution of $\partial_x \phi(x) = 0$. Hence there is only one contributing ray at each value of μ . For $\alpha > 1$, the lens is “strong:” for a finite range of μ values centred on zero, there are three real solutions of $\partial_x \phi(x) = 0$ hence three contributing rays. Correspondingly, one finds three images of the source in this range of μ . The values of μ bounding this range mark a transition from three contributing rays (*i.e.* three real saddles

¹²See https://github.com/jfeldbrugge/Picard_Lefschetz_Integrator.

in the phase) to one. At these values of μ , a maximum and a minimum of $\phi(x)$ merge into a *cubic* stationary point (*i.e.* a point of inflexion), creating the simplest “fold” catastrophe. If we now decrease the strength of the lens α towards unity, the two “fold” catastrophes approach the point $\mu = 0$ where they merge to form a “cusp” catastrophe, in which there is a *quartic* stationary point in the phase $\phi(x)$. Since the phase (viewed as a function of x) is flatter in the vicinity of higher order stationary points, there is less destructive interference. The intensity of light grows more rapidly as ν is increased as compared to the intensity from a quadratic saddle, so that “folds” become increasingly bright compared to the unlensed image and “cusps” become even brighter. While higher order catastrophes are rarer, their brightness makes them easier to detect. This has encouraged the conjecture, yet to be verified [14, 8, 12, 33, 25], that the brightest sources seen may be those which happen to be lensed into high order catastrophes.

In order to emphasize the foundational character of the physics at play and by way of a pedagogical introduction, we show how the standard Fresnel-Kirchhoff integral (see, *e.g.*, Ref. [34], Chapter 8), central to the description of lensing in radioastronomy and in optics¹³, can be derived directly from Feynman’s path integral for a massless particle propagating through a refractive medium, *i.e.*, one in which the speed of light varies across space. Our main focus in this paper is on dispersive but non-dissipative lensing, in which the lensing phase factor always has modulus unity. However, the methods we use may equally well be applied to dissipative (lossy) lensing, in which the plasma dispersion relation is complex (for a review of dispersion relations, for example in water or in the ionosphere, see, *e.g.*, [35] Ch. 7). In this more general circumstance, the “phase factor” over which the Fresnel-Kirchhoff integral is taken has a varying modulus.

As an illustration of such a case, as well as to provide a foretaste of the use of our method in describing quantum mechanical interference, in Appendix 11.10 we examine Young’s famous double-slit experiment. We consider a thin, flat one-dimensional lens which modulates the intensity rather than the phase of the light passing through it. We model the lens with a smooth function which allows very little light through except in two narrow regions comprising the slits. We calculate the resulting interference pattern by deforming the contour onto the relevant Lefschetz thimbles numerically, observing how different real and complex saddle points become relevant and irrelevant, as one moves across the observational screen, through an intricate sequence of Stokes phenomena. Using this smooth lens model, we can also study in detail the emergence of the classical limit as Planck’s constant \hbar is taken to zero, so that the de Broglie wavelength becomes small. In this limit we find as expected that only the real, classical saddles contribute and all interference effects disappear.

Finally, as an aside, we remark that the work presented here represents a step in a larger program, involving two of us [36] and our collaborators, seeking to rigorously define, calculate and interpret real time (Lorentzian) path integrals, with diverse applications in quantum physics, both nonrelativistic and relativistic, including quantum gravity and cosmology. We expect to report further on this work in the near future. Recently, Dunne, Unsal and collaborators have pursued a very interesting (and closely related) program in quantum field theory and quantum mechanics, based upon Euclidean path integrals [37, 38, 39, 40, 41]. See also, the closely related work of [42], and earlier work of [43].

The outline of this paper is as follows. In Section 11.2 we show how the Fresnel-Kirchhoff integral and Fermat’s principle follow from the relativistic path integral for a massless particle, *i.e.*, a spinless photon, moving in a medium with a variable speed of light. In Section 11.3 we

¹³The integral formula has a fascinating history of successive approximate derivations and subsequent critiques, reviewed in Ref. [34], Chapter 8. Exact solutions of Maxwell’s equations (or their scalar version) representing quasi-realistic interference patterns created by diffraction around physical obstacles of various types are still few in number, and are reviewed in Chapter 11 of the same work. It would be interesting to revisit these solutions and, perhaps to find others, using the ideas we develop here.

discuss the Fresnel-Kirchhoff integral for thin astrophysical lenses, putting the answer into a canonical dimensionless form. We then discuss the intensity in the geometric optics limit, along with the occurrence of critical points and caustics. We introduce catastrophe theory, describing the “normal form” of critical points of increasing complexity and their relation to observable parameters. In Section 11.4 we discuss Picard-Lefschetz theory for a one-dimensional lens – first in the geometric optics limit and then beyond, to include diffraction. We introduce the key concept of “flowing” the integration contour into the complex plane, in order to find the set of relevant Lefschetz thimbles upon which the integral becomes absolutely convergent. We describe a simple and powerful numerical code which implements this idea. In section 11.5 we numerically compute the interference patterns of the seven elementary catastrophes, giving a comprehensive analysis of their “unfoldings.” In section 11.6 we turn to localized lens models, which are analytically intractable. In section 11.7 we anticipate possible applications to the study of Fast Radio Bursts, which is an exciting current prospect. Section 11.8 concludes. Appendix 11.9 provides some instructive background on the simplest (Gaussian) oscillatory integrals - both one- and two-dimensional, and Appendix 11.10 tackles Young’s famous double-slit experiment.

10.2 From Feynman to Fermat to Fresnel-Kirchhoff

(WORK OUT THIS!!!!)

Imagine a bright source emitting coherent electromagnetic waves which traverse an astrophysical plasma on their way to our telescopes on earth. Let us describe the propagation in terms of the elementary *quanta* of such waves, considered to be relativistic particles. The Feynman path integral over these particle’s trajectories in spacetime yields the quantum mechanical amplitude to propagate from the source to any particular location. The square of the amplitude yields the intensity, determining the interference pattern in position and frequency. As we shall see, one or more classical trajectories dominate the amplitude: these dominant trajectories obey Fermat’s “principle of least time.” For simplicity, we shall ignore polarization effects, taking the elementary quanta to be spinless. We shall furthermore study only the simplest dispersion relation for astrophysical plasmas, valid in the high frequency regime – generalizations to more complex and realistic dispersion relations should be straightforward. Our derivation emphasizes the fundamental nature of the physics involved - as we shall show, the Fresnel-Kirchhoff integral (see, *e.g.*, Ref. [34], Chapter 8, 8.3.3 (28)) follows directly from the Feynman path integral. We hope the reader will enjoy the directness and economy of this approach compared to more standard (and cumbersome) derivations based on Maxwell’s equations, or their scalar counterpart.

We start from the dispersion relation in a tenuous plasma (see, *e.g.*, [35] Section 7.9)

$$\omega^2 = k^2 c^2 + \omega_p^2(\mathbf{x}). \quad (10.2)$$

Here, ω and k are the angular frequency and wavenumber of the waves, c is the speed of light and $\omega_p(\mathbf{x})$ is the plasma frequency at position \mathbf{x} , determined by the local density of electrons, assumed to vary across space on scales much larger than the wavelength of the electromagnetic waves. Notice that (11.2) takes exactly the same form as the dispersion relation for a relativistic particle whose mass varies with spatial position.

The dispersion relation (11.2) yields a phase propagation speed

$$v_p(\mathbf{x}) \equiv \frac{\omega}{k} = c \sqrt{1 + \frac{\omega_p^2(\mathbf{x})}{k^2 c^2}}, \quad (10.3)$$

which is greater than the speed of light. This should be no cause for concern, as the analogy with a massive particle assures us, since information only propagates at the group velocity, $\mathbf{v}_g \equiv \nabla_{\mathbf{k}} \omega$ whose magnitude $c_g = c^2/c_p$ is always less than the speed of light.

The quanta of these waves may be described as relativistic particles, following parameterized worldlines in spacetime: $x^\mu(\lambda) = (ct(\lambda), \mathbf{x}(\lambda))$. Reparameterizations $\lambda \rightarrow \tilde{\lambda}(\lambda)$ are generated by a Hamiltonian, and reparameterization invariance corresponds to the constraint that the Hamiltonian vanishes, $\mathcal{H} = 0$. The correct expression for the Hamiltonian \mathcal{H} may be read off from the dispersion relation (11.2), using the correspondence $p_\mu = (p_0, \mathbf{p}) \leftrightarrow \hat{p}_\mu = -i\hbar\partial_\mu = \hbar(-\omega/c, \mathbf{k})$:

$$\mathcal{H} = -p_0^2 c^2 + \mathbf{p}^2 c^2 + \hbar^2 \omega_p^2(\mathbf{x}). \quad (10.4)$$

The first order (phase space) action, with the initial and final spacetime locations of the particle held fixed, is:

$$S[x; x^\mu(0), x^\mu(1)] = \int_0^1 d\lambda (p_0 \dot{x}^0 + \mathbf{p} \cdot \dot{\mathbf{x}} - \tau(\lambda)\mathcal{H}). \quad (10.5)$$

where dots denote derivatives with respect to λ , taken to run from 0 to 1 as the the particle trajectory runs from the initial spacetime point $x^\mu(0) \equiv (ct_i, \mathbf{x}_i)$ to the final point $x^\mu(1) \equiv (ct_f, \mathbf{x}_f)$. The ‘einbein’ $\tau(\lambda)$ serves as a Lagrange multiplier enforcing the Hamiltonian constraint and ensuring the action is reparameterization invariant (it transforms under reparameterization so that $d\lambda\tau(\lambda)$ is invariant). Varying the action with respect to the momenta yields Hamilton’s equations for the momenta $p_0 c^2 = -\dot{x}^0/(2\tau)$ and $\mathbf{p} c^2 = \dot{\mathbf{x}}/(2\tau)$. Varying with respect to τ yields the constraint $\mathcal{H} = 0$. The energy $E = -p_0 c$ is conserved because the action is invariant under constant translations of the time x^0 .

In seeking to derive Fermat’s principle, we face a conundrum. If the initial and final times t_i and t_f are held fixed, how can the total time $t_f - t_i$ possibly vary? The resolution is that, for a monochromatic beam, we should fix the initial energy E , not the initial time t_i . We cannot fix both because of the time-energy uncertainty relation (which follows from the commutator $[\hat{p}_0, \hat{x}^0] = -i\hbar$). The action appropriate to fixing the initial energy is obtained by adding a boundary term. The latter must be chosen to ensure that the variation of the action is zero when the initial energy and the final time, as well as the initial and final spatial positions, are held fixed and the equations of motion are satisfied. The initial time is then free to vary, which is how Fermat’s principle can arise. The required total action is:

$$S[\mathbf{x}; E, \mathbf{x}_i, t_f, \mathbf{x}_f] = p_0 x^0(0) + \int_0^1 d\lambda (p_0 \dot{x}^0 + \mathbf{p} \dot{\mathbf{x}} - \tau(\lambda)\mathcal{H}), \quad (10.6)$$

with \mathcal{H} given in (11.4).

Since the action (11.6) is quadratic in the momenta and linear in τ , we can integrate out those variables. At the relevant saddle, we may use Hamilton’s equations for the momenta, and the constraint, to obtain a reduced action expressed purely in terms of reparameterization-invariant quantities:

$$S_r[\mathbf{x}] = -Et_i - \int_{t_i}^{t_f} dt \frac{\hbar^2 \omega_p^2(\mathbf{x}(t))}{E}. \quad (10.7)$$

(HOW IS IT DONE???????)

Writing $t_i = t_f - \int_{t_i}^{t_f} dt$ where the final time t_f is held fixed, we find, up to an irrelevant constant phase,

$$S_r[\mathbf{x}] = E \int_{t_i}^{t_f} dt \left(1 - \frac{\hbar^2 \omega_p^2(\mathbf{x})}{E^2} \right) = E \int_{\mathbf{x}_i}^{\mathbf{x}_f} \frac{|d\mathbf{x}|}{c} \left(1 - \frac{\hbar^2 \omega_p^2(\mathbf{x})}{E^2} \right)^{\frac{1}{2}} = \int_{\mathbf{x}_i}^{\mathbf{x}_f} |d\mathbf{x}| |\mathbf{p}|, \quad (10.8)$$

where, again, we used Hamilton’s equations for the momenta and the Hamiltonian constraint. Finally, we express the result in terms of the phase velocity (11.3), obtaining

$$S_r[\mathbf{x}] = E \int_{\mathbf{x}_i}^{\mathbf{x}_f} \frac{|d\mathbf{x}|}{c_p(\mathbf{x})}. \quad (10.9)$$

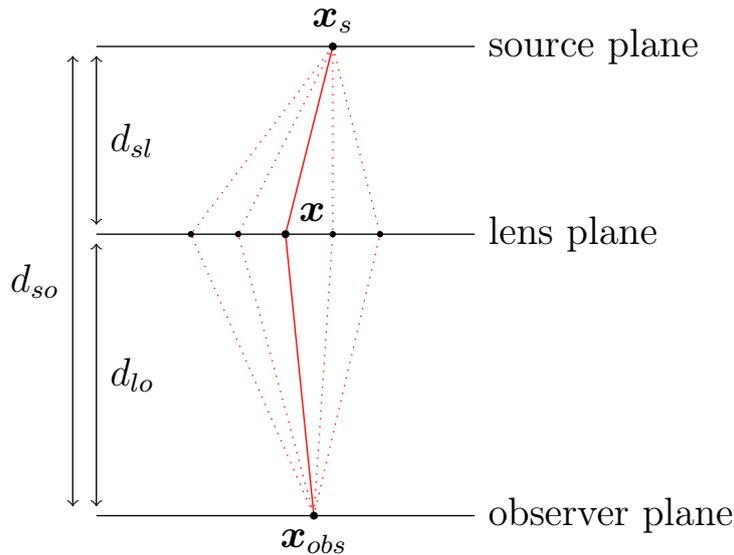


Figure 10.1: The geometry of interfering paths passing through a thin lens.

Note that, although the phase velocity $c_p(\mathbf{x})$ appearing here is always greater than the speed of light, nowhere in our derivation does any on-shell particle actually travel faster than light.

The reduced action (11.9) embodies Fermat's principle of least time or, more correctly, the principle that the time taken is stationary on dominant classical trajectories. The path integral over all paths, weighted by $e^{iS_r[\mathbf{x}]/\hbar}$, is the Fresnel-Kirchhoff integral we seek.

As an aside, note that one may, equally well, obtain the result (11.9) starting from the square root (Nambu-type) action for a particle with a spatially dependent mass $m(\mathbf{x})$, by making use of the correspondence $m(\mathbf{x})c^2 \leftrightarrow \hbar\omega_p(\mathbf{x})$, namely

$$\mathcal{S}[\mathbf{x}] = - \int_{t_i}^{t_f} dt \hbar\omega_p(\mathbf{x}(t)) \left(1 - \frac{\dot{\mathbf{x}}(t)^2}{c^2} \right)^{\frac{1}{2}}, \quad (10.10)$$

where the dot now denotes a t derivative. This action is explicitly reparameterization invariant from the start. However, it is the action appropriate to fixing the initial time t_i whereas we need to fix the initial energy E . As before, we must supplement the action (11.10) by a boundary term, which turns out to be $+E(t_f - t_i)$. One can easily check that the identity $\partial S_{cl}/\partial t_i = E$ for Hamilton's principal function S_{cl} implies the total action is stationary, provided the desired boundary conditions and the equations of motion are fulfilled. Using $E = \hbar\omega_p/(1 - \dot{\mathbf{x}}^2/c^2)^{\frac{1}{2}}$, the total action reduces (again, up to a constant phase) to (11.8) as before.

10.3 Evaluating the Fresnel-Kirchhoff integral

(WORK OUT THIS!!!!)

Consider now a radio wave quantum, as described above, traversing an astrophysical plasma from its initial position at the source $\mathbf{x}_i = \mathbf{x}_s$ to its final position at the observer $\mathbf{x}_f = \mathbf{x}_{obs}$. For simplicity we assume the plasma takes the form of a thin, flat lens, with the phase velocity $c_p(\mathbf{x}) = c$, the speed of light *in vacuo*, everywhere except on the lens (see Fig. 11.1). Let us redefine the spatial coordinates $\mathbf{x} \rightarrow (\mathbf{x}, z)$ to separate out the coordinates in the lens plane \mathbf{x} from the normal coordinate z . The real classical paths are piecewise linear, with a possible bend at the lens, and the integral over these paths reduces to an ordinary integral over the lens plane [1]. The path integral amplitude for a (spinless) photon is obtained by integrating over

all paths weighted by the phase factor $e^{iS_r[\mathbf{x}]/\hbar}$ obtained from (11.9):

$$\Psi(\mathbf{x}_{obs}, \mathbf{x}_s) = \int d\mathbf{x} \exp \left[i\omega \int_{\mathbf{x}_s}^{\mathbf{x}_{obs}} |d\mathbf{x}| \frac{n(\mathbf{x})}{c} \right]. \quad (10.11)$$

where we replaced E with $\hbar\omega$, ω is the angular frequency of the light, and the phase velocity $c_p(\mathbf{x})$ with $c/n(\mathbf{x})$ where c is the speed of light *in vacuo* and $n(\mathbf{x})$ is the refractive index. For an astrophysical plasma, as mentioned above, at high frequency we have $n(\mathbf{x}) \approx 1 - \omega_p^2(\mathbf{x})/\omega^2$ where ω_p is the plasma frequency, given by $\omega_p^2(\mathbf{x}) \approx n_e(\mathbf{x})e^2/(m_e\epsilon_0)$, with $n_e(\mathbf{x})$, e , and m_e respectively the local electron density and the charge and mass of the electron in SI units (see *e.g.* [13]). We explicitly exhibit the \mathbf{x} -dependence since it governs the structure of the lens.

In the thin lens approximation, variations in the phase arise in part geometrically, from variations in the length of the straight line segments on either side of the lens, and in part from the passage through the lens. The former are straightforwardly computed using the Pythagorean theorem in the approximation that the relative horizontal displacements in Fig. 11.1, $|\mathbf{x}_{obs} - \mathbf{x}|$ and $|\mathbf{x}_s - \mathbf{x}|$ are much smaller than the vertical distances d_{lo} and d_{sl} . The latter are likewise computed approximately, noting that, to lowest order, the paths pass vertically through the lens so we may replace $\int n_e(\mathbf{x}, z) dz$ with $\Sigma_e(\mathbf{x})$, the electron surface density.

The path integral amplitude then becomes

$$\Psi(\mathbf{x}_{obs}, \mathbf{x}_s; \nu) \propto \int d\mathbf{x} \exp \left[i\frac{\omega}{2c} \left(\frac{(\mathbf{x} - \boldsymbol{\mu})^2}{\bar{d}} - \frac{\Sigma_e(\mathbf{x})e^2}{m_e\epsilon_0\omega^2} \right) \right], \quad (10.12)$$

where $\bar{d} = d_{sl}d_{lo}/(d_{sl} + d_{lo})$ is the reduced distance and $\boldsymbol{\mu} = (\mathbf{x}_s d_{lo} + \mathbf{x}_{obs} d_{sl})/d_{so}$ is a weighted average of the transverse displacements of the source and the observer. Notice that $\mathbf{x} - \boldsymbol{\mu}$ depends only on the *relative* displacements of the source, the lens and the observer, so that the answer is independent of the choice of origin for the transverse coordinates.

It is convenient to normalize the amplitude by dividing it by the amplitude obtained with the same geometry but no lens present. We may then write the resulting normalized amplitude as a dimensionless integral. Redefining $\mathbf{x} \rightarrow a\mathbf{x}$, $\boldsymbol{\mu} \rightarrow a\boldsymbol{\mu}$ where a is some convenient physical scale associated with the lens, we set $\nu = \omega a^2/(2c\bar{d}) = a^2/(2R_F^2)$ where $R_F = (\lambda\bar{d})^{1/2}$ is the Fresnel scale [6]. Notice that, because lensing alters the *angle* of propagation, the fringe spacing grows with the distance. Hence, it is the Fresnel scale – the geometric mean of the distance and the wavelength – rather than the wavelength which should be compared with the source dimensions to determine whether the interference pattern is observed in the heavily diffracted (low ν) or eikonal (high ν) regime. Finally, we define $\varphi(\mathbf{x}) = -\Sigma_e(\mathbf{x})e^2\bar{d}/(m_e\epsilon_0a^2\omega^2)$ to obtain the normalized, dimensionless amplitude,

$$\Psi(\boldsymbol{\mu}; \nu) = \left(\frac{\nu}{\pi} \right)^{N/2} \int_{\mathbb{R}^N} d\mathbf{x} \exp [i\phi(\mathbf{x}; \boldsymbol{\mu})\nu], \quad \text{with} \quad \phi(\mathbf{x}; \boldsymbol{\mu}) = (\mathbf{x} - \boldsymbol{\mu})^2 + \varphi(\mathbf{x}), \quad (10.13)$$

for an N -dimensional lens. Since $\nu \propto \omega$, we see that the eikonal limit is high frequency limit. However, the strength of the lens is controlled by φ which is proportional to ω^{-2} . Therefore the lens becomes stronger at lower frequencies where, of course, diffraction becomes important. The highest magnifications attained involve a payoff between strong lensing, creating effects like caustics and catastrophes, and diffraction which tends to smear out intensity peaks. Hence, to model the most interesting regime for astrophysical plasma lenses, one must go beyond geometrical optics and include diffractive effects.

The intensity corresponding to the amplitude (11.13) is proportional to the probability for a photon to be detected at $\boldsymbol{\mu}$:

$$I(\boldsymbol{\mu}; \nu) \propto |\Psi(\boldsymbol{\mu}; \nu)|^2. \quad (10.14)$$

The observed intensity should be normalized to the energy flux received by the detector, at each frequency, integrated over all observed $\boldsymbol{\mu}$. For a more detailed analysis see [44].

Except in special cases, the Fresnel-Kirchhoff integral (11.13) is not possible to evaluate analytically. At large ν (and with the dimensionless form of the lens $\varphi(\boldsymbol{x})$ held fixed) and in the geometric optics limit, one can easily model the intensity, as we shall explain. However, the most interesting regime for astrophysical plasma lenses occurs in the intermediate regime, where focusing and caustic catastrophes generate bright features whose peak intensity is controlled by diffraction. In this regime, there are characteristic patterns in the intensity, controlled by the topological character of the lens. In this intermediate- ν regime, conventional integration techniques typically fail, and it is hard to capture the complex, oscillatory interference pattern numerically. For example, G. Grillo and J. M. Cordes [25] implemented a procedure based on Fourier methods but found this technique to generate numerical artifacts. Here, motivated by our earlier work on Picard-Lefschetz theory, we instead employ analytic continuation and Cauchy's theorem to unambiguously define and to evaluate the relevant oscillatory integrals. We have developed a custom numerical scheme (made available [online](#)¹⁴.) which is fast and accurate, and applicable to a generic one dimensional oscillatory integral. A two dimensional version will be made available shortly. A nice feature of our method is that it typically becomes more efficient, *i.e.*, its convergence is improved, as the integrand becomes more oscillatory and difficult to handle via conventional techniques.

10.3.1 The geometric optics limit

In the limit of large ν , the Fresnel-Kirchhoff integral is dominated by real stationary points of the phase function ϕ which, except at special values of $\boldsymbol{\mu}$, are well-approximated by Gaussians. Furthermore, any interference between different stationary points leads to oscillations in the intensity which, in the limit $\nu \rightarrow \infty$, become increasingly rapid. In the geometric optics approximation, one averages these oscillations away. Physically this averaging occurs through the incoherence of any realistic extended source, as we explain later. Although this paper is devoted to the study of interference phenomena, it proves useful to begin by studying the geometric optics limit.

In the large ν (eikonal) limit, we focus on real critical points of the exponent, *i.e.*, those values of \boldsymbol{x} for which

$$\nabla_{\boldsymbol{x}}\phi(\boldsymbol{x}; \boldsymbol{\mu}) = 0, \quad (10.15)$$

considered as a function of the parameter $\boldsymbol{\mu}$. The critical points are generally smooth complex-valued functions of $\boldsymbol{\mu}$. In the eikonal limit, only the real critical points contribute because contributions from complex saddle points are exponentially suppressed. The critical points can be described in terms of the Lagrangian map $\boldsymbol{\xi} : X \rightarrow M$, mapping the points in the base space $\boldsymbol{x} \in X = \mathbb{R}^N$ to points in the parameter space $\boldsymbol{\mu} \in M$ according to the critical point condition

$$\nabla_{\boldsymbol{x}}\phi(\boldsymbol{x}; \boldsymbol{\mu})|_{\boldsymbol{\mu}=\boldsymbol{\xi}(\boldsymbol{x})} = 0. \quad (10.16)$$

The Lagrangian map is determined by the gradient of the phase of the lens:

$$\nabla_{\boldsymbol{x}}\phi(\boldsymbol{x}; \boldsymbol{\mu}) = 2(\boldsymbol{x} - \boldsymbol{\mu}) + \nabla\varphi(\boldsymbol{x}) \quad \implies \quad \boldsymbol{\mu} = \boldsymbol{\xi}(\boldsymbol{x}) = \boldsymbol{x} + \frac{1}{2}\nabla\varphi(\boldsymbol{x}). \quad (10.17)$$

The Lagrangian map $\boldsymbol{\xi}$ determines the optical rays, giving a purely geometric description of the lens. Every point \boldsymbol{x} is mapped to a point $\boldsymbol{\mu}$ in the space of observational parameters. In general, a point $\boldsymbol{\mu} \in M$ might be obtained from several points in $\boldsymbol{x} \in X$, *i.e.*, the Lagrangian

¹⁴See https://github.com/jfeldbrugge/Picard_Lefschetz_Integrator.

Distant point source

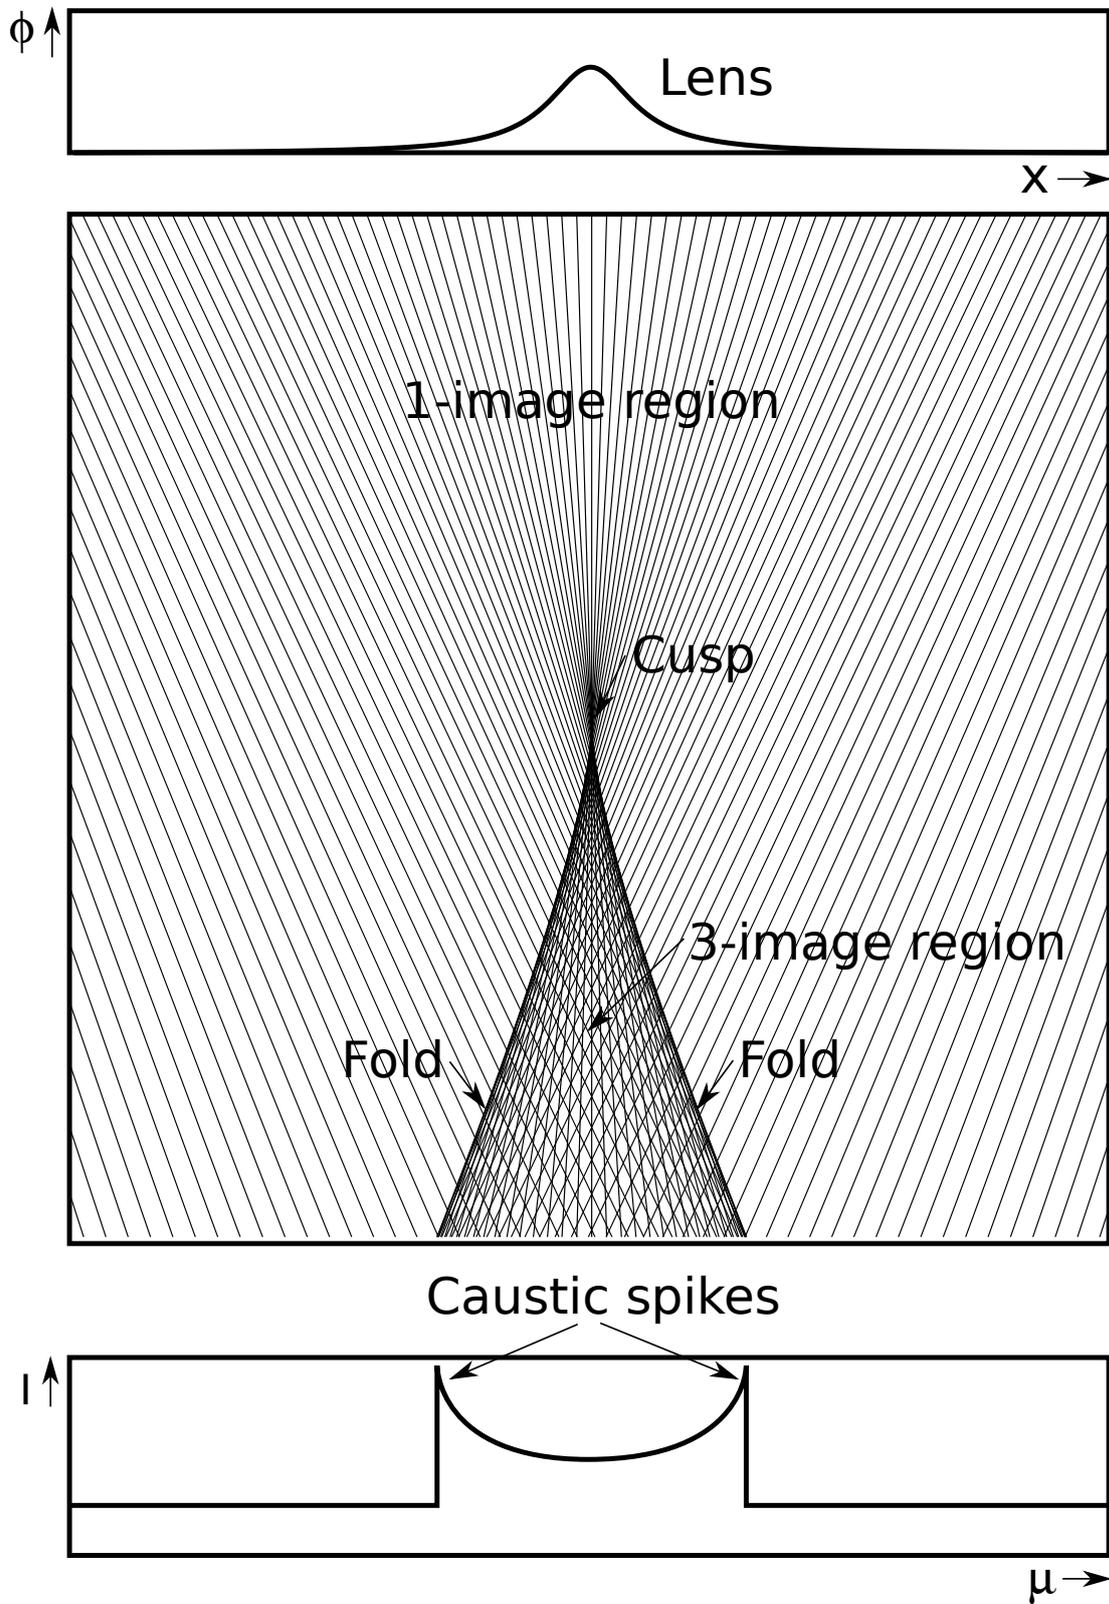


Figure 10.2: The Lagrangian map in geometrical optics. The image consists of two single and one triple image regions separated by a fold caustic at which the normalized intensity spikes.

map can be many-to-one. The regions in $\boldsymbol{\mu}$ where each point is obtained from n points in X are known as n -image regions. In multi-image regions, one adds the intensities due to each of the contributing paths: performing the relevant Gaussian integrals one finds for the normalized intensity

$$I(\boldsymbol{\mu}; \infty) = \sum_{\boldsymbol{x} \in \boldsymbol{\xi}^{-1}(\boldsymbol{\mu})} \frac{2}{|\lambda_1(\boldsymbol{x})| \dots |\lambda_N(\boldsymbol{x})|}, \quad (10.18)$$

with $\boldsymbol{\xi}^{-1}$ the pre-image of the Lagrangian map and $\lambda_1, \dots, \lambda_N$ the eigenvalues of the deformation tensor,

$$\mathcal{M}_{ij}(\boldsymbol{x}) = \frac{\partial^2 \phi(\boldsymbol{x}; \boldsymbol{\mu})}{\partial x_i \partial x_j} = 2 \frac{\partial \xi_i(\boldsymbol{x})}{\partial x_j}, \quad (10.19)$$

evaluated at the relevant critical points $\boldsymbol{x} = (x_1, x_2, \dots, x_N)$. Below, we shall study these multi-image regions in detail, at finite ν , where they exhibit intricate interference patterns.

At the boundaries between regions with a different number of images, at least one of the eigenvalue fields λ_i must vanish. At infinite ν this leads to an infinite spike in the normalized intensity map, signalling a caustic. See Fig. 11.2 for an illustration of the Lagrangian map corresponding to a one-dimensional lens with a single- and a triple-image region. The triple-image region is separated from the single-image regions by two fold caustics. At the fold caustic the normalized intensity profile diverges.

Formally, the Lagrangian map $\boldsymbol{\xi}$ forms a caustic at $\boldsymbol{x}_c \in X$ when the deformation tensor becomes singular, *i.e.*, its determinant vanishes. However, the X space is generally not observed. In the space M of observable parameters, the caustic at \boldsymbol{x}_c appears at the point $\boldsymbol{\mu}_c = \boldsymbol{\xi}(\boldsymbol{x}_c)$. For one-dimensional lenses caustics occur at isolated points. For higher-dimensional cases, the determinant of the deformation tensor vanishes on a manifold $X_c = \{\boldsymbol{x} \in X \mid |\mathcal{M}(\boldsymbol{x})| = 0\}$ which is mapped to a caustic set $M_c = \boldsymbol{\xi}(X_c)$ in the parameter space. Note that the set M_c is generally *not* a manifold, as it includes higher order caustics, such as cusps and swallowtails, at which the variety is non-differentiable and therefore M_c is not locally homeomorphic to Euclidean space. We shall discuss examples of this kind later, but note here that they are exactly the points at which the lensing integral exhibits the most interesting behaviour.

The geometric optics limit is attained in two stages: at short wavelengths, each real stationary point corresponds to a distinct image. As the wavelength is increased, each image itself forms an interference pattern, as illustrated in the Young's double slit experiment examined in Appendix 11.10. The limit of short wavelengths, in which phase coherence is maintained, is often called the eikonal approximation. However, when phase coherence is lost - for example, when the source size becomes larger than the spacing of its fringes, interference effects disappear altogether. This assumption of loss of coherence is implicit in the geometric optics limit. However, objects smaller than the Fresnel scale are still seen to scintillate, as a result of coherent interference effects on unresolved scales. This is reflected in the expression *stars twinkle, planets don't*. Interstellar scintillation typically occurs for sources smaller than about a micro arcsecond, corresponding to the Fresnel angle $\theta_F = \sqrt{\lambda/d}$ (with λ the wavelength and d the distance from the lens) on the sky. This condition is true for most FRBs and pulsars. Interplanetary scintillation due to the solar wind is commonly seen for many compact extragalactic radio sources at low frequencies [45]. In this case, the characteristic Fresnel angle for wavelengths of a few meters and distances of an astronomical unit is a fraction of an arcsecond. Ionospheric scintillation is strongest at the lowest frequencies, and is commonly seen at solar maximum or at equatorial locations near sunrise or sunset [46], and causes loss of lock in GPS. The Fresnel angular scale for a screen at a distance of 200km at wavelengths of a meter is 8 arc minutes, causing all celestial sources except the sun and the moon to scintillate.

Name	Symbol	K	N	$\phi(\mathbf{x}; \boldsymbol{\mu})$
Maximum/minimum	A_1^\pm	0	1	$\pm x^2$
Fold	A_2	1	1	$x^3/3 + \mu x$
Cusp	A_3	2	1	$x^4/4 + \mu_2 x^2/2 + \mu_1 x$
Swallowtail	A_4	3	1	$x^5/5 + \mu_3 x^3/3 + \mu_2 x^2/2 + \mu_1 x$
Butterfly	A_5	4	1	$x^6/6 + \mu_4 x^4/4 + \mu_3 x^3/3 + \mu_2 x^2/2 + \mu_1 x$
Elliptic umbilic	D_4^-	3	2	$x_1^3 - 3x_1 x_2^2 - \mu_3(x_1^2 + x_2^2) - \mu_2 x_2 - \mu_1 x_1$
Hyperbolic umbilic	D_4^+	3	2	$x_1^3 + x_2^3 - \mu_3 x_1 x_2 - \mu_2 x_2 - \mu_1 x_1$
Parabolic umbilic	D_5	4	2	$x_1^4 + x_1 x_2^2 + \mu_4 x_2^2 + \mu_3 x_1^2 + \mu_2 x_2 + \mu_1 x_1$

Table 1: The unfoldings of the seven elementary catastrophes with codimension $K \leq 4$, with $\mathbf{x} = (x_1, x_2, \dots, x_N)$ and $\boldsymbol{\mu} = (\mu_1, \mu_2, \dots, \mu_K)$. The normal forms are defined as the unfolding at parameter $\boldsymbol{\mu} = \mathbf{0}$, *i.e.*, $\phi(\mathbf{x}; \mathbf{0})$.

10.3.2 Catastrophe theory

Catastrophe theory is the mathematical classification of stable critical points. Caustics are classified by Lagrangian catastrophe theory [47, 48], which is a special application of the general theory. Given the definition of the Lagrangian map $\boldsymbol{\xi}$, the connection between caustics in optical systems and critical points is not surprising. For one-dimensional functions, the classification consists only of minima and maxima. The local minima and maxima of a one-dimensional function are stable, *i.e.*, the addition of a small perturbation merely leads to a displacement of the critical point. Degenerate critical points are not included, as they are not stable in one dimension. For example, a cubic critical point decomposes into a minimum and a maximum, or no critical point at all, when perturbed.

In the catastrophe theory of higher-dimensional functions, degenerate critical points are included because they are stable. René Thom (1972) proved [49] that the stable critical points with co-dimension¹⁵ K less than or equal to 4 are classified by the seven “elementary catastrophes.” These seven singularities suffice to classify the full range of caustics emerging in three-dimensional lenses. Thom named the seven catastrophes: *the fold, cusp, swallowtail, butterfly, and the elliptic, hyperbolic and parabolic umbilic*. The caustics were in the subsequent years connected and labeled by the Coxeter reflection groups (Arnol’d [50, 51]). The theory was subsequently applied to optical interference patterns by Berry and collaborators, and beautiful experiments were performed [6]. For a more recent theoretical investigation of catastrophe theory and caustics with applications to large-scale structure formation see [52]. Here we briefly review catastrophe theory and its application to oscillatory integrals.

Table 1 lists the seven “elementary catastrophes” and their unfoldings $\phi(\mathbf{x}; \boldsymbol{\mu})$. The unfolding $\phi(\mathbf{x}; \boldsymbol{\mu})$ evaluated at $\boldsymbol{\mu} = \mathbf{0}$ is the normal form of the catastrophe, representing the archetypical form of the critical point near $\mathbf{x} = \mathbf{0}$. We observe that the fold and the cusp respectively correspond to a cubic and quartic critical point of x . The parameter $\boldsymbol{\mu}$ represents the ways in which the caustic can decompose into lower-order caustics. In the case of the fold, we see that a linear perturbation decomposes the fold into a minimum and a maximum for $\mu < 0$ and no critical point at all for $\mu > 0$. The seven catastrophes belong to two families, classified by their co-rank¹⁶. The A -family is of co-rank $N = 1$, while the D -family is of co-rank $N = 2$. Critical points with higher co-rank have a co-dimension higher than 4, and for this reason are not included here. The co-rank N and the co-dimension K characterize the critical point. It generally takes N variables to describe the critical point, and it takes K

¹⁵The co-dimension of a caustic is roughly the dimensionality of the singularity. The stable critical points of a n -dimensional function are completely classified by the caustics with co-dimension smaller or equal to n .

¹⁶The co-rank is the number vanishing eigenvalues of the Hessian matrix.

Catastrophe	Symbol	I_0	β	σ_j
Fold	A_2	1.584	1/6	$\sigma_1 = 2/3$
Cusp	A_3	2.092	1/4	$\sigma_1 = 3/4, \sigma_2 = 1/2$
Swallowtail	A_4	1.848	3/10	$\sigma_1 = 4/5, \sigma_2 = 3/5, \sigma_3 = 2/5$
Butterfly	A_5	1.991	1/3	$\sigma_1 = 5/6, \sigma_2 = 2/3, \sigma_3 = 1/2, \sigma_4 = 1/3$
Elliptic umbilic	D_4^-	1.096	1/3	$\sigma_1 = 2/3, \sigma_2 = 2/3, \sigma_3 = 1/3$
Hyperbolic umbilic	D_4^+	0.580	1/3	$\sigma_1 = 2/3, \sigma_2 = 2/3, \sigma_3 = 1/3$
Parabolic umbilic	D_5	2.258	3/8	$\sigma_1 = 5/8, \sigma_2 = 3/4, \sigma_3 = 1/2, \sigma_4 = 1/4$

Table 2: The intensity and fringe separation scaling relations for the catastrophes shown shown in Table 1. At large ν the maximum intensity (11.14) is given by $I_0\nu^{2\beta}$ (see the discussion following Eq. (11.21)) and the fringe scaling exponents are defined in (11.22).

parameters to describe its unfolding. In more prosaic terms, N is the dimension of the space of \mathbf{x} 's and K is the dimension of the space of $\boldsymbol{\mu}$'s.

For each of the normal forms listed in Table 1, the normalized amplitude

$$\Psi(\boldsymbol{\mu}; \nu) = \left(\frac{\nu}{\pi}\right)^{N/2} \int e^{i\phi(\mathbf{x}; \boldsymbol{\mu})\nu} d\mathbf{x}, \quad (10.20)$$

forms a caustic at the critical point $\boldsymbol{\mu} = 0$. For a detailed analysis including illustrations of the intensities obtained in each case, see chapter 36 of [53]. As ν is increased to large values, the normalized intensity $I(\boldsymbol{\mu}; \nu) = |\Psi(\boldsymbol{\mu}; \nu)|^2$ diverges and the scale of the associated diffraction fringes shrinks to zero according to scaling laws which are specific for each catastrophe. At large ν , the maximum of the intensity is attained near $\boldsymbol{\mu} = 0$ as illustrated, for example, by the fold singularity shown in Fig. 11.13. The maximum intensity obeys the following scaling law at large ν :

$$I(\mathbf{0}, \nu) = I_0\nu^{2\beta}. \quad (10.21)$$

The constant β , termed the *singularity index* by Arnold (Arnold [51] and Varchenko [54]), is universal, being invariant under diffeomorphisms and depending only on the topological class of the catastrophe. It is given, for each case, in the fourth column of Table 2. The scaling with ν is easily seen by examining the corresponding normal form. Setting the unfolding parameter to zero, *i.e.*, $\boldsymbol{\mu} = \mathbf{0}$, in the phases listed in Table 1, one can render the phase of the integrand independent of ν by rescaling the integration variables \mathbf{x} . For example, for A_2 we set $x = \nu^{-\frac{1}{3}}y$. Taking into account the ν -dependence arising from the Jacobian in the integration measure as well from the prefactor in (11.20), one infers that the amplitude at the caustic scales as $\nu^{\frac{1}{6}}$ for A_2 and hence that $\beta = \frac{1}{6}$ for this case. For the two dimensional lenses, one has to rescale both x_1 and x_2 in order to remove ν from the exponent but the argument is otherwise the same.

For each of the normal forms of the phase listed in Table 1, one may also analytically compute the constant I_0 , and we provide its numerical value in Table 2. When considering a class of lens models for modeling purposes (such as the localised models we consider later), it may be helpful to notice an additional scaling property. At large ν , the amplitude is determined by the form of the phase near the critical point. Indeed, this is how universality arises. For any lens model which includes a given catastrophe, the leading terms in the Taylor expansion of the phase about the associated critical point will, after coordinate redefinitions, take the form of one of the "normal forms" listed in Table 1. Coordinate rescalings are one of the simplest such transformations, which have a simple effect on any model and on its Taylor expansion about any critical point. One can easily derive the scaling behavior of the intensity under such transformations of the lens model. For an A_n catastrophe, for example, we may consider a set

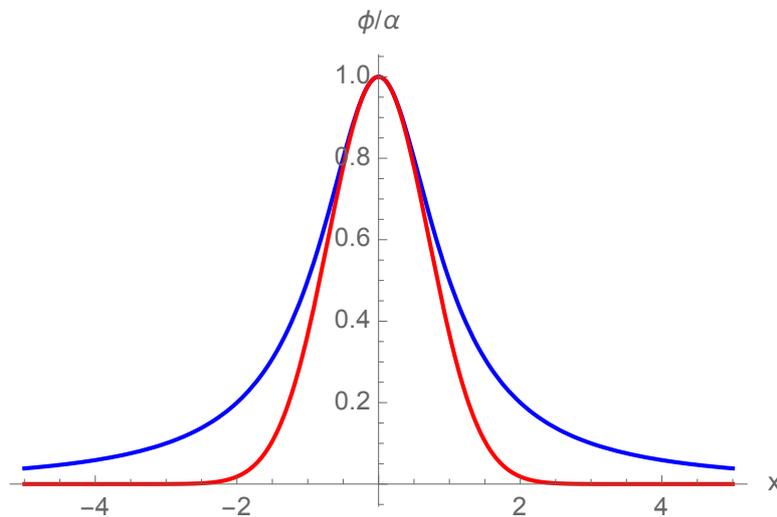


Figure 10.4: A comparison of the Gaussian lens (red) and the rational approximation (blue).

The exponent $i\phi(\mathbf{x}; \boldsymbol{\mu})\nu$ evaluated along a steepest descent contour \mathcal{J}_i has a constant imaginary part while its real part $h = \text{Re}[i\phi(\mathbf{x}; \boldsymbol{\mu})\nu]$ is monotonically decreasing. As a consequence, the conditionally convergent oscillatory integral is transformed into a sum of convex integrals. This is remarkable, as the originally conditionally convergent integral is generally sensitive to regularization and also, if the integral is performed iteratively, to the order in which the partial integrals are taken (see Appendix 11.9 for an instructive example). The integrals over Lefschetz thimbles have none of these ambiguities. It is for this reason that we will interpret the integral over the sum of Lefschetz thimbles as the *definition* of the integral taken over the real integration domain. Once we have identified the correct set of thimbles, we can use conventional numerical methods to evaluate the integral on each thimble.

We shall describe two distinct methods to obtain the sum of Lefschetz thimbles corresponding to the Fresnel-Kirchhoff integral for a one-dimensional lens. In the first method, we follow the techniques explained, for example, in Ref. [27]. We start by computing all the saddle points of the exponent $i\phi(\mathbf{x}; \boldsymbol{\mu})\nu$ and their corresponding steepest descent and ascent contours. We subsequently study the intersections of the steepest ascent contours with the original integration domain to find the relevant saddle points and associated steepest descent contours \mathcal{J}_i . This method is well suited to the Picard-Lefschetz analysis of one-dimensional integrals, for which we can plot the steepest descent and ascent contours in the complex plane \mathbb{C} .

In the second method, we instead flow the real integration domain along the downward flow of the real part of the exponent, $h = \text{Re}[i\phi(\mathbf{x}; \boldsymbol{\mu})\nu]$. We show that this downward flow terminates on the correct sum of Lefschetz thimbles $\sum_i \mathcal{J}_i$. The relevant saddle points are given by the local maxima of h restricted to this thimble sum \mathcal{J} . Note that this second scheme is completely determined by the gradient of the h with respect to the real and imaginary parts of the complexified coordinates \mathbf{x} . We do not need to find all the saddle points nor evaluate the corresponding steepest ascent and descent contours. We moreover do not need to study the intersections of the steepest ascent contours with the original integration domain. Any Stokes transitions are automatically taken care of. This method is furthermore ideally suited to higher-dimensional oscillatory integrals, where the steepest ascent and descent contours are expensive to evaluate and the intersections are computationally difficult to find.

10.4.1 Geometric optics approximation

In the introduction to this paper, we discussed a one-parameter family of one-dimensional localized lenses

$$\varphi(x) = \frac{\alpha}{1+x^2}, \quad (10.24)$$

with $\alpha \in \mathbb{R}$. For plasma lenses, the parameter α follows the dispersion relation $\alpha \propto \omega^{-2}$ with ω the angular frequency of the source. The longer the wavelength, the stronger the lens. We restrict our analysis to rational lenses for two reasons:

1. their analytic continuation into the complex x -plane does not contain branch-cuts and consists of only a finite number of poles,
2. the phase ϕ has only a finite number of saddle points and corresponding steepest-descent contours.

Picard-Lefschetz theory, however, applies to analytic lenses in general. The lens (11.24) is a rational approximation to the Gaussian lens

$$\varphi(x) = \alpha e^{-x^2}, \quad (10.25)$$

which has an essential singularity at infinity on the Riemann sphere and an infinite number of saddle points in the complex plane. See Fig. 11.4 for a comparison between the two lenses. It is a wonderful fact that many real-valued functions with intricate structure in the complex plane, can be well-approximated with a Padé approximation, whose analytic continuation possesses only a finite number of poles.

As we derived in Section 11.3.1, the Lagrangian map ξ of the rational lens φ ,

$$\xi(x) = x - \frac{\alpha x}{(1+x^2)^2}, \quad (10.26)$$

forms caustics at the real roots of the second order derivative of the exponent

$$\frac{\partial^2 \phi(x)}{\partial x^2} = 2 \frac{\partial \xi(x)}{\partial x} = 2 + 2\alpha \frac{3x^2 - 1}{(1+x^2)^3} = 0. \quad (10.27)$$

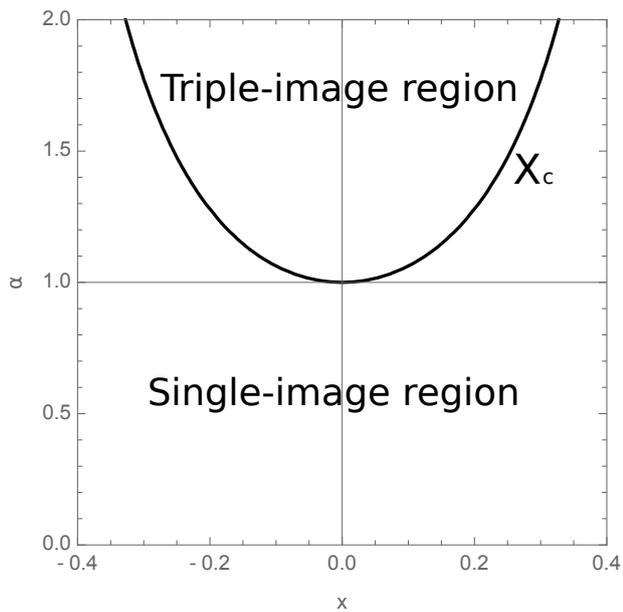
See Fig. 11.5 for the caustic surface in the x - α and the μ - α planes. For $\alpha < 1$ no such real root exists. The lensed image consists of a single-image region. For $\alpha = 1$ there is a single real-valued root at $x_c = 0$ with multiplicity two and the corresponding point $\mu_c = \xi(x_c) = 0$ in the parameter space M . In the μ - α plane, this point is non-differentiable on the caustic set. This is an example of a cusp caustic. For $\alpha > 1$ there are two symmetric real roots. These are examples of fold caustics. For further reference, for $\alpha = 2$, the two caustics are located at

$$X_c = \{-0.327334 \dots, 0.327334 \dots\}, \quad (10.28)$$

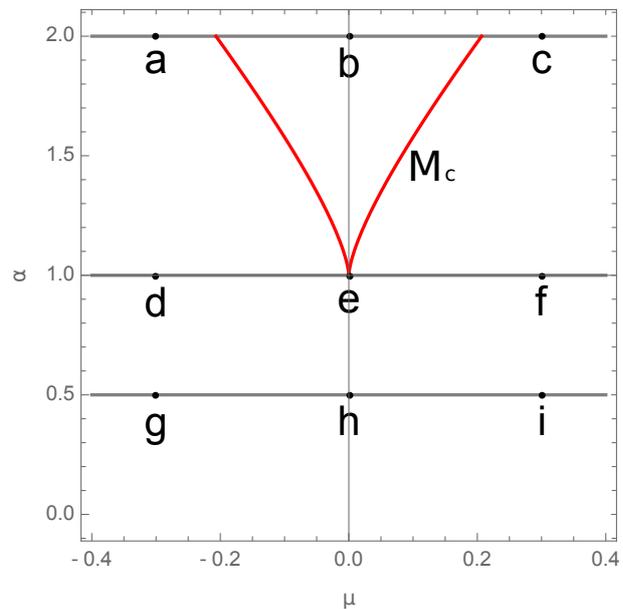
in the base space $X = \mathbb{R}$. In the parameter space M the caustic appears at

$$M_c = \xi(X_c) = \{+0.206751 \dots, -0.206751 \dots\}. \quad (10.29)$$

The relative normalized intensity of the lens in the geometric optics limit (see Section 11.3.1) is plotted in figure 11.6. For $\alpha = 1/2$, the lens does not form a caustic. The normalized intensity map is finite. For $\alpha = 1$, we see a cusp caustic at $\mu = 0$. For $\alpha = 2$, we observe two fold caustics at M_c enclosing a triple-image region. The black curves in the triple-image region are the three contributions corresponding to the three images. The red curve is the sum over the multi-image regions.

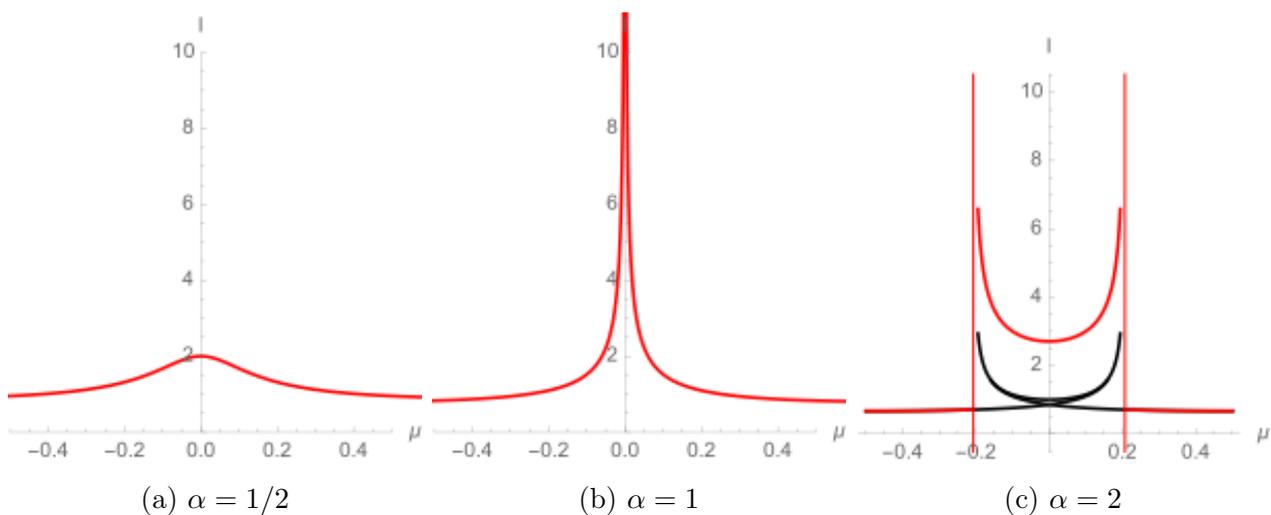


(a) Caustic in the x - α plane.



(b) Caustic in the μ - α plane.

Figure 10.5: The caustics of the one-dimensional rational lens in the x - α and the μ - α plane. The points in the *left panel* correspond to the Picard-Lefschetz diagrams in Fig. 11.7. The lines correspond to the panels in Fig. 11.11.



(a) $\alpha = 1/2$

(b) $\alpha = 1$

(c) $\alpha = 2$

Figure 10.6: The normalized intensity in geometric optics, $I(\mu; \infty)$ is plotted as a function of μ .

10.4.2 Finding the thimbles

We now turn to evaluating the full expression (11.23). First, we need to determine which Lefschetz thimbles contribute. We shall describe two distinct methods, the second of which is more efficient for numerical purposes.

Method 1: following steepest ascent contours

The exponent

$$i\phi(x; \mu)\nu = i \left[(x - \mu)^2 + \frac{\alpha}{1 + x^2} \right] \nu \quad (10.30)$$

is imaginary for real μ and x . Its analytic continuation has two poles at $x = \pm i$, and five saddle points in the complex x -plane, satisfying

$$\frac{\partial\phi(x; \mu)}{\partial x} = 2(x - \mu) - \frac{2\alpha x}{(1 + x^2)^2} = 0. \quad (10.31)$$

For the Picard-Lefschetz analysis we start by writing the analytic continuation of the exponent in terms of its real and imaginary part

$$i\phi(\mathbf{x}; \boldsymbol{\mu})\nu = h(\mathbf{u} + i\mathbf{v}; \mu) + iH(\mathbf{u} + i\mathbf{v}; \mu), \quad (10.32)$$

with the complex expansion $\mathbf{x} = \mathbf{u} + i\mathbf{v}$ and the real-valued functions h, H . For generality, we describe the flow of the integration contour in N dimensions. The real part h is, in the Picard-Lefschetz analysis, known as the h -function. The downward flow of the h -function $\gamma_\lambda : \mathbb{C}^N \rightarrow \mathbb{C}^N$ is defined by

$$\frac{\partial\gamma_\lambda(\mathbf{z})}{\partial\lambda} = -\nabla_{\mathbf{u}+i\mathbf{v}}h[\gamma_\lambda(\mathbf{z})] \quad (10.33)$$

with the boundary condition $\gamma_0(\mathbf{z}) = \mathbf{z} \in \mathbb{C}^N$, the parameter λ in a subset of \mathbb{R} , and the complex gradient defined as

$$\nabla_{\mathbf{u}+i\mathbf{v}}h = \nabla_{\mathbf{u}}h + i\nabla_{\mathbf{v}}h. \quad (10.34)$$

Note that in defining the gradient, we have assumed a corresponding metric on the space \mathbb{C}^N . In this paper we will always associate \mathbb{C}^N with \mathbb{R}^{2N} and use the corresponding Euclidean metric. We are of course free to consider different metrics. Given the saddle points we can compute the steepest ascent and descent contours and intersect the ascent contours with the real axis, to obtain the relevant saddle points.

Depending on μ and α either one or three of the saddle points are real-valued. The lens thus has both single- and triple-image regions. See Fig. 11.7 for the five saddle points \bar{x}_i and the corresponding steepest ascent and descent contours. By intersecting the steepest ascent contours with the real line, we obtain the Lefschetz thimble (plotted in blue). The thimbles run from $x = -\infty$ to $x = +\infty$, while passing through the poles at $x = \pm i$.

From the caustic structure in Fig. 11.5 we can distinguish three regimes:

- In the regime $\alpha < 1$, the lens forms a single image. The corresponding Picard-Lefschetz analysis yields a single real-valued saddle point. For large $|\mu|$ there is, in addition, a relevant complex saddle point. When $|\mu|$ decreases to 0, the complex saddle point becomes irrelevant due to a Stokes transition. This phenomenon is discussed in detail in the next section. For $\mu = 0$, only the real saddle point is relevant. Note that the thimble can for all μ be deformed to the original integration domain \mathbb{R} . See the lower panels of Fig. 11.5.

- For $\alpha = 1$, the lens contains a cusp caustic at $\mu_c = 0$. For $\mu \neq \mu_c$, the Picard-Lefschetz analysis is similar to the $\alpha < 1$ regime. The thimble passes through one real-valued and one complex-valued saddle point. At the caustic $\mu = \mu_c$, three non-degenerate saddle points merge forming a degenerate saddle point. This is the signature of the cusp caustic, whose normal form is the quartic function x^4 . See the middle panels of Fig. 11.5.
- In the regime $\alpha > 1$, the Picard-Lefschetz analysis splits into three intervals. In the single-image region, *i.e.*, μ in $(-\infty, -\mu_c)$ or (μ_c, ∞) , the Picard-Lefschetz analysis consists of two relevant saddle points; one real and one complex. At the caustic, the complex saddle point approaches the real line and merges with its complex conjugate saddle point. This is the signature of the fold caustic. In the triple-image region, *i.e.*, $\mu \in (-\mu_c, \mu_c)$, the analysis consists of three real-valued relevant saddle points. See the upper panels of Fig. 11.5.

Method 2: flowing the integration domain

We can alternatively obtain the Lefschetz thimble \mathcal{J} by flowing the original integration domain \mathbb{R} along the downward flow of the real part h .

Given the downward flow γ_λ for general points $z \in \mathbb{C}$, we flow the original integration domain X to

$$X_\lambda = \gamma_\lambda(X) \subset \mathbb{C}. \quad (10.35)$$

The steepest descent contours \mathcal{J}_i corresponding to the saddle points \bar{x}_i are the fixed points of the flow, *i.e.*,

$$\gamma_\lambda(\mathcal{J}_i) = \mathcal{J}_i \quad (10.36)$$

for all λ . When the h -function has saddle point in the complex plane, it follows from Morse-Smale theory [55, 56] that the flowed contour X_λ will converge to a set of steepest descent contours \mathcal{J}_i as $\lambda \rightarrow \infty$. Since X_λ is a continuous deformation of the original integration domain X , it follows that X_λ converges to the Lefschetz contour, *i.e.*,

$$\lim_{\lambda \rightarrow \infty} X_\lambda = \mathcal{J}. \quad (10.37)$$

When we perform the flow of the original integration domain as a function of the parameter μ , we obtain a family of thimbles. The thimble generally changes smoothly as a function of μ . There are however two ways in which the Picard-Lefschetz structure of the integral can abruptly change its geometry:

1. If for some μ , a few non-degenerate saddle points merge to form a higher order saddle point, the number of relevant critical points will change. At these points, the integral $\Psi(\mu; \nu)$ forms a caustic. This phenomenon can be described by catastrophe theory (see sections 11.3.2 and 11.5).
2. When the imaginary part H evaluated in two saddle points coincides for some parameter μ_s , the two corresponding steepest-descent contours can coincide. At such a parameter μ_s , the Lefschetz thimbles flip changing the number of relevant saddle points (see Fig. 11.8 for an illustration). This is known as a Stokes transition. The parameters μ_s for which this happens form so-called Stokes lines.

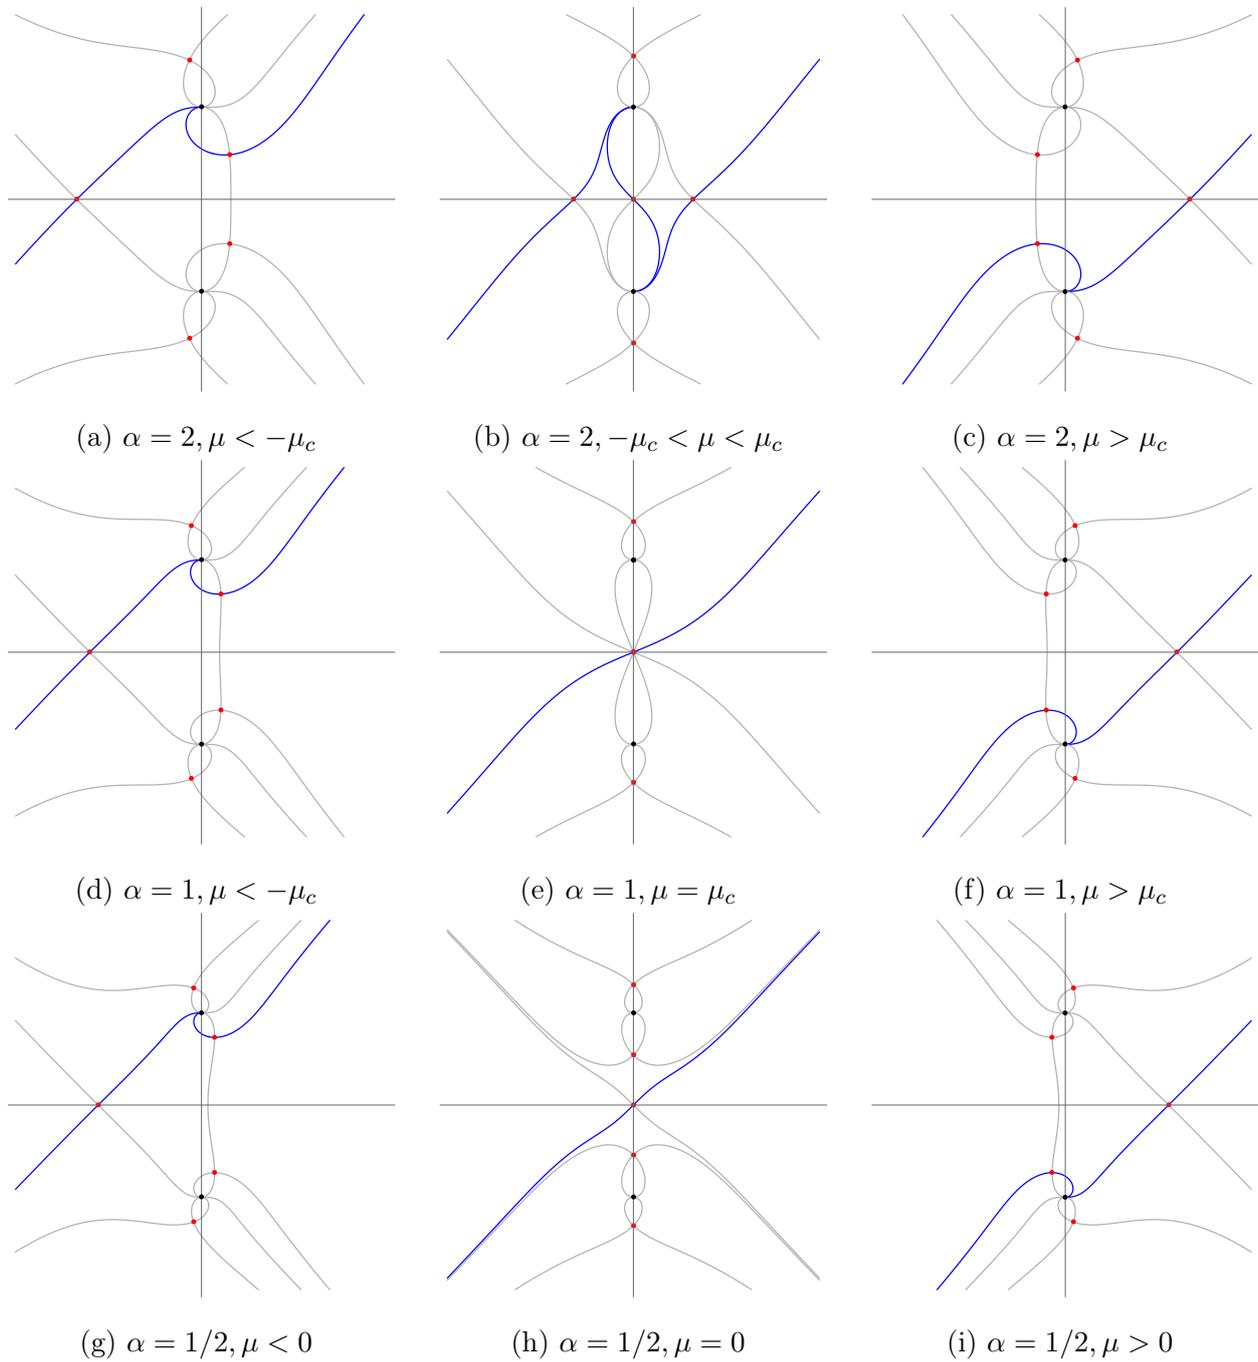


Figure 10.7: The Picard-Lefschetz thimbles for $\alpha = 1/2, 1, 2$ as a function of μ . The red and black points are the saddle points and poles. The curves are paths of steepest descent and ascent. The blue ones are relevant, while the grey ones are irrelevant.

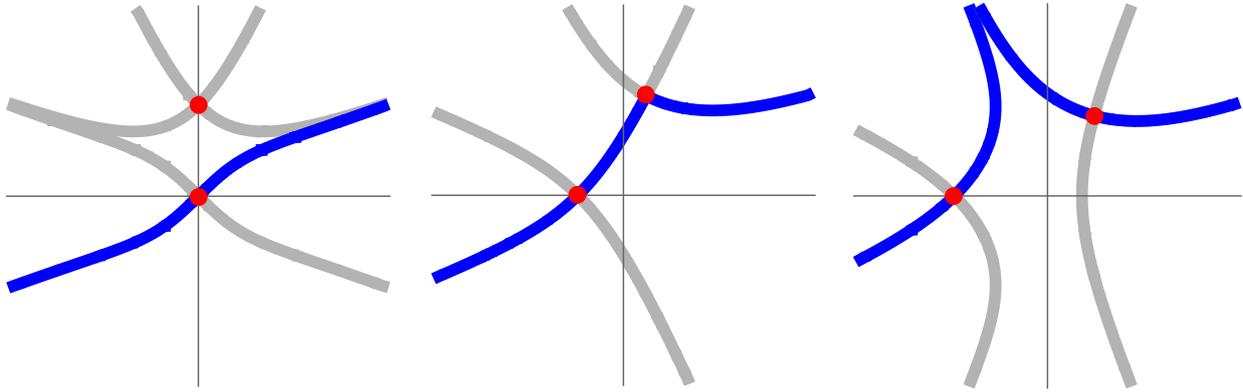


Figure 10.8: The Stokes phenomenon at which the steepest descent contours coincide and the relevant saddle points change. In the *left* panel, we have one real relevant saddle point. In the *central* panel, we see a Stokes phenomenon. The Lefschetz thimble passes to both a real and a complex saddle point. In the *right* pane, we have one real and one complex relevant saddle point. The Stokes phenomenon occurs when the steepest descent contour of a saddle point connects with another saddle point.

We study both phenomena in detail in the Section 11.5.

We numerically evaluate the flow X_λ by approximating X by a set of line-segments and flowing the endpoints. Since the real part of the analytic continuation of an analytic function does not have local extrema (this follows from the Cauchy-Riemann equation), all points $z \in \mathbb{C}$ flow to poles as $\lambda \rightarrow \infty$. The limit $\lim_{\lambda \rightarrow \infty} X_\lambda$ should not be interpreted as a pointwise limit. We, for this reason, trace the length of the line-segments and add points when neighboring points move too far apart. We moreover remove line-segments in the neighborhoods of the poles of the h -function. The contour X_λ has converged to the thimble when the imaginary part H is approximately constant along the line-segments.

This idea is implemented by the algorithm:

Algorithm 1 The flow of the contour of one-dimensional oscillatory integrals.

Require: Represent a subset $[a, b]$ of the original integration domain $X = \mathbb{R}$ by the regular lattice $p_i = a + i\Delta x$ with $\Delta x = \frac{b-a}{n}$ for some $n \in \mathbb{Z}_{>0}$, and the line-segments $(p_0, p_1), (p_1, p_2), \dots, (p_{n-1}, p_n)$.

while the variance of the imaginary part H on the points p_i exceeds threshold T_1 **do**
 flow the points: $p_i \mapsto p_i - \nabla h(p_i)\Delta t$
 if the h -function evaluated in the point p_i is smaller than the threshold T_2 **then**
 remove the corresponding line segments
 end if
 if the length of the line-segments (p_i, p_{i+1}) exceeds the threshold T_3 **then**
 split the line segment into the two lines $(p_i, \frac{p_i+p_{i+1}}{2}), (\frac{p_i+p_{i+1}}{2}, p_{i+1})$.
 end if
end while

with the parameters $a, b, T_1, T_2, T_3 \in \mathbb{R}$, and $n \in \mathbb{Z}_{>0}$.

See Fig. 11.9 for the flow of the original integration domain corresponding the rational lens for $\alpha = 2$ and $\mu = 0$. For $\lambda = 0$ the contour X_λ coincides with the real line. As λ is increased to 1, the original integration domain smoothly flows to the Lefschetz thimble \mathcal{J} consisting of three steepest descent contours \mathcal{J}_i corresponding to three relevant saddle points \bar{x}_i . By evaluating the flow for varying α and μ , we obtain the Picard-Lefschetz analysis of the lens.

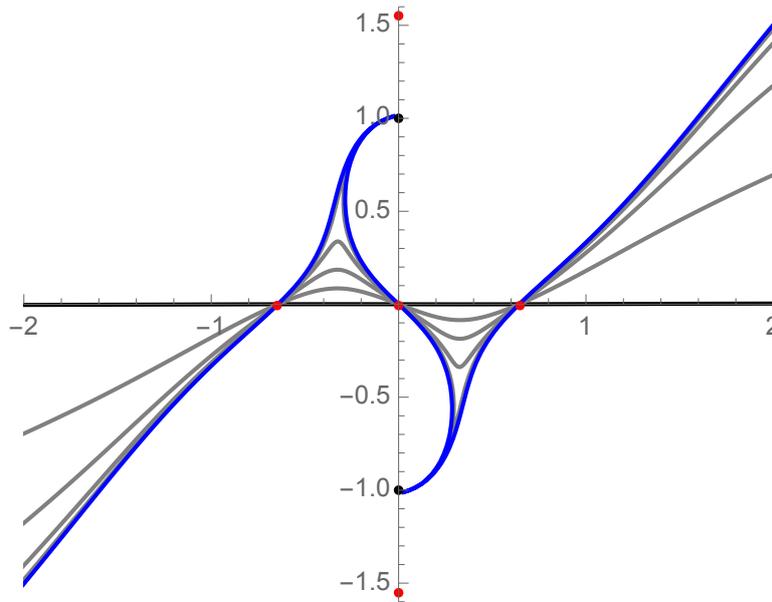


Figure 10.9: The downward flow of the integration domain. The contour X_λ for $\lambda = 0, 0.2, 0.4, 0.6, 0.8, 1$ is plotted by the black, grey and blue curves. The five saddle points are plotted in red and the poles are plotted in black.

For multi-dimensional oscillatory integrals, the flow algorithm can be generalized by flowing the cells of a tessellation of the original integration domain. In this paper, we start our calculations with the tessellation of a rectilinear lattice. For an two-dimensional illustration see Fig. 11.10.

Algorithm 2 The flow of the contour of multi-dimensional oscillatory integrals.

Require: Represent a subset of the original integration domain X with a regular tessellation consisting of cells V_i spanned by the points $\mathbf{p}_{i,1}, \mathbf{p}_{i,2}, \dots$

while the variance of the imaginary part H on the points $p_{i,j}$ exceeds threshold T_1

do

flow the points: $\mathbf{p}_{i,j} \mapsto \mathbf{p}_{i,j} - \nabla h(\mathbf{p}_{i,j}) \Delta t$

if the h -function evaluated in the point p_i is smaller than the threshold T_2 **then**

remove the corresponding cells

end if

if the volume of a cell V_i exceeds the threshold T_3 **then**

subdivide the cell into smaller cells

end if

end while

There are various possible implementations of this algorithm. However, it follows from Cauchy's theorem that the integral is insensitive to the details of the tessellation employed. For all reasonable tessellations, the algorithm terminates in a polynomial number of steps as it scales roughly linearly with the number of simplices. Remarkably, this cost scaling is no worse than that required by the geometric optics approximation.

10.4.3 Integrating along the thimbles

Given a Lefschetz thimble \mathcal{J} for a range of α and μ , obtained with either one of the above-described methods, we perform the resulting integral along the thimble with the trapezium rule. Given a thimble \mathcal{J} represented as a set of line-segments $l_i = (p_{i,1}, p_{i,2})$, the integral is

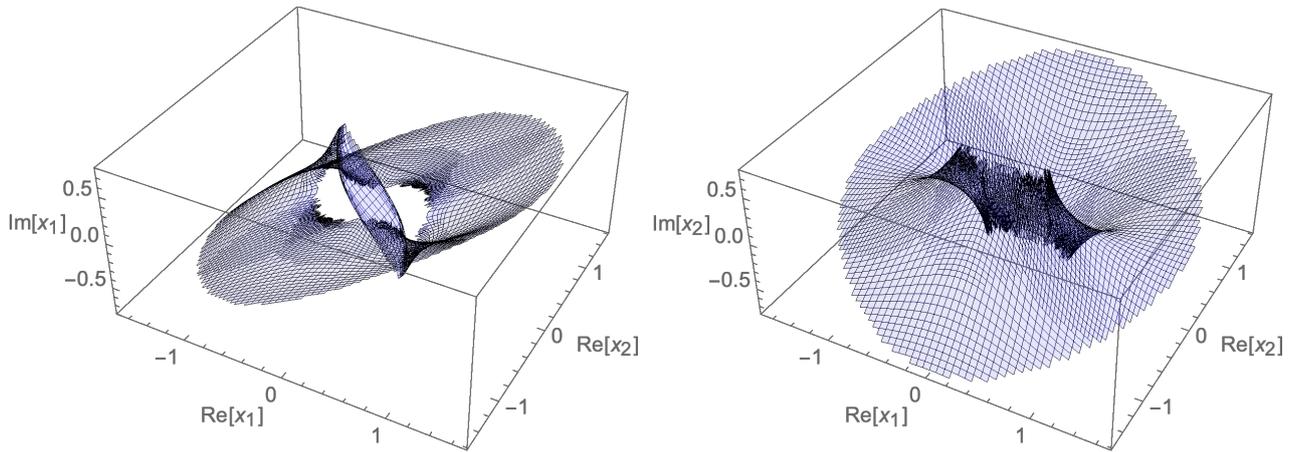


Figure 10.10: Two projections of the numerically obtained two-dimensional thimble \mathcal{J} in \mathbb{C}^2 for a two-dimensional oscillatory integral.

approximated by

$$\Psi(\mu; \nu) \approx \sum_i \frac{e^{i\phi(p_{i,1}; \mu)\nu} + e^{i\phi(p_{i,2}; \mu)\nu}}{2} (p_{i,2} - p_{i,1}) \quad (10.38)$$

summed over the line segments. For multi-dimensional oscillatory integrals, we evaluate the integral on a linear approximation of the integrand on the tessellation of the thimble. Naively, one might expect to have to compute the Lefschetz contour \mathcal{J} for every μ for which one wishes to perform the integral. However, since the thimble is a smooth function of μ , it suffices to compute the thimble for a range of μ . When integrating, we instead evaluate the integral on the thimble corresponding to the closest μ for which we have evaluated the thimble. Finally, it should be noted that for increasing ν , the support of the integral is increasingly concentrated around the relevant saddle points. As a consequence we can, for large ν , restrict the integral to the line segments close to the saddle points. It follows from this that the numerical evaluation of the integral along the thimble becomes more and more efficient as the frequency is increased. This is in sharp contrast with conventional integration techniques which need to trace many oscillations of the integrand along the real line.

See Fig. 11.11 for the normalized intensity profiles of the lens evaluated along the thimble for frequencies $\nu = 50, 100, 500$. We observe the following properties of the normalized intensity profiles:

- In the regime $\alpha < 1$, the lens leads to a single-image region. The normalized intensity profile does not oscillate and is moreover independent of the frequency ν . See the lower panels of Fig. 11.11.
- For $\alpha = 1$, the lens forms a cusp caustic. The caustic corresponds to the peak at $\mu_c = 0$. For increasing frequency, ν , the peak is enhanced and becomes increasingly narrow. In the eikonal limit $\nu \rightarrow \infty$, the normalized intensity diverges as $\nu^{1/2}$ at the caustic μ_c (see the scaling relations in Table 2). See the middle panels of Fig. 11.11.
- In the regime $\alpha > 1$, the lens forms a triple-image region which is bounded by two-fold caustics. We see that the triple-image region $(-\mu_c, \mu_c)$, with $\mu_c = 0.206751 \dots$ for $\alpha = 2$, consists of an interference pattern bounded by two peaks at $\mu = \pm\mu_c$. The interference pattern in the triple-image region is the result of the three real saddle points. The oscillations in the single-image region result from the interplay between the relevant real and the complex saddle point. For increasing ν , the fringes of the interference pattern

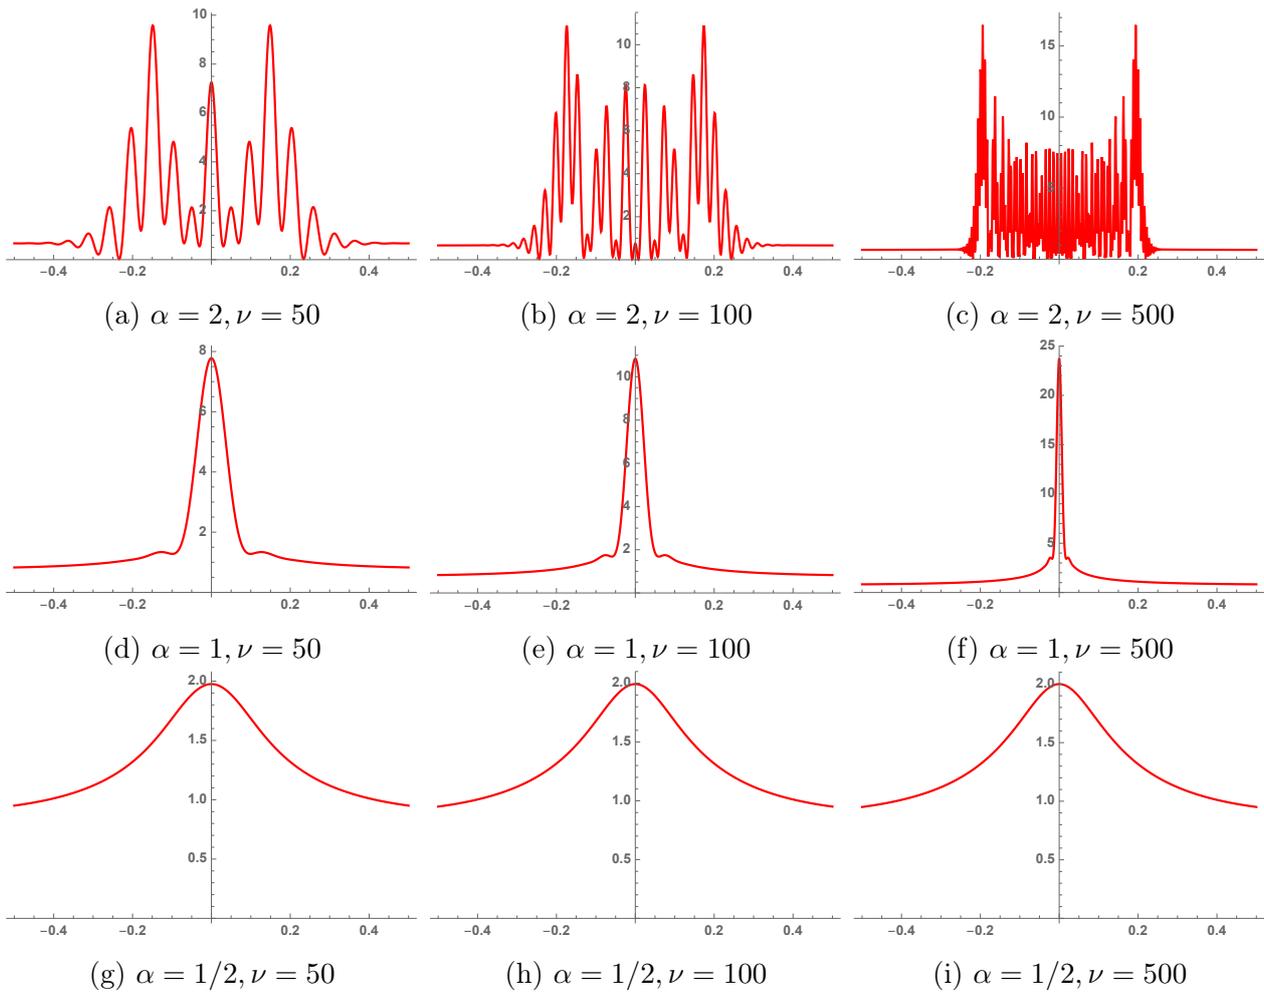


Figure 10.11: The normalized intensity $I(\mu; \nu) = |\Psi(\mu; \nu)|^2$ for $\alpha = 1/2, 1, 2$ as a function of μ for $\nu = 50, 100, 500$.

shrink and spikes corresponding to the fold get sharper and are increasingly enhanced. For the relevant scalings see Table 2. See the upper panels of Fig. 11.11.

Note that the normalized intensity in the cusp exceeds the normalized intensity in the fold caustic. This related to the co-dimension of the caustic as described in Section 11.3.2. Moreover remark that the cusp caustic only exists at a single α for the one-dimensional lens, while the fold caustic appears for a range of α . Table 2 shows the frequency dependence of the pattern. Furthermore, note that the normalized intensity profiles at frequency $\nu = 500$, for $\alpha = 1/2, 1$, and 2, are close to the normalized intensity maps predicted by geometric optics (see Fig. 11.6).

In the context of astronomical radio sources, the signal is dramatically enhanced when the relative position of the observer and the source move through the fold or the cusp caustic of the lens. One would in this context interpret the μ axis as the line traced by the source on the sky, i.e. $\mu = vt + \mu_0$ with μ_0 the initial position, v the speed of the source in parameter space and t the time. This amplification of the signal may be relevant as an selection effect for the recently observed Fast Radio Bursts. Note that if the observed FRBs are indeed the result of caustics in plasma lenses, we expect the peaks to evolve in a characteristic way and satisfy specific scaling relations in frequency space. See Section 11.7 for a more detailed discussion.

10.5 The elementary catastrophes

The unfoldings of the seven *elementary singularities* (see Table 1), form a local description of lenses near the caustics. We here study the Picard-Lefschetz analysis of the elementary catastrophes appearing in two-dimensional lenses and evaluate the corresponding normalized intensity maps using the flow algorithm described above. This analysis is complementary to the asymptotic analysis described in chapter 36 of [53].

10.5.1 The fold A_2

The fold singularity is the simplest degenerate critical point and can be viewed as the superposition of two non-degenerate saddle points. The Picard-Lefschetz analysis of the unfolding of the fold singularity is illustrated in Fig. 11.12. For negative μ , there are two relevant real saddle points (see Fig. 11.12a). As μ approaches the caustic at $\mu_c = 0$, the two saddle points merge and form the fold singularity (see Fig. 11.12b). Note that the fold saddle point emanates three steepest ascent and three descent curves. The thimble is non-differentiable at the degenerate saddle point. When μ is increased passed the caustic μ_c , the two saddle points move off the real axis and into the complex plane (see Fig. 11.12c). In this regime only one of them remains relevant.

The Fresnel-Kirchhoff integral for the fold singularity can be related to the Airy function

$$\Psi(\mu; \nu) = \sqrt{\frac{\nu}{\pi}} \int_{-\infty}^{\infty} e^{i\left(\frac{x^3}{3} + \mu x\right)\nu} dx = 2\sqrt{\pi}\nu^{1/6} \text{Ai}[\nu^{2/3}\mu]. \quad (10.39)$$

Note the appearance of the singularity and fringe indices $1/6$ and $2/3$ as listed in Table 2. It straightforward to derive the scaling of the amplitude and the fringes, with the change of coordinates $z = \nu^{1/3}x$. The other scaling relations are derived analogously. The Airy function is a good illustration of the interference phenomenon present in multi-image regions (seen in Fig. 11.13). The range $\mu < 0$, for which the two relevant saddle points reside on the real line corresponds to a double-image region, where two saddle points lead to an interference pattern. The range $\mu > 0$ corresponds to a zero-image region in which the amplitude asymptotes to zero as $\nu \rightarrow \infty$.

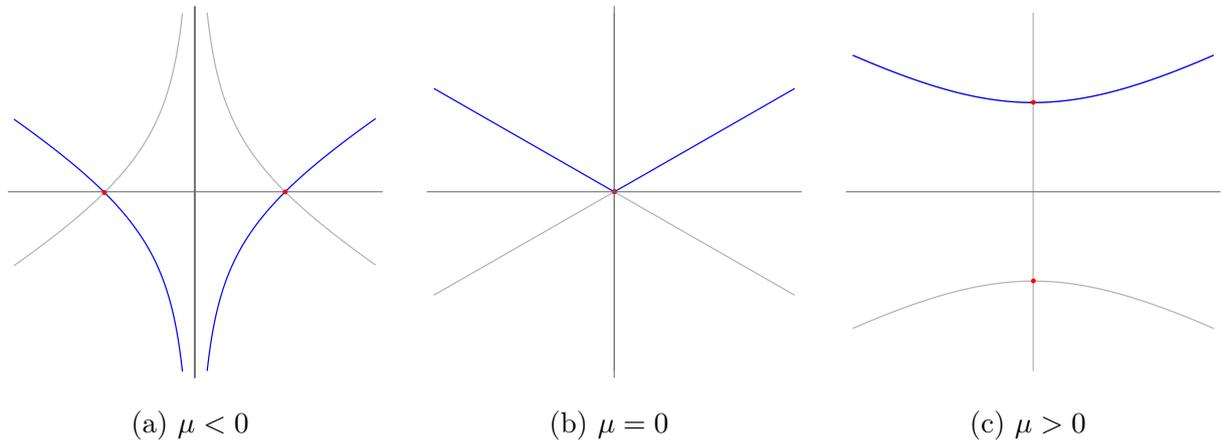


Figure 10.12: The saddle points and the Lefschetz thimbles in the complex plane \mathbb{C} of the unfolding of the fold singularity A_2 .

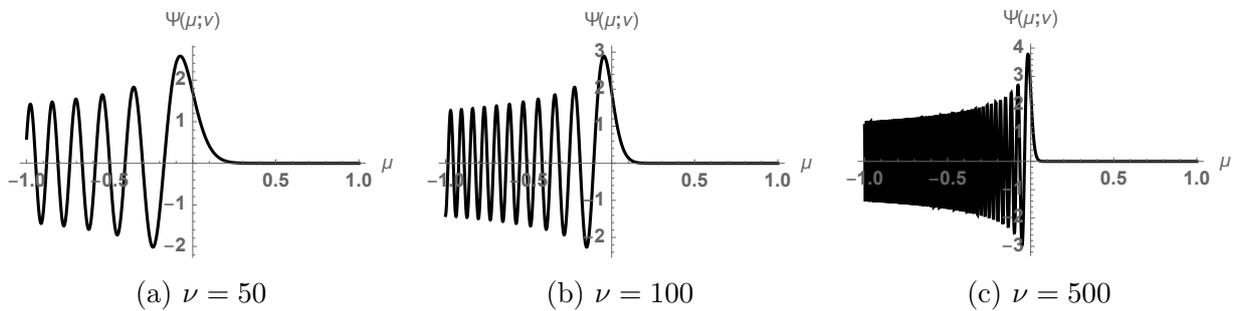


Figure 10.13: The integral $\Psi(\mu, \nu)$ for the fold singularity as a function of μ for $\nu = 50, 100, 500$.

Asymptotics

Using the Picard-Lefschetz diagrams (Fig. 11.12) we can derive asymptotics for the integral $\Psi(\mu; \nu)$. For $\mu > 0$ the Picard-Lefschetz analysis consists of a single relevant saddle point located at $i\sqrt{\mu}$. The exponent can be approximated around the saddle point by

$$\phi(x; \mu) = i\frac{2}{3}\mu^{3/2} + i\sqrt{\mu}(x - i\sqrt{\mu})^2 + \mathcal{O}((x - i\sqrt{\mu})^3). \quad (10.40)$$

The saddle point approximation for this point gives an exponential falloff

$$\Psi(\mu; \nu) \approx \frac{e^{-\frac{2}{3}\mu^{3/2}\nu}}{\sqrt{2}\mu^{1/4}}. \quad (10.41)$$

This matches the the behaviour in Fig. 11.13.

For $\mu < 0$, the Picard-Lefschetz analysis consists of two real relevant saddle points located at $x = \pm\sqrt{|\mu|}$. A saddle point approximation around these points gives us the oscillatory behaviour

$$\Psi(\mu; \nu) \approx \frac{e^{-\frac{2i}{3}\mu^{3/2}\nu} + ie^{\frac{2i}{3}\mu^{3/2}\nu}}{\mu^{1/4}}, \quad (10.42)$$

seen in Fig. 11.13. Observe that wave function becomes increasingly oscillatory and falls off as a power law $\Psi(\mu; \nu) \propto \frac{1}{|\mu|^{1/4}}$ in the geometric limit $\nu \rightarrow \infty$.

10.5.2 The cusp A_3

The cusp singularity consists of the superposition of three non-degenerate saddle points. The singularity is of co-dimension $K = 2$ and has two unfolding parameters μ_1 and μ_2 , *i.e.*,

$$\Psi(\boldsymbol{\mu}; \nu) = \sqrt{\frac{\nu}{\pi}} \int_{-\infty}^{\infty} e^{i\left(\frac{x^4}{4} + \mu_2 \frac{x^2}{2} + \mu_1 x\right)\nu} dx. \quad (10.43)$$

See Fig. 11.14 for an illustration of unfolding of the cusp caustic and the Picard-Lefschetz analysis in the (μ_1, μ_2) -plane.

The the exponent $\phi(x; \boldsymbol{\mu})$ has three saddle points \bar{x}_i , given by the roots of the cubic equation

$$x^3 + \mu_2 x + \mu_1 = 0. \quad (10.44)$$

Depending on μ , either one or three of the saddle points are real-valued. The complex-valued saddle points always come in conjugate pairs since $\phi(x; \boldsymbol{\mu})$ is a real-valued function, *i.e.* real for real x .

Geometric optics applied to this integral shows that the cusp caustic at $(\mu_1, \mu_2) = (0, 0)$ emanates two fold-lines $A_2 \subset M$, given by cubic root

$$\mu_2 = -\frac{3}{2^{2/3}} |\mu_1|^{2/3}. \quad (10.45)$$

The fold lines are non-differentiable at the cusp singularity $(\mu_1, \mu_2) = (0, 0)$.

In the triple-image region enclosed by the two fold-lines, the thimble passes through three real-valued saddle points. When approaching one of the fold lines, we see that two of the real saddle points merge and move in the complex plane. Only one of the two complex saddle points remains relevant to the integral. This is analogous to the behavior observed in the analysis of the fold caustic. At the cusp saddle point at $(\mu_1, \mu_2) = (0, 0)$ all three saddle points merge at the origin. Finally, note that the single-image region consists of three subregions, for which the Picard-Lefschetz analysis either consists of one or two relevant saddle points. These subregions are separated by two Stokes lines (red dashed lines in Fig. 11.14). Along these lines, the Lefschetz thimbles flip while the saddle points remain separated. The Stokes lines can be found by equating the imaginary parts of the exponents evaluated at the saddle points, *i.e.*,

$$\text{Im}[i\phi(\bar{x}_i; \boldsymbol{\mu})\nu] = \text{Im}[i\phi(\bar{x}_j; \boldsymbol{\mu})\nu] \quad (10.46)$$

for $i \neq j$. For the unfolding of the cusp, we see that the Stokes lines are described by

$$\mu_2 = 3 \sqrt[3]{\frac{3\sqrt{3} - 5}{2}} |\mu_1|^{2/3}, \quad (10.47)$$

for $\mu_1 < 0$ and μ_2 . Note that the amplitude across a Stokes line is smooth, even though the saddle point structure changes abruptly. The Stokes lines can be interpreted as the points for which the saddle point approximation of the integral fails.

Numerics

Given the Lefschetz thimble, we can numerically evaluate the amplitude (see Fig. 11.15). In the eikonal limit $\nu \rightarrow \infty$ we observe the emergence of a fold-line (A_2) with a sharp exponential falloff in most of the single-image region and the power-law falloff in the triple-image region. We also see the emergence of a cusp caustic at the origin with a power-law falloff along the line $\{\mu_1 = 0\}$.

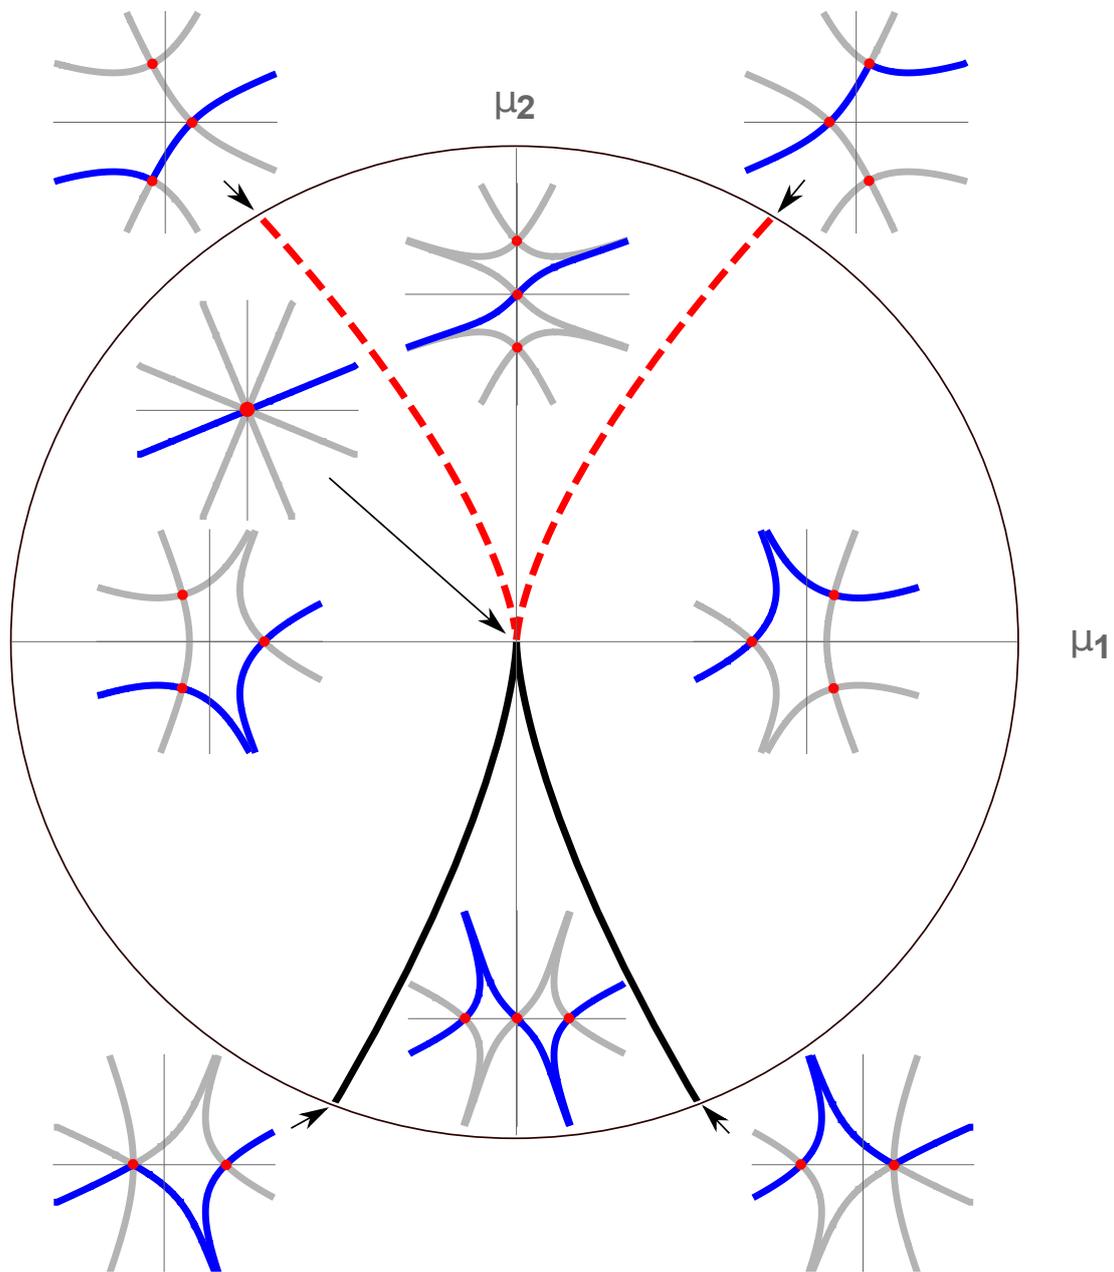


Figure 10.14: The cusp singularity in the unfolding μ -plane. The black curve is the fold line separating the single-image region (upper region) from the triple-image region (lower region). The dashed red line is the Stokes line. We observe that the Stokes transition corresponds to one complex saddle point becoming (ir)relevant, and that the caustics correspond to the superposition of real non-degenerate saddle points. The upper region is a single-image region since there is only one real relevant saddle point. The lower region is a triple-image region since there are three real relevant saddle points.

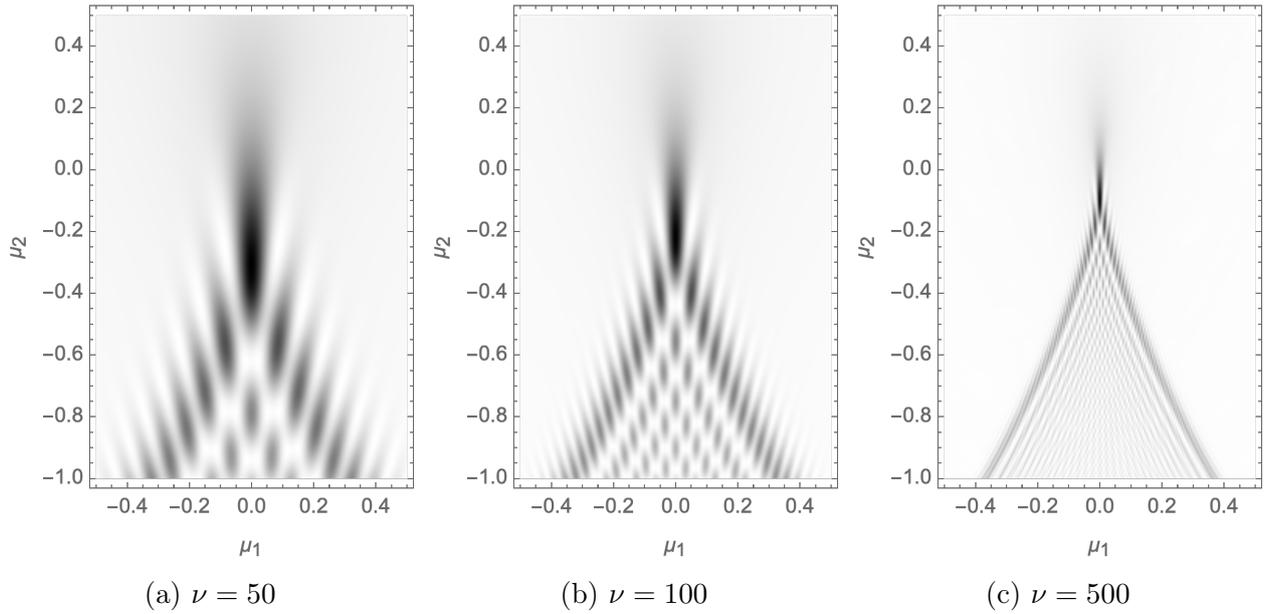


Figure 10.15: The normalized intensity $I(\boldsymbol{\mu}; \nu)$ for the cusp caustic for $\nu = 50, 100$ and 500 .

Asymptotics

The Picard-Lefschetz diagrams (Fig. 11.14) allow us to derive limiting behaviour for $\Psi(\boldsymbol{\mu}; \nu)$. From the functional form of $\phi(x; \boldsymbol{\mu})$ along the line $\{\mu_1 = 0\}$, *i.e.*,

$$\phi(x; \boldsymbol{\mu}) = \frac{\mu_2 x^2}{2} + \frac{x^4}{4}, \quad (10.48)$$

we observe that one of the relevant saddle points is located at the origin $x = 0$. For $\mu_2 > 0$ this is the only relevant saddle point, whereas for $\mu_2 < 0$ it is one of three real relevant saddle points.

In the case $\mu_2 > 0$, we find that the single saddle leads to a power-law

$$\Psi(x; \boldsymbol{\mu}) \approx \sqrt{\frac{\nu}{\pi}} \int_{-\infty}^{\infty} e^{i\frac{\mu_2 \nu}{2} x^2} dx = \sqrt{2} (-i\mu_2)^{-1/2} \quad (10.49)$$

which in the normalized intensity corresponds to the falloff

$$I(x; \boldsymbol{\mu}) = |\Psi(x; \boldsymbol{\mu})|^2 \approx \frac{2}{\mu_2}, \quad (10.50)$$

independent of the frequency. This feature is absent in the unfolding of the fold caustic.

For $\mu_2 < 0$, the Picard-Lefschetz analysis consists of three real relevant saddle points located at $\pm\sqrt{-\mu_2}$ and 0 . The exponent at the saddle point $\pm\sqrt{-\mu_2}$ can be approximated by

$$\phi(x; \boldsymbol{\mu}) = -\frac{\mu_2^2}{4} - \mu_2(x \pm \sqrt{-\mu_2})^2 + \mathcal{O}((x \pm \sqrt{-\mu_2})^3). \quad (10.51)$$

In the saddle point approximation,

$$\Psi(\boldsymbol{\mu}; \nu) \approx \sqrt{2} \frac{-(-1)^{3/4} + (1+i)e^{-\frac{i}{4}\mu_2^2\nu}}{\sqrt{-\mu_2}}. \quad (10.52)$$

The normalized intensity $I(\boldsymbol{\mu}; \nu)$ thus oscillates in μ_2 with increasing frequency with power-law suppression

$$I(\boldsymbol{\mu}; \nu) \propto \frac{2}{\mu_2}. \quad (10.53)$$

Along the line $\{\mu_2 = 0\}$, the Picard-Lefschetz analysis consists of two relevant saddle points, one real and one complex. The real relevant saddle point is located at $\bar{x} = -\mu_1^{1/3}$, giving the oscillatory behaviour

$$\Psi(\boldsymbol{\mu}; \nu) \approx \sqrt{2} \frac{(-1)^{1/4} e^{-i\mu_1^{4/3}\nu}}{-\sqrt{3}\mu_1^{1/3}}, \quad (10.54)$$

so that the normalized intensity again falls off as a power-law

$$I(\boldsymbol{\mu}; \nu) = |\Psi(\boldsymbol{\mu}; \nu)|^2 \approx 2\mu_1^{-2/3}. \quad (10.55)$$

10.5.3 The swallowtail A_4

The swallowtail singularity is more complicated, as it consists of the superposition of four non-degenerate saddle points. The singularity is of co-dimension $K = 3$ and has three unfolding parameters μ_1, μ_2 and μ_3 , *i.e.*,

$$\Psi(\boldsymbol{\mu}; \nu) = \sqrt{\frac{\nu}{\pi}} \int_{-\infty}^{\infty} e^{i\left(\frac{x^5}{5} + \mu_3 \frac{x^3}{3} + \mu_2 \frac{x^2}{2} + \mu_1 x\right)\nu} dx. \quad (10.56)$$

See figures 11.17 and 11.18 for an illustration of unfolding of the swallowtail caustic and the Picard-Lefschetz analysis in the (μ_1, μ_2, μ_3) -space.

The analytic continuation of the exponent $i\phi(x; \boldsymbol{\mu})\nu$ has four saddle points \bar{x}_i , given by the roots of the quartic equation

$$x^4 + \mu_3 x^2 + \mu_2 x + \mu_1 = 0. \quad (10.57)$$

Depending on $\boldsymbol{\mu}$, either zero, two or four of the saddle points are real-valued. The complex-valued saddle points always come in conjugate pairs since $\phi(x; \boldsymbol{\mu})$ is a real-valued function.

Geometric optics applied to this integral shows that the swallowtail caustic at $\boldsymbol{\mu} = (0, 0, 0)$ emanates a cusp-line and a fold-surface (see Fig. 11.16). The fold-surface (the yellow surface in Fig. 11.16) is given by

$$A_2 = \{(3u^4 + u^2v, -4u^3 - 2uv, v) | (u, v) \in \mathbb{R}^2\} \subset M \quad (10.58)$$

satisfying the two constraints

$$\frac{d\phi(x; \boldsymbol{\mu})}{dx} = 0, \quad \frac{d^2\phi(x; \boldsymbol{\mu})}{dx^2} = 0. \quad (10.59)$$

The cusp-line (the black curve in Fig. 11.16) lays on the fold-surface and is given by

$$A_3 = \{(-3t^4, 8t^3, -6t^2) | t \in \mathbb{R}\} \subset M \quad (10.60)$$

satisfying the three constraints

$$\frac{d\phi(x; \boldsymbol{\mu})}{dx} = 0, \quad \frac{d^2\phi(x; \boldsymbol{\mu})}{dx^2} = 0, \quad \frac{d^3\phi(x; \boldsymbol{\mu})}{dx^3} = 0. \quad (10.61)$$

Note that the caustics are symmetric in the (μ_1, μ_3) -plane and that caustics only appear for negative μ_1 . This aids our analysis, since we can consider the three-dimensional swallowtail unfolding as a one parameter family of unfoldings in the $\{\mu_3 = \text{const}\}$ planes.

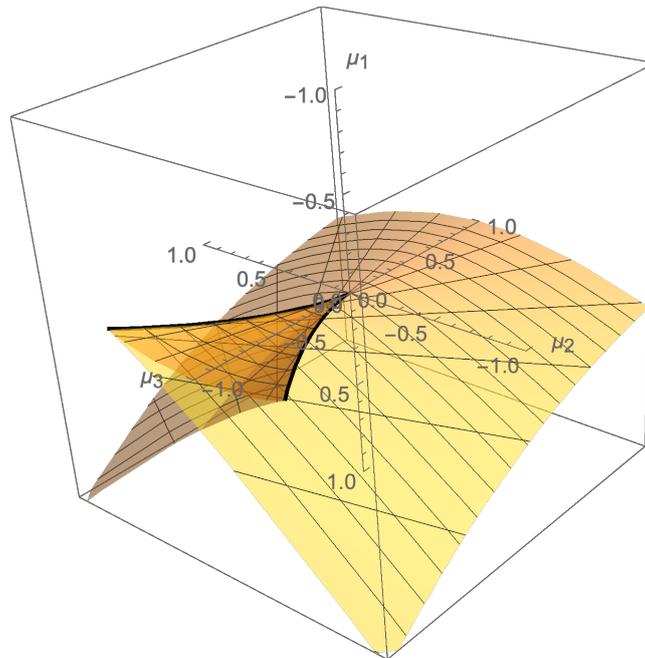


Figure 10.16: The swallowtail singularity in the unfolding space (μ_1, μ_2, μ_3) . The yellow surface is the fold surface separating the single- double- and triple-image regions. The black line is the cusp line, along which we find the cusp saddle points.

In figures 11.17 and 11.18 we plot three slices of the fold-surface and cusp-line for $\mu_3 = -1, 0$ and $+1$. For $\mu_3 = -1$ we obtain the characteristic swallowtail shape in the fold-surface with the cusp-line intersecting at the tips, which gives the singularity its name. For $\mu_3 = 0$ we see the actual swallow caustic. The slice $\mu_3 = +1$, is simpler as it does not contain intersections with the cusp-line and only consists of the fold-surface separating two regions.

Given the caustics of geometric optics, we can evaluate the Lefschetz thimble. It again suffices to study the three cases $\mu_3 < 0, \mu_3 = 0$ and $\mu_3 > 0$:

- We start by analysing the saddle points in the $\mu_3 = -1$ plane (Fig. 11.17). The Picard-Lefschetz analysis for the enclosed region in the middle of the circle consists of four relevant real saddle points. This is a quadruple-image region. Note that multiple-image regions for localized lenses always consist of an odd number of images. In such lenses, the swallowtail will in practice always appear near another caustic such as a fold.

Starting from the quadruple-image region and moving through the fold-line on the left, we observe that the two central saddle points merge to form a fold saddle point. The two saddle points subsequently move in the complex plane, one remaining relevant. Since this region corresponds to two real saddle points it is a double-image region.

Again, starting from the quadruple-image region and moving in the vertical direction, we observe that two of the outer saddle points merge to form a fold saddle point and subsequently move into the complex plane. The resulting Picard-Lefschetz analysis again consists of three relevant saddle points; two real and one complex. This again is a double-image region. If we, however, move from this double-image region to the double-image region on the left of the quadruple-image region, we pass through two Stokes lines, at which the complex saddle point switches from relevant to irrelevant. The Stokes lines are defined by

$$\text{Im}[i\phi(\bar{x}_i; \boldsymbol{\mu})\nu] = \text{Im}[i\phi(\bar{x}_j; \boldsymbol{\mu})\nu] \quad (10.62)$$

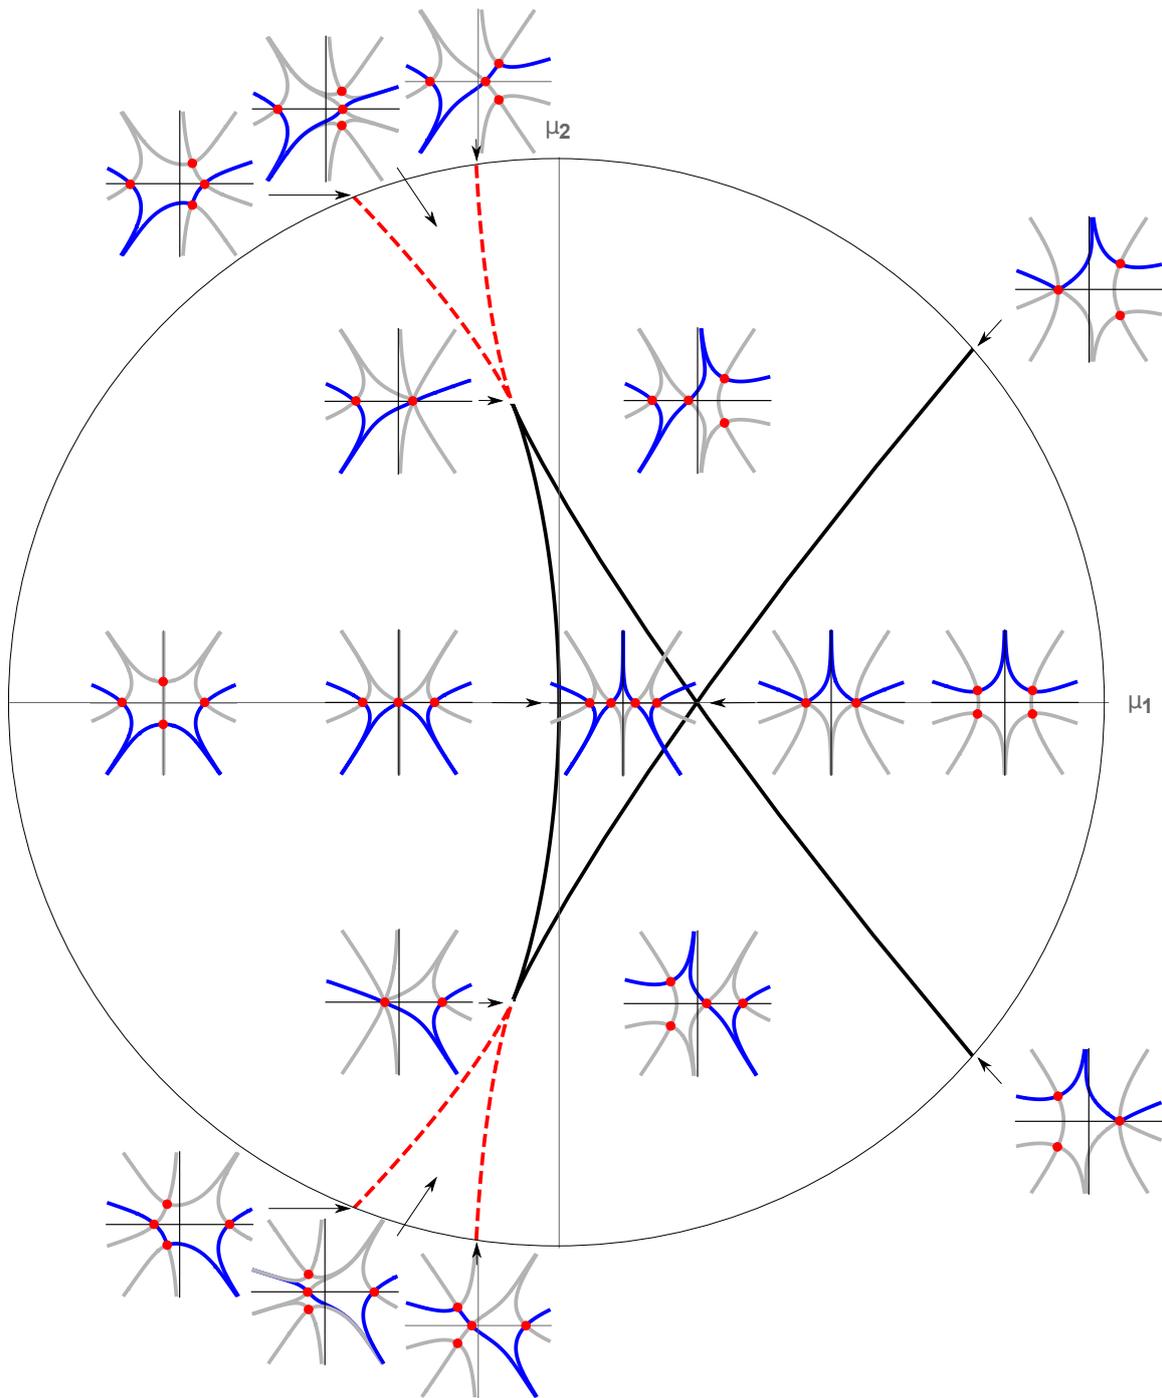


Figure 10.17: The Picard-Lefschetz analysis of the unfolding of the swallowtail (A_4) singularity at $\mu_3 = -1$. The number of real saddles gives the number of images in geometric optics.

for $i \neq j$. Note that the Stokes lines can be associated with the cusp caustic at the tips of the fold-line. Note that the three relevant saddle points merge at these tips, to form a cusp saddle point.

Finally, if we move from the quadruple-image region along the line $\mu_2 = 0$ to the right, we pass through the intersection of the fold lines. At this point, both the left and right two real saddle points merge to form a ‘double’ fold caustic. After passing this point, the four saddle points move in the complex plane. The Picard-Lefschetz analysis consists here of two relevant complex saddle points. This is a zero-image region (which will not be realized in localized lenses). If we pass from the double-image region to the zero-image region, we again observe a fold caustic in which two relevant real saddle points merge and move in the complex plane. This completes the analysis of the unfolding of the swallowtail caustic at $\mu_3 = -1$.

- For $\mu_3 = 0$, the geometry of the fold-line is simpler as the quadruple-image region has merged into the swallowtail caustic at the origin (see Fig. 11.18a). The Picard-Lefschetz analysis of this slice is largely similar to the one at $\mu_3 = -1$. The double-image region (including the Stokes lines) has been deformed but is otherwise the same. The zero-image region is also unchanged. However, the intersection of the two fold-lines is replaced by the swallowtail saddle point at the origin of in the (μ_1, μ_2) -plane. Since this saddle point is the superposition of four non-degenerate saddle points, the amplitude integral is enhanced.
- For $\mu_3 = +1$, the geometry of the caustics is depicted in Fig. 11.18b. The fold-line separates the zero-image region on the right from the double-image region on the left. Since the Picard-Lefschetz diagram in the zero-image region consists of four complex saddle points – two of them being relevant – there exist two distinct ways in which we can pass to the double-image region; either by merging the two saddle points on the left or on the right (see upper and lower diagram). The transition between these two takes place at the origin, where the four saddle points are located on the imaginary axis. The double-image region consists of three subregions. The rightmost Stokes lines at $\mu_3 = 0$ (see Fig. 11.18a) have partly moved into the zero-image region.

By patching the Picard-Lefschetz analysis at $\mu_3 = -1, 0$ and $+1$ together, we obtain a complete description of the unfolding of the swallowtail singularity in the (μ_1, μ_2, μ_3) -space. Note that the Stokes lines obtained in figures 11.17 and 11.18 are intersections of Stokes-surfaces, which together with the fold-surface partition the μ -space.

Numerics

Given the Lefschetz thimble, we can numerically compute the normalized intensity map of the lens (see Fig. 11.19). The left, central and right panels depict the normalized intensity $I(\boldsymbol{\mu}; \nu)$ for $\mu_3 = -1, 0$ and $+1$. The upper, middle and lower panels depict the different frequencies $\nu = 50, 100$ and 500 .

We observe that for $\nu = 50$, interference is a dominant feature of the geometry of the caustic. The images are blurry and the geometry of the swallowtail is not resolved (Fig. 11.19a). We do observe the power-law falloff associated with the cusp singularities, which contrasts with the exponential falloff of the fold singularities.

In the eikonal limit $\nu \rightarrow \infty$ we observe the emergence of a fold-line (A_2) with cusps (A_3). For $\nu = 500$ the swallowtail structure at $\mu_3 = -1$ is fully resolved. Note the difference in normalized intensity between the double- and quadruple-image regions. As the frequency ν is increased we observe that the enhanced flares, in the double-image regions, corresponding to the cusp caustics get thinner. However, note that they are independent of the frequency ν .

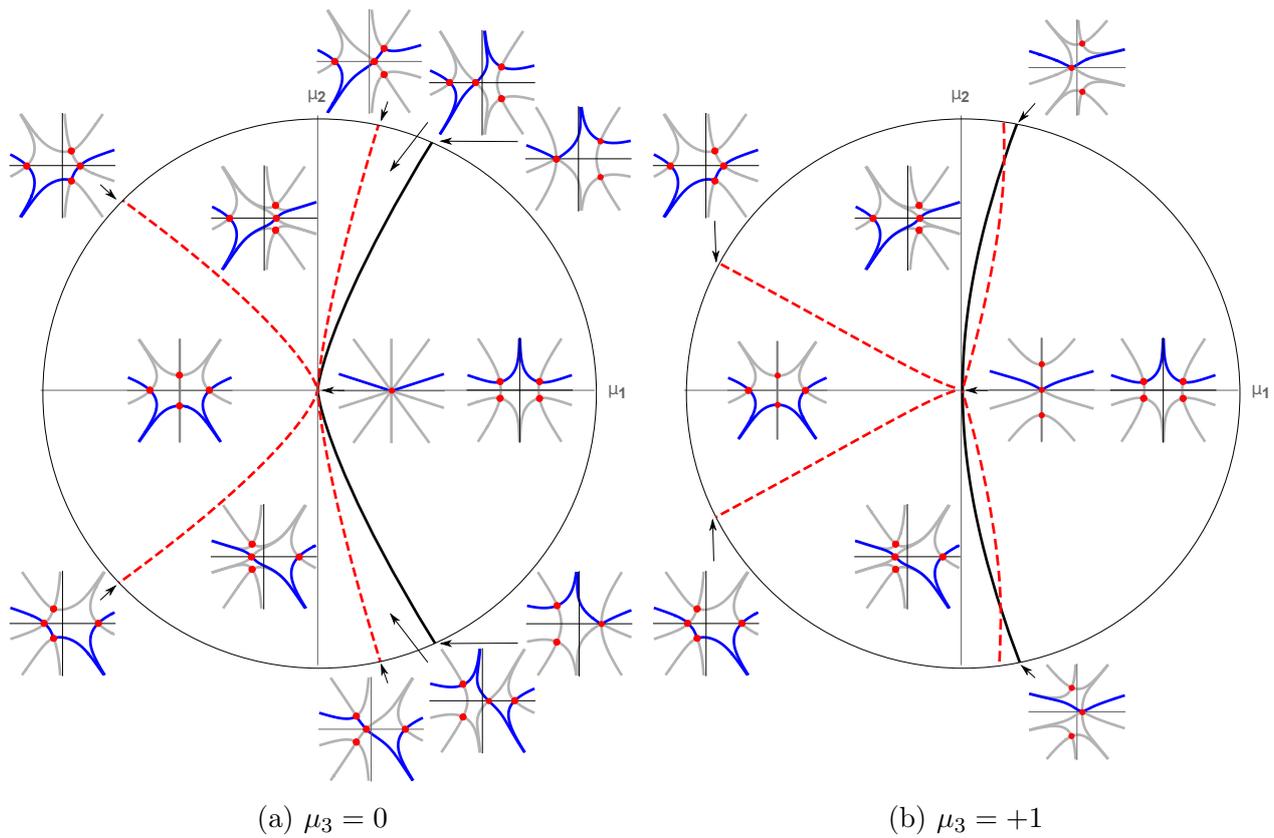


Figure 10.18: The Picard-Lefschetz analysis of the unfolding of the swallowtail (A_4) singularity at $\mu_3 = 0$ and $+1$. The number of real saddles gives the number of images in the geometric optics approximation.

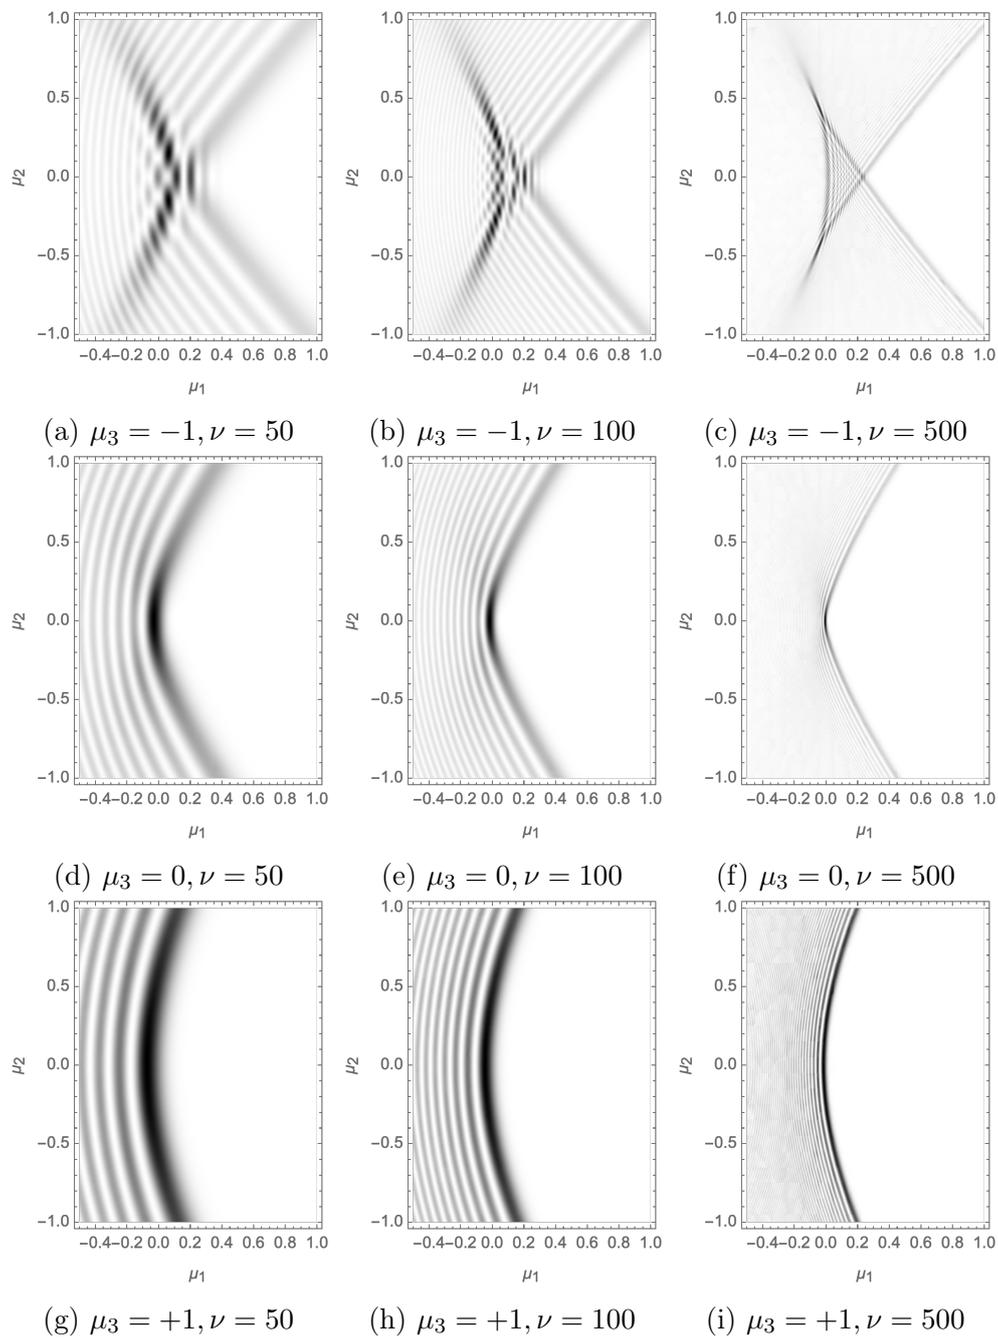


Figure 10.19: The normalized intensity, $I(\boldsymbol{\mu}; \nu)$, of the unfolding of the swallowtail caustic (A_4) sliced by the surfaces $\{\mu_3 = -1\}$, $\{\mu_3 = 0\}$, $\{\mu_3 = +1\}$ (respectively the left, central and right panels) for the frequencies $\nu = 50, 100$ and 500 (respectively the upper, the middle and lower panels).

10.5.4 The elliptic umbilic D_4^-

The caustics described above were part of the A -family. They are of co-rank 1 and can be described by a one-dimensional integral. This should be contrasted by the D family which is of co-rank 2 and can only be studied in two-dimensional integrals.

The elliptic umbilic D_4^- is a singularity with co-rank 2 and co-dimension $K = 3$. The unfolding is described in terms of the three unfolding parameters (μ_1, μ_2, μ_3) . We consider the interference pattern emerging from the integral

$$\Psi(\boldsymbol{\mu}; \nu) = \frac{\nu}{\pi} \int_{\mathbb{R}^2} e^{i(x_1^3 - 3x_1x_2^2 - \mu_3(x_1^2 + x_2^2) - \mu_2x_2 - \mu_1x_1)\nu} dx_1 dx_2. \quad (10.63)$$

The analytic continuation of the exponent $i\phi(\mathbf{x}; \boldsymbol{\mu})\nu$ has four saddle points \bar{x}_i , given by the roots of the two quadratic equations

$$3x_1^2 - 3x_2^2 - 2\mu_3x_1 - \mu_1 = 0 \quad (10.64)$$

$$-6x_1x_2 - 2\mu_3x_2 - \mu_2 = 0. \quad (10.65)$$

Depending on $\boldsymbol{\mu}$, either two or four of the saddle points are real-valued. The complex-valued saddle points always come in conjugate pairs since $\phi(\mathbf{x}; \boldsymbol{\mu})$ is real-valued for real x . Solving this set of equations for μ_1 and μ_2 we obtain the Lagrangian map as a function of μ_3 ,

$$\xi_{\mu_3}(x_1, x_2) = (3x_1^2 - 3x_2^2 - 2x_1\mu_3, -2x_2(3x_1 + \mu_3), \mu_3). \quad (10.66)$$

In the geometric limit, we form a fold-surface and three cusp lines. The fold-surface in base space $X = \mathbb{R}^2$ is given by

$$A_2^X(\mu_3) = \left\{ \left(\frac{\mu_3}{3} \cos \theta, \frac{\mu_3}{3} \sin \theta \right) \mid \theta \in [0, 2\pi) \right\} \quad (10.67)$$

which is a cylinder with radius $\frac{\mu_3}{3}$, satisfying the equation

$$|\mathcal{M}| = 0, \quad (10.68)$$

where the deformation tensor is given by

$$\mathcal{M} = \left[\frac{\partial^2 \phi(\mathbf{x}; \boldsymbol{\mu})}{\partial x_i \partial x_j} \right]_{i,j=1,2} \quad (10.69)$$

$$= \begin{pmatrix} 6x_1 - 2\mu_3 & -6x_2 \\ -6x_2 & -6x_1 - 2\mu_3 \end{pmatrix}. \quad (10.70)$$

The three cusp-lines are straight lines and lay on the fold-surface,

$$A_3^X(\mu_3) = \left\{ \left(\frac{\mu_3}{3}, 0 \right), \left(\frac{\mu_3}{3} \cos \frac{2\pi}{3}, \frac{\mu_3}{3} \sin \frac{2\pi}{3} \right), \left(\frac{\mu_3}{3} \cos \frac{4\pi}{3}, \frac{\mu_3}{3} \sin \frac{4\pi}{3} \right) \right\} \quad (10.71)$$

in the X space.

In M space, after being mapped by ξ_{μ_3} , the elliptic umbilic point is located at the origin. The fold-surface is given by

$$A_2 = \left\{ \left(\frac{\mu_3^2}{3} (\mp 2 \cos \theta + \cos(2\theta)), -\frac{2\mu_3^2}{3} (\pm 1 + \cos(\theta)) \sin(\theta), \pm \mu_3 \right) \mid \theta \in [0, 2\pi), \mu_3 \in \mathbb{R} \right\} \quad (10.72)$$

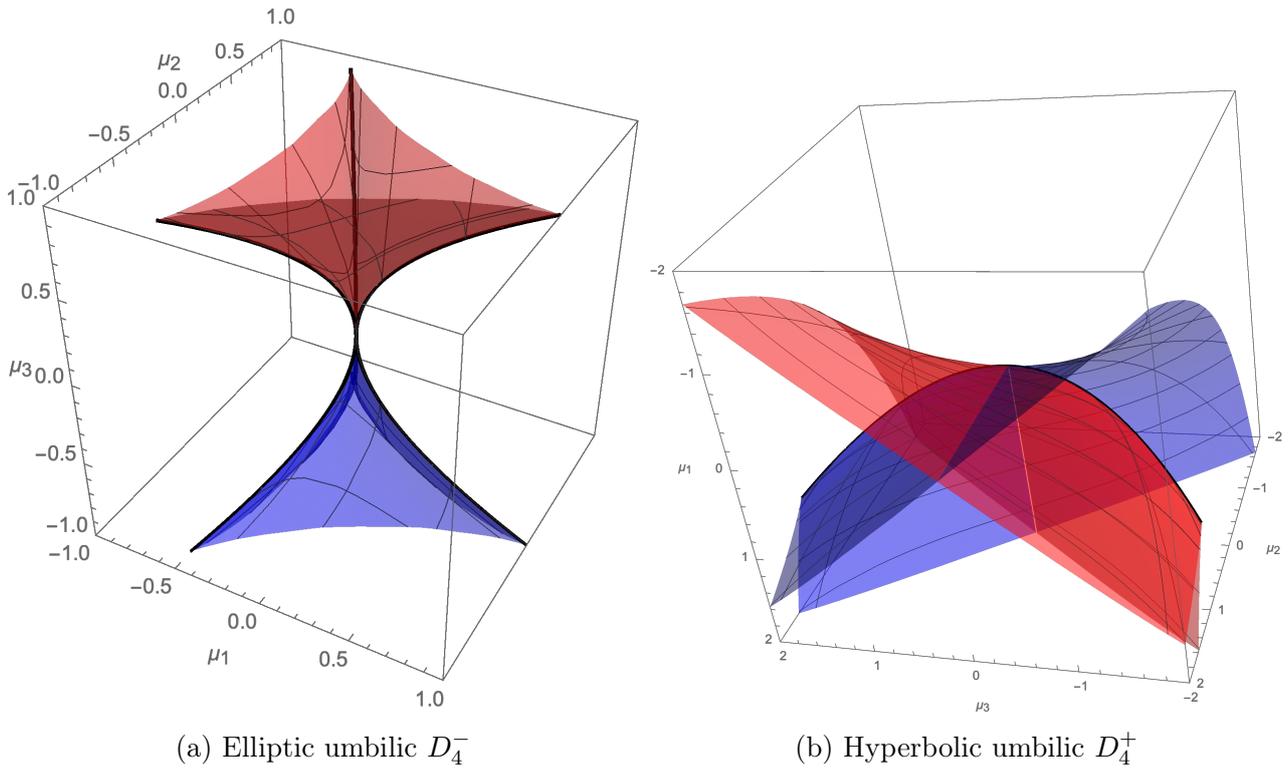


Figure 10.20: The elliptic (D_4^-) and hyperbolic umbilic (D_4^+) singularity in the unfolding space (μ_1, μ_2, μ_3) . The red and blue surfaces are the fold-surfaces corresponding to the eigenvalue field λ_1 and λ_2 separating the single- double- and triple-image regions. The black lines are the cusp-lines, along which we find the cusp saddle points.

where the two branches corresponding to \pm correspond to two disconnected pieces corresponding to the two eigenvalue fields of \mathcal{M} . The cusp lines are given by

$$A_3 = \left\{ (t^2, 0, t), (-t^2/2, \sqrt{3}t^2/2, t), (-t^2/2, -\sqrt{3}t^2/2, t) \mid t \in \mathbb{R} \right\}. \quad (10.73)$$

The fold-surface and cusp lines are illustrated in Fig. 11.20a. The red and the blue surfaces denote the fold surfaces corresponding to the eigenvalue fields λ_1 and λ_2 . The fold surface has a sharp edge at the cusp lines (in black).

Note the symmetry of the triangular singularity and point symmetry of the caustic. By performing the Picard-Lefschetz analysis for the two slices $\mu_3 = \pm 1$ and $\mu_3 = 0$ we can obtain the Picard-Lefschetz diagram of the unfolding of the singularity. See Fig. 11.21 for the Picard-Lefschetz analysis of the two slices. The small diagrams are the real parts of the four saddle points in the (x_1, x_2) -plane. The black circle is the caustic in the base space at the corresponding μ_3 .

- We first consider the case $\mu_3 \neq 0$. At the origin, the four saddle points are real-valued (see Fig. 11.21a). As a consequence, we conclude that they are all relevant. This is a quadruple-image region. One of the four saddle points is located inside the black circle. The other three are symmetrically distributed around the circle.

When crossing the fold-line, the saddle point in the circle merges with one of the outer saddle points on the circle. After passing the fold-line, the two saddle points become complex. The saddle point with the smallest real part of the exponent $i\phi(\bar{x}; \mu)$ will remain relevant whereas the other saddle point becomes irrelevant. The outside of the triangle is a double-image region. Note that the real parts of the two complex saddle points always coincides with the black circle.

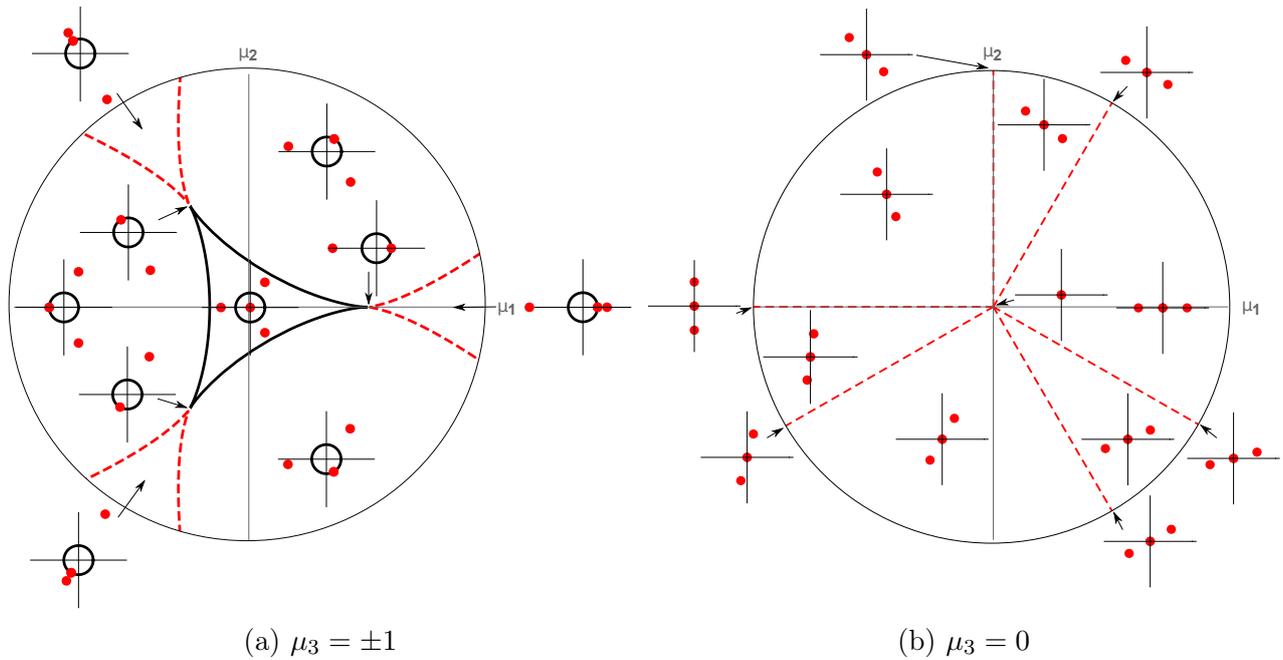


Figure 10.21: Intersection of the elliptic umbilic D_4^- by the surface $\{\mu_3 = -1\}$. The black line is the fold-line and the red line is the Stokes line. The number of real saddles gives the number of images in the geometric optics approximation.

When approaching one of the three the cusp points, three of the four saddle points merge at a single point on the circle. Note that the four saddle points are collinear in the cusps.

Finally, note that the double-image region consists of six subregions divided by six Stokes lines. In the regions on the left, the upper right and the lower right, the Lefschetz thimble passes through two real and one complex saddle point. In the regions to the right, upper left and lower right, the Picard-Lefschetz analysis consist of only two real saddle points (the ones outside the circle. The Stokes lines are again associated with the three cusps.

- In the case, $\mu_3 = 0$, the central region is replaced by the elliptic umbilic saddle point (see Fig. 11.21b). The rest of the μ_1 - μ_2 -plane is divided into six distinct regions by the six Stokes lines. The upper left, lower left, and the upper right regions consist of two relevant real saddle points. These regions correspond to the upper left, lower left, and right region in Fig. 11.21a. The three remaining regions consist of two real and one complex relevant saddle points.

These slices form a complete description of the Lefschetz thimble of the unfolding of the elliptic umbilic in the μ -space.

Numerics

Given the Lefschetz thimble, we can numerically evaluate the normalized intensity map of the lens (Fig. 11.22). The upper and lower panels depict the normalized intensity $I(\boldsymbol{\mu}; \nu)$ for $\mu_3 = \pm 1$ and 0. The left, central and right panels depict the frequencies $\nu = 50, 100$ and 500.

The normalized intensity map corresponding to the unfolding of the elliptic umbilic (D_4^-) has a triangular symmetry. As the frequency increases, the normalized intensity profile steepest and increases in amplitude. In the plane $\mu_3 = \pm 1$ we observe a fold-line in a triangular configuration with three cusp caustics at the corners. For the frequency, $\nu = 50$ the fold-line is relatively blurry. We again observe outward stripes emanating from the cusp caustics. These again follow a power-law falloff independent of the frequency. As the frequency is raised to $\nu = 100$ and

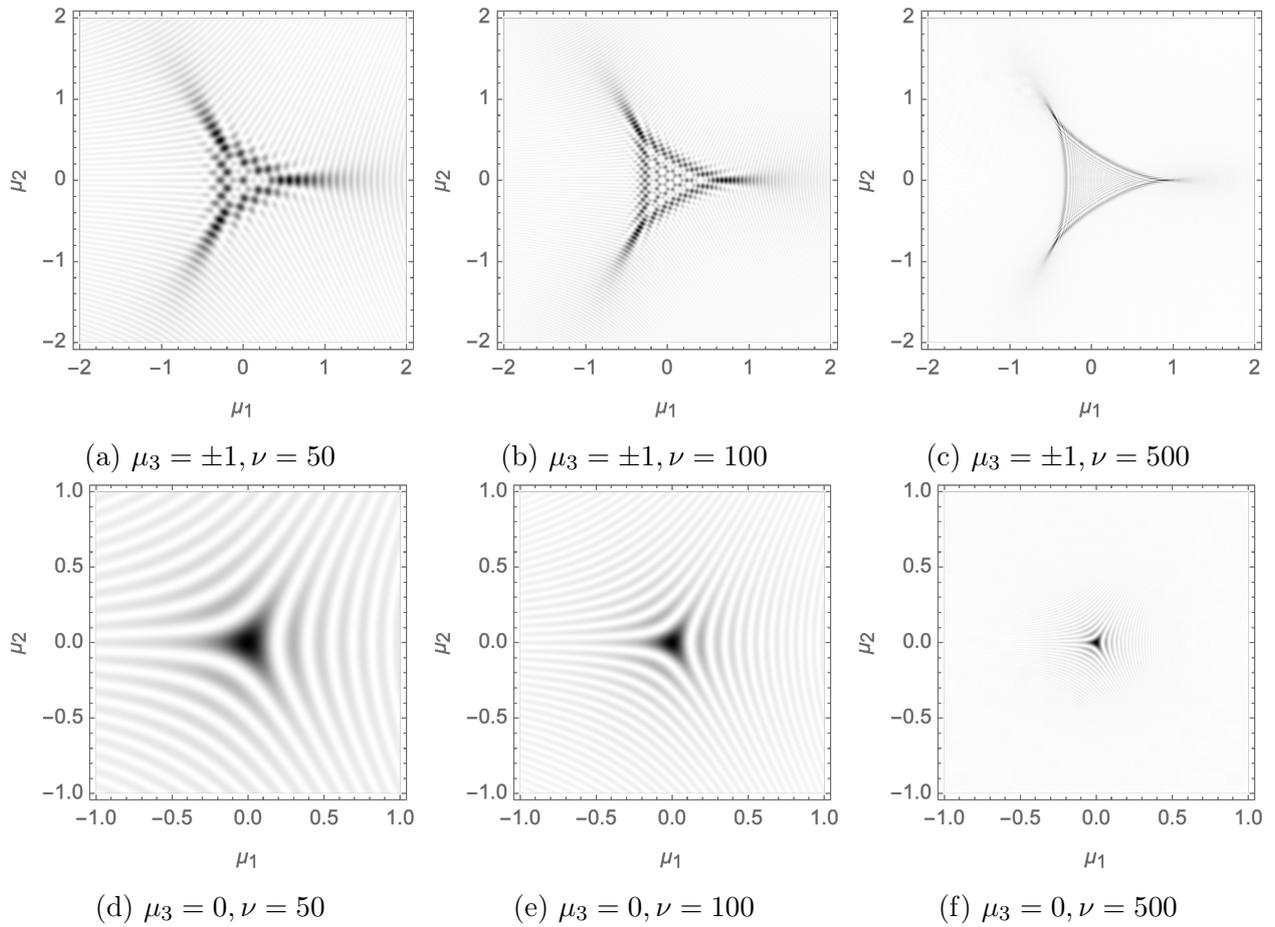


Figure 10.22: The normalized intensity, $I(\boldsymbol{\mu}; \nu)$, of the unfolding of the elliptic umbilic caustic (D_4^-) sliced by the surfaces $\{\mu_3 = \pm 1\}$, $\{\mu_3 = 0\}$ (respectively the upper and lower panels) for the frequencies $\nu = 50, 100$ and 500 (respectively the left, the centre and right panels).

$\nu = 500$ we observe that the fold lines become sharper and the fringes in the quadruple image region shrink. The normalized intensity at frequency $\nu = 500$ is very close to the normalized intensity map predicted by geometric optics.

10.5.5 The hyperbolic umbilic D_4^+

The hyperbolic umbilic D_4^+ completes the set of caustics appearing in two-dimensional lenses. It is again a singularity with co-rank 2 and co-dimension $K = 3$. The unfolding is described in terms of the three unfolding parameters (μ_1, μ_2, μ_3) . We consider the integral

$$\Psi(\boldsymbol{\mu}; \nu) = \frac{\nu}{\pi} \int_{\mathbb{R}^2} e^{i(x_1^3 + x_2^3 - \mu_3 x_1 x_2 - \mu_2 x_2 - \mu_1 x_1)} \nu dx_1 dx_2. \quad (10.74)$$

The analytic continuation of the exponent $i\phi(\mathbf{x}; \boldsymbol{\mu})\nu$ has four saddle points \bar{x}_i , given by the roots of the two quadratic equations

$$3x_1^2 - \mu_3 x_2 - \mu_1 = 0, \quad (10.75)$$

$$3x_2^2 - \mu_3 x_1 - \mu_2 = 0. \quad (10.76)$$

Depending on $\boldsymbol{\mu}$, either zero, two or four of the saddle points are real-valued. The complex-valued saddle points always come in conjugate pairs since $\phi(\mathbf{x}; \boldsymbol{\mu})$ is real-valued for real x . Solving this set of equations for μ_1 and μ_2 we obtain the Lagrangian map as a function of μ_3 ,

$$\xi_{\mu_3}(x_1, x_2) = (3x_1^2 - x_2\mu_3, 3x_2^2 - x_1\mu_3). \quad (10.77)$$

In the geometric limit, we form a fold-surface and a cusp lines. The fold-surface in X space is given by

$$A_2^X(\mu_3) = \left\{ \left(\pm \frac{\mu_3^2}{36t}, t \right) \mid t \in \mathbb{R} \right\} \quad (10.78)$$

which is a cylinder with radius $\frac{\mu_3}{3}$, satisfying the equation

$$|\mathcal{M}| = 0 \quad (10.79)$$

where the deformation tensor is given by

$$\mathcal{M} = \left[\frac{\partial^2 \phi(\mathbf{x}; \boldsymbol{\mu})}{\partial x_i \partial x_j} \right]_{i,j=1,2} \quad (10.80)$$

$$= \begin{pmatrix} 6x_1 & -\mu_3 \\ -\mu_3 & 6x_2 \end{pmatrix}. \quad (10.81)$$

The three cusp-lines are linear lines laying on the fold-surface,

$$A_3^X(\mu_3) = \{(-\mu_3/6, -\mu_3/6)\} \quad (10.82)$$

in the $X = \mathbb{R}^2$ space.

In the parameter space M , the elliptic umbilic point is located at the origin. The fold-surface is given by

$$A_2 = \{(3u^4 \pm 6uv^3, \pm 6u^3v + 3v^4, \mp 6uv) \mid u, t \in \mathbb{R}\} \quad (10.83)$$

$$A_2 = \{(3u^4 \mp 6uv^3, \mp 6u^3v + 3v^4, \mp 6uv) \mid u, t \in \mathbb{R}\} \quad (10.84)$$

where the two solutions correspond to two disconnected pieces corresponding to the two eigenvalue fields of \mathcal{M} . The cusp line in the parameter space is given by

$$A_3 = \{(t^2/4, t^2/4, t) \mid t \in \mathbb{R}\}. \quad (10.85)$$

The fold-surface and cusp-line are illustrated in Fig. 11.20b. The red and the blue surfaces denote the fold surfaces corresponding to the eigenvalue fields λ_1 and λ_2 of the deformation tensor \mathcal{M} . The fold surface has a harp edge at the cusp lines (in black).

Note the symmetry of the triangular singularity and point symmetry of the caustic. By performing the Picard-Lefschetz analysis for the two slices $\mu_3 = \pm 1$ and $\mu_3 = 0$ we can obtain an understanding of the relevant saddle points.

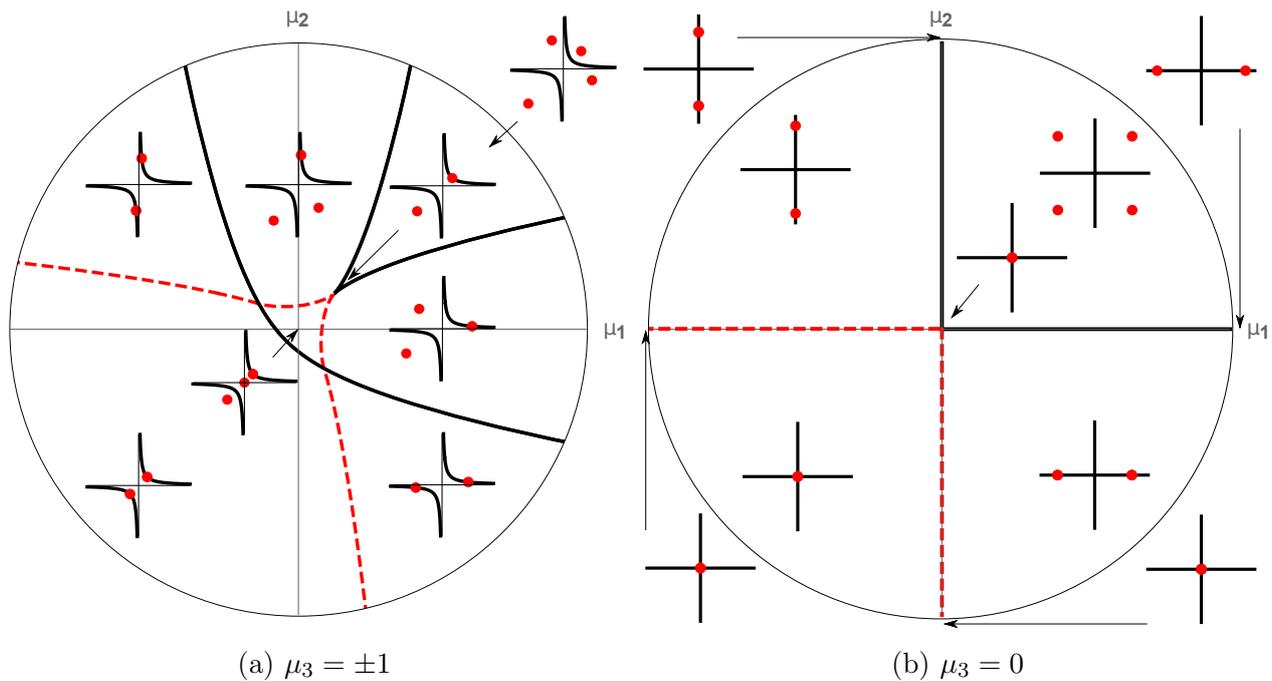


Figure 10.23: Intersection of the hyperbolic umbilic D_4^+ with the surface $\{\mu_3 = -1\}$. The black line is the fold-line and the red line is the Stokes line. The number of real saddles gives the number of images in the geometric optics approximation.

See Fig. 11.23 for the Picard-Lefschetz analysis of the two slices. The small diagrams are the real parts of the four saddle points in the (x_1, x_2) plane. The black circle is caustic in X space at $\mu_3 = \pm 1$ and 0.

- Consider the slice $\mu_3 = \pm 1$ (see Fig. 11.23a). In the upper right corner, the four saddle points are real-valued. In the corresponding Picard-Lefschetz analysis, they are all relevant. This is a quadruple-image region. When we pass the left or lower fold-line, two of the four saddle points merge at the hyperbola in X space, to form a fold singularity. Afterward, both saddle points become complex. The one with the smallest real part of the exponent $i\phi(\bar{\mathbf{x}}; \boldsymbol{\mu})\nu$ remains relevant whereas the other saddle point becomes irrelevant. Just like in the elliptic umbilic, the real part of the complex saddle points remains on the hyperboloid. This is a double-image region. Depending on whether we cross the fold line to the left or below the quadruple region, two different saddle points merge.

If we move from the quadruple-image region to the cusp, we obtain a singularity due to the merger of three saddle points. After passing through the cusp, only the two real saddle points will be relevant. The two complex saddle points are irrelevant.

From the double-image region, we can pass the second fold-line. At this fold-line, the two remaining real saddle points merge to form a fold saddle point after which they move in the complex plane. Note that the real parts of these two saddle points remain on the second branch of the hyperbolic. Since the Picard-Lefschetz analysis does not contain any real-valued saddle points after passing the second fold-line, this is a zero-image region. The zero-image region is again subdivided into three subregions. In the upper left and lower right regions, the Picard-Lefschetz analysis consists of two relevant complex saddle points. In the lower left region, the Picard-Lefschetz analysis consists of one relevant complex saddle point.

- In the case, $\mu_3 = 0$ the analysis is similar to the one obtained for $\mu_3 = \pm 1$, since the regions

are trivially deformed (see Fig. 11.23b). In the upper right region, again four saddle points are real. All of them are relevant. This is still a quadruple-image region. The fold line along the positive μ_1 and μ_2 axis is double fold lines, as the two fold lines at $\mu_3 = \pm 1$ have merged. The left and lower right regions are zero-image regions. In the upper left and lower right regions, the Picard-Lefschetz analysis consists of two relevant complex saddle points. In the lower left region, the Picard-Lefschetz analysis again consists of one relevant complex saddle point. This concludes the Picard-Lefschetz analysis.

Numerics

Given the Lefschetz thimble, we can numerically compute the normalized intensity map (see Fig. 11.24). The upper and lower panels depict the normalized intensity $I(\boldsymbol{\mu}; \nu)$ for $\mu_3 = \pm 1$ and 0. The left, central and right panels depict the different frequencies $\nu = 50, 100$ and 500.

For both unfoldings at $\mu_3 = \pm 1$ and $\mu_3 = 0$, the normalized intensity map closely follows the caustics structure represented in Fig. 11.23. In the zero-image region, the normalized intensity vanishes. In the double-image regions, for $\mu_3 = \pm 1$, the normalized intensity oscillates forming lines of equal normalized intensity as should be expected from the presence of the left fold line. In the quadruple-image regions, the normalized intensity oscillates in two directions, for $\mu_3 = \pm 1$ forming the structure we observed for the cusp caustic, and for $\mu_3 = 0$ forming an interference pattern with rectangular symmetry.

In the eikonal limit $\nu \rightarrow \infty$, the normalized intensity becomes sharper and the caustics become more pronounced. It should, in particular, be noted that the normalized intensity at the hyperbolic umbilic (in the origin in the plots corresponding to $\mu_3 = 0$), the normalized intensity rises rapidly with ν . This is in correspondence with the scaling relations we found above.

10.6 Two dimensional localized lenses

The seven elementary singularities form a dictionary of the local behavior of the lens integral

$$\Psi(\boldsymbol{\mu}; \nu) = \left(\frac{\nu}{\pi}\right)^{N/2} \int_{\mathbb{R}^N} e^{i\phi(\boldsymbol{x}; \boldsymbol{\mu})\nu} d\boldsymbol{x}, \quad (10.86)$$

$$\phi(\boldsymbol{x}; \boldsymbol{\mu}) = (\boldsymbol{x} - \boldsymbol{\mu})^2 + \varphi(\boldsymbol{x}), \quad (10.87)$$

near caustics. Their corresponding normalized intensity map completely describes the local properties of lensed images. However, the global structure of the caustic is in general different. Since the normal forms of the elementary singularities are polynomials, the corresponding phase φ has support throughout the base space $X = \mathbb{R}^N$. The catastrophes with an even co-dimension K , lead to an image with an even number of images. In contrast, localized lenses lead to n -image regions with n an odd integer. We now turn to the study of interference patterns appearing in localized lenses near caustics. We evaluate three two-dimensional lenses, which simulate the behavior of a localized lens and include the five elementary catastrophes appearing in two-dimensional lenses. In the process, we also demonstrate the accuracy of the integration scheme along the Lefschetz thimble.

10.6.1 A generic peak

In general, lensing effects are strongest near the extrema of the variation of the phase φ . It is for this reason natural to study the effect of an asymmetric peak in the phase variation φ , with

$$\varphi(\boldsymbol{x}) = \frac{\alpha}{1 + x_1^2 + 2x_2^2}, \quad (10.88)$$

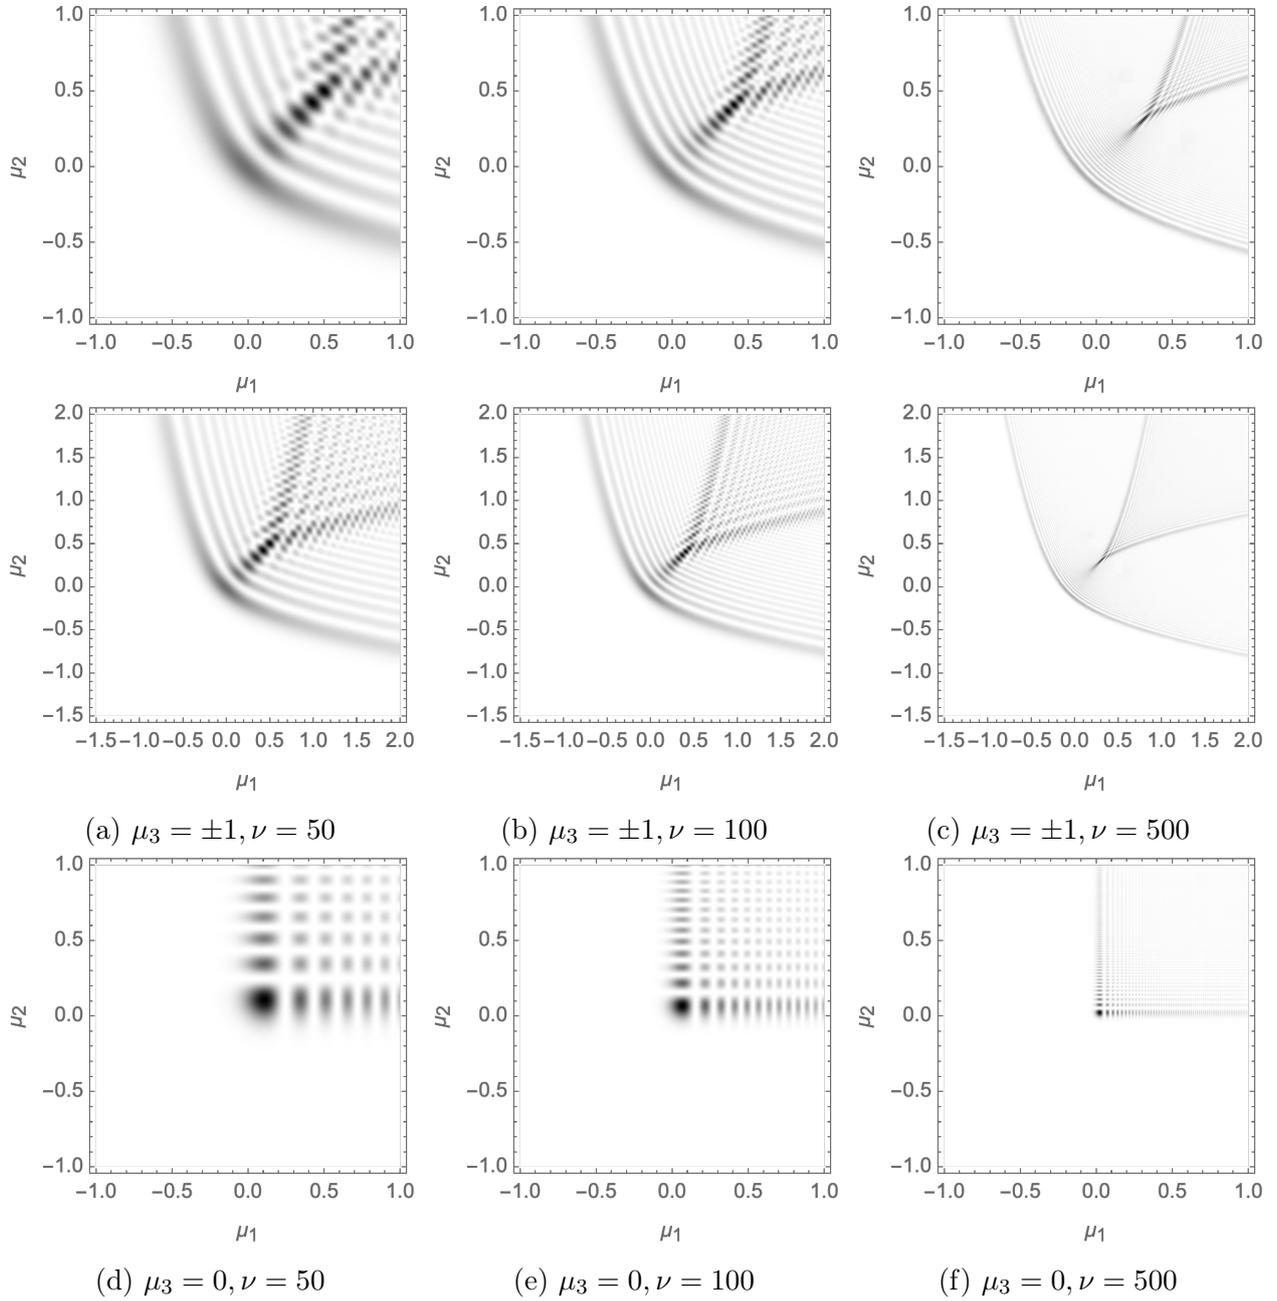


Figure 10.24: The normalized intensity, $I(\boldsymbol{\mu}; \nu)$, of the unfolding of the hyperbolic umbilic caustic (D_4^+) sliced by the surfaces $\{\mu_3 = \pm 1\}, \{\mu_3 = 0\}$ (respectively the upper and lower panels) for the frequencies $\nu = 50, 100$ and 500 (respectively the left, the centre and right panels).

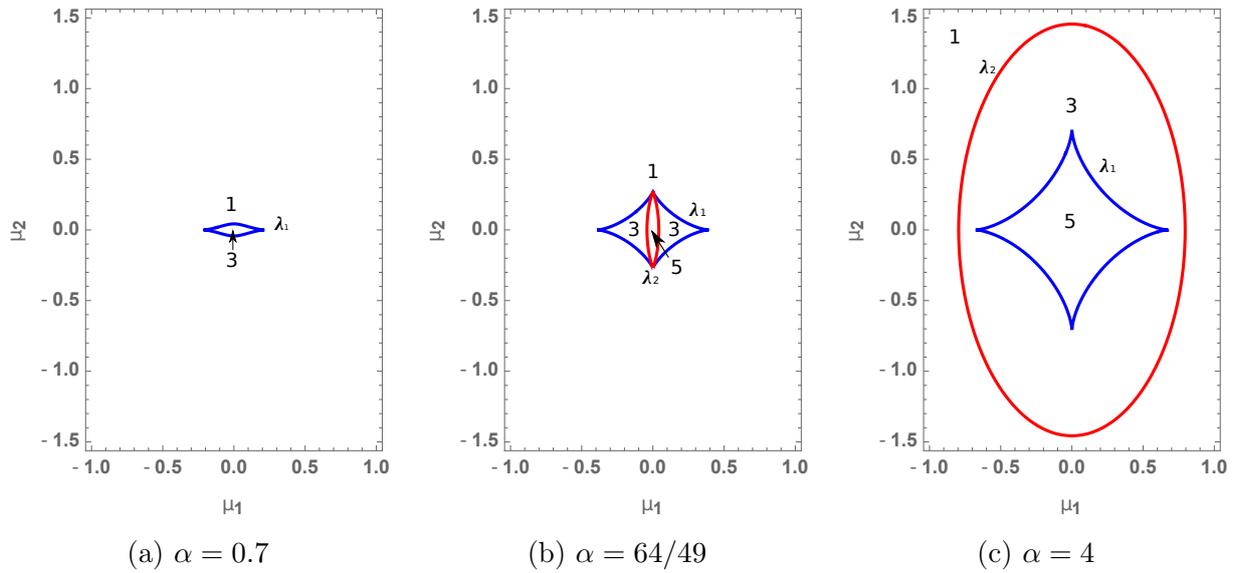


Figure 10.25: The caustics corresponding to λ_1 (blue) and λ_2 (red) as a function of α . We observe the formation of a triple- and a five-image region.

the two-dimensional generalization of the one-dimensional lens studied in Section 11.4. For astrophysical plasma lenses, the parameter α scales according to the dispersion relation $\alpha \propto \nu^{-2}$.

The Lagrangian map is given by

$$\xi(\mathbf{x}) = \mathbf{x} + \frac{1}{2} \nabla \varphi(\mathbf{x}) \quad (10.89)$$

$$= \mathbf{x} - \frac{\alpha}{(1 + x_1^2 + 2x_2^2)^2} (x_1, 2x_2). \quad (10.90)$$

The map forms a caustic where the deformation tensor

$$\mathcal{M}_{ij} = \frac{\partial^2 \phi(\mathbf{x}; \boldsymbol{\mu})}{\partial x_i \partial x_j}, \quad (10.91)$$

with the eigenvalue and eigenvector fields $\lambda_i(\mathbf{x}), v_i(\mathbf{x})$, is singular, *i.e.*,

$$|\mathcal{M}(\mathbf{x})| = \lambda_1(\mathbf{x})\lambda_2(\mathbf{x}) = 0. \quad (10.92)$$

For convenience, we order the eigenvalue and eigenvector fields by $\lambda_1(\mathbf{x}) \leq \lambda_2(\mathbf{x})$.

The first caustic forms at the origin $(\mu_1, \mu_2) = (0, 0)$ for the parameter $\alpha = \frac{1}{2}$ (see Fig. 11.25). This is a cusp singularity. Note that by construction this caustic corresponds to the eigenvalue field λ_1 . For $\frac{1}{2} < \alpha < \frac{64}{49}$ the A_3 point forms an outgoing fold-line (A_2) with two cusps (A_3) on the left and the right. At $\alpha = 1$, a new A_3 point is created, this time corresponding to the second eigenvalue field λ_2 . For $1 < \alpha < \frac{64}{49}$ the A_3 point forms a fold-line (A_2) with two cusps (A_3) at the top and the bottom. At $\alpha = \frac{64}{49}$ the two fold lines merge in a hyperbolic umbilic (D_4^+) at $(\mu_1, \mu_2) = (0, \pm 1/\sqrt{14})$. For $\alpha > \frac{64}{49}$ the two fold lines continue to move outwards, where the fold-line corresponding to λ_1 has four cusps while the fold-line corresponding to λ_2 does not contain a cusp. Outside the fold-line of the caustics, the image consists of a single-image region. Inside the blue fold line, we find a triple-image and a five-image region enclosed by the red fold line.

The analytic continuation of the exponent, $\phi(\mathbf{x}; \boldsymbol{\mu})$, possesses a pole on the two-dimensional surface $x_1^2 + 2x_2^2 + 1 = 0$. Note that poles are never isolated in multi-dimensional complex analysis [57]. The exponent has nine saddle points $\bar{\mathbf{x}}_i$. By evaluating the gradient of the h -function and

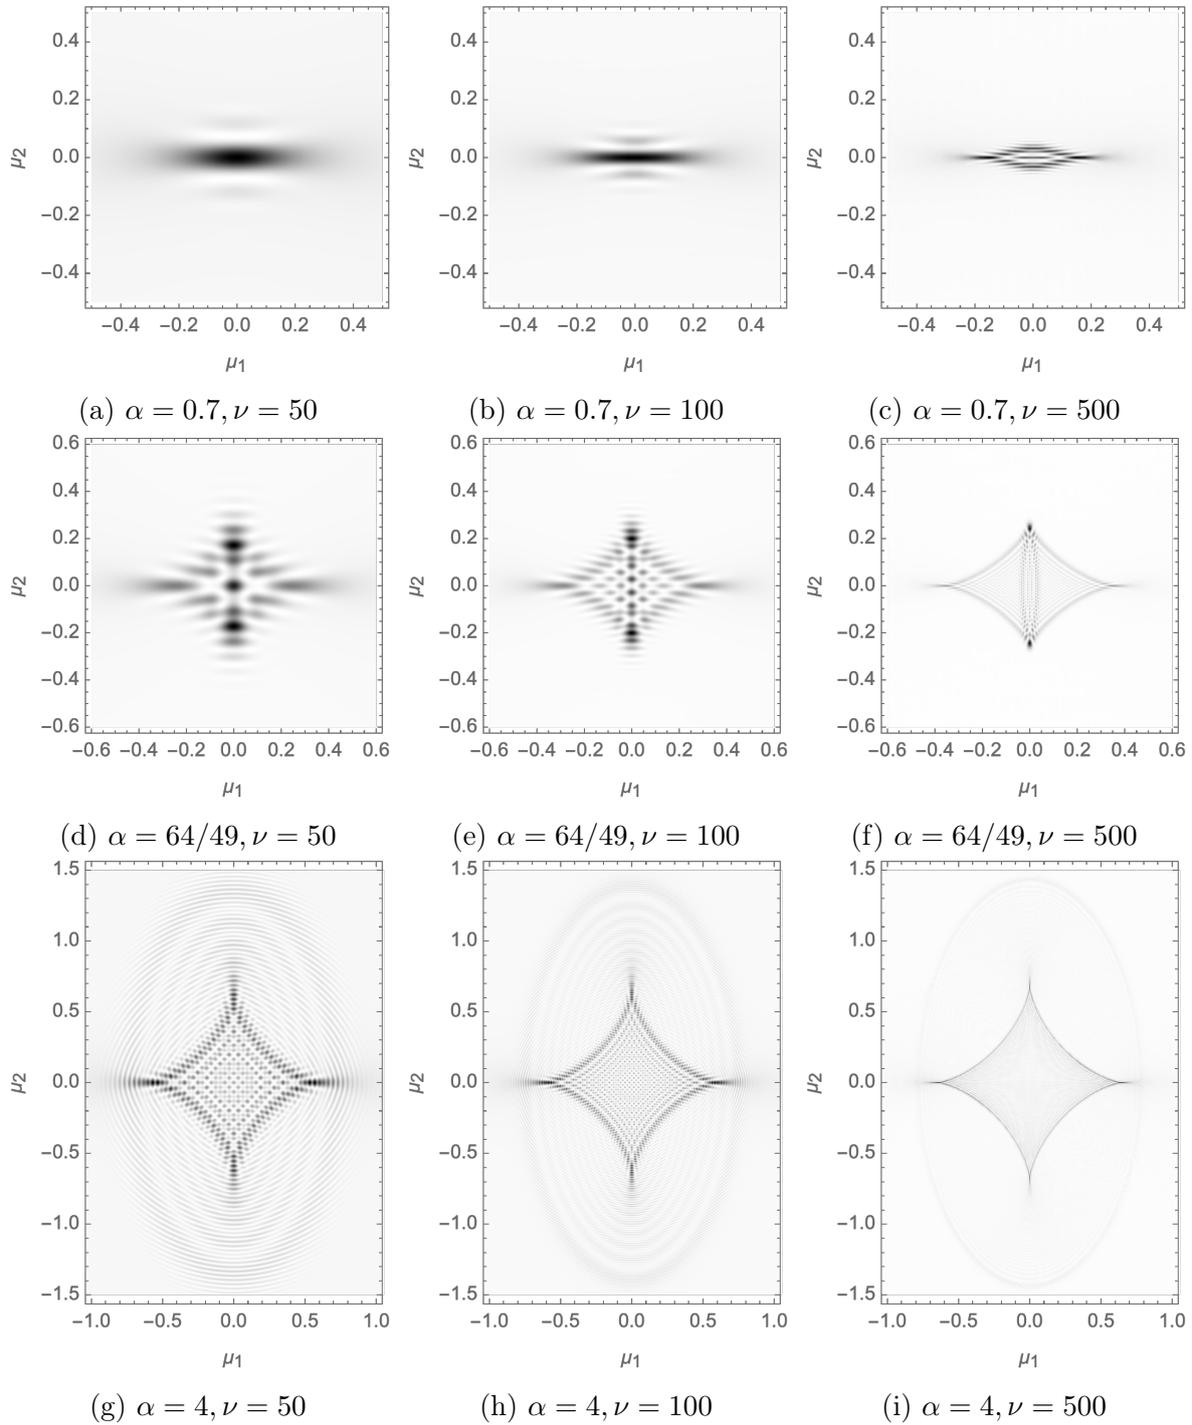


Figure 10.26: Intensity $I(\boldsymbol{\mu}; \nu)$ of the local lens at $\alpha = 0.7, 64/49, 4$ for $\nu = 50, 100$ and 500 .

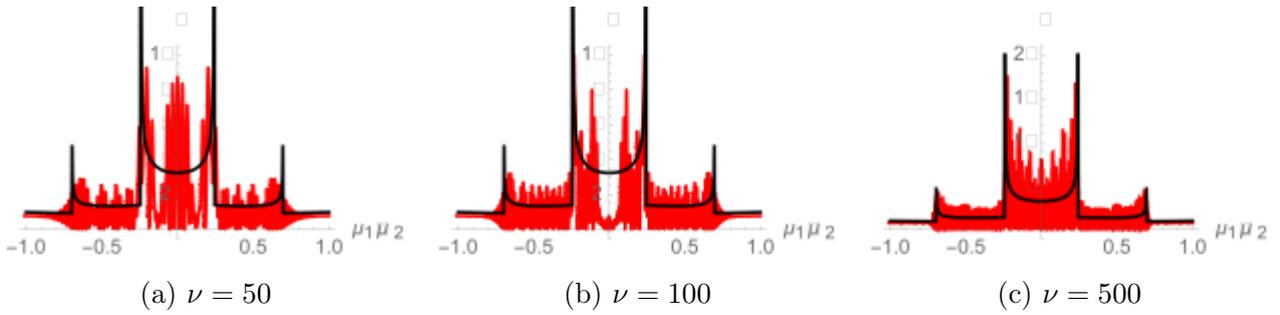


Figure 10.27: The normalized intensity evaluate along the diagonal in the (μ_1, μ_2) -plane for $\alpha = 4$ for $\nu = 100$ and 500 . The black curve is the envelope predicted by geometric optics (see equation (11.18)).

flowing the original integration domain, we obtain a numerical representation of the thimble $\mathcal{J} \subset \mathbb{C}^2$.

Given the two-dimensional thimble \mathcal{J} , we numerically evaluate the integral $\Psi(\boldsymbol{\mu}; \nu)$. In Fig. 11.26, we plotted the normalized intensity of the sensed signal for $\alpha = 7/10, 64/49$ and 4 as a function of the frequency. Observe that when the wavelength is comparable to the size of the caustic structure, the normalized intensity is blurred. The caustics emerge when the wavelength becomes shorter. At the frequency $\nu = 500$, we accurately recover the image corresponding to geometric optics. Remark the stripes emanating from the cusp singularities. This is the frequency independent power-law falloff we observed in the elementary singularities.

In Fig. 11.27, we plot the cross-section of the normalized intensity map along the diagonal $\mu_1 = \mu_2$ for the lens with $\alpha = 4$ for $\nu = 50, 100$ and 500 . Observe the four spikes while passing through the fold catastrophe. Note that the spikes increase in magnitude as ν is raised. In the astronomical context, these spikes correspond to amplification in the light-curve of the lensed source.

10.6.2 A degenerate peak

A more intricate structure arises for the lens corresponding to the degenerate peak in the phase,

$$\varphi(\mathbf{x}) = \frac{\alpha}{1 + x_1^4 + x_2^2}, \quad (10.93)$$

with the Lagrangian map

$$\xi(\mathbf{x}) = \mathbf{x} + \frac{1}{2} \nabla \varphi(\mathbf{x}) \quad (10.94)$$

$$= \mathbf{x} - \frac{\alpha}{(1 + x_1^4 + x_2^2)^2} (2x_1^3, x_2). \quad (10.95)$$

The caustics structure of the Lagrangian map for varying α is plotted in Fig. 11.28. For $\alpha = 1$ we find two disconnected components, which are joined at $\alpha = 1.5$ and form an intricate pattern at $\alpha = 2$ and $\alpha = 2.5$. At $\alpha = 2$ we again find a hyperbolic umbilic caustic (D_4^+) at the two points where the cusps corresponding to the first and second eigenvalue fields λ_1, λ_2 coincide. We thus see that not only the structure at the peak but also the falloff of the variation in the phase φ is important in the study of caustics in lensed images. The caustic structure is generally sensitive to the Hessian of the phase φ , *i.e.*, the second order derivatives.

After flowing the original integration contour to the Lefschetz thimble \mathcal{J} , we numerically evaluate the amplitude $\Psi(\boldsymbol{\mu}; \nu)$ and the corresponding normalized intensity $I(\boldsymbol{\mu}; \nu)$. The resulting normalized intensity maps are plotted in figures 11.29 and 11.30. For the frequency

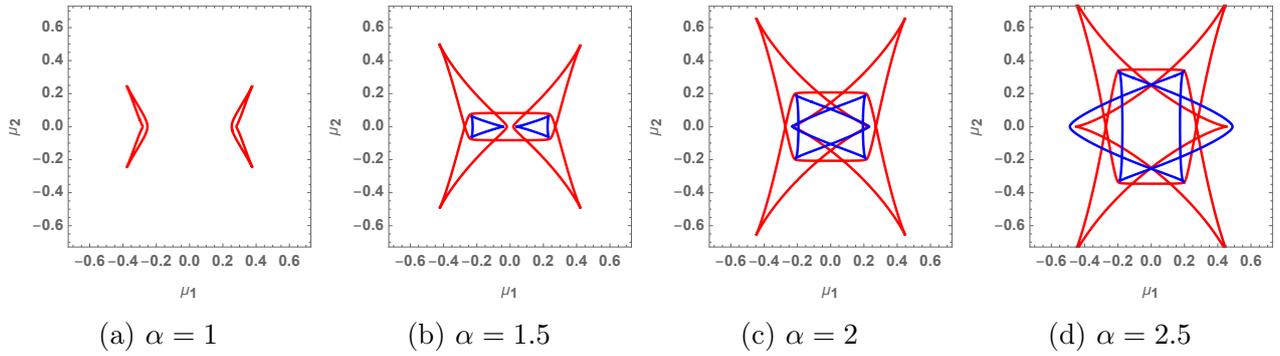


Figure 10.28: The caustics of the Lagrangian map for varying α . The caustics corresponding to the first and second eigenvalue fields λ_1, λ_2 in red and blue.

$\nu = 50$, the image is again rather blurry. We can see the general shape, but cannot distinguish the detailed line structure. For the frequency $\nu = 100$, the lines are better resolved. However, the length scale of the caustics is comparable to the length scales of the interference patterns in the multi-image regions. For $\nu = 500$, we see the complete geometric structure of the caustics. The oscillations in the multi-image regions are now very fine. For this frequency, we are very close to the geometric optics approximation.

10.6.3 The swallowtail caustic

In the previous two examples of lenses corresponding to the simple peaks, we found both fold (A_2) and cusp caustics (A_3) corresponding to a single eigenvalue field, and the interaction between two eigenvalue fields via the hyperbolic umbilic (D_4^+). The two remaining caustics, *i.e.*, the swallowtail (A_4) and the elliptic umbilic (D_4^-), appear in slightly more involved lenses. For the swallowtail caustic, consider the lens

$$\varphi(\mathbf{x}) = \frac{\alpha x_1}{1 + x_1^4 + x_2^2}. \quad (10.96)$$

Again, in the astrophysical context, α follows the dispersion relation $\alpha \propto \nu^{-2}$.

The corresponding integrand $i\phi(\mathbf{x}; \boldsymbol{\mu})\nu$, consists of 23 saddle point in the complex plane. By deforming the integration domain to the thimble, we evaluate the two-dimensional lens integral numerically. See figures 11.31 and 11.32 for the caustics obtained from geometric optics and the corresponding normalized intensity maps for the frequency $\nu = 50, 100, 500$.

- For $\alpha = 2$, the lens forms a caustic corresponding to a single eigenvalue field (see the upper panels of figure 11.31). The profile a pancake with two cusps at the tips. In the corresponding normalized intensity field, we see an interference pattern in the triple-image region, two stripes emanating from the cusps and more strikingly two diagonal stripes going to the left in the single image region. These stripes are a precursor of the swallowtail caustic emerging at later α .
- As α is raised to 3, a second caustic emerges in the triple-image region (see the lower panels of figure 11.31). This caustic corresponds to the second eigenvalue field of the deformation tensor. At $\alpha = 3$ one of the two cusps of the second fold line merges with the outer fold line and transfers the cusp singularity via an elliptic umbilic caustic (D_4^-). For larger α , the blue line will thus have three cusps whereas the red line has only one.

However, more importantly, the lens forms a swallowtail caustic (A_4) in the blue line at $\alpha = 2$. This phenomenon cannot be observed in the blue fold-line but is apparent in the

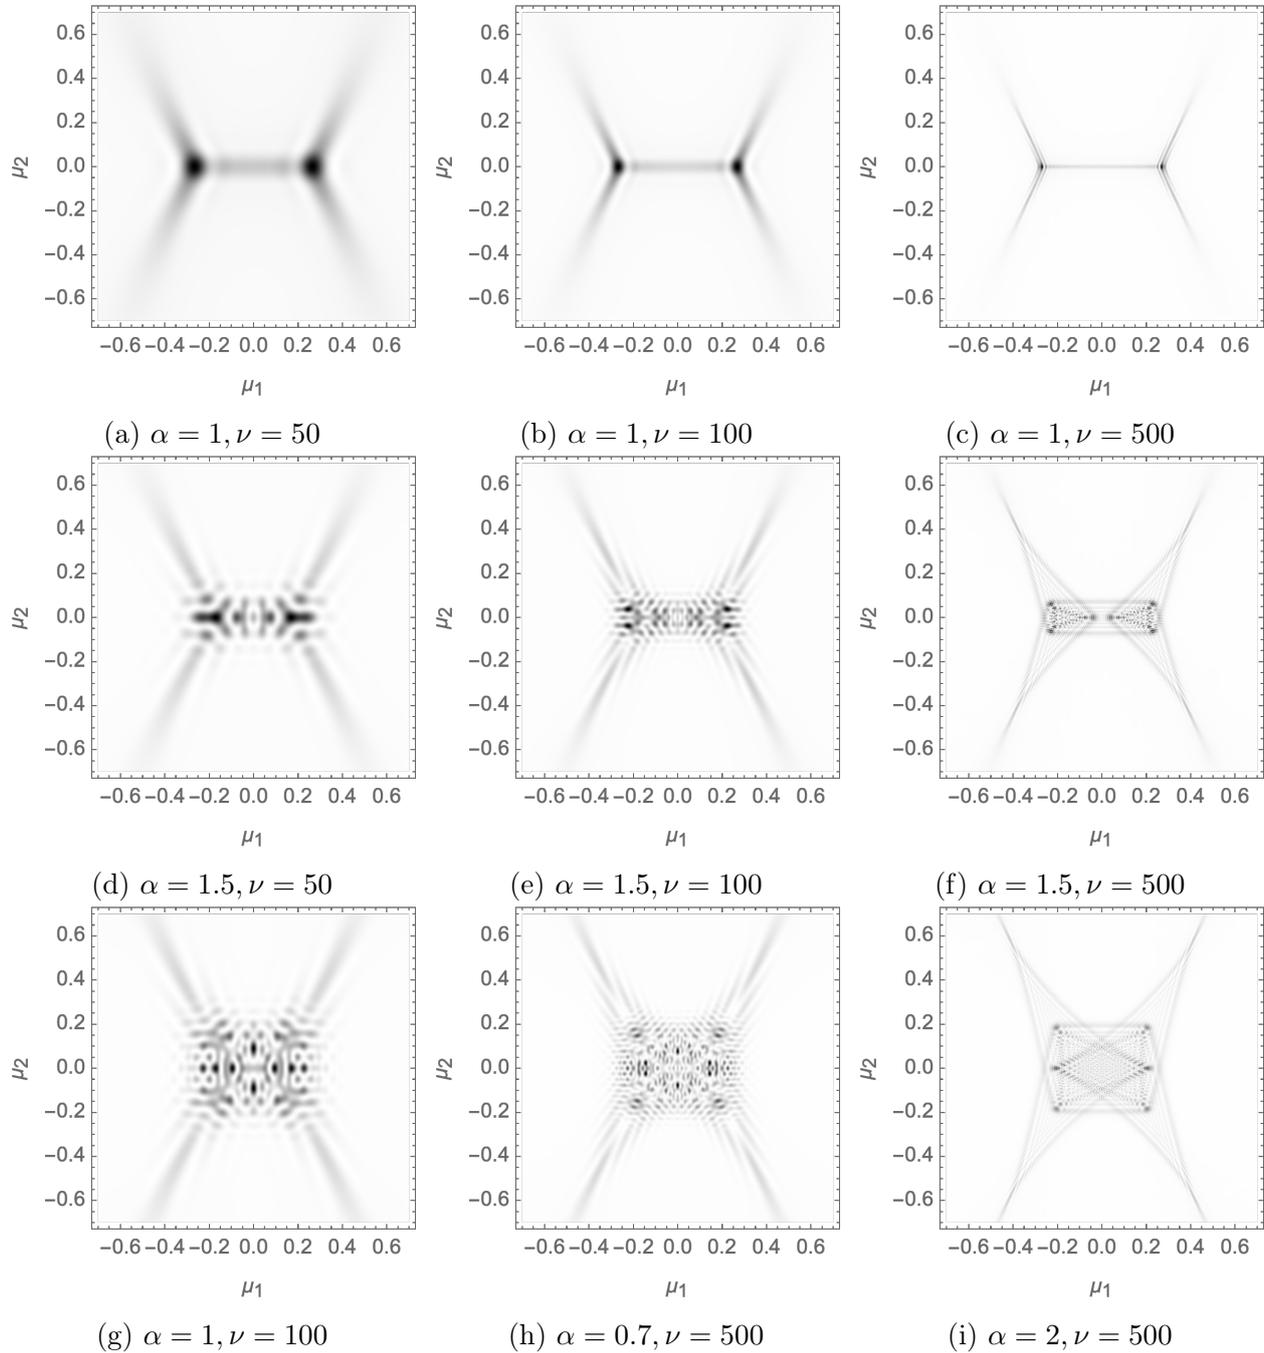


Figure 10.29: The normalized intensity map, $I(\boldsymbol{\mu}; \nu)$, for different frequencies.

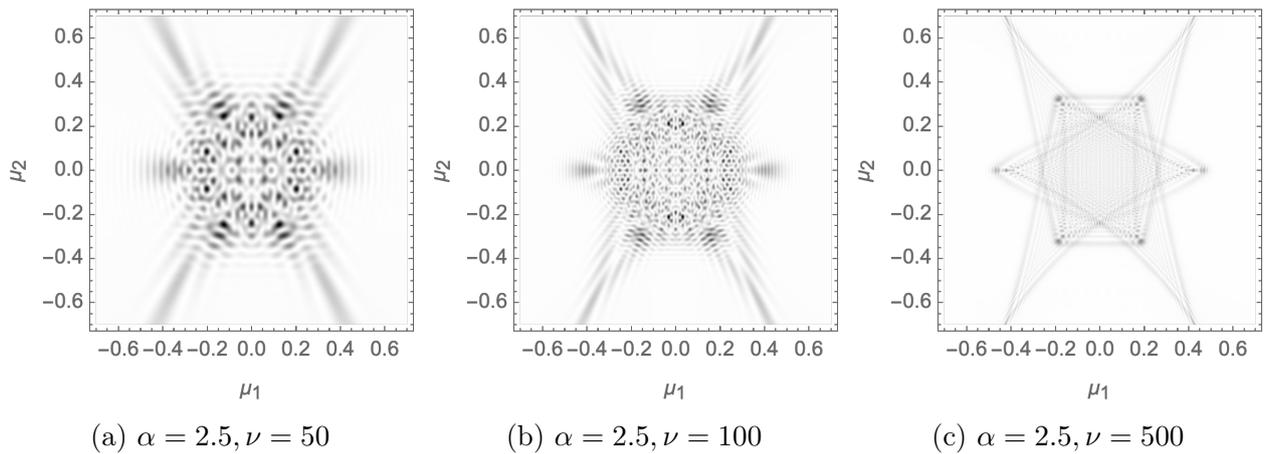


Figure 10.30: The normalized intensity map, $I(\boldsymbol{\mu}; \nu)$, for different frequencies.

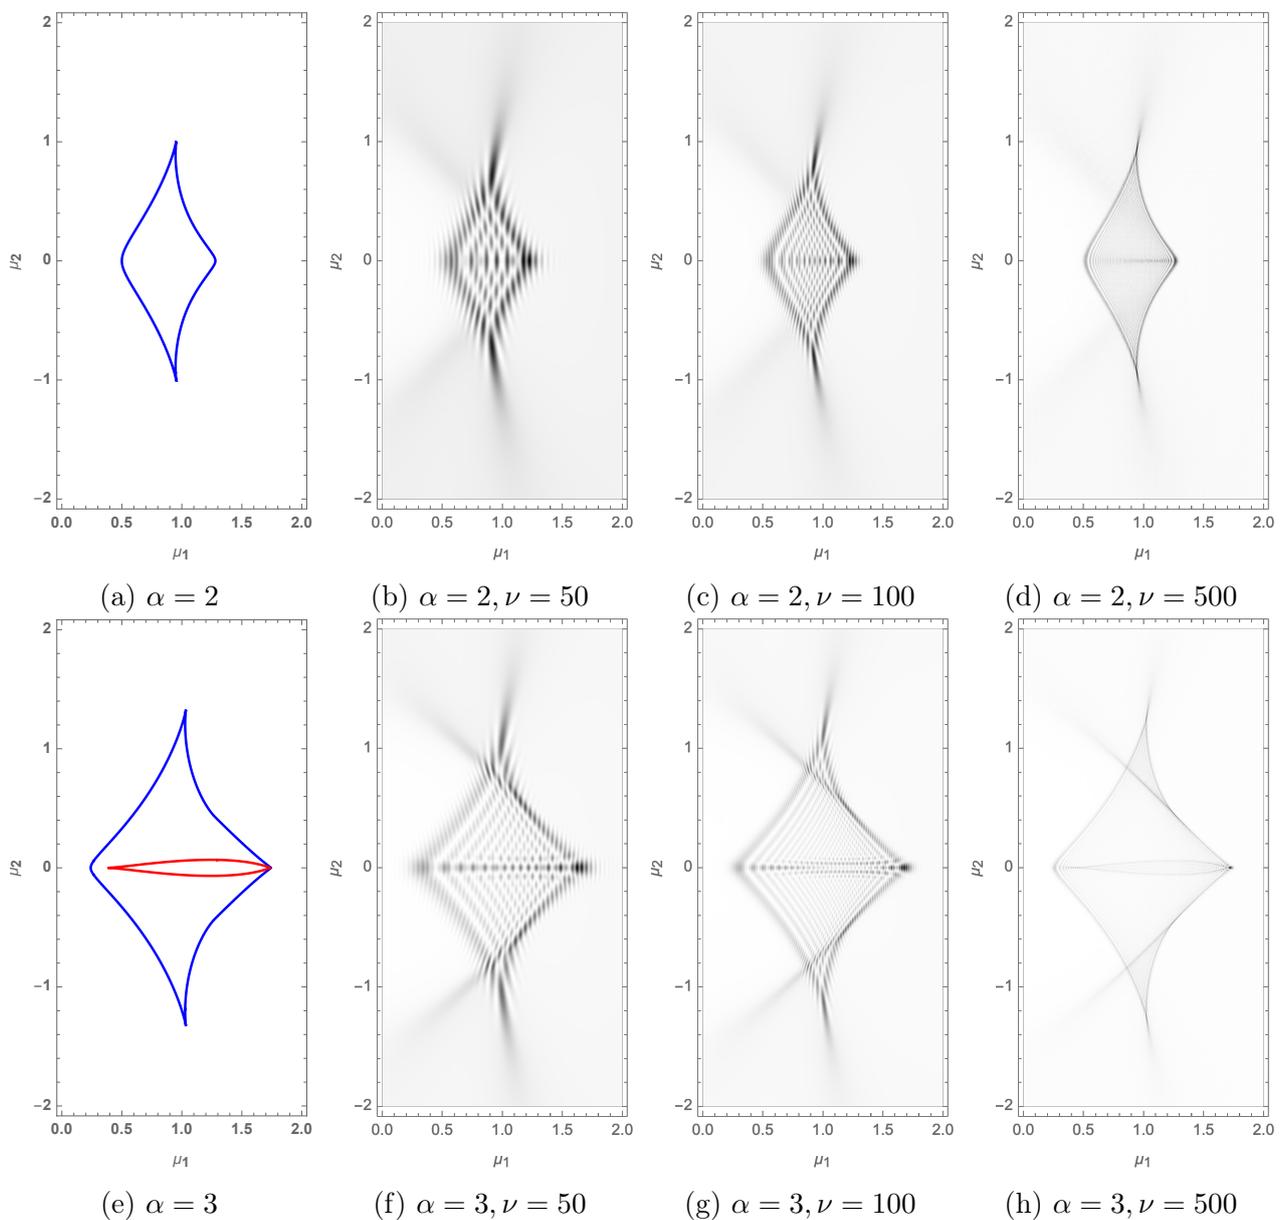


Figure 10.31: The normalized intensity map, $I(\boldsymbol{\mu}; \nu)$, for $\alpha = 2, 3$ and frequencies $\nu = 50, 100, 500$.

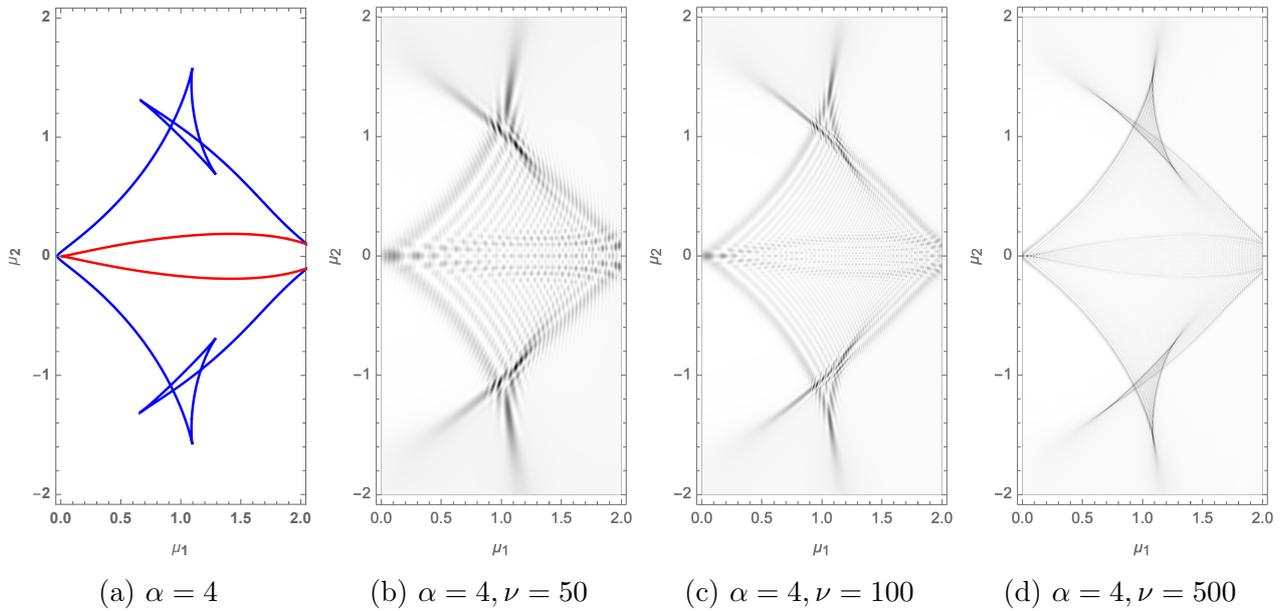


Figure 10.32: The normalized intensity map, $I(\boldsymbol{\mu}; \nu)$, for $\alpha = 4$ and frequencies $\nu = 50, 100, 500$.

normalized intensity map. The two stripes already visible for $\alpha = 2$ are amplified. At the location where the swallowtail stripe coincides with the fold-line, we see an amplification of the normalized intensity in the swallowtail point.

In the normalized intensity map, we see that the geometry becomes sharper and sharper as we increase the frequency and approach the geometric optics limit. Note that the normalized intensity of the hyperbolic umbilic (D_4^+) outshines the other caustics at frequency $\nu = 500$.

- Finally, for $\alpha = 4$, we see that the swallowtail caustic has unfolded into its characteristic shape in the blue fold-line (see Fig. 11.32). We see the same structure emerge in the normalized intensity map. However, in addition, we how to obtain a large number of stripes emanating from the cusp caustics.

We also see that the lens at $\alpha = 4$, consists of a second hyperbolic umbilic (D_4^+) appearing at the origin, where the blue and the red fold-lines meet. As the frequency is raised, we again see that the normalized intensity spikes for this caustic.

10.6.4 The elliptic umbilic caustic

We conclude this section by studying the elliptic umbilici (D_4^-) caustic in a localized lens. The elliptic umbilic forms when the deformation tensor is singular due to two eigenvalues vanishing simultaneously. The geometry of the caustic however differs from the hyperbolic umbilic (D_4^+), in that it includes the merger of three cusp caustics. We here study the localized lens

$$\varphi(\mathbf{x}) = \frac{\alpha(x_1^3 - 3x_1x_2^2)}{1 + x_1^2 + x_2^2}. \quad (10.97)$$

From geometric optics, we observe the caustic structure of the lens (see Fig. 11.33).

- For small, $\alpha < 1.4$, the lens consists of three Zel'dovich pancakes with a triangular symmetry. Three of the cusp caustic point to the origin of the parameter space.

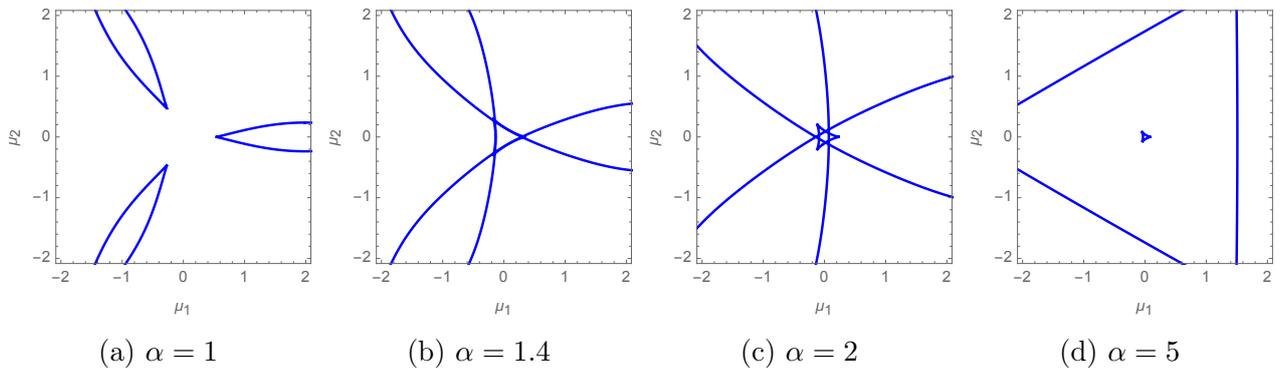


Figure 10.33: The caustics of the Lagrangian map for varying α .

- At $\alpha = 1.4$, we observe that the three Zel'dovich pancakes are joined by three fold-lines forming a triangular structure.
- As $\alpha > 1.4$, the triangle decouples from the three Zel'dovich pancakes. The three resulting fold lines move away from the origin and the triangle shrinks to a point. The point at which the triangle is contracted to a point is the elliptic umbilic caustic. The region enclosed by the large triangle is a 5-image region. The region enclosed by the small triangle is a 7-image region.

Note that since the elliptic umbilic caustic only forms after three cusp caustics have formed a triangular fold line, the caustic will be rare in simple lenses. It is nonetheless a stable configuration, as a small deformation of the lens preserves the structure.

Using the Picard-Lefschetz analysis, we evaluate the normalized intensity map for the configurations $\alpha = 1, 1.4$, and 5 for the frequencies $\nu = 50, 100$ (see Fig. 11.34).

- For $\alpha = 1$, we observe that even though the triangular structure is not yet present in the geometric optics analysis, it is present in the normalized intensity map at finite frequency (see the left panels of Fig. 11.34). That is to say, the normalized intensity is enhanced at the triangle, however as $\nu \rightarrow \infty$ the normalized intensity at the triangle will remain finite.
- At $\alpha = 1.4$, the triangle has formed in the geometric optics analysis (see the central panels of Fig. 11.34). In the normalized intensity maps, the triangle is enhanced. The normalized intensity will now diverge in the geometric optics limit.
- As α is further increased to $\alpha = 5$, the triangle shrinks to a point and interference effects between the different fold lines start to appear (see the right panels of Fig. 11.34). At $\alpha = 5$, we do no longer observe the fold lines but rather observe a triangular blob at the origin of the parameter space. This closely resembles the normalized intensity map of the elementary elliptic umbilic catastrophe. It is however a bit more intricate as a close inspection demonstrates that caustic structure oscillates at a high frequency due to the interference of the elliptic umbilic with the surrounding multi-image region.

Note that there are a few small numerical artefacts present in the normalized intensity map for the lens at $\alpha = 5$. The lens outside of the triangle, is a 5-image region in which some of the real saddle points are located far away from the origin in the lens plane. The inside of the triangular region is a 7-image region. The Lefschetz thimble has a complicated shape and the tessellation of the thimble can occasionally miss a few points.

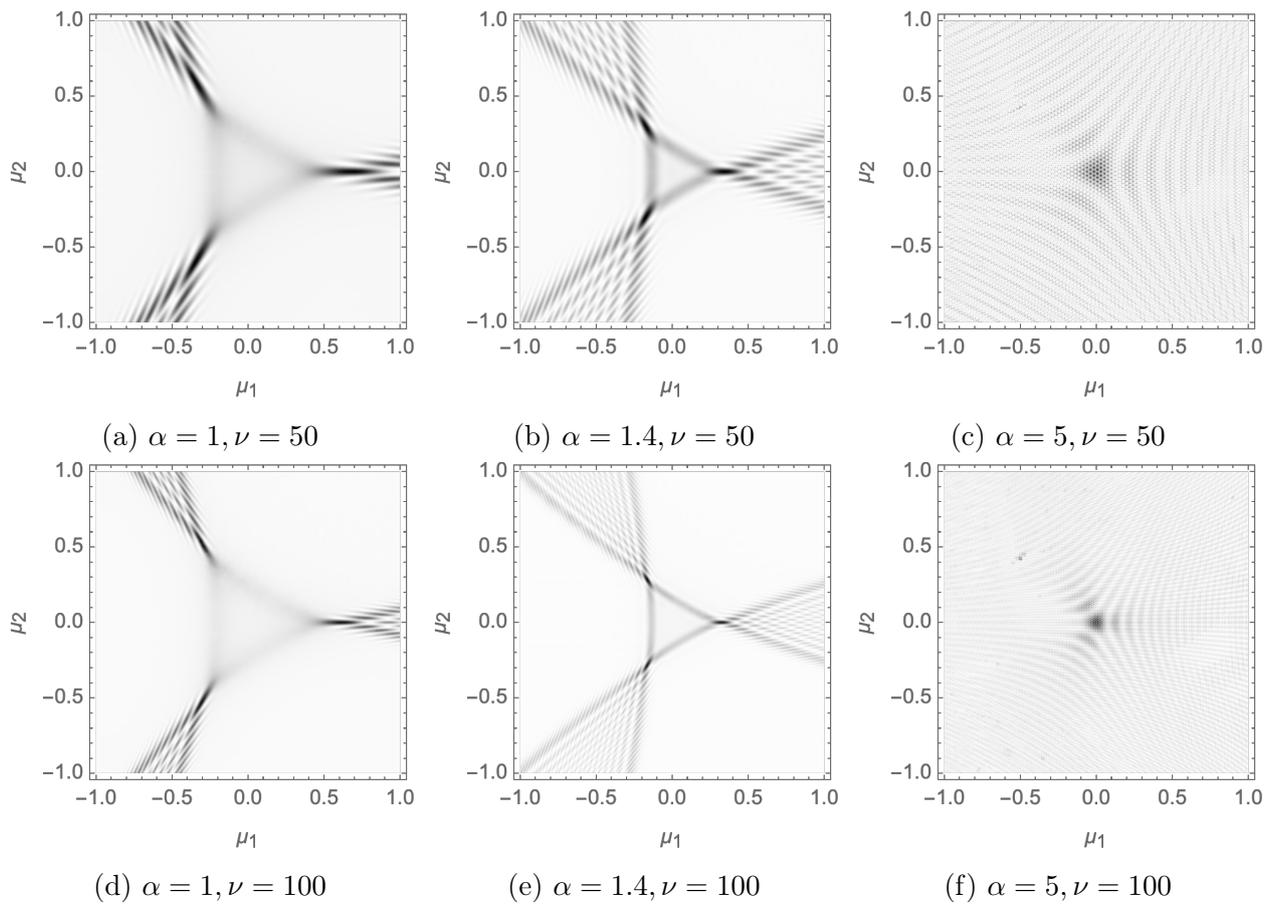


Figure 10.34: The normalized intensity map, $I(\boldsymbol{\mu}; \nu)$, for $\alpha = 1, 1.4, 5$ and frequencies $\nu = 50, 100$.

10.7 Signatures of caustics in fast radio bursts

A Fast Radio Burst (FRB) is a millisecond transient radio pulse, caused by some yet to be identified high-energy astrophysical process. The first burst was found by Duncan Lorimer and his student David Narkevic in 2007 while scanning through archival pulsar survey data [58]. The burst in question had been detected in 2001 by the Parkes Observatory in Australia. In subsequent years, several other bursts were observed, among which the first repeating source (named FRB 121102) [8] was detected in 2012 by the Arecibo Observatory in Puerto Rico. In the last few months, several new detections have been announced by the Canadian Hydrogen Intensity Mapping Experiment (CHIME) collaboration, including the second repeating FRB source (named FRB 180814) [59]. FRBs are now known to be relatively common, with approximately 10,000 bright fast radio bursts occurring per day over the entire sky. Telescopes capable of detecting a significant fraction of these bursts should become possible in coming decades, an exciting prospect indeed.

The source of fast radio bursts is yet to be identified. Many different models have been proposed but none is yet compelling. They range from rapidly spinning neutron stars or black holes and regions of very high electromagnetic fields, to more exotic sources [60, 61]. It seems likely that the bursts are extragalactic in origin, as the first observed repeater, FRB 121102, has been identified with a galaxy at a distance of approximately 3 billion light years [62, 63, 64]. As mentioned in the introduction, it is likely that the phenomenology of fast radio bursts is strongly affected by astrophysical plasma lensing. They have a characteristic time-frequency profile, their frequency typically falling during the pulse, or series of pulses. This profile is probably due to the fact that lower frequencies are more strongly lensed and thus follow longer geometrical paths, and also because they propagate more slowly.

The methods and results we have reported here should be helpful in modeling the effects of plasma lensing on observed FRBs. The lensing may take place in a variety of places – near the source, near the observer or in between. If the line of sight encounters a caustic due to a plasma lens, the FRB may be amplified, enhancing the chances of detection. For reasons we have explained, caustics are likely to be localized in frequency, leading to the observed spectral shape. The “marching down” features could also be due to asymmetric structures in the lens, leading to angled caustics. This requires a preferred time asymmetry, which could in turn provide hints about the structure of the lens itself. In the lensing example of B1957+20 [13], the lens is due to a companion wind. In this specimen, the time-frequency caustics march both up and down. This symmetry could be broken if the wind contained shock waves, which could preferentially move retrograde in the rotating frame. Quantitative lens modeling can be tested on the pulsar binary system, and then applied to FRB data. This could be the scope of a future paper.

Since the observed radio waves have a relatively long wavelength, the corresponding diffraction catastrophes are likely to fill a significant volume in the parameter space of the normalized intensity maps. Therefore it is important to study the complete interference pattern. It follows from Table 2 that the elliptic (D_4^-), the hyperbolic (D_4^+) umbilic and to a lesser extent the swallowtail (A_4) caustic lead to the largest spikes in the normalized intensity map. Of these three caustics, the swallowtail (A_4) and the hyperbolic (D_4^+) umbilic caustics are most likely to be realized in simple lenses, of which the hyperbolic caustic gives the greatest amplification. However, these caustics will not generically occur in time-frequency data, as they are formed at point in three-dimensional functions. The line of sight, is, however, reasonably likely to pass close to them, as they fill a finite volume of the parameter space. In principle, we do expect to see the cusp (A_3) points and the fold (A_2) lines caustics, in the data. However, note that these caustics lead to a lesser amplification of the source.

As we observed in the previous sections, caustics due to multi-dimensional lenses never occur as isolated events. The caustics of co-dimension four, *i.e.*, the umbilics D_4^\pm , and the

swallowtail A_4 caustics, are always accompanied by cusp (A_3) points and fold (A_2) lines. It thus follows that when a fast radio burst is indeed amplified by a lens, that the corresponding peak in time-frequency space will be of characteristic shape. More concretely, after identifying the time and the frequency with the two of the unfolding parameters μ , we expect the peak to resample the normalized intensity map of the corresponding elementary catastrophe computed in Section 11.5. That is to say, the peak corresponding to elliptic (D_4^-) umbilic caustic should exhibit a triangular symmetry and the peak corresponding to the swallow (A_4) caustic will exhibit the characteristic swallowtail geometry in the fold-line and two cusps caustics.

Further investigation is required to estimate the number density of the different caustics for generic two-dimensional lenses and the most likely normalized intensity profiles along the line of sight.

10.8 Conclusions

Conditionally convergent oscillatory integrals play a central role in modern physics. However, these integrals are often difficult to define as their definition, in the multi-dimensional case, can depend on the order of integration or the regularization scheme. They are, moreover, generically impossible to evaluate analytically and too expensive to evaluate with conventional numerical methods. In this paper we have brought Picard-Lefschetz theory to bear. We have shown how in a multi-dimensional oscillatory integral, the integrand generically defines a set of relevant Lefschetz thimbles in the complexified integration domain, along which the integral is absolutely convergent. These thimbles can be thought of as an ‘integrand-dependent Wick rotation’. The integral evaluated along the set of relevant thimbles in fact provides an unambiguous definition of the original integral itself. We moreover have presented a new, efficient numerical scheme both to find the thimbles and to efficiently evaluate the integral along them in polynomial time. The virtue of this new method that the efficiency actually *increases* as the integrand becomes more oscillatory.

In particular, we have studied the Lefschetz thimbles for caustic catastrophes and the Stokes phenomenon occurring in two-dimensional lenses. Given the thimbles, we numerically evaluate the normalized intensity maps over all frequencies study the resulting interference patterns. We have shown that the normalized intensity maps smoothly converge to the caustics predicted by geometric optics, without introducing numerical artifacts.

Our method renders feasible the calculation of interference patterns in a wide variety of interesting astrophysical contexts, in particular to model the effect of plasma lenses on radio sources. So far, such modeling has been restricted to the simplest examples of fold and cusp singularities, produced by one dimensional lenses. More realistic, two-dimensional models, including the swallowtail, elliptic umbilic and hyperbolic umbilic caustics are now accessible. We have computed the normalized intensity maps for a few representative examples, and commented briefly on likely observational signatures. A statistical analysis of the normalized intensity profiles for the diffraction catastrophes generated by a realistic plasma lens ensembles will be the subject of further investigations.

Finally, we analyzed a simple model of Young’s double slit experiment, representing an initial exploration of the use of these methods for describing interference in quantum mechanics.

10.9 Defining oscillatory integrals

Oscillatory integrals, which do not converge absolutely, are sometimes claimed to be ill-defined since the key theorems of measure theory, *e.g.* the dominated convergence theorem and Fubini’s theorem, do not apply [65]. We here study conditionally convergent oscillatory integrals for the one- and multi-dimensional case and propose a definition using Picard-Lefschetz theory

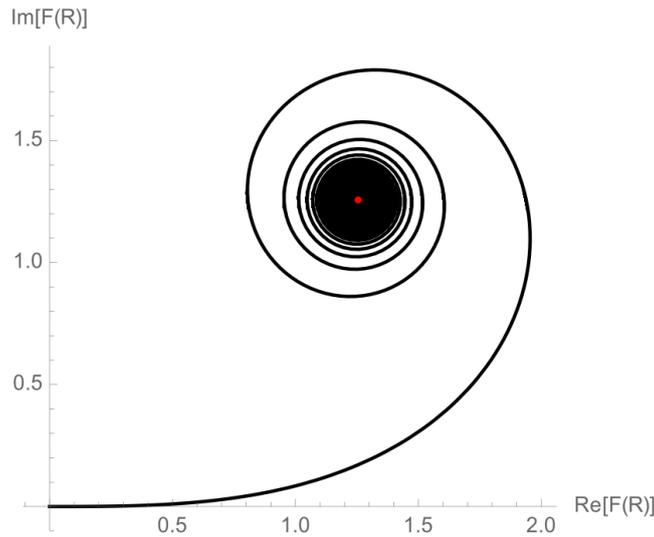


Figure 10.35: The Euler or Cornu spiral. The black line is the real and imaginary part of $F(R)$ as a function of R . The black point is the real and imaginary part of the limit $\lim_{R \rightarrow \infty} F(R) = (1 + i)\sqrt{\frac{\pi}{2}}$.

in terms of absolutely convergent ones.

10.9.1 One-dimensional integral

The Fresnel integral

$$F(\infty) = \int_{-\infty}^{\infty} e^{ix^2} dx = (1 + i)\sqrt{\frac{\pi}{2}} \quad (10.98)$$

exists, even though the integral is only conditionally convergent. The integral is usually defined as a limit of the partial integral

$$F(R) = \int_{-R}^R e^{ix^2} dx, \quad (10.99)$$

i.e., $\lim_{R \rightarrow \infty} F(R) = (1 + i)\sqrt{\frac{\pi}{2}}$ following the Euler or Cornu spiral (see Fig. 11.35). This definition is as important to the integral as the integrand, as different regularization schemes – which do not approach the real line by adding points incrementally – lead to different answers.

The definition of the conditionally convergent integral in terms of the limit $R \rightarrow \infty$ is equivalent to the assumption of analyticity, since Cauchy's integral theorem only applies to integrals over \mathbb{R} defined this way. We can alternatively define the integral by deforming the integration contour \mathbb{R} in the complex plane \mathbb{C} to the Lefschetz thimble

$$\mathcal{J} = \{(1 + i)u | u \in \mathbb{R}\} = (1 + i)\mathbb{R}, \quad (10.100)$$

for which the integrand is convex and the integral is absolutely convergent, *i.e.*,

$$\int_{\mathbb{R}} e^{ix^2} dx = \int_{\mathcal{J}} e^{ix^2} dx \quad (10.101)$$

$$= (1 + i) \int_{\mathbb{R}} e^{-2u^2} du \quad (10.102)$$

$$= (1 + i)\sqrt{\frac{\pi}{2}}. \quad (10.103)$$

Note that this definition does not depend on a limit. The regularization is completely determined by the assumption of analyticity.

10.9.2 Multi-dimensional integrals

Multi-dimensional conditionally convergent oscillatory integrals such as

$$\int_{\mathbb{R}^N} e^{if(x_1, \dots, x_N)} dx_1 \dots dx_N \quad (10.104)$$

for $N \in \mathbb{N}$ and appropriate functions f , play an important role in optics but cannot be uniquely defined using an extension of the regularization scheme described above for the one-dimensional case. To show this, let's consider the two-dimensional analogue of the Fresnel integral

$$\int_{\mathbb{R}^2} e^{i(x^2+y^2)} dx dy. \quad (10.105)$$

Since this integral factorizes, it is reasonable to require the integral to converge to

$$F(\infty)^2 = \left((1+i) \sqrt{\frac{\pi}{2}} \right)^2 = i\pi. \quad (10.106)$$

However, for general $f(x_1, \dots, x_n)$ we are not able to write the integral as a product of one-dimensional integrals. This thus should not be considered as a desirable definition of the integral.

To see the dependence on the regularization scheme, consider the integral in polar coordinates. We write

$$I(R) = \int_{\mathbb{D}_R} e^{i(x^2+y^2)} dx dy \quad (10.107)$$

$$= 2\pi \int_0^R r e^{ir^2} dr \quad (10.108)$$

$$= i\pi \left(1 - e^{iR^2} \right), \quad (10.109)$$

with \mathbb{D}_R the disk of radius R centred at the origin. We thus find that the limit $\lim_{R \rightarrow \infty} I(R)$ does not exist! The function $I(R)$ instead circles the 'correct answer' $i\pi$ with increasing angular velocity.

It is instead appropriate to define the integral in terms of the Lefschetz thimble

$$\mathcal{J} = \{(1+i)(u, v) \mid (u, v) \in \mathbb{R}^2\} \quad (10.110)$$

$$= (1+i)\mathbb{R}^2. \quad (10.111)$$

Along the thimble, the integral is absolutely convergent

$$\int_{\mathbb{R}^2} e^{i(x^2+y^2)} dx dy = \int_{\mathcal{J}} e^{i(x^2+y^2)} dx dy \quad (10.112)$$

$$= (1+i)^2 \int_{\mathbb{R}^2} e^{-2(u^2+v^2)} du dv \quad (10.113)$$

$$= 2i \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} e^{-2(u^2+v^2)} du dv \quad (10.114)$$

$$= i\pi. \quad (10.115)$$

On the thimble we can safely convert the integral over the real plane \mathbb{R}^2 into the iterative integral using Fubini's theorem, since the integral in u and v over \mathbb{R}^2 is absolutely convergent. This definition straightforwardly generalizes to general multi-dimensional conditionally convergent integrals.

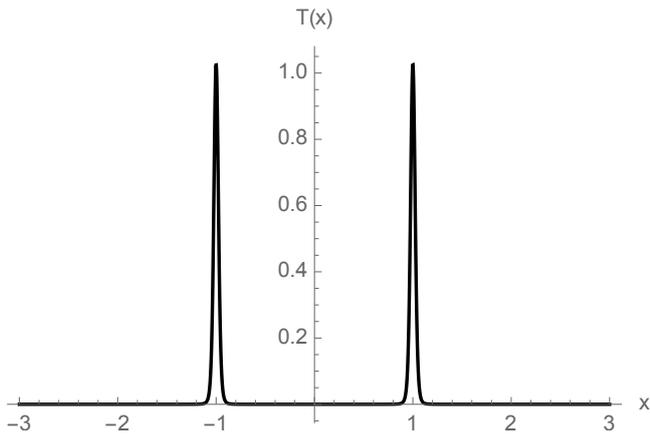


Figure 10.36: The transition amplitude of the wall $T(x)$ with two slits at $s_1 = -1$, $s_2 = 1$ with width $\epsilon = 0.1$.

10.10 Young's double-slit experiment

In this appendix, we generalize our treatment of interference in order to tackle Young's famous double slit experiment. In spite of the extreme simplicity of this example, and its centrality to introductory discussions of quantum physics, detailed interference patterns are surprisingly hard to compute. By generalizing our treatment of the Fresnel-Kirchhoff integral we shall be able to efficiently study the pattern created by a pair of smooth, finite size slits in detail. In particular, we shall see how quantum interference effects disappear in the classical limit, as \hbar is taken to zero.

The generalization required is to make the interference "phase" complex in order to damp out the amplitude away from two narrow slits. Modeling this complex phase with a simple rational function, our numerical techniques allow us to efficiently find the relevant Lefschetz thimbles and compute the detailed interference pattern at all values of the parameters.

Consider a distant point source emitting particles towards a screen, with a thin barrier separating the screen from the source. The barrier is opaque to the particles except in the neighbourhood of two slits. In dimensionless coordinates (which we shall define below), the transmission amplitude takes the form

$$T(x) \propto \exp \left[\frac{\epsilon}{\epsilon^2 + (x - s_1)^2} + \frac{\epsilon}{\epsilon^2 + (x - s_2)^2} - \frac{1}{\epsilon} \right], \quad (10.116)$$

consisting of two peaks each of strength unity, centered respectively at $x = s_1$ and $x = s_2$. Here, $\epsilon > 0$ is a small number representing both the width of the slits (see Fig. 11.36) and the opacity of the barrier: away from the slits, the latter is given by $T \sim \exp(-1/\epsilon)$.

Assuming the incident amplitude for the particles to be coherent and constant across the slits, we may then compute the path integral amplitude just as in Section 11.2. Here, however, we deal with a particle of fixed mass m , energy E and momentum $p = \sqrt{2mE}$. The last formula in (11.8), in the same small displacement-approximations made in Eq. (11.12) above, yields a Pythagorean contribution to the phase, $p(x - \mu)^2/(2d\hbar)$ where d is the distance from the slits to the screen. Setting $x \rightarrow xa$ where a is the characteristic dimension of the slits and x is dimensionless, we take the quantity $\hbar 2d/(pa^2)$ to be our new, dimensionless \hbar . In terms of these dimensionless quantities, the amplitude for the particle to arrive at position μ on the screen is therefore given by the oscillatory integral

$$\Psi(\mu) = \mathcal{N} \int e^{\frac{i}{\hbar}(x-\mu)^2} T(x) dx \quad (10.117)$$

with the normalization constant \mathcal{N} , ensuring unitarity $\int |\Psi(\mu)|^2 d\mu = 1$. The probability for the particle to arrive at μ on the screen is given by the absolute square of the wavefunction

$$I(\mu) = |\Psi(\mu)|^2. \quad (10.118)$$

Note that the dimensionless version of Planck's constant \hbar appears in this nonrelativistic problem, whereas it cancelled out of our earlier formulae for a massless particle, as a result of the latter's scale covariance.

We evaluate the wavefunction (equation (11.117)), by analytically continuing the exponent

$$\phi(x; \mu) = \frac{i}{\hbar}(x - \mu)^2 + \frac{\epsilon}{\epsilon^2 + (x - s_1)^2} + \frac{\epsilon}{\epsilon^2 + (x - s_2)^2} - \frac{1}{\epsilon}, \quad (10.119)$$

in the complex plane and evaluating the Lefschetz thimble. The exponent, ϕ , has four poles and nine saddle points. The poles at $x = s_i \pm i\epsilon$ correspond to the slit centered at s_i . The saddle points are roots of a ninth order polynomial. We can associate four saddle points to each slit. The remaining saddle point is shared and moves between the the poles corresponding to the two slits as a function of the position on the screen μ .

Fig. 11.37 shows the corresponding Picard-Lefschetz diagrams for various positions μ for $\hbar = 1$. In the description we will for simplicity assume the left slit to be at s_1 and the right slit to be at s_2 , *i.e.*, $s_1 < s_2$:

- For positions on the screen far to the left of the slits, $\mu \ll s_1$, the thimble consists of five steepest descent contours. The thimble runs from the lower left to the upper right via a complex saddle point. The thimble subsequently loops around the upper left and the upper right poles. For positions $\mu \leq \frac{s_1 + s_2}{2}$, the wavefunction is dominated by the left slit. It is for this reason not surprising to see that the thimble corresponding to the right slit is representative in this regime.
- As μ approaches s_1 , we observe a Stokes transition after which only four saddle points are relevant. The thimble moves from the lower-left via a saddle point to the upper left pole, after which it passes through the saddle point between the two left poles. The right part of the thimble is largely unchanged.
- For μ near s_1 , we observe yet another Stokes transition after which only three saddle points remain relevant. The thimble runs from the lower left via the saddle point between the to left poles to the upper right.
- When μ approaches the mid-point $\frac{\mu_1 + \mu_2}{2}$, we observe that a complex saddle point becomes relevant after a Stokes phenomenon. The thimble now consists of four steepest descent contours.
- For μ near the mid-point $\mu = \frac{\mu_1 + \mu_2}{2}$, we observe that after yet another Stokes transition, we obtain a thimble consisting of five steepest descent contours. Note that the middle saddle point has moved to the origin $x = 0$. When the position μ is increased further, this saddle point will move to the poles corresponding to the right slit. The corresponding Picard-Lefschetz diagrams are mirror images of the ones discussed above.

In the semi-classical limit $\hbar \rightarrow 0$, the geometry of the Lefschetz thimble is to an increasing extent determined by the Pythagorean term in equation (11.119). As a consequence, after a few Stokes transitions, the eight saddle points which can be associated to the poles corresponding to the two slits become tighter bound to the poles representing the geometry of the right part of the thimble in figure 11.37. The remaining saddle point still moves between the poles corresponding

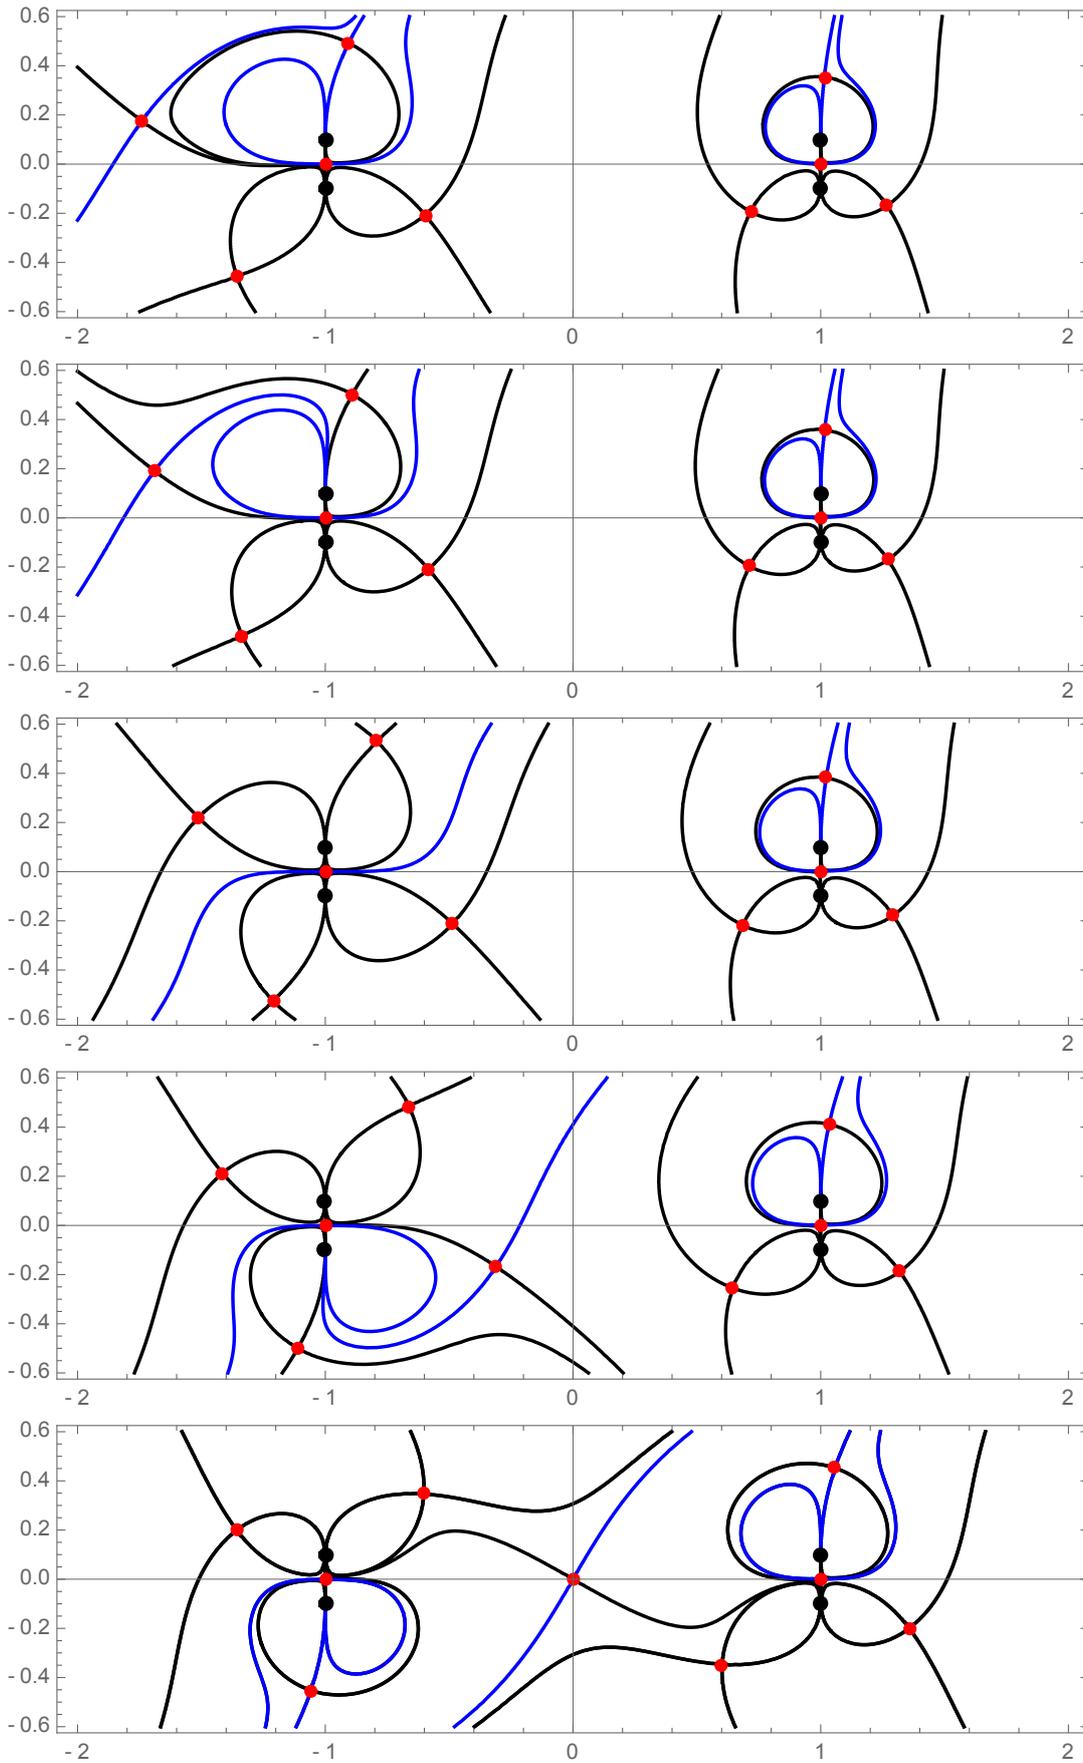
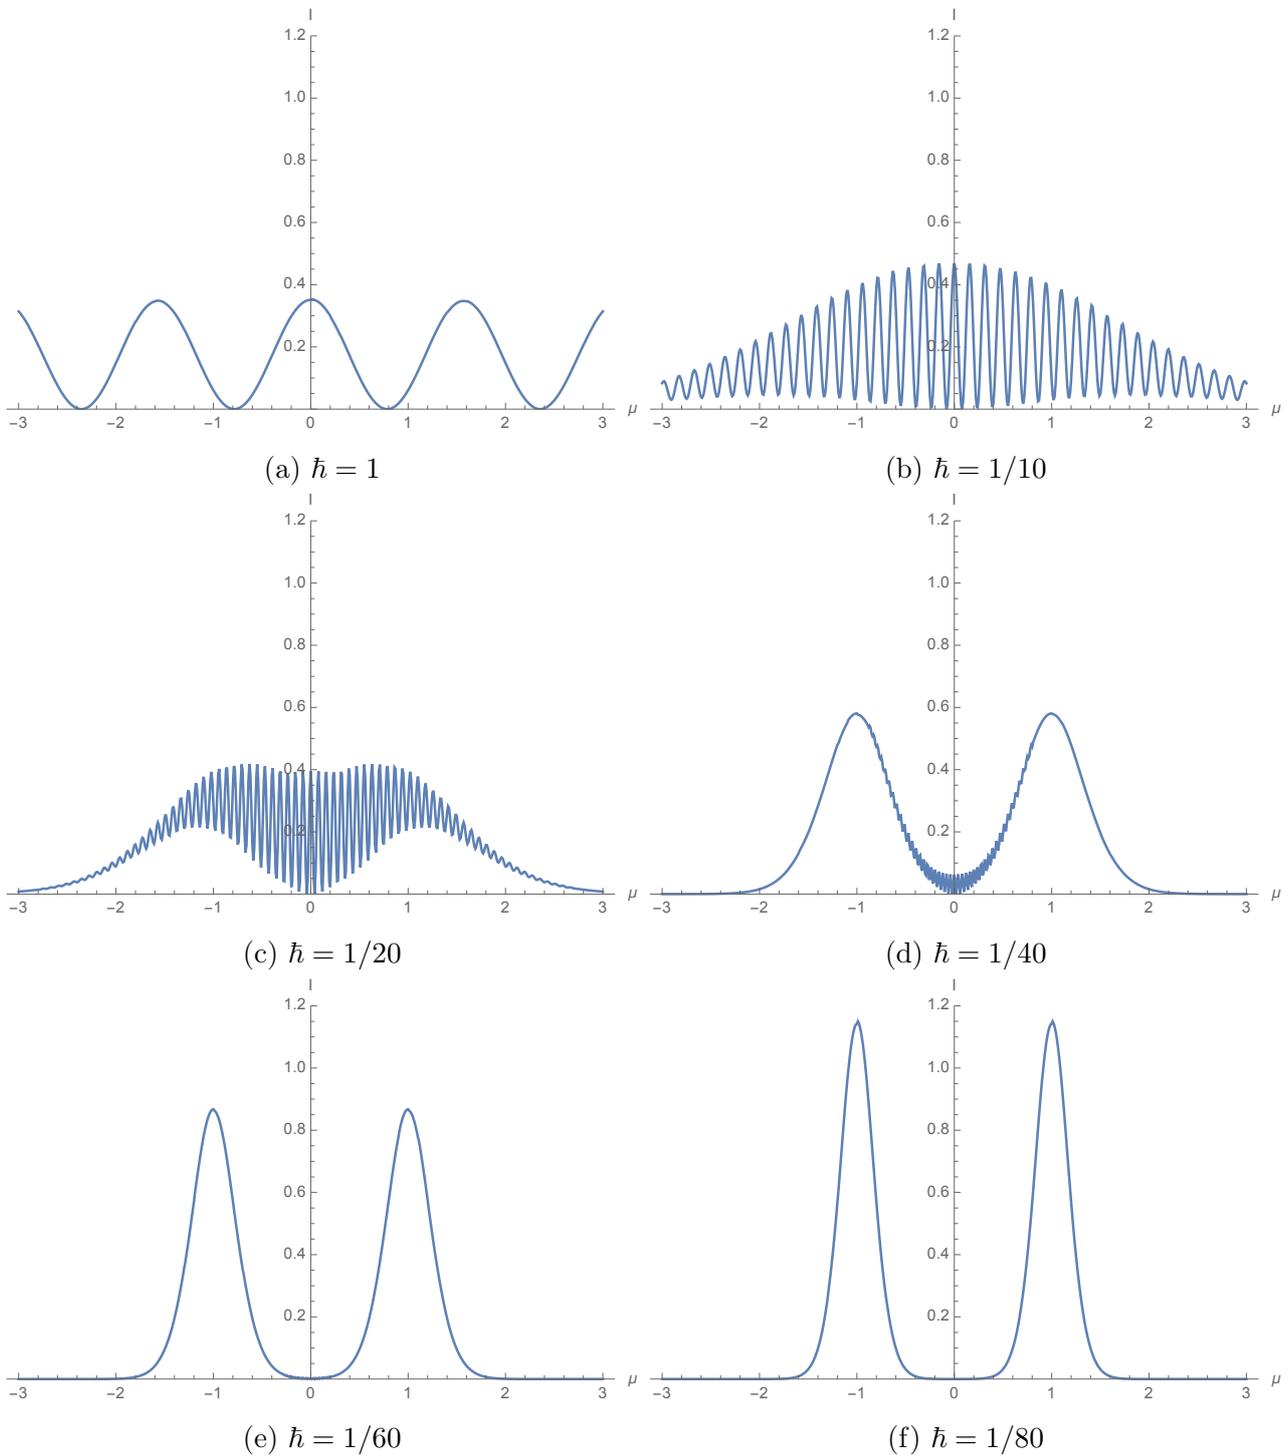


Figure 10.37: The Picard-Lefschetz diagram for the Young experiment for $\epsilon = 0.1$ from $\mu = -1.6, -1.5, -1, -0.5$, and $\mu = 0$, with the steepest ascent, descent contours (black) and thimbles (blue) corresponding to the saddle points (red) for $\hbar = 1$.

to the two slits. However, note that the integral is increasingly dominated by the two saddle points between the four poles. These two saddle points approach the real line at $x = s_1$ and $x = s_2$ in this limit.

Given the thimble, we can efficiently evaluate the oscillatory integral for various \hbar (see figure 11.38). For relatively large \hbar , the intensity on the screen is dominated by interference effects. For both $\hbar = 1$ and $\hbar = 1/2$ we do not observe the classical intensity peaks corresponding to the two slits. In the semi-classical limit, $\hbar \rightarrow 0$, the interference pattern is slowly replaced by the classical peaks. Note that this transition from the quantum to the classical regime cannot be studied in the traditional thin slit approximation.

Observe that, while the behavior of strong lenses is dominated by caustics where the saddle points become degenerate, the qualitative behavior of the double-slit experiment is completely determined by the Stokes transitions. The saddle points are everywhere non-degenerate and the h -function is a Morse function. In both instances, the saddle point approximation fails and the integral should be evaluated along the complete Lefschetz thimble. We expect this to be a generic feature in quantum mechanical interference phenomena.

Figure 10.38: The intensity I as a function of position μ for various \hbar

11 Proper time path integrals for gravitational waves: an improved wave optics framework by Braga, Ricciardone, et al.

I am very not sure, how useful this article is. It is very new, 2024, and looks like a summary of books and other articles.

When gravitational waves travel from their source to an observer, they interact with matter structures along their path, causing distinct deformations in their waveforms. In this study we introduce a novel theoretical framework for wave optics effects in gravitational lensing, addressing the limitations of existing approaches. We achieve this by incorporating the *proper time* technique, typically used in field theory studies, into gravitational lensing. This approach allows us to extend the standard formalism beyond the eikonal and paraxial approximations, which are traditionally assumed, and to account for polarization effects, which are typically neglected in the literature. We demonstrate that our method provides a robust generalization of conventional approaches, including them as special cases. Our findings enhance our understanding of gravitational wave propagation, which is crucial for accurately interpreting gravitational wave observations and extracting unbiased information about the lenses from the gravitational wave waveforms.

11.1 Introduction

Currently, I don't believe in detection of such effects.

below there is a very good literature review, but mostly about GL of GW. Still, the articles are maybe useful. Gravitational lensing of gravitational waves (GWs) is a rich and complex phenomenon that can provide valuable insights into the matter structures in the Universe. One of its intriguing aspects is the possible manifestation of *wave optics* (WO) effects: frequency-dependent modulations in the observed waveforms due to interference and diffraction, occurring when the wavelength λ of the GW is comparable to the Schwarzschild radius \mathcal{R}_s of the gravitational lens. For GWs in the LISA frequency band, these effects are expected to become important for lenses with mass smaller than $\lesssim 10^8 M_\odot$ and for sources at redshift $z_s \sim 1$, with event rates of $(0.1 - 1.6)\%$ for black-hole binaries in the mass range of $(10^5 - 10^{6.5}) M_\odot$ [18, 2]. For ground-based interferometers, the possibility of detecting wave-related effects as been assessed in [3, 4], confirming that these modulations can be detected for events with high enough signal-to-noise ratio. Given the broadness of the GW frequency spectrum, and the expected increasing in the number of detections with future interferometers [5, 6, 7], wave optics effects in GW lensing will become an important investigation tool. The body of literature exploring the possibility of using WO signatures in GW waveforms for various purposes is growing, given the importance of having more accurate waveforms [8]. These purposes include lens parameter estimation [18, 32, 10], uncovering intermediate-mass black holes [11], investigating the matter power spectrum on small scales [10, 12, 13, 14, 15, 16, 17, 18], and constraining the primordial black hole abundance [19, 20, 3, 22, 23, 24, 25].

The current body of literature about wave optics effects is based on the so-called *diffraction integral*, derived in the context of gravitational wave's lensing in [26] and further developed in subsequent works [27, 18, 23, 28, 29, 30]. this is not quite true. [27] was before and others are only about GW. The formalism is established also for electromagnetic radiation [31, 32], and its characteristic trait is to describe WO in a similar language as quantum mechanics. Indeed, in these works the wave equation is recast into a Schrödinger-like one, whose solution is given in terms of a path integral where the inverse frequency of the wave, $1/\omega$, takes the role of \hbar [33, 34]. Wave optics phenomena are thus reinterpreted under the light of interference effects between the

multiple paths that the wave can take, and, the laws of optics in the high-frequency regime ($\omega \rightarrow \infty$), correspond to the classical trajectories found in the $\hbar \rightarrow 0$ limit. This analogy between optics and quantum mechanics is very well known in the literature of electromagnetic waves, and is called *Hamilton analogy* [34, 35]. Another avenue to recover the diffraction integral is the one pursued in [12], where the authors derive its expression from assigning a particle Hamiltonian, coming from the knowledge of a dispersion relation, and building a path integral description from it. Both derivations hinge on the fundamental assumption of being able to describe the wave effectively as a collection of (quantum) particles: the path integral is not performed over all possible field's configuration as in quantum field theory, rather over trajectories as in quantum mechanics. This fact is reflected also in the type of assumptions the diffraction integral relies on. Two of these are the *eikonal* and *paraxial* approximations [37, 38, 39], which resemble the Wentzel–Kramers–Brillouin (WKB) semiclassical limit of quantum mechanics. In order, the eikonal approximation requires that the modifications of the waveform induced by the lens occur on scales larger than the wavelength of the wave. In a similar spirit, the paraxial approximation assumes that any orthogonal displacement from the direction of propagation is negligible within a period. The direct consequence of both of these two assumptions, is to individuate a direction of propagation of the wave, and consequently, to define the wavefronts, at the price of introducing a lower bound in frequency on the validity of the diffraction integral. Note that the knowledge of wavefronts is equivalent to the assignment of a dispersion relation. Given the significant cosmological and astrophysical insights that can be gained from studying the interaction between waves and matter structures, it is important to move beyond the limitations of the diffraction integral and produce a formalism that remains flexible across various length scales, so that we can include a greater range of events in the analysis.

The first aim of this paper is to generalize the treatment of the diffraction integral, going beyond the eikonal and paraxial assumptions. We do so by introducing the *proper time technique* or *worldline formalism* in the field of gravitational wave's lensing. This approach, sometimes referred to as *first quantised worldline*, was introduced by Feynman and has been developed to formally address the path-integral solution of the Helmholtz equation [40, 37, 38] and it has been used in various fields: ranging from quantum electrodynamics [41, 42, 43, 44], quantum field theory [45, 46, 47, 48, 49, 50, 51, 52, 53, 54] and quantum cosmology [55, 56, 57, 58] under the name of *Schwinger proper time*, to light optics with the *Feynman-Fradkin* path integral representation [39, 59, 60, 61, 62, 63], and also in seismology for acoustic wave propagation [64, 65]. The proper time technique sees the addition of a new time-like variable, in terms of which the wave equation can be recast *exactly* into a Schrödinger-like one, with an initial-value constraint, and without the need of any approximation. The advantage of this technique, therefore, is to provide an effective description of the fully relativistic fields in terms of associated particles, with their trajectories and associated momenta. The additional temporal parameter is removed at the end upon integration. The first main result of this paper is to provide the worldline representation of the gravitational wave's propagator describing their dynamics during a lensing event. This is given as the sum over all possible values of the proper time of particle-like path integrals, generalizing the diffraction integral of [27]. The reason behind the need of a relativistic quantum mechanical description in terms of particles, rather than fields, lies in the way the high-frequency limit is taken. Working only at the field level, it is not straightforward to see the emergence of the laws of geometric optics (GO) in the high-frequency limit, were the waves can be very well described as a collection of particles. By assuming eikonal and paraxial approximation *a posteriori*, we show that our formalism encompasses the diffraction integral, proving its solidity as its generalization, as well as reproducing the classical laws of geometric optics, such as Snell's laws, the Fermat principle, and the geodesic equations. Not only, having a particle-like description for GWs in any optical regime is even more compelling after the first detection of the stochastic gravitational wave background (SGWB) from the pulsar timing ar-

rays collaboration [66, 67, 68, 69, 70]. The description of the incoherent superposition of many gravitational waves requires statistical tools, most of which are staged in the phase-space, either through a Boltzmann equation or line-of-sight techniques, where the waves are categorized by their positions and momenta [71, 72, 73, 74, 75, 76, 77, 78, 79, 80], rather than their field configurations. All of these methods rely on the geodesic equation, a typical feature of the high-frequency regime [81, 82] where the waves have well defined wavefronts and direction of propagation. The inclusion of wave optics effects via the introduction of a collisional term in the Boltzmann equation has been considered in [83, 84]. Our formalism aims at providing a comprehensive description of the anisotropies of the SGWB, while being truly agnostic about the optical regime.

The second goal of this work is to address another limitation of the diffraction integral, namely the neglecting of polarization effects. Indeed, in this approach GWs are treated as scalar fields, assuming that their amplitude will suffer the largest effects from the lensing events, while its polarization will be unaffected by it. This is a well known high-frequency result [81, 82], where the GW's polarization is parallel transported along the paths, and it is assumed to remain valid also in the WO regime. However, it is known that finite size effects can both modify the dynamics of tensor modes and also generate additional polarization content in the GW sector [85, 86, 87], questioning the validity of the scalar wave assumption. Geometric optics works in the far field zone ($\lambda/\mathcal{R}_s \ll 1$), where only radiative degrees of freedom are present, and neglects the near zone dynamics which also encompasses non-radiative components of the metric perturbation. Accounting for both of these two sectors could lead to additional interaction channels, producing a richer phenomenology in wave optics phenomena. To address the inclusion of polarization effects, similarly to [88, 89, 90], we model the lenses as black holes (BH) so that we can borrow directly the results of the long-standing field of BH perturbation [91] in our discussion. This research line aims at understanding the behavior of BH under an external perturbation predicting, among other results, their quasi normal modes spectra [92, 93] and the behavior of their tidal deformations encoded in the Love numbers [94, 95, 96]. In the case of binary systems, and especially for extreme mass ratio inspirals, BH perturbation tools are also employed to compute strong field correction to post-Newtonian results [97]. The master equation dictating the dynamics of fields perturbations of different spin on a Kerr background is the so-called Teukolsky equation [98], written in terms of Newman-Penrose (NP) scalars [99], proxies for the field perturbations. For instance, when studying $s = 2$ metric perturbations, the NP scalars are related to components of the perturbed Weyl tensor, while $s = 1$ perturbations to the perturbed Maxwell tensor. At linear order in perturbation theory, NP scalars are also gauge invariant [91], so that they are readily related to observables. In the last part of this work, we will apply the proper time formalism also to Teukolsky equation, showing how to include spin effects into the field of wave optics. Our result, therefore, constitute a first step into understanding the role of the interaction with matter on the polarization content of GWs in the context of lensing in the WO regime.

This paper is organized as follows: in Section 10.2 we review the derivation of the diffraction integral as done in [27], illustrating the role of the eikonal and paraxial approximations. In Section 10.3 we introduce the proper time technique for a Klein-Gordon field, and derive the worldline representation of the wave's propagator. Section 10.4 is dedicated to recover all the standard results of paraxial- and geometric optics, proving that our formalism encompasses them as sub cases. In Section 10.5 we set up the perturbative expansion of the Green function. In Section 10.6 we extend our formalism to massive scalar fields and we provide explicit computations for a Coulomb-like gravitational potential. Finally, Section 10.7 addresses the inclusion of spin effects for BH lenses.

11.2 The diffraction integral

In this section, we review the standard formalism of gravitational lensing of gravitational waves in the wave optics limits, as laid out in the pioneering work of [26] (see [27, 18, 23, 28, 29, 30] for other references), bringing into the context of gravitational lensing previous works addressing the analogy between optics and quantum mechanics under the paraxial approximation [100, 101, 102].

We consider the propagation of a wave through a Universe containing a lens in the Newtonian approximation

$$ds^2 = -(1 + 2\alpha U(\mathbf{x}))dt^2 + (1 - 2\alpha U(\mathbf{x}))d\mathbf{x}^2, \quad (11.1)$$

where U is the gravitational potential of the lens, assumed to be static $U = U(\mathbf{x})$. In the expression above we introduced α as the bookkeeping parameter representing the strength of the gravitational potential. In our conventions, therefore, U is of order ~ 1 and $\alpha \ll 1$. It is a

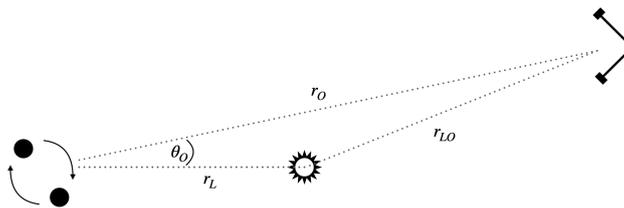


Figure 11.1: Lensing situation considered.

known fact that, in General Relativity, the polarization content of a GW is parallel transported in the geometric optics limit [81, 82]. In the standard wave optics literature, this result is carried over also in the WO regime, where the GW is treated as a scalar field, namely $\Psi(\mathbf{x}, t)$. Then, all the physics regarding the interaction between the GWs and the lens is accounted for in the GW's equations of motion on the curved background, namely $\bar{\square}\Psi(\mathbf{x}, t) = 0$, where the covariant derivative are taken with respect to the metric in Eq. (10.1). Solving the latter equation is non-trivial due to the non-linearity of the wave operator. One can use the small strength of the gravitational potential ($\alpha \ll 1$) to set up a perturbative expansion, and to study the interaction up to first order in α . This approximation is usually called *Born approximation* [30, 29]. Since we are considering a static background, we perform a Fourier transform in frequency domain with conventions

$$\Psi(\mathbf{x}, t) = \int \frac{d\omega}{2\pi} \tilde{\Psi}_\omega(\mathbf{x}) e^{-i\omega t}, \quad (11.2)$$

below there is a rewritten theory from Nakamura's original article. and obtain, at first order in α , the associated Helmholtz equation

$$[\nabla^2 + \omega^2(1 - 4\alpha U)] \tilde{\Psi}_\omega(\mathbf{x}) = 0, \quad (11.3)$$

where $\nabla^2 \equiv \partial_i \partial_i$. This equation is valid outside the wave's source, which we assume to be localized at some point \mathbf{x}_S , and generically address it as $\mathcal{S}(\mathbf{x}_S)$. Eq. (10.3) is solved by introducing the amplification factor F

$$F(\mathbf{x}) \equiv \frac{\tilde{\Psi}_\omega(\mathbf{x})}{\tilde{\Psi}_\omega^{NL}(\mathbf{x})}, \quad (11.4)$$

defined as the ratio between the lensed $\tilde{\Psi}_\omega$ and unlensed $\tilde{\Psi}_\omega^{NL}$ waves, namely the solution of Eq. (10.3) with $U = 0$. The Helmholtz equation, then, turns into an equation for the amplification factor

$$\frac{1}{r^2} \nabla_\theta^2 F + \frac{\partial^2 F}{\partial r^2} + 2i\omega \frac{\partial}{\partial r} F = 4\alpha\omega^2 U F, \quad (11.5)$$

in spherical coordinates $\mathbf{x} = \{r, \theta, \varphi\}$, where we have chosen $\tilde{\Psi}_\omega^{NL}(\mathbf{x}) = e^{-i\omega r}/r$, and ∇_θ^2 is the Laplacian operator on the 2-D sphere. It is important to point out that the latter equation has been derived without any assumption on the frequency of the GW. Therefore, in the high-frequency regime, one should recover the standard GO results, while in the opposite limit we expect to have interference and diffraction patterns before reaching the fully diffractive regime [31] where the amplification is suppressed. Eq. (10.5) allows setting up the previously mentioned analogy between geometric optics (GO) and wave optics (WO) on one side, and classical and quantum mechanics on the other. Indeed, if one is able to remove the second derivative term from it, then the latter becomes a Schrödinger-like equation where the radius takes the role of the time variable. This way, one could solve Eq. (10.5) via a path integral description, which nicely accounts for wave-effects as interference between the multiple propagation paths taken by the GW, in a similar fashion of quantum effects in quantum mechanics. Throughout the entire paper, we will make such analogy more and more explicit, and exploit it as much as possible to borrow tools from known literature, and bring them in the realm of wave optics effects in gravitational lensing. In this review section, we proceed as in [26], and work out the particle-like path integral solution of Eq. (10.5). First of all, we have to reduce a second order partial differential equation, to a first order, Schrödinger-like, one. This can be achieved by neglecting the second order derivative term, with respect to the one proportional to the first order derivative,

$$|\partial_r^2 F| \ll |2i\omega \partial_r F|. \quad (11.6)$$

If we call \mathcal{R} the typical variation scale of F (which is related to the typical length scale of the lens' potential) in the direction of motion, then Eq. (10.6) would imply $\mathcal{R} \ll \omega$. In [26], the authors call this assumption the *eikonal* approximation, as it effectively prescribes a hierarchy between the length scales associated to the GW and those of the lens. It shares also similarities with the non-relativistic limit, except that the role of the time is now assumed by the radial coordinate, and the one of the particle's mass by the frequency. We also require that the wave's amplitude does not change much in the θ direction, namely the angle spanning from the line of sight, and that the observer is located at $\theta_O \ll 1$ (see Fig. 10.1). This assumption, combined with Eq. (10.6), gives us the *paraxial approximation*. It allows us to consider the two vector $\boldsymbol{\theta} = \{\theta, \varphi\}$ as a two dimensional vector on a flat plane and $\sin \theta \approx \theta$ and to treat the system as a 2-dimensional one, evolved in "time" with the radial coordinate. As a consequence of Eq. (10.6), Eq. (10.5) takes the following form

$$i\partial_r F = -\frac{1}{2\omega} \partial_\theta^2 F + 2\alpha\omega U F, \quad (11.7)$$

where $\partial_\theta^2 = \partial^2/\partial\theta^2 + \theta^{-1}\partial/\partial\theta + \theta^{-2}\partial^2/\partial\varphi^2$.

We just obtained a Schrödinger-like equation, where r plays the role of time and ω the one of the particle mass. The corresponding Lagrangian, which yields to the classical equation of motion, is given by:

$$L[\boldsymbol{\theta}, \dot{\boldsymbol{\theta}}, r] = \left[\frac{r^2}{2} |\dot{\boldsymbol{\theta}}|^2 - 2\alpha U(r, \boldsymbol{\theta}) \right], \quad (11.8)$$

where $\dot{\boldsymbol{\theta}} = d\boldsymbol{\theta}/dr$. We point out that other analogous assumptions have been made in literature in order to reduce the order of the differential equation of F . For instance in [33], the authors introduce the first Fresnel radius $R_F = \sqrt{\lambda D}$, where D is the distance between the source and a given plane, while λ is the wavelength of the wave with $\lambda \ll D$. Then, Eq. (10.5) turns into a first order differential equation by requiring $R_F \gg \lambda$. The *diffraction integral* is the path integral solution of Eq. (10.7)

$$F(\vec{r}_O) = \int \mathcal{D}\boldsymbol{\theta}(r) \exp \left\{ i\omega \int_0^{r_O} dr L[\boldsymbol{\theta}(r), \dot{\boldsymbol{\theta}}(r), r] \right\}, \quad (11.9)$$

as done in quantum mechanics [103], and where \mathbf{r}_O is the position of the observer. Therefore, the amplification factor is described in terms of a path integral, where the action weighting each path is built from the particle Lagrangian in Eq. (10.8). The diffraction integral in Eq. (10.9) shows manifestly that the inverse frequency $1/\omega$ is the parameter taking the place of \hbar and formally establishes the analogy between WO effects and quantum mechanical ones. The laws of GO, which are valid in the $\omega \rightarrow \infty$ limit, are interpreted as the counterparts of the classical limit $\hbar \rightarrow 0$, where one or more classical trajectories dominate the path integral solution. These trajectories, namely the *rays* along which the wave propagates, are the one that satisfies the Euler-Lagrange equations descending from the Lagrangian in Eq. (10.8). Conversely, for smaller values of the GW's frequency, multiple path contribute to the final amplitude, producing interference patterns. One can bring Eq. (10.9) in a more familiar form under the *thin lens* assumption [31], namely that the lens extension in the direction of propagation is small compared to the one on the plane orthogonal to it. In this case, $\alpha U(r, \boldsymbol{\theta}) = \frac{1}{2}\delta(r - r_L)\hat{\psi}(\boldsymbol{\theta})$, and the diffraction integral takes the form

$$F(\vec{r}_O) = \int d^2\boldsymbol{\theta}_L \exp \left\{ i\omega \left[\frac{r_L r_O}{2r_{LO}} |\boldsymbol{\theta}_L - \boldsymbol{\theta}_O|^2 - \hat{\psi}(\boldsymbol{\theta}_L) \right] \right\}, \quad (11.10)$$

where r_L is the distance to the lens and $\boldsymbol{\theta}_L$ is the two vector in the lens plane. Note that the path integral has been reduced to a standard 2-dimensional integral under the assumption that the paths contributing the most to the diffraction integral are constant $\boldsymbol{\theta}(t) \approx \boldsymbol{\theta}_L$ (see [27]). The integrand of Eq. (10.10) is the *time delay* function

$$t_d \equiv \frac{r_L r_O}{2r_{LO}} |\boldsymbol{\theta}_L - \boldsymbol{\theta}_O|^2 - \hat{\psi}(\boldsymbol{\theta}_L), \quad (11.11)$$

where $\hat{\psi}(\boldsymbol{\theta}_L)$ is the deflection potential associated to the lens [33, 104, 105]. In the high frequency regime, only the stationary points of Eq. (10.11) contribute to the diffraction integral. These points are obtained by solving the differential equation [31]

$$\nabla_{\boldsymbol{\theta}_L} t_d = 0, \quad (11.12)$$

which can be recast in the more familiar form:

$$\boldsymbol{\theta}_O = \boldsymbol{\theta}_L - \frac{r_{LO}}{r_L r_O} \nabla_{\boldsymbol{\theta}_L} \hat{\psi}. \quad (11.13)$$

Eq. (10.13) is the *lens equation* [31, 106, 11, 107], and determines the position of images that can be produced by a lensing event with an associated strong enough gravitational potential.

11.2.1 Field vs particle description

Before delving into the main body of this work, we aim to address in more details the significance of the approximation in Eq. (10.6) and the paraxial assumption. Firstly let us point out that there is a disagreement in literature concerning what authors means for *eikonal* and *geometric optics* limits. Some consider the two description equivalent, while others attribute to the first coherent events and incoherent lensing to the second [108, 33]. In the analogy between quantum- and wave- effects, this differentiation divides interference events, occurring only if the trajectories in the path integral "sum" their probability amplitude in phase, from incoherent events where only the classical trajectory remains. In our work, we take Eq. (10.6) as the definition of the eikonal approximation, and reserve the geometric optics limit, for the saddle point approximation of the path integral. Therefore, the eikonal approximation is a half way step: it does not fully coincide with GO (formally valid when $\omega \rightarrow \infty$), but it also requires a

lower bound on the possible frequency values. This is because, the interference effects allowed in this intermediate regimes may still be represented in terms of the dynamics of de-phased wavefronts. However, to have wavefronts in the first place, a field must be describable in terms of associated particles (wavefronts define dispersion relations [12, 34]), regardless of whether their probability amplitudes of transitioning from one spacetime point to another, are summed coherently (as waves), or incoherently (as particles).

The diffraction integral Eq. (10.10) also hinges on the paraxial assumption, by considering the propagation occurring mainly along a given direction, with slow (compared to the wavelength of the wave) deviations from it [109]. Therefore, we can summarize the physical content of the assumptions behind the diffraction integral as the requirements of having a i) slowly changing direction of propagation and ii) wavefronts.

Our goal is to go beyond the diffraction integral, namely to understand the behavior of GWs when these are treated as *fields*, rather than *quantum particles*, in the Hamilton analogy.¹⁷ We expect these two descriptions not to be entirely equivalent: transitioning from a relativistic wave equation to a Schrödinger-like one (e.g., assuming Eq. (10.6)) requires certain approximations, excluding possible interesting phenomenology. After considering these points, one might argue for solving the Helmholtz equation directly using field theory methods, namely in terms of a path integral over field's configurations rather than trajectories in configuration space. Indeed, the analogy between quantum effects and wave optics extends to the field representation of the wave, and it does not exclusively belong to the particle description. However, it's important to emphasize the merit and significant advantages of the particle-like description (in terms of trajectories and momenta rather than field variables). Two reasons stand out prominently. The first one regards the limit toward the GO regime, where the wave's description naturally transitions to one in terms of effective associated particles propagating along the *rays* [81, 82]. Much of the standard lensing literature in GO is derived from solving the ray equation, thereby describing waves in terms of coordinates and momenta, compatibly with the fact that in the classical limit the uncertainty principle is trivial. Hence, having a generalized particle-like description of the field allows a direct high frequency limit toward the well-known geometric optics laws, contrary to a field approach, where this connection is more obscure. Secondly, as point out in the Introduction, a particle-like description allows setting up a phase-space where the statistical study of the stochastic gravitational wave background is usually staged.

11.3 Proper time path integral for a scalar wave

In this section we present our new formalism to generalize the path integral solution of the amplification factor without imposing the eikonal and paraxial approximations of Eq. (10.6). Firstly, we do so in the case of massless scalar fields, and then we generalize to the case of a massive scalar field, and a massless field with spin in the next Sections. Our starting point is once again Helmholtz Eq. (10.3), which is completely general. We solve it in terms of the associated Green's function $G_\omega(\mathbf{x}_f, \mathbf{x}_i)$ satisfying

$$[\nabla^2 + \omega^2(1 - 4\alpha U)] G_\omega(\mathbf{x}_f, \mathbf{x}_i) = \delta^{(3)}(\mathbf{x}_f - \mathbf{x}_i) . \quad (11.14)$$

In order to continue having a particle-like description of the GW, we have to reduce Eq. (10.14) to a Schrödinger-like equation as well. To achieve this description, we introduce the proper time technique in the realm of gravitational lensing, instead of taking the route illustrated in Section 10.2. This is a well-known method in the literature, which prescribes the introduction of a new time-like parameter, τ , called *proper time* in a specific way, so that Eq. (10.14) becomes

¹⁷We stress that all the phenomena considered in this work are strictly classical, we only rely on $1/\omega \leftrightarrow \hbar$ as an analogy to borrow tools from known literature.

a first order equation in τ , hence allowing an effective description of the GWs through an associated particle in a higher dimensional spacetime.

11.3.1 Adding a new time parameter

We apply the proper time method in the case of waves propagating through a spacetime with a lens. The main advantage of this methodology is that it allows a direct interpretation of fields in terms of associated quanta, hence providing a description of the field in terms of paths rather than field configurations. Here, we propose it for the first time to our knowledge, in a classical gravitational lensing context. Among the references cited in the Introduction, we mainly follow [40, 63, 39].

In the proper time method, the system's dynamics is studied in a 5-dimensional (???) spacetime and it is captured by the Green function $\tilde{G}_\omega(\mathbf{x}_f, \mathbf{x}_i, \tau)$. This is related to the 4-dimensional propagator through its definition

$$G_\omega(\mathbf{x}_f, \mathbf{x}_i) \equiv -\frac{i}{\omega} \int_0^\infty d\tau e^{i\omega\tau} \tilde{G}_\omega(\mathbf{x}_f, \mathbf{x}_i, \tau). \quad (11.15)$$

By inserting this form of G_ω into its equation of motion, namely Eq. (10.14), and performing an integration by part, one finds

$$\tilde{G}_\omega(\mathbf{x}_f, \mathbf{x}_i, 0) + \int_0^\infty d\tau e^{i\omega\tau} \left[\frac{\partial}{\partial\tau} - \frac{i}{\omega} \nabla^2 + 4i\alpha\omega U \right] \tilde{G}_\omega(\mathbf{x}_f, \mathbf{x}_i, \tau) = \delta^{(3)}(\mathbf{x}_f - \mathbf{x}_i). \quad (11.16)$$

Finally, imposing the initial condition $\tilde{G}_\omega(\mathbf{x}_f, \mathbf{x}_i, 0) = \delta^{(3)}(\mathbf{x}_f - \mathbf{x}_i)$, Eq. (10.14) reduces to a first-order differential equation in terms of the proper time τ

$$\frac{i}{\omega} \frac{\partial}{\partial\tau} \tilde{G}_\omega(\mathbf{x}_f, \mathbf{x}_i, \tau) = -\frac{1}{\omega^2} \nabla^2 \tilde{G}_\omega(\mathbf{x}_f, \mathbf{x}_i, \tau) + V(\mathbf{x}) \tilde{G}_\omega(\mathbf{x}_f, \mathbf{x}_i, \tau), \quad (11.17)$$

where we have defined

$$V(\mathbf{x}) \equiv 4\alpha U(\mathbf{x}). \quad (11.18)$$

Eq. (10.17) is equivalent to a Schrödinger equation for a particle moving in a potential $V(\mathbf{x})$. Note that we have obtained it without the eikonal and paraxial approximations, hence Eq. (10.17) is completely general and valid for every frequency. The time evolving parameter is now played by the proper time τ , differently from the diffraction integral where the trajectories were parametrized by the radial coordinate. As in Section 10.2, also in Eq. (10.17) $1/\omega$ plays the role of \hbar , implying that the limit of high frequency $\omega \rightarrow \infty$ can be treated as the semi-classical limit in quantum mechanics, namely the limit of $\hbar \rightarrow 0$. The fact that the Hamiltonian analogy is valid also in our formalism, tells us that we expect to recover the GO description [81, 82] in the high-frequency limit also from Eq. (10.17). In Section 10.4 we will show that, by taking the eikonal and paraxial approximation *a posteriori*, our formalism is the proper generalization of the WO analysis summarized in Section 10.2, proving its solidity as a generalization of the standard GO and WO formalism.

Having recovered a Schrödinger-like equation, we can now express the associated solution in the path integral representation [40]

$$\tilde{G}_\omega(\mathbf{x}_f, \mathbf{x}_i, \tau) = \int_{\mathbf{x}(\tau'=0)=\mathbf{x}_i}^{\mathbf{x}(\tau'=\tau)=\mathbf{x}_f} \mathcal{D}\mathbf{x}(\tau') e^{i\omega S}, \quad (11.19)$$

where S is interpreted as the associated particle action, taking the form

$$S[\mathbf{x}, \dot{\mathbf{x}}] \equiv \int_0^\tau d\tau' L[\mathbf{x}(\tau'), \dot{\mathbf{x}}(\tau')], \quad \text{with} \quad L = \frac{1}{4} \left(\frac{d\mathbf{x}(\tau')}{d\tau'} \right)^2 - V(\mathbf{x}(\tau')). \quad (11.20)$$

According to Eq. (10.15), the original propagator is then given by

$$G_\omega(\mathbf{x}_f, \mathbf{x}_i) = -\frac{i}{\omega} \int_0^\infty d\tau e^{i\omega\tau} \int_{\mathbf{x}(\tau'=0)=\mathbf{x}_i}^{\mathbf{x}(\tau'=\tau)=\mathbf{x}_f} \mathcal{D}\mathbf{x}(\tau') e^{i\omega S}. \quad (11.21)$$

Being the proper time τ integrated over, it does not enter the system Green's function, as expected since this was introduced as a mathematical device. The equivalent form of $G_\omega(\mathbf{x}_f, \mathbf{x}_i)$ in quantum field theory is known as the *worldline representation* of the propagator, and it has been used as an alternative and equivalent avenue to formalize perturbation theory. In this formalism, one focuses on a 1-dimensional field theory, where the particle position x^μ is seen as a set of four fields living on the one dimensional space of the proper time τ , called *worldline* [47, 51, 45]. Namely, one sees the worldline description as the (relativistic) quantum mechanics of the particles that are the quanta of the associated fields in the field theory. Eq. (10.21) suggests the following interpretation for the propagator $G_\omega(\mathbf{x}_f, \mathbf{x}_i)$: it describes the probability for the particle associated to $\tilde{\Psi}_\omega(\mathbf{x})$ to propagate from a given initial point located at \mathbf{x}_i , to the final one at \mathbf{x}_f , in a fictitious time τ . The integral over all possible values of τ means that we have to consider all paths from \mathbf{x}_i to \mathbf{x}_f with duration $0 < \tau < +\infty$. In other words, it allows for the motion to occur at different “velocities”.

We also point out that the Green function in Eq. (10.21) is defined up to a dimensionless constant, which can be fixed by physical arguments, such that the propagator in the absence of the potential reduces to the standard Green function of the wave operator. Also the boundary conditions are yet fixed. Since we are working in frequency space, one can pick for instance retarded boundary conditions by continuing $\omega \rightarrow \omega + i\epsilon$ upon performing the Fourier integral.

Although the Feynman-Fradkin representation of the propagator is an exact form, much of its physics is embodied in the last τ - integration, and before performing it, it is hard to draw conclusion about the dynamics of the system [37]. Because of this, it is sometimes referred to as an *indirect* representation. On the contrary, direct representations of the Helmholtz propagator, where the integral over the proper time is removed, are usually approximate in nature. Indeed, to do so one would have to exchange the integration order between the proper time and the path integral, but this is not possible in an exact fashion because of the boundary conditions. We shall see that one can exchange the integration order, or in any case perform the integral over τ , in the high frequency regime, where ω becomes large [110, 60]. Since the aim of this paper is to introduce in the context of GW lensing the proper time technique for the first time, we only show its connection to the existing literature, as its limit toward the high-frequency regime. Nonetheless, our propagator encompasses much more physics than that. Because Eq. (10.21) is an exact solution, it can be used as starting point to investigate the behavior of GWs near caustics, tunneling and many more effects [39, 12, 104, 111]. Understanding these phenomena will be the object of future work.

11.3.2 Hamiltonian derivation

The proper time path integral representation of the Green function Eq. (10.21) can also be recovered from an Hamiltonian point of view. We can identify $G_\omega(\mathbf{x}_f, \mathbf{x}_i)$, obeying Eq. (10.14), as the matrix element

$$G_\omega(\mathbf{x}_f, \mathbf{x}_i) = \langle \mathbf{x}_f | [\nabla^2 + \omega^2(1 - 4\alpha U)]^{-1} | \mathbf{x}_i \rangle, \quad (11.22)$$

taking the form

$$G_\omega(\mathbf{x}_f, \mathbf{x}_i) = -\frac{i}{\omega} \int_0^\infty d\tau \langle \mathbf{x}_f | e^{i\omega\tau\hat{\mathcal{H}}} | \mathbf{x}_i \rangle, \quad (11.23)$$

where the particle Hamiltonian is given by

$$\hat{\mathcal{H}} \equiv -p^2 + (1 - 4\alpha U), \quad (11.24)$$

with $\mathbf{p} \equiv i\omega^{-1}\nabla$ and $p^2 = \mathbf{p} \cdot \mathbf{p}$. As usual (see e.g. [31]), the propagation of a wave through the background curved by the lens can also be seen as the problem of a wave propagating through a medium with refractive index

$$n^2 \equiv 1 - 4\alpha U, \quad \rightarrow \quad n \simeq 1 - 2\alpha U \quad (11.25)$$

hence the particle Hamiltonian would be given by $\hat{\mathcal{H}} = -p^2 + n^2$. Eq. (10.23) clearly shows that the proper time allows for a particle-like description of a field. Indeed, the matrix element inside the integral is particle propagator, given in terms of a particle path integral, while the field propagator is $G_\omega(\mathbf{x}_f, \mathbf{x}_i)$, expressed as an integral over the proper time. Despite that a field description and a particle-one are not equivalent, one can achieve a particle-like description of the waves with the aid of the proper time. The Hamiltonian derivation of the field propagator shows that one can introduce the additional time-like variable exploiting first quantization tools and not just embedding techniques as in Eq. (10.15). Moreover, it nicely shows that the gravitational potential U can be interpreted as the *self-energy*, dressing the free propagator in momentum space, and supports its identification with the index of refraction. In Section 10.5.4 we will also show that it satisfies Dyson equations.

Note that this Hamiltonian derivation of the worldline representation of the propagator is also used in [12] to derive the diffraction integral. However, the proper time is fixed by considering fixed average energy configurations. Indeed, the Green function in Eq. (10.22) can be interpreted as the *fixed energy propagator* [39].

11.4 The high-frequency limit

In this section we prove that the formalism we set up is able to recover both the results of WO literature (as described in Section 10.2) and the GO laws of propagation by taking the high-frequency limit of Eq. (10.21).

11.4.1 Geometric Optics: the ray equation

In the limit in which $\omega \rightarrow \infty$, the complex exponential in Eq. (10.21) becomes highly oscillating and the main contribution to the path integral comes from the paths that are close to the minima of the phase. Therefore, the rays of the GO description emerge naturally in the high frequency limit.

We start by considering the total action

$$W[\mathbf{x}, \dot{\mathbf{x}}, \tau] \equiv \tau + \int_0^\tau d\tau' L[\mathbf{x}(\tau'), \dot{\mathbf{x}}(\tau')], \quad (11.26)$$

which is a function of both the proper time and the trajectories. The function W is also known as the *Hamilton characteristic function* [34, 35]. The equation of motion of τ gives the constraint

$$\frac{\partial W}{\partial \tau} = 1 - \left[\frac{1}{4} \left(\frac{d\mathbf{x}(\tau')}{d\tau'} \right)^2 + V(\mathbf{x}(\tau')) \right] = 0, \quad (11.27)$$

selecting a specific value of τ as a function of the trajectory and the boundary conditions, which we call τ_{cl} (this enters in the boundary condition of the trajectory). Considering, for instance, a free particle, the only path must satisfy

$$\ddot{\mathbf{x}}_{free}(\tau') = 0, \quad 1 - \frac{1}{4}|\dot{\mathbf{x}}_{free}|^2 = 0, \quad (11.28)$$

with boundary conditions $\mathbf{x}_{free}(\tau' = 0) = \mathbf{x}_i$ and $\mathbf{x}_{free}(\tau' = \tau) = \mathbf{x}_f$, implying $\dot{\mathbf{x}}_{free} \equiv (\mathbf{x}_f - \mathbf{x}_i)/\tau$, and

$$\tau_{free} = \frac{|\mathbf{x}_f - \mathbf{x}_i|}{2}. \quad (11.29)$$

Eq. (10.27) is related to the conservation of energy. One might have thought that computing the variation of the phase with respect to the proper time τ , would have lead to a constraint on the associated Lagrangian L in Eq. (10.20). However, this is not the case as in changing τ , one also changes the path along which the action S is evaluated. Specifically, it can be shown [103, 39] that

$$\frac{\partial W}{\partial \tau} = -\hat{\mathcal{H}}, \quad (11.30)$$

where $\hat{\mathcal{H}}$ is the associated Hamiltonian of Eq. (10.24). Indeed, one can derive Eq. (10.27) also from $\hat{\mathcal{H}} = 0$ by plugging the expression of the velocity $\dot{\mathbf{x}}$, related to the momenta through Hamilton equations $\dot{\mathbf{x}} = \partial\hat{\mathcal{H}}/\partial\mathbf{p}$, in $\hat{\mathcal{H}}$. Therefore, the Hamiltonian constraint above selects those trajectories where the total energy is conserved (see Figure 10.2), among all the possible paths. The fact that it is equivalent to the vanishing of the Hamiltonian, can be also interpreted as the assignment of a dispersion relation, which defines the wavefronts, since $\hat{\mathcal{H}} = 0$ implies $p^2 = n^2$. This is also supported by the fact that Eq. (10.27) is the Hamilton-Jacobi equation for the particle action S . We can also give an interpretation of this result in light of the eikonal approximation described in Section 10.2. The diffraction integral has among its basic assumptions the one of transforming a second order (in r) partial differential equation, to a first order one. We also commented that such a procedure is similar to taking the non-relativistic limit from a Klein-Gordon to a Schrödinger one. Hence, we can read the fact that removing the proper time fixes one particular τ parametrization under the light of the non-relativistic approximation.

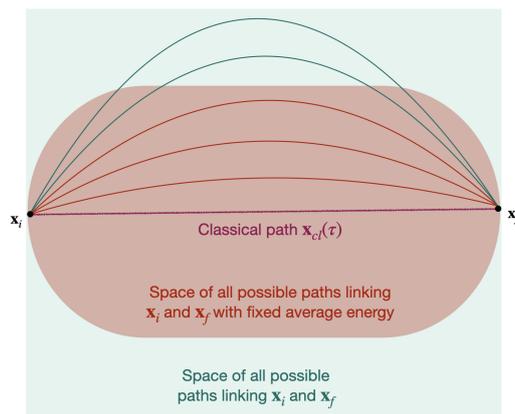


Figure 11.2: Among all the possible paths (green region), the constraint in Eq. (10.27) selects only those paths connecting \mathbf{x}_i and \mathbf{x}_f with fixed average energy. The classical path satisfies $\delta\hat{S}/\delta x = 0$.

Note that the constraint in Eq. (10.27), is satisfied on the classical trajectory up to a constant. The Euler-Lagrange equations of the particle's action (10.20), pick out the path

$$\ddot{\mathbf{x}}_{cl} = -2\nabla V[\mathbf{x}_{cl}], \quad (11.31)$$

and the time derivative of the Hamiltonian constraint

$$\frac{\partial^2 W}{\partial \tau^2} = \dot{\mathbf{x}} \cdot \left[\frac{\ddot{\mathbf{x}}}{2} + \nabla V[\mathbf{x}] \right], \quad (11.32)$$

vanishes on the classical trajectory in virtue of the Euler-Lagrange Eq. (10.31). Therefore, on \mathbf{x}_{cl} the constraint Eq. (10.27) is already implemented up to a constant. This is the usual result stating that only for the on-shell motion energy is well defined (begin constant in time). However, it also means that the the knowledge of \mathbf{x}_{cl} does not fix τ_{cl} because the constraint equation is only recovered up to an arbitrary constant.

We can rewrite the total action W in a more familiar form by chaining the path's parametrization from the proper time to the path's length $dl = \sqrt{\dot{\mathbf{x}} \cdot \dot{\mathbf{x}}} d\tau$, leading to

$$W = \int_0^{l_{cl}} dl' \sqrt{1 - 4\alpha U(\mathbf{x}(l'))} = \int_0^{l_{cl}} dl' n(\mathbf{x}(l')), \quad (11.33)$$

where n is the refractive index defined in Eq. (10.25) and where we have used the constraint in Eq. (10.27). We can use this result to find the equations of the trajectories, namely the rays of GO. These follow from the variational principle applied to Eq. (10.33)

$$\delta \left(\int_0^{l_{cl}} dl' n(\mathbf{x}(l')) \right) = 0, \quad (11.34)$$

leading to the known Fermat principle [31, 32, 34], a fundamental part of optics in the high frequency regime. Note that the knowledge of a ray, gives automatically the phase space variables: a unique trajectory has a well defined position (\mathbf{x}^μ) and momentum ($d\mathbf{x}^\mu/dl$) vectors at all points.

11.4.2 WKB form of the propagator

The results just found concern the stationary points of the path integral action. We can do better and study the contributions of the first non-null fluctuations around the classical configuration, which will dictate how the field's amplitude changes during the propagation. As we will show, these contributions are not suppressed by a factor $1/\omega$ in the final results, they are an integrating part of the GO result. Following [39, 60, 111], we start the analysis from the unconstrained particle action S (rather than W) in Eq. (10.20) and perform a saddle-point approximation first in the trajectories, and then in the proper time. The procedure leads to

$$G_\omega^{\text{WKB}}(\mathbf{x}_f, \mathbf{x}_i) \approx \mathcal{N} \sum_{IJ} \sqrt{-\frac{\det[\partial^2 S_{cl}/\partial x_i^a \partial x_f^b]}{\partial^2 S_{cl}/\partial \tau^2}} \Big|_I e^{i\omega W[\mathbf{x}_{cl}^I, \tau_{cl}^{IJ}] - i\pi n^I/2}, \quad (11.35)$$

where we have used the Van Vleck-Morette determinant [39]. The sum over I , corresponds to the sum over different classical configurations, i.e., different images, if Eq. (10.31) has multiple solutions, while the one over J accounts for possible different values of proper time fixed by the constraint Eq. (10.27) for each image. \mathcal{N} is a dimensionless normalization constant that has to be fixed, and the factor n^I is the number of negative eigenvalues of the Hessian matrix of the image I , also related to the number of focal points lying on the classical trajectory [39], as stated by Morse theory. The amplitude of the WKB propagator can, in principle, diverge. These can occur on caustics (infinities of the Van Vleck-Morette determinant) or at the zeros of the denominator. Concerning the first, these are lines of infinite magnification, signaling the breaking of the saddle-point approximation when two classical trajectories merge into each

other. On the other hand, the zeros of $\partial^2 S_{cl}/\partial\tau^2$ are found in correspondence of the turning points of the trajectory, namely when $\dot{\mathbf{x}} = 0$ and the system changes direction of motion [39].

As an example, we compute the square root in Eq. (10.35) in the absence of the potential, and show that this leads to the decay with distance $\propto 1/|\mathbf{x}_f - \mathbf{x}_i|$ typical of the free propagation. Indeed, if one assumes a GW source localized at a certain spacetime point \mathbf{x}_s , then the field solution takes the form

$$\Psi_\omega(\mathbf{x}) = \mathcal{S}(\mathbf{x}_s) G_\omega^{\text{WKB}}(\mathbf{x}, \mathbf{x}_s), \quad (11.36)$$

where $\mathcal{S}(\mathbf{x}_s)$ is the value of the source's amplitude at its location. Hence, the decay with distance typical of freely propagating fields must be contained in the Green function. The first ingredient we need is the classical trajectory. In the absence of the potential, this satisfies Eq. (10.28), and the velocity is given by $\dot{\mathbf{x}}_{cl} \equiv (\mathbf{x}_f - \mathbf{x}_i)/\tau$. Therefore

$$S_{cl}[\tau] = \frac{|\mathbf{x}_f - \mathbf{x}_i|^2}{4\tau}. \quad (11.37)$$

The derivatives entering the square root in Eq. (10.35) are then easily computed

$$\frac{\partial^2 S_{cl}}{\partial x_f^a \partial x_i^b} = -\frac{\delta_{ab}}{2\tau}, \quad \frac{\partial^2 S_{cl}}{\partial \tau^2} = \frac{|\mathbf{x}_f - \mathbf{x}_i|^2}{2\tau^3}, \quad (11.38)$$

giving

$$G_\omega^{\text{WKB}}(\mathbf{x}_f, \mathbf{x}_i) = -\frac{1}{4\pi} \frac{e^{i\omega|\mathbf{x}_f - \mathbf{x}_i|}}{|\mathbf{x}_f - \mathbf{x}_i|}, \quad (11.39)$$

since the Morse index is null in this case, and where we have used Eq (10.29) to compute the phase. We have recovered the free propagator (see later Eq. (10.57) for its explicit computation), and fixed the normalization constant accordingly, choosing $\mathcal{N} = -1/2\pi$. The same computation can be performed again accounting also for the gravitational potential, if one knows how to solve the equations of motion of the classical trajectory. To keep the potential generic, here we take a perturbative approach considering that its effects on the trajectory are small. Recalling that $V = 4\alpha U$, we use the smallness of α and solve Eq. (10.31) as $\mathbf{x}_{cl}(\tau') = \mathbf{x}_{free}(\tau') + \alpha \mathbf{q}(\tau')$, where $\ddot{\mathbf{x}}_{free} = 0$ while

$$\frac{d^2 \bar{\mathbf{q}}(\tau')}{d\tau'^2} = -2\nabla V[\bar{\mathbf{x}}_{free}(\tau')], \quad (11.40)$$

with boundary conditions $\mathbf{q}(\tau' = \tau) = \mathbf{q}(\tau' = 0) = 0$. Therefore the action evaluated on classical path is

$$\begin{aligned} S_{cl}[\tau] &= \int_0^\tau d\tau' \left[\frac{1}{4} \left(\frac{d\mathbf{x}_{cl}}{d\tau'} \right)^2 - V[\mathbf{x}_{cl}(\tau')] \right] = \\ &= \frac{|\mathbf{x}_f - \mathbf{x}_i|^2}{4\tau} + \alpha \int_0^\tau d\tau' \left[\frac{\mathbf{x}_f - \mathbf{x}_i}{2\tau} \cdot \frac{d\bar{\mathbf{q}}}{d\tau'} - 4U[\bar{\mathbf{x}}_{free}(\tau')] \right] = \\ &= \frac{|\mathbf{x}_f - \mathbf{x}_i|^2}{4\tau} + \alpha \mathcal{V}, \end{aligned} \quad (11.41)$$

where in the last line we defined \mathcal{V} collecting all together the first order contributions. From this expression, we can take again the second derivatives of the action evaluated on the classical trajectory with respect to the initial and final points, and the proper time

$$\frac{\partial^2 S_{cl}[\tau]}{\partial x_f^a \partial x_i^b} = -\frac{\delta_{ab}}{2\tau} + \alpha \frac{\partial^2 \mathcal{V}}{\partial x_f^a \partial x_i^b}, \quad \frac{\partial^2 S_{cl}[\tau]}{\partial \tau^2} = \frac{|\mathbf{x}_f - \mathbf{x}_i|^2}{2\tau^3} + \alpha \frac{\partial^2 \mathcal{V}}{\partial \tau^2} \quad (11.42)$$

yielding the WKB form of the propagator

$$G_{\omega}^{\text{WKB}}(\mathbf{x}_f, \mathbf{x}_i) \approx -\frac{1}{4\pi} \frac{e^{i\omega W[\mathbf{x}_{cl}, \tau_{cl}]}}{|\mathbf{x}_f - \mathbf{x}_i|} \left[1 - \alpha \frac{|\mathbf{x}_f - \mathbf{x}_i|}{2} \left(\text{Tr} \left[\frac{\partial^2 \mathcal{V}}{\partial x_f^a \partial x_i^b} \right] + \frac{1}{4} \frac{\partial^2 \mathcal{V}}{\partial \tau^2} \right) \Big|_{\tau_{cl}} \right], \quad (11.43)$$

where we have fixed the normalization constant in the limit $\alpha \rightarrow 0$ using the free propagator. Since we performed the expansion in $\alpha \ll 1$, this result shows the decay with distance typical of the wave's propagation, at zero-th order in the parameter. The presence of the lens generates both a modification of the amplitude (the second term in the parenthesis) and in the phase ($W[\mathbf{x}_{cl}, \tau_{cl}]$ has a term $\mathcal{O}(\alpha)$). Both of these two terms are frequency independent, as it should since we are considering the high-frequency limit and the GO laws are achromatic. Note also that the assumption $\alpha \ll 1$ implies that there is only one classical trajectory, reason why we removed the summation over the different images.

11.4.3 Recovering the diffraction integral under the paraxial approximation

In this section we want to show that Eq. (10.21) reduces to the WO literature presented in Section 10.2 starting from our propagator given in the worldline description. Given all that we discussed, few points are clear: we should recover the diffraction integral by imposing the paraxial and eikonal approximations on our Green function and, in doing so, the integral over the proper time should be removed. Indeed, the main difference between our approach and the diffraction integral is the presence of the integral over all possible values of the proper time. If one was able to remove this last step, then the two description should match. In general, though, it is not possible to exchange the integration order between the proper time and path integrals and perform the former first, because the latter has integration boundaries which depend on τ . Nonetheless, exchanging the two integrals can be done in the high-frequency limit [37], as shown by the *Feynman-Garrod* representation of the propagator [37, 38, 112, 60]. To achieve it, we assume to pass the proper time integral through and perform a half-step compared to the derivation of the WKB form of the propagator, fixing the value of the proper time, while leaving free to vary the trajectories. This approach can be carried out in two ways: performing a saddle point approximation considering variations in τ around its classical configuration, or by strictly imposing the vanishing of the Hamiltonian, hence requiring that Eq. (10.27) holds exactly and simply evaluating the integral in τ_{cl} . Practically, both of the approaches imply a series in $\delta\tau = \tau - \tau_{cl}$ of the phase $e^{i\omega W}$, and in the first case the series is truncated at second order, while in the second case at first order. We illustrate both of the two options, showing that the first leads to the Feynman-Garrod form of the propagator while the second one to the diffraction integral in the paraxial limit. Note that this derivation is similar to the one displayed in [104] where the diffraction integral is found starting from the dispersion relation.

The diffraction integral

We start with the approach that leads to the diffraction integral, namely requiring that the Hamiltonian constraint is valid exactly. Therefore we simply evaluate the integral in τ as a suitable normalization (taking care of the dimensions) times the path integral evaluated in the classical configuration τ_{cl} . Enforcing Eq. (10.27) and writing the path integral in terms of the arc length, instead of the proper time, leads us to

$$G_{\omega}(\mathbf{x}_f, \mathbf{x}_i) \approx \mathcal{N} \int_{\mathbf{x}(l'=0)=\mathbf{x}_i}^{\mathbf{x}(l'=l_{cl})=\mathbf{x}_f} \mathcal{D}\mathbf{x}(l') e^{i\omega \int_0^{l_{cl}} dl' n(\mathbf{x}(l'))}, \quad (11.44)$$

showing that the action weighting the paths is given by the refractive index, as it known in optics [12, 31, 32, 113]. The normalization factor \mathcal{N} will be subsequently determined by requiring that in the limit $\omega R_s \rightarrow 0$ we find the free propagator of Eq. (10.57). Hence we expect a behaviour like $\mathcal{N} \propto 1/|\mathbf{x}_f - \mathbf{x}_i|$. Having removed the integral over the proper time, we proceed in evaluating the Green function under the paraxial approximation by considering waves propagating along a main direction, with small deviations from it. In full generality, we choose the propagation direction to be along the z axis and that the lens is thin. We redefine the spatial coordinates as $\mathbf{x} = (\vec{x}, z)$, where \vec{x} is the 2-dimensional vector $\vec{x} = (x^1, x^2)$. The constrained total action receives two contributions

$$W = \int_0^{l_{cl}} dl' n(\mathbf{x}(l')) = \int_0^{l_{cl}} dl' [1 - 2\alpha U(\mathbf{x}(l'))] = l_{cl} - 2\alpha \int_0^{l_{cl}} dl' U(\mathbf{x}(l')), \quad (11.45)$$

where the first term represents the geometric time delay, t_{geo} , while the second one is the gravitational time delay t_{grav} experienced by the paths immersed in the gravitational field of the lens. The geometric time delay can be integrated using the Pythagorean theorem in the paraxial approximation. We implement this by requiring that the rays bend instantly as they hit the thin lens plane, and the horizontal displacement has to be much smaller than the distances r_O and r_L separating the observer from the lens and the source from the lens (see Figure 10.1). After some manipulations, we obtain

$$t_{geom} = l_{cl} = \frac{1}{2} \frac{r_O}{r_{LO} r_L} (\vec{x} - \vec{\mu})^2, \quad \text{where} \quad \vec{\mu} \equiv \frac{\vec{x}_O r_L + \vec{x}_S r_{LO}}{r_O} \quad (11.46)$$

where \vec{x}_O and \vec{x}_S are the 2-D vectors describing the observer and the source position in the lens' plane, as in [12]. As for the gravitational time delay t_{grav} , the thin lens approximation allows us to approximate the integral by collapsing the gravitational field over the z axis. In this way we can re-express t_{grav} as a function of the impact parameter $\vec{x}(l)$ in the following way

$$t_{grav} = -2\alpha \int_0^{l_{cl}} dl' U(\mathbf{x}(l')) = -2\alpha \int U(\vec{x} + z\hat{z}) dz \equiv -\hat{\psi}(\vec{x}). \quad (11.47)$$

As explained in [26], the result of the thin-lens approximation is that the paths that contribute the most to the phase integral are well approximated by $\vec{x}(l) = \vec{x}_L$, namely by a constant vector. In this case the path integral collapses to a 2D integral on the planes orthogonal to the z direction

$$G_\omega(\mathbf{x}_O, \mathbf{x}_S) = \mathcal{N} \int_{\vec{x}_S}^{\vec{x}_O} d\vec{x} \exp \left\{ i\omega \left[\frac{1}{2} \frac{r_O}{r_{LO} r_L} (\vec{x} - \vec{\mu})^2 - \hat{\psi}(\vec{x}) \right] \right\}, \quad (11.48)$$

where we have considered the source located in \mathbf{x}_i and the observer in \mathbf{x}_f . This result matches the diffraction integral found in [12], and it is equivalent to the one described in Section 10.2. Thus we have proved that our formalism encompasses the current WO literature in the eikonal and paraxial limits.

The Feynman-Garrod propagator

Instead of enforcing the Hamiltonian constraint strictly, one can also expand the total phase up to second order in $\tau - \tau_{cl}$, and perform the Gaussian integral over the proper time. This requires to exchange the order of the proper time and the path integral, procedure which is allowed in the high frequency regime [110, 38, 112, 60]. Doing so, leads to the Feynman-Garrod representation of the propagator

$$G_\omega(\mathbf{x}_f, \mathbf{x}_i) \approx \mathcal{N} \int_{\mathbf{x}(\tau'=0)=\mathbf{x}_i}^{\mathbf{x}(\tau'=\tau_{cl})=\mathbf{x}_f} \mathcal{D}\mathbf{x}(\tau') \frac{e^{i\omega W[\mathbf{x}, \dot{\mathbf{x}}, \tau_{cl}]}}{\sqrt{[\partial^2 W[\mathbf{x}, \dot{\mathbf{x}}, \tau_{cl}]/\partial \tau^2]_{\tau_{cl}}}}, \quad (11.49)$$

where the normalization still needs to be fixed suitably. The square root at the denominator comes from the Gaussian integral and it can develop poles, corresponding points where the approximation performed breaks down. Indeed, the enforcing of the Hamiltonian constraint leads to the existence of caustics [102, 109, 39] where the saddle point approximation breaks down, as it is unable to account for nontrivial wave contributions. As it can be understood from Eq. (10.32), the zeros of the denominator in Eq. (10.49) are found in correspondence of: the classical trajectories, turning points where $\dot{\mathbf{x}} = 0$, or points where the velocity becomes orthogonal to the acceleration and the force exerted on the particle. Regarding the first case, from the WKB analysis carried out previously, we know that this divergence is not there: when the classical path dominates the path integral, the second saddle point approximation performed on the paths brings down an additional amplitude factor canceling out the vanishing $\partial^2 W / \partial \tau^2$. Before examining the other two possibilities generating a pole in Eq. (10.49), we stress that we expect to recover a description exclusively in terms of particle trajectories under the two assumptions of eikonal and paraxial. The second one in particular, requires that the wave propagates along a main optical axis, with small deviations away from it. Therefore, the other two options for the vanishing of the denominator, are excluded from the range of validity of the paraxial assumption: when the velocity vanishes, the particle's trajectory hits a turning point inverting the direction of motion, while when the velocity and the acceleration (or force) are orthogonal, the particle changes direction of propagation. In both cases, the path changes direction considerably, hence invalidating the paraxial assumption. Intuitively, to transition toward the paraxial approximation we demand that the trajectories included in the path integral always *move forward* and we further restrict the trajectories considered for the path integral in Eq. (10.49), already constrained by Eq. (10.27), to those that satisfy the paraxial approximation.

In the standard literature of WO, the diffraction integral is also used to compute the *beyond geometric optics* corrections. These are corrections to the $\omega \rightarrow \infty$ limit of the amplification factor and are computed by expanding up to fourth order in $\mathbf{x} - \mathbf{x}_{cl}$ the phase of the highly oscillating exponential of the diffraction integral. However, from our derivation it is clear that the derivation of Eq. (10.10) is approximated in nature, and especially the comparison with the Feynman-Garrod representation in Eq. (10.49) questions the validity of the beyond geometric optics expansion. Indeed, in such expansions the proper time τ and the path $\mathbf{x}(\tau')$ are treated on unequal footings. One should first make sure that the assumption of $\hat{\mathcal{H}} = 0$ is consistent with keeping higher orders in the perturbation series around the classical trajectory.

11.5 Perturbation theory

In the previous Section, we derived the laws of GO and the diffraction integral starting from our main result in Eq. (10.21). In this part of the work, we set up the perturbative expansion of the Green function in the parameter $\alpha\omega$, investigating its prediction away from the high-frequency regime.

11.5.1 Setting up the the expansion

Identifying the free action as $S_0 \equiv \int_0^\tau d\tau' \frac{1}{4} \left(\frac{d\mathbf{x}}{d\tau'} \right)^2$, we can rewrite the total Green function of Eq (10.21) as

$$G_\omega(\mathbf{x}_f, \mathbf{x}_i) = -\frac{i}{\omega} \int_0^{+\infty} d\tau e^{i\omega\tau} \int_{\mathbf{x}(\tau'=0)=\mathbf{x}_i}^{\mathbf{x}(\tau'=\tau)=\mathbf{x}_f} \mathcal{D}\mathbf{x}(\tau') e^{i\omega S_0} \exp \left[-i\omega \int_0^\tau d\tau' V(\mathbf{x}(\tau')) \right], \quad (11.50)$$

and expand in series the last exponential if the combination of the parameter $\alpha\omega$ is much smaller than one (α is contained in the definition of V),

$$\exp\left[-i\omega\int_0^\tau d\tau' V(\mathbf{x}(\tau'))\right] \simeq 1 - i\omega\int_0^\tau d\tau' V(\mathbf{x}(\tau')) - \frac{\omega^2}{2}\left(\int_0^\tau d\tau' V(\mathbf{x}(\tau'))\right)^2 + \dots \quad (11.51)$$

We can use this expression to define the Green function at various perturbation orders

$$G_\omega(\mathbf{x}_f, \mathbf{x}_i) = -\frac{i}{\omega}\int_0^{+\infty} d\tau e^{i\omega\tau} \left[\tilde{G}_\omega^{(0)}(\mathbf{x}_f, \mathbf{x}_i, \tau) - i\omega\tilde{G}_\omega^{(1)}(\mathbf{x}_f, \mathbf{x}_i, \tau) - \frac{\omega^2}{2}\tilde{G}_\omega^{(2)}(\mathbf{x}_f, \mathbf{x}_i, \tau) + \dots \right] \quad (11.52)$$

with

$$\tilde{G}_\omega^{(0)}(\mathbf{x}_f, \mathbf{x}_i, \tau) \equiv \int_{\mathbf{x}(\tau'=0)=\mathbf{x}_i}^{\mathbf{x}(\tau'=\tau)=\mathbf{x}_f} \mathcal{D}\mathbf{x}(\tau') e^{i\omega S_0}, \quad (11.53)$$

$$\tilde{G}_\omega^{(1)}(\mathbf{x}_f, \mathbf{x}_i, \tau) \equiv \int_{\mathbf{x}(\tau'=0)=\mathbf{x}_i}^{\mathbf{x}(\tau'=\tau)=\mathbf{x}_f} \mathcal{D}\mathbf{x}(\tau') e^{i\omega S_0} \int_0^\tau d\tau' V(\mathbf{x}(\tau')), \quad (11.54)$$

$$\tilde{G}_\omega^{(2)}(\mathbf{x}_f, \mathbf{x}_i, \tau) \equiv \int_{\mathbf{x}(\tau'=0)=\mathbf{x}_i}^{\mathbf{x}(\tau'=\tau)=\mathbf{x}_f} \mathcal{D}\mathbf{x}(\tau') e^{i\omega S_0} \left[\int_0^\tau d\tau_1 V(\mathbf{x}(\tau_1)) \int_0^{\tau_1} d\tau_2 V(\mathbf{x}(\tau_2)) \right] \quad (11.55)$$

Given a specific form of the gravitational potential V , one can then use these equations to compute its effects at various order. Note that, because $V = 4\alpha U$, each order in the expansion is proportional to $(\omega\alpha)^n$. In the next two Sections, we will provide further details about the latter equations.

11.5.2 The free relativistic particle propagator

We start by computing the free propagator, namely Eq. (10.53). In this case, the path integral can be evaluated (see Appendix 10.9) and $G_\omega^{(0)}(\mathbf{x}_f, \mathbf{x}_i)$ takes the explicit form

$$\tilde{G}_\omega^{(0)}(\mathbf{x}_f, \mathbf{x}_i, \tau) = \left(\frac{\omega}{4i\pi\tau}\right)^{3/2} e^{\frac{i\omega}{4\tau}|\mathbf{x}_f-\mathbf{x}_i|^2}, \quad (11.56)$$

while the integration over proper time gives

$$G_\omega^{(0)}(\mathbf{x}_f, \mathbf{x}_i) = -\frac{i}{\omega}\int_0^{+\infty} d\tau e^{i\omega\tau} \tilde{G}_\omega^{(0)}(\mathbf{x}_f, \mathbf{x}_i, \tau) = -\frac{1}{4\pi} \frac{e^{i\omega|\mathbf{x}_f-\mathbf{x}_i|}}{|\mathbf{x}_f-\mathbf{x}_i|}, \quad (11.57)$$

which is the standard form of the Green function of the wave operator in frequency space. In order to get a more familiar form, one can move to momentum space by performing a 3D Fourier transform

$$G_\omega^{(0)}(\mathbf{p}) = -\frac{1}{4\pi} \int d^3x \frac{e^{i\omega|\mathbf{x}_f-\mathbf{x}_i|-i\mathbf{p}\cdot(\mathbf{x}_f-\mathbf{x}_i)}}{|\mathbf{x}_f-\mathbf{x}_i|} = \frac{1}{p^2-\omega^2} \quad (11.58)$$

Drawing an analogy with the propagator of a massive scalar field, we see that $i\omega$ plays the role of the mass of the associated particle [33].

If the propagation does not start at $\tau' = 0$ (we will see an example later), the free result is generalized to

$$\begin{aligned} \tilde{G}_\omega^{(0)}(\mathbf{x}_a, \mathbf{x}_b, \tau_a - \tau_b) &= \int_{\mathbf{x}(\tau'=\tau_b)=\mathbf{x}_b}^{\mathbf{x}(\tau'=\tau_a)=\mathbf{x}_a} \mathcal{D}\mathbf{x}(\tau') e^{i\omega\int_{\tau_b}^{\tau_a} d\tau' L_0[\mathbf{x}(\tau'), \dot{\mathbf{x}}(\tau'), \tau']} = \\ &= \left(\frac{\omega}{4i\pi(\tau_a - \tau_b)}\right)^{3/2} e^{\frac{i\omega}{4} \frac{|\mathbf{x}_a-\mathbf{x}_b|^2}{(\tau_a-\tau_b)}}. \end{aligned} \quad (11.59)$$

11.5.3 First order

We now turn to the first order contribution, namely Eq. (10.54). In this expression the only information about the specific path $\mathbf{x}(\tau)$ that the gravitational potential can give is the position of the path at the given proper time τ' . Before and after that moment, the scalar field is evolving freely. With this consideration, we proceed following Feynman's argument in [103], and divide each path in two parts (one before $\tau = \tau'$ and one after), and require that, at the given time τ_1 , each path goes through a specific point $\mathbf{x}_1^* = \mathbf{x}(\tau_1)$. By doing so, and integrating over all possible values of this division point \mathbf{x}_1^* , one gets the following expression

$$\tilde{G}_\omega^{(1)}(\mathbf{x}_f, \mathbf{x}_i, \tau) = \int_0^\tau d\tau_1 \int_{-\infty}^{+\infty} d\mathbf{x}_1^* \tilde{G}_\omega^{(0)}(\mathbf{x}_f, \mathbf{x}_1^*, \tau - \tau_1) V(\mathbf{x}_1^*) \tilde{G}_\omega^{(0)}(\mathbf{x}_1^*, \mathbf{x}_i, \tau_1), \quad (11.60)$$

which suggests the interpretation of a particle propagating freely except for when it interacts with V . Such interaction can be seen as a “scattering event” occurring at a specific point in space. The full probability amplitude is then built from marginalizing over all possible scattering centers.

This representation is widely known in literature. Particularly in the context of cosmology and gravity, it has been used to describe the contributions due to self-forces and tail effects [114, 115, 116, 117].

11.5.4 Higher orders

Now we work Eq. (10.55). As in [103], we assume $\tau_1 < \tau_2$ and double the integral. To avoid clutter in the notation, we implement this condition by assuming causal propagation, namely that $\tilde{G}_\omega^{(2)}(\mathbf{x}_a, \mathbf{x}_b, \tau_a - \tau_b) = 0$ if $\tau_a - \tau_b < 0$. Then, as done for the first order, we divide the propagation path into three pieces: before τ_1 , between τ_1 and τ_2 and after τ_2 . This reflects the fact that the particle scatters with the potential at the two proper time instants of τ_1 and τ_2 and propagates freely otherwise. Proceeding as before, one obtains

$$\begin{aligned} \tilde{G}_\omega^{(2)}(\mathbf{x}_f, \mathbf{x}_i, \tau) = & 2 \int_0^\tau d\tau_1 \int_0^\tau d\tau_2 \int_{-\infty}^{+\infty} d\mathbf{x}_1^* \int_{-\infty}^{+\infty} d\mathbf{x}_2^* \times \\ & \times \left[\tilde{G}_\omega^{(0)}(\mathbf{x}_f, \mathbf{x}_1^*, \tau - \tau_1) V(\mathbf{x}_1^*) \tilde{G}_\omega^{(0)}(\mathbf{x}_1^*, \mathbf{x}_2^*, \tau_1 - \tau_2) V(\mathbf{x}_2^*) \tilde{G}_\omega^{(0)}(\mathbf{x}_2^*, \mathbf{x}_i, \tau_2) \right]. \end{aligned} \quad (11.61)$$

Going to higher orders in perturbation theory in the worldline description means adding more and more scattering points. Indeed, organizing the perturbation theory in this fashion has the advantage of being straightforwardly iterated to higher orders: $\tilde{G}_\omega^{(3)}(\mathbf{x}_f, \mathbf{x}_i, \tau)$ will have a similar structure with one more gravitational potential insertion and so on. This way, at the n -th order in perturbation, the probability of going from \mathbf{x}_i to \mathbf{x}_f is built as the sum of the probabilities of doing so with i scattering, where $i \in [0, n]$. It is important to point out that no assumption has been made on the *shape* of the gravitational potential, such as the thin lens approximation. Indeed, here $V(\mathbf{x}(\tau))$ is extended and the integration over all the division points \mathbf{x}^* spans all spacetime, also the points where the potential is non-zero. Nonetheless, the perturbation theory considers free propagation before and after each scattering event, *even if* the propagation is occurring in a region where $V \neq 0$. This assumption is called *Born approximation* [103]. The way that the perturbation theory reaches higher levels of accuracy is by considering more and more interactions points between the wave and the potential. At higher perturbative orders, the number of scattering points grows, accounting for a less coarse exploration of the potential. This structure is represented in Figure 10.3. In the $n \rightarrow \infty$ limit, Born approximation becomes exact as there is no more free propagation inside the gravitational

potential. This suggest that, in this limit, one can resum the perturbation theory and obtain an exact result. Indeed, by considering infinite scattering points, the expansion series can be organized in such a way that, by factoring out the last scattering point, the Green function takes the form

$$\tilde{G}_\omega(\mathbf{x}_f, \mathbf{x}_i, \tau) = \tilde{G}_\omega^{(0)}(\mathbf{x}_f, \mathbf{x}_i, \tau) - i\omega \int_0^\tau d\tau_{LS} \tilde{G}_\omega^{(0)}(\mathbf{x}_f, \mathbf{x}_{LS}, \tau - \tau_{LS}) V(\mathbf{x}_{LS}) \tilde{G}_\omega(\mathbf{x}_{LS}, \mathbf{x}_i, \tau_{LS}). \quad (11.62)$$

This equation states that the full probability amplitude for the associate particle to propagate form the initial \mathbf{x}_i to the final \mathbf{x}_f point (the l.h.s), can be obtained as the sum of the free propagation (first term on the r.h.s) and the conditional probability of interacting on the last scattering point (LS) after having propagated with probability amplitude dictated by the full theory. Eq. (10.62) is an integral equation for the exact Green function. Interpreting the gravitational potential $V(\mathbf{x})$ as the proper self-energy, then Eq. (10.62) becomes equivalent to *Dyson equation*, found in quantum field theory [118] and in non-relativistic many body theory [119]. The interpretation of the gravitational potential as self-energy, dressing the free propagator, is also clear from the Hamiltonian description laid out in Section 10.3.2.

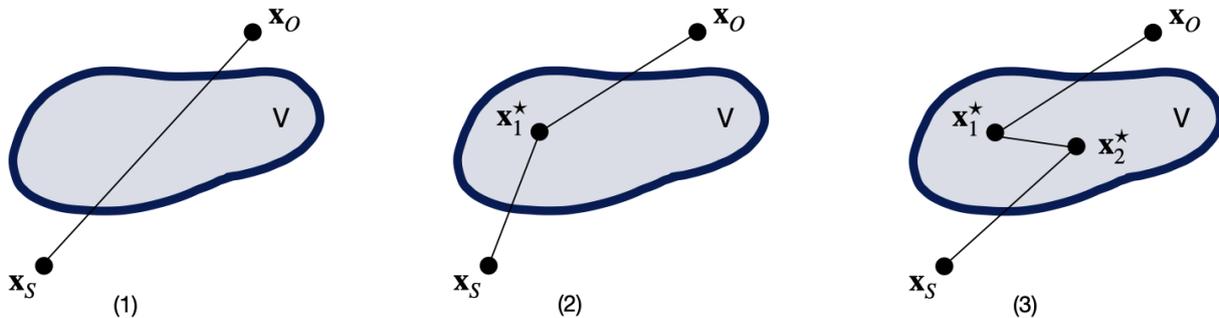


Figure 11.3: Depiction of the perturbation series we set up, showing three orders in the Born approximation. The solid line represent the free propagator $\tilde{G}_\omega^{(0)}(\mathbf{x}_f, \mathbf{x}_i, \tau)$ connecting various initial and final points. In figure (1), the particles travels through the region where $V \neq 0$ freely, hence representing the lowest perturbative order. Adding more scattering points means going to higher orders in perturbation theory. Figures (2) and (3) represent $\tilde{G}_\omega^{(1)}(\mathbf{x}_f, \mathbf{x}_i, \tau)$ and $\tilde{G}_\omega^{(2)}(\mathbf{x}_f, \mathbf{x}_i, \tau)$ respectively, where the wave scatters once and twice with the potential at the points \mathbf{x}_1^* and \mathbf{x}_2^* . The propagation occurs from the source $\mathbf{x}_i = \mathbf{x}_S$, to the observer $\mathbf{x}_f = \mathbf{x}_O$. In the ideal limit of infinite scattering points, there is no more free propagation inside the potential.

11.5.5 Perturbation theory with the Generating Functional

A key feature of the path integral description, is its organization in terms of the generating functional [45, 120, 121]. By adding a source to the phase weighting the paths, one can obtain expectation values of operators by taking functional derivatives of the generating functional with respect to it. In the context of the worldline representation, the generating functional is defined as

$$Z[\mathbf{J}, \tau] = \int_{\mathbf{x}(\tau'=0)=\mathbf{x}_i}^{\mathbf{x}(\tau'=\tau)=\mathbf{x}_f} \mathcal{D}\mathbf{x}(\tau') \exp \left[i\omega \int_0^\tau d\tau' (L + \mathbf{J}(\tau') \cdot \mathbf{x}(\tau')) \right], \quad (11.63)$$

where L is the associated Lagrangian, given by Eq. (10.20) and $\mathbf{J}(\tau')$ is the current acting as the source of \mathbf{x} . Note that, since we are working at the particle level through the worldline

approach, operators are functions of $\mathbf{x}(\tau)$, rather than fields. For instance, in Eq. (10.55), we have two gravitational potentials evaluated in $\mathbf{x}(\tau_1)$ and $\mathbf{x}(\tau_2)$. Expectation values of operators can be obtained by taking functional derivatives with respect to the current \mathbf{J} and then setting it to zero, namely

$$\frac{\delta^n Z[\mathbf{J}, \tau]}{\delta \mathbf{J}(\tau_1) \dots \delta \mathbf{J}(\tau_n)} \Big|_{\mathbf{J}=\mathbf{0}} = \int_{\mathbf{x}(\tau'=0)=\mathbf{x}_i}^{\mathbf{x}(\tau'=\tau)=\mathbf{x}_f} \mathcal{D}\mathbf{x}(\tau') e^{i\omega S} \mathbf{x}(\tau_1) \dots \mathbf{x}(\tau_n), \quad (11.64)$$

since the action of the functional derivative is analogous to its counterpart in quantum field theory [118], namely $\delta J_i(\tau_1)/\delta J_j(\tau_2) = \delta_i^j \delta(\tau_1 - \tau_2)$, thus

$$\frac{1}{i\omega} \frac{\delta}{\delta \mathbf{J}(\tau_1)} \exp \left[i\omega \int_0^\tau d\tau' \mathbf{J}(\tau') \cdot \mathbf{x}(\tau') \right] = \mathbf{x}(\tau_1) \exp \left[i\omega \int_0^\tau d\tau' \mathbf{J}(\tau') \cdot \mathbf{x}(\tau') \right]. \quad (11.65)$$

As usual, one can produce a perturbative expansion also through the generating functional by writing the latter as

$$Z[\mathbf{J}, \tau] = \exp \left[i\omega \int_0^\tau d\tau' V \left(\frac{1}{i\omega} \frac{\delta}{\delta \mathbf{J}} \right) (\tau') \right] Z_0[\mathbf{J}, \tau], \quad (11.66)$$

where $Z_0[\mathbf{J}, \tau]$ is the generating functional of the free theory,

$$Z_0[\mathbf{J}, \tau] \equiv \int_{\mathbf{x}(\tau'=0)=\mathbf{x}_i}^{\mathbf{x}(\tau'=\tau)=\mathbf{x}_f} \mathcal{D}\mathbf{x}(\tau') \exp \left[i\omega \int_0^\tau d\tau' \left(\frac{\dot{\mathbf{x}}^2}{4} + \mathbf{J}(\tau') \cdot \mathbf{x}(\tau') \right) \right]. \quad (11.67)$$

In Eq. (10.66) the operator V evaluated in the derivative with respect to the current, is to be intended through its Taylor expansion. The free generating functional is quadratic in the velocity, hence it can be computed explicitly

$$\begin{aligned} Z_0[\mathbf{J}, \tau] &= \exp \left[i\omega \int_0^\tau d\tau' \left(-\frac{1}{4} \mathbf{x} \frac{d^2}{d\tau'^2} \mathbf{x} + \mathbf{J}(\tau') \cdot \mathbf{x}(\tau') \right) \right] = \\ &= \exp \left[i\omega \int_0^\tau d\tau' \mathcal{A} \left(-\frac{1}{2} \mathbf{x} - \mathbf{J}(\tau') \mathcal{A}^{-1} \right)^2 \right] \exp \left[i\omega \int_0^\tau d\tau' d\tau'' \mathbf{J}(\tau') \Delta(\tau' - \tau'') \mathbf{J}(\tau'') \right], \end{aligned} \quad (11.68)$$

where $\mathcal{A} \equiv \frac{d^2}{d\tau^2}$ and $\Delta(\tau' - \tau'')$ can be interpreted as a *proper time propagator*, describing the probability to go from τ to τ' and obeying

$$\frac{d^2}{d\tau^2} \Delta(\tau' - \tau'') = \delta(\tau' - \tau''). \quad (11.69)$$

Note that the first factor in Eq. (10.68) gives an overall normalization constant, which can be reabsorbed in the overall one.

11.6 Applications

In this Section we provide two applications of our formalism. First we show how it can be implemented also for massive scalar fields. Then, we provide an explicit example of the first order correction to the propagator, i.e. we compute Eq. (10.60), for a specific choice of gravitational potential.

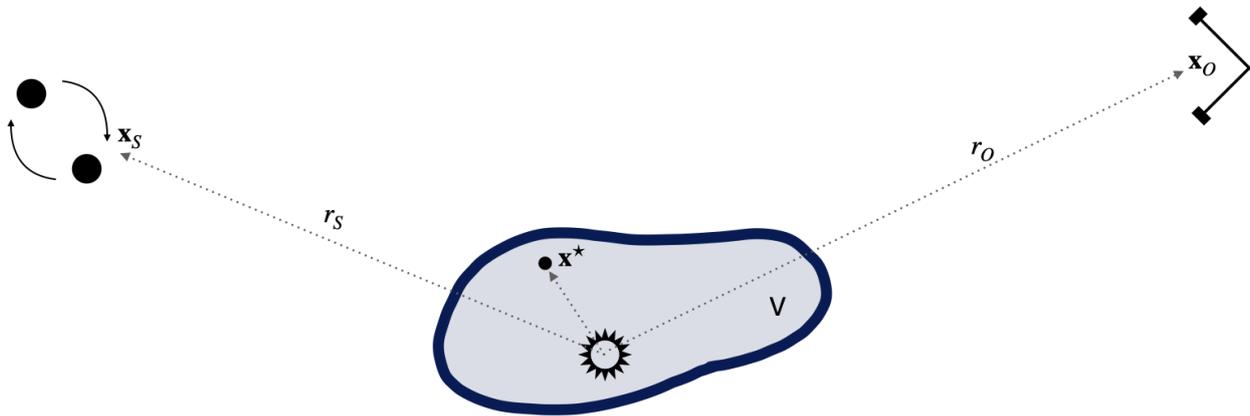


Figure 11.4: Lensing situation considered. The coordinate system is now centered on the lens and the distances to the source and observer are large.

11.6.1 Coulomb potential: an explicit example

In this section, we obtain an explicit result for the case of a spherically symmetric potential, describing the standard scenario of a gravitational potential generated by a compact object. In particular, we shall compute $\tilde{G}_\omega^{(1)}(\mathbf{x}_f, \mathbf{x}_i, \tau)$ of Eq. (10.60) for a Coulomb-like potential, i.e. $U \sim 1/r$, and the total first order Green function in frequency space, $G_\omega^{(1)}(\mathbf{x}_f, \mathbf{x}_i)$, by integrating over the proper time. To do so, we will first consider a Yukawa potential,

$$U = -\frac{MG}{r}e^{-r/a}, \quad (11.70)$$

where a sets the scale of the range of the interaction, and then take the $a \rightarrow \infty$ limit. Introducing this length-scale, allows us to regularize the divergent integral in Eq. (10.60), and perform it. We also recall that, in the unit where $c = 1$, then MG has dimension of a length and it coincides with the Schwarzschild radius of the lens, and the gravitational potential is dimensionless. In order to have a more clear picture in mind, let us identify the initial propagation point with the position of the source emitting the GW (i.e. $\mathbf{x}_i = \mathbf{x}_S$) and the final point with the position of the observer (i.e. $\mathbf{x}_f = \mathbf{x}_O$). It is convenient to center the coordinate frame on the lens, and consider that r_S and r_O are much larger than the distance to the scattering point r_* . Namely, we consider that the points where the waves are produced and observed are distant from the point where the interaction occurs. This is a good approximation when $V(\mathbf{x})$ has a compact support, or decays fast enough at large distances, hence allowing us to stop at first order in perturbation theory. In particular, we consider (see Figure 10.4)

$$\mathbf{x}_S = \{r_S, -\boldsymbol{\theta}_S\}, \quad \mathbf{x}_O = \{r_O, \boldsymbol{\theta}_O\}, \quad \mathbf{x}^* = \{r^*, \boldsymbol{\theta}^*\}. \quad (11.71)$$

The integral over time of Eq. (10.60) can be performed

$$\begin{aligned} \int_0^\tau d\tau_1 \tilde{G}_\omega^{(0)}(\mathbf{x}_O, \mathbf{x}^*, \tau - \tau_1) V(\mathbf{x}^*) \tilde{G}_\omega^{(0)}(\mathbf{x}^*, \mathbf{x}_S, \tau_1) &= \\ &= \left(\frac{\omega}{4i\pi}\right)^{5/2} \tau^{-3/2} \left(\frac{1}{|\mathbf{x}_O - \mathbf{x}^*|} + \frac{1}{|\mathbf{x}_S - \mathbf{x}^*|} \right) e^{\frac{i\omega}{4r}(|\mathbf{x}_O - \mathbf{x}^*| + |\mathbf{x}_S - \mathbf{x}^*|)^2} V(\mathbf{x}^*), \end{aligned} \quad (11.72)$$

where we have used that $V(\mathbf{x}^*)$ does not depend on τ' at this stage (see Appendix 10.10 for the explicit computation). To proceed with the spatial integration we consider that the source and observer are located far away compared to the lens. Imposing $r_S \gg r^*$ and $r_O \gg r^*$ gives

$$|\mathbf{x}^* - \mathbf{x}_S| \simeq r_S \left(1 + \frac{r^*}{r_S} \boldsymbol{\theta}_S \cdot \boldsymbol{\theta}^* \right), \quad |\mathbf{x}^* - \mathbf{x}_O| \simeq r_O \left(1 - \frac{r^*}{r_O} \boldsymbol{\theta}_O \cdot \boldsymbol{\theta}^* \right). \quad (11.73)$$

which we keep at lowest order in the first parenthesis of Eq. (10.72), and at first order in the exponential since this is quite sensitive to small relative changes in the phase. As prescribed in Eq. (10.60), we compute the spatial integral in the case of the Yukawa potential in Eq. (10.70). Defining the vector $\mathbf{R} \equiv (r_S + r_O)(\boldsymbol{\theta}_S - \boldsymbol{\theta}_O)$, with modulus $R \equiv \sqrt{|\mathbf{R}|^2}$ the first order worldline propagator takes the form (see Appendix 10.10 for the explicit steps)

$$\begin{aligned} \tilde{G}_\omega^{(1)}(\mathbf{x}_O, \mathbf{x}_S, \tau) &= \int d^3\mathbf{x}^* \int_0^\tau d\tau' \tilde{G}_\omega^{(0)}(\mathbf{x}_O, \mathbf{x}^*, \tau - \tau') V(\mathbf{x}^*) \tilde{G}_\omega^{(0)}(\mathbf{x}^*, \mathbf{x}_S, \tau') \\ &\approx 16\pi\alpha MG \left(\frac{\omega}{4i\pi}\right)^{5/2} \left(\frac{1}{r_O} + \frac{1}{r_S}\right) \times \frac{\tau^{-3/2} e^{\frac{i\omega}{4\tau}(r_S+r_O)^2}}{\left(\frac{i\omega R}{2\tau}\right)^2 - \frac{1}{a^2}}. \end{aligned} \quad (11.74)$$

We see that the scale of the interaction a , can be removed at this point by taking $a \rightarrow \infty$, as the divergent integral has now been performed. We take this limit and produce the result for the full propagator, by integrating out the proper time (see Appendix 10.10 for the explicit steps)

$$\begin{aligned} G_\omega^{(1)}(\mathbf{x}_O, \mathbf{x}_S) &\equiv -\frac{i}{\omega} \int_0^{+\infty} d\tau e^{i\omega\tau} \tilde{G}_\omega^{(1)}(\mathbf{x}, \mathbf{x}_0, \tau) = \\ &= \frac{\alpha MG}{\pi\omega^2} \left(\frac{1}{r_O} + \frac{1}{r_S}\right) \frac{e^{i\omega(r_S+r_O)}(1 - i\omega(r_S + r_O))}{R^2}. \end{aligned} \quad (11.75)$$

Hence we have arrived to the first order Born approximation of the Green function describing the dynamics of the system, which considers that the wave propagates freely except when the one interaction takes place. We can compose the full Green function by summing the zero-th and first order results for the Coulomb potential

$$\begin{aligned} G_\omega(\mathbf{x}_O, \mathbf{x}_S) &= G_\omega^{(0)}(\mathbf{x}_O, \mathbf{x}_S) - i\omega G_\omega^{(1)}(\mathbf{x}_O, \mathbf{x}_S) \approx \\ &\approx -\frac{1}{4\pi} \frac{e^{i\omega(r_O+r_S)}}{r_O + r_S} \left[1 + 4i\alpha(\omega MG) \frac{1 - i\omega(r_S + r_O)}{r_S r_O |\boldsymbol{\theta}_S - \boldsymbol{\theta}_O|^2 \omega^2} \right] \end{aligned} \quad (11.76)$$

where we have used Eq. (10.57) and in the last line we have written the parenthesis in a way that clearly shows that the first order contribution is dimensionless (MG is the Schwarzschild radius of the lens). We point out some characteristic of Eq. (10.76). First, the divergence in the forward limit $\boldsymbol{\theta}_S \approx \boldsymbol{\theta}_O$ is typical of an interaction mediated by a long range force, inducing an interaction between the wave and the lens also at very long distances, and it is mitigated by the smallness of the interaction (the coupling constant is G). Another remarkable feature of Eq. (10.76) regards the static limit, which seems not to be possible for the $1/\omega$ factor. However, in our treatment we had to assume $\omega \neq 0$ in order to introduce the proper time propagator with the correct $-i/\omega$ normalization. Because of this reason, taking the static limit is subtle but one can do so by considering $\omega \rightarrow 0$ while keeping $\omega r_S, \omega r_O$ finite. This prescription is necessary to remain within the realm of our assumption, and it leads to the correct results as well. Another possibility, is to take the limit $\omega MG \rightarrow 0$ as proxy for the static limit. In this case, the wavelength of the wave ($\lambda \sim 1/\omega$) becomes much larger than the Schwarzschild radius of the lens and the interaction is suppressed as expected. Indeed, we see that the interaction (encoded in $G_\omega^{(1)}(\mathbf{x}_O, \mathbf{x}_S)$) vanishes either if $\alpha = 0$, i.e. when there is no lens, or in the $MG/\lambda \rightarrow 0$ limit which corresponds to the fully diffractive regime. We also point out that any frequency dependent modification to the Green function appears in combination $i\omega$ or ω^2 compatibly with the fact that the Green function in time domain is real, implying $G_\omega = G_{-\omega}^*$.

Finally, we use our result to connect our theoretical framework to an observable, namely the distance to the source. This can be obtained from the waveform through an inverse proportionality relation in GO regime [81, 82]. The presence of the lens along the propagation will

introduce modifications in the inferred distance, which are also direction dependent [122]. We consider a source \mathcal{S} localized at \mathbf{x}_s , so that the scalar field at the observer position takes the form $\tilde{\Psi}_\omega(\mathbf{x}_O) = \mathcal{A}_{in} G_\omega(\mathbf{x}_O, \mathbf{x}_S)$, where \mathcal{A}_{in} is the initial amplitude. From the form of the free propagator, we recognize in $r_S + r_O$ the distance from the source to the observer, since we set the center of the coordinate system on the lens. We call this distance $\bar{d}_{SO} \equiv r_S + r_O$, which is the unperturbed distance in the sense that it is the correct one in the absence of the lens, and it does not depend on the direction of arrival of the wave. We can then use the fact $\alpha\omega \ll 1$ to rewrite the propagator in Eq. (10.76) as

$$\tilde{\Psi}_\omega(\mathbf{x}_O) \approx -\frac{\mathcal{A}_{in}}{4\pi} \frac{e^{i\omega(r_S+r_O)}}{\bar{d}_{SO} \left[1 - 4i\alpha(\omega MG) \frac{1-i\omega(r_S+r_O)}{r_S r_O |\boldsymbol{\theta}_S - \boldsymbol{\theta}_O|^2 \omega^2} \right]} = -\frac{\mathcal{A}_{in}}{4\pi} \frac{e^{i\omega(r_S+r_O)}}{d_{SO}}, \quad (11.77)$$

where in the last step we identified the denominator as the corrected distance in presence of the lens, d_{SO} . While the unperturbed distance depends only on the radial coordinates, the presence of the lens introduces a dependence of d_{SO} also on the deviation angle. We can compute the source-observer distance fluctuation as

$$\frac{d_{SO} - \bar{d}_{SO}}{\bar{d}_{SO}} \equiv \frac{\Delta d_{SO}}{\bar{d}_{SO}} = -4i\alpha(\omega MG) \frac{1 - i\omega(r_S + r_O)}{r_S r_O |\boldsymbol{\theta}_S - \boldsymbol{\theta}_O|^2 \omega^2}. \quad (11.78)$$

To facilitate the comparison with other works in literature, we also report the expression of the amplification factor which can be understood from Eq. (10.77). This is given by

$$F_\omega(\mathbf{x}_O) = \frac{\tilde{\Psi}_\omega(\mathbf{x}_O)}{\tilde{\Psi}_\omega^{NL}(\mathbf{x}_O)} = 1 + 4i\alpha(\omega MG) \frac{1 - i\omega(r_S + r_O)}{r_S r_O |\boldsymbol{\theta}_S - \boldsymbol{\theta}_O|^2 \omega^2} \quad (11.79)$$

where the unlensed solutions can be found by setting $M = 0$ in Eq. (10.77). One can compare this expression for the amplification factor to its analogous result obtained from the diffraction integral (see e.g. [123] Eq. 22). Both amplification factors show that deviations from the unlensed solution are proportional to ωMG , indicating the vanishing of the effects in the fully diffractive regime, and their structure in the complex frequency plane is such that the real time counterpart is real. Both expressions show a divergence in the $\omega \rightarrow 0$ limit, in Eq. (10.79) with a $1/\omega$ falloff, and in [123] through a logarithmic term. However, the radial dependencies are different, especially concerning the divergence in the forward limit, due to the different approaches and approximations employed.

11.6.2 Massive scalar field

We now consider a massive scalar field as prototype of massive gravity theory. This case was already considered in [124] to investigate the possibility of constraining the field's mass through lensing observations of combined GW and electromagnetic signals. In a standard lensing scenario, electromagnetic waves are usually in the GO regime, while GWs can have longer wavelengths and display WO effects. In these situations, the two lensing patterns can be used to break degeneracies and constraint both the properties of the lens and of the two messengers. Also [124] relies on the diffraction integral for its prediction, hence here we generalize their description within our framework, which we can easily re-propose with suitable modifications.

We start from a KG equation for a massive scalar field $(\square - m^2)\Psi = 0$ in the curved background. The Green function Eq. (10.14) then becomes

$$\left[\nabla^2 + \omega^2 \left(1 - \frac{m^2}{\omega^2} \right) [1 - 4\alpha U] \right] G_\omega(\mathbf{x}_f, \mathbf{x}_i) = \delta^{(3)}(\mathbf{x}_f - \mathbf{x}_i). \quad (11.80)$$

We can make the frequency redefinition

$$\omega_m \equiv \omega \sqrt{1 - \frac{m^2}{\omega^2}}, \quad (11.81)$$

and obtain an equation which is exactly equal to Eq. (10.14) upon replacing ω with ω_m . Consequently, the exact solution of Eq. (10.80) is given by

$$G_{\omega_m}(\mathbf{x}_f, \mathbf{x}_i) = -\frac{i}{\omega_m} \int_0^\infty d\tau e^{i\omega_m \tau} \int_{\mathbf{x}(\tau'=0)=\mathbf{x}_i}^{\mathbf{x}(\tau'=\tau)=\mathbf{x}_f} \mathcal{D}\mathbf{x}(\tau') e^{i\omega_m S}, \quad (11.82)$$

and the discussion of the previous sections can be reproduced straightforwardly but, in this case, the optical regime of the wave is regulated by the parameter ω_m instead of ω . When the frequency of the wave is much larger than the field's mass, i.e. $\omega^2 \gg m^2$, then $\omega_m \rightarrow \omega$. This regime corresponds to Fourier modes whose associated length-scale ($\sim 1/\omega$) is much smaller than the Compton wavelength of the field, hence they are not influenced by the presence of the new scale. Vice versa, when $\omega^2 \ll m^2$, then the frequency ω_m becomes purely imaginary and the dynamic is suppressed.

11.7 Lensing by a Black Hole: adding the polarization

An interesting example!!!

Having set the formalism for a scalar field, in this section we generalize our results for perturbations of higher spin. This task would require to solve Maxwell's equations for fields of spin $s = 1$, or the linearized Einstein's equations for $s = 2$, over the background described in Eq. (10.1). This task is highly complicated and not analytically solved for generic gravitational potentials yet. Nonetheless, for specific forms of V , the solutions of the mentioned equations are known in literature. One example are Kerr and Schwarzschild BHs, which we consider in this Section. To be consistent with the literature, in this Section we consider units where $c = G = 1$. To restore the dependence upon Newton's gravitational constant, it is sufficient to perform $M \rightarrow MG$.

11.7.1 Preliminary about Black Hole perturbation theory

Studying the equations of motion of test fields of different spin is the heart of the literature concerning BH perturbation theory. The field has started with the pioneering works of Regge and Wheeler [125] and Zerilli [126] for Schwarzschild lenses, and it was then expanded to Kerr BHs by Teukolsky in [98], providing a master equations for perturbations of different spins over such background. The key feature of Teukolsky equation is that, rather than being written in terms of the field variables, it is an equation for a specific Newman-Penrose (NP) scalar [99]. Such variables are related to derivatives of the test fields and, at linear order in perturbation theory, they are gauge invariant [91]. For instance, for $s = 2$, Teukolsky equation is an equation for $\psi = \{\psi_0, \rho^{-4}\psi_4\}$ where $\rho = -1/(r - ia \cos \theta)$, a being the Kerr parameter of Eq. (10.84), and

$$\psi_0 = -\delta (C_{\mu\alpha\nu\beta} l^\mu m^\alpha l^\nu \bar{m}^\beta), \quad \psi_4 = -\delta (C_{\mu\alpha\nu\beta} n^\mu \bar{m}^\alpha n^\nu \bar{m}^\beta), \quad (11.83)$$

where $C_{\mu\alpha\nu\beta}$ is the Weyl tensor and $\{l^\mu, n^\mu, m^\mu, \bar{m}^\mu\}$ are the elements of the NP null tetrad. As an example, the two scalars ψ_0 and ψ_4 on a Minkowski background, are related to gravitational waves as $\{\psi_0, \psi_4\} \propto \{\dot{h}_+ + i\dot{h}_\times, \dot{h}_+ - i\dot{h}_\times\}$ [127], namely the second time derivative of the two GW polarizations. For a vector field, the NP scalars composing the doublet ψ are related to NP components of the Maxwell tensor, while for $s = 0$, ψ is proportional to the scalar field itself. In this section, therefore, we borrow the results of BH perturbation theory studies, and generalize

the formalism presented in the previous Sections, for perturbations of spin $|s| = 0, 1/2, 1, 2$. Hence, here we consider as background metric the one of a Kerr BH

$$d\bar{s}^2 = \left(1 - \frac{2Mr}{\Sigma}\right) dt^2 + \frac{4Mar \sin^2 \theta}{\Sigma} dt d\varphi - \frac{\Sigma}{\Delta} d\theta^2 - \sin^2 \theta \left(r^2 + a^2 + \frac{2Ma^2 r \sin^2 \theta}{\Sigma}\right) d\varphi^2, \quad (11.84)$$

with

$$\Sigma = r^2 + a^2 \cos^2 \theta, \quad \Delta = r^2 - 2Mr + a^2, \quad (11.85)$$

instead of the one in Eq. (10.1). Teukolsky master equation then reads [98, 91]

$$\left\{ \left[\frac{(r^2 + a^2)^2}{\Delta} - a^2 \sin^2 \theta \right] \frac{\partial^2}{\partial t^2} + \frac{4Mar}{\Delta} \frac{\partial^2}{\partial t \partial \varphi} + \left[\frac{a^2}{\Delta} - \frac{1}{\sin^2 \theta} \right] \frac{\partial^2}{\partial \varphi^2} - \frac{1}{\sin \theta} \frac{\partial}{\partial \theta} \left(\sin \theta \frac{\partial}{\partial \theta} \right) - \right. \\ \left. - 2s \left[\frac{a(r-M)}{\Delta} + \frac{i \cos \theta}{\sin^2 \theta} \right] \frac{\partial}{\partial \varphi} - 2s \left[\frac{M(r^2 - a^2)}{\Delta} - r - ia \cos \theta \right] \frac{\partial}{\partial t} + \right. \\ \left. + (s^2 \cot^2 \theta - s) - \Delta^{-s} \frac{\partial}{\partial r} \left(\Delta^{s+1} \frac{\partial}{\partial r} \right) \right\} \psi(t, r, \theta, \varphi) = 0, \quad (11.86)$$

where ψ is the NP scalar, as explained. It is straightforward to show that the latter equation reduces to Eq. (10.3) for a spherically symmetric ($a = 0$), static lens in Newtonian regime ($r \gg 2M$), by taking the suitable limit. Indeed, the metric in Eq. (10.1) is included in the Kerr family if one restricts the study to radial potentials of the form, i.e. $U(r, \theta, \varphi) = -M/r$. However, we emphasize that Eq. (10.93) is exact while Helmholtz Eq. (10.3) for the scalar problem was found by linearizing to first order in α , namely the strength of the gravitational potential.

11.7.2 Proper time solution of Teukolsky equation

The goal of this Section is to show how to recast Eq. (10.86) into a form similar to Eq. (10.3), so that one can re-apply all the previous arguments in the case of higher spin fields. Following the standard procedure, we exploit the symmetries of the Kerr metric and decompose the scalar variable ψ as

$$\psi = e^{-i\omega t} e^{im\varphi} S(\theta) R(r), \quad (11.87)$$

so that Teukolsky equation decouples into two, one for the radial and one for the angular variable

$$\Delta^{-s} \frac{d}{dr} \left(\Delta^{s+1} \frac{dR}{dr} \right) + V_R(r) R = 0, \quad (11.88)$$

$$\frac{1}{\sin \theta} \frac{d}{d\theta} \left(\sin \theta \frac{dS}{d\theta} \right) + V_S(\theta) S = 0, \quad (11.89)$$

with the potentials given by

$$V_R(r) \equiv \left(\frac{K^2 - 2is(r-M)K}{\Delta} + 4is\omega r - a^2\omega^2 + 2am\omega - A \right), \quad (11.90)$$

$$V_S(\theta) \equiv \left(a^2\omega^2 \cos^2 \theta - \frac{m^2}{\sin^2 \theta} - 2a\omega s \cos \theta - \frac{2ms \cos \theta}{\sin^2 \theta} - s^2 \cot^2 \theta + s + A \right), \quad (11.91)$$

where $K \equiv (r^2 + a^2)\omega - am$ and A is the separation constant [128], found as the eigenvalue problem relative to Eq. (10.89). For non-rotating BH ($a = 0$), the angular Teukolsky equation

admits as solution spin-weighted spherical harmonics, ${}_s Y_{\ell m} = {}_s S_{\ell m}(\theta) e^{im\varphi}$, and $A = \ell(\ell + 1) - s(s + 1)$. For rotating BH, the solution of the angular equation are “spin-weighted spheroidal harmonics”, and the separation constant A are known perturbatively for $a\omega \ll 1$, and are given by $A = \ell(\ell + 1) - s(s + 1) - 2a\omega s^2 m / (\ell^2 + \ell) + \mathcal{O}(a^2\omega^2)$ [128]. Because the solutions of the angular problem are known, we treat the radial Eq. (10.88). This can be turned into a 1 dimensional Helmholtz equation, similar to Eq. (10.3), by making the variable redefinition

$$R(r) = \Delta^{-\frac{s+1}{2}} \tilde{R}(r). \quad (11.92)$$

In terms of \tilde{R} , Eq. (10.88) becomes

$$\frac{d^2 \tilde{R}}{dr^2} + \omega^2 \left[1 - 4\tilde{U}_{\ell m}^s(\omega, r) \right] \tilde{R} = 0, \quad (11.93)$$

with

$$4\omega^2 \tilde{U}_{\ell m}^s(\omega, r) = \omega^2 + \frac{(1+s)(a^2 - M^2 + (M-r)^2 s)}{\Delta^2} - \frac{V_R(r)}{\Delta}. \quad (11.94)$$

Note that to write the equation above we collected out a factor of ω^2 , hence we have implicitly assumed $\omega \neq 0$. From the expression in Eq. (10.94) we can already point out one interesting new feature of lensing by BH, compared to a Newtonian lens: differently from the potential $V(r)$, $\tilde{U}_{\ell m}^s(\omega, r)$ depends also on the frequency ω , and not simply on the radial direction. This important difference is introduced both by near-horizon effects, or by the rotation of the BH or by the spin of the perturbation, phenomena that were neglected in the previous discussion but were already predicted in literature [90, 89, 129]. Writing the Teukolsky equation as in Eq. (10.93), allows us to apply immediately all the tools developed in the previous Sections also in this context, and find a solution of the BH perturbation equation in the worldline representation. In particular, the Green function of Eq. (10.93) will be written as

$$G_\omega(r_f, r_i) = -\frac{i}{\omega} \int_0^\infty d\tau e^{i\omega\tau} \int_{r(\tau'=0)=r_i}^{r(\tau'=\tau)=r_f} \mathcal{D}r(\tau') e^{i\omega \int_0^\tau d\tau' \left[\frac{1}{4} \left(\frac{dr}{d\tau'} \right)^2 - V_{\ell m}^s(\omega, r) \right]}, \quad (11.95)$$

with potential given by

$$V_{\ell m}^s(\omega, r) = 1 - \left(\frac{-(1+s)(a^2 - M^2 + (M-r)^2 s)}{\omega^2 \Delta^2} + \frac{V_R(r)}{\omega^2 \Delta} \right). \quad (11.96)$$

With these two equations, we have successfully generalized the standard wave optics treatment for scalar waves to a field with generic spin s , and without requiring the paraxial approximation, nor the Newtonian expansion. Eq. (10.95) is valid for any frequency ω and field spin $|s| = 0, 1/2, 1, 2$, when the lens is a BH of the Kerr family also close to the horizon.

11.7.3 Weak field limit and spin effects

To compare the results of this Section with the previous ones, we take the weak field limit ($r \gg 2M$) for non-rotating ($a = 0$) BH. This will allow us to understand the role of the spin, which was not accounted for in the standard description of the diffraction integral. In the Newtonian limit $\Delta \approx r^2$ and it can be showed that

$$4\tilde{U}_{\ell m}^s(\omega, r) \approx -4\frac{M}{r} + \frac{\ell(\ell + 1) + s(s + 1)}{\omega^2 r^2} - \frac{2is}{\omega r}. \quad (11.97)$$

Therefore, the dynamics of the system is encapsulated in Eq. (10.93), equipped with the potential given in Eq. (10.97). We compare this result against Eq. (10.3) with the gravitational

potential of a point mass lens, namely $U(r) = -M/r$. In order, in Eq. (10.97) we recognize the gravitational potential of the same lens (first term) and the angular momentum contributions. The latter is explicitly present in Eq. (10.97) because we have performed a separation of variables and decomposed the field on a spherical harmonics basis. A similar term can be made manifest also in Eq. (10.3) by performing the same manipulations. Furthermore the potential above, contains two spin-dependent terms, absent in the previous description valid for scalar waves. Among these two, we recognize in the first a contribution to the angular momentum, building up to the total angular momentum $J^2 = L^2 + S^2$. Even though the general theory has these additional spin-dependent contributions, we claim that these two are negligible in the high frequency limit. This is already evident from Eq. (10.97) since the spin dependent terms are always suppressed by different powers of the frequency. Indeed, s is bounded to be smaller than 2, while the angular momentum can take arbitrarily large numbers to keep $l(l+1)/(\omega^2 r^2)$ potentially finite even in the $\omega \gg 1$ limit. Note that this behavior might change close to the horizon. For a non-rotating BH, the potential (10.96) can be written as

$$V_{\ell m}^s(\omega, r) = \frac{4Mr^2(M-r)}{(r^2-2Mr)^2} + \frac{1}{\omega} \frac{2is(r-3M)}{(r-2M)^2} + \frac{1}{\omega^2} \left(\frac{l(l+1)(r^2-2Mr) + M(s^2-1)}{r^2(r-2M)^2} \right), \quad (11.98)$$

where we organized the terms in powers of $1/\omega$. Sufficiently far away from the horizon, where Newtonian theory is valid, the spin-dependent contributions vanish in the limit $\omega \rightarrow \infty$. However, the horizon corrections can boost the spin dependent terms (second and fourth). This fact also tells us that neglecting spin effects in the diffraction integral, as described in Section 10.2, is not an inconsistent assumption, since this relies on the eikonal and paraxial approximations (which pose a lower bound on ω) and it is derived in the weak field limit. We point out that this fact has a clear interpretation in the context outlined in the previous Sections: in the limit where a dispersion relation can be assigned to a wave, these can be described in terms of effective particles propagating on geodesics of the background spacetime regardless of their spin.

11.8 Conclusions

Gravitational lensing of gravitational waves is a rich field of research at the intersection of astrophysics and cosmology. While geometric optics offers a powerful tool for analyzing gravitational waves in the high frequency regime ($\lambda \ll \mathcal{R}_s$), recent years have seen a growing interest in the opposite optical regime, the wave optics one. Its current description relies on the diffraction integral, which we summarized in Section 10.2. While this derivation has the advantage to make manifest the interpretation of the field in terms of the associated particle, and smoothly transition to geometric optics in the high-frequency limit, it is largely based on the *paraxial* and *eikonal* approximations, effectively posing a lower frequency limit on the validity range of such a description. Moving beyond these approximations is crucial. It not only deepens our theoretical understanding of the behavior of gravitational waves in curved spacetimes, but it also becomes vital as efforts intensify to probe the Universe's matter content using low-frequency gravitational lensing.

In this work, we proposed a generalization of the diffraction integral based on the introduction of a fifth coordinate, called *proper time* (see Section 10.3). The introduction of this new variable allows us to seamlessly transition from the standard Helmholtz equation for wave propagation, to a Schrödinger-like one, without the need of any approximation. This procedure led us directly to the *Feynman-Fradkin* exact representation of the Green function in Eq. (10.21) for a scalar field propagating through a spacetime containing a gravitational lens in the Newtonian regime. Just as in the diffraction integral, also in our formalism the wave's frequency plays

the role of $1/\hbar$, allowing us to set up the Hamilton analogy and compare wave effects to the quantum ones. The worldline representation of the Green function, $\tilde{G}_\omega(\mathbf{x}_f, \mathbf{x}_i, \tau)$, is once again constructed out of particle-like path integrals over the paths, rather than one over different field's configurations, making possible a direct link to the laws of high-frequency optics, since the proper time is removed upon integration at the end.

In Section 10.4, we investigated the high-frequency limit ($\omega \rightarrow \infty$) using three approaches. First, we analyze the WKB form of the propagator through saddle point approximations for both proper time and trajectories. This recovers the geometric optics (GO) ansatz in Eq. (10.35). Next, aiming to obtain the diffraction integral and eliminate the proper time integration while preserving the path integral over trajectories, we assume the ability to interchange the integration order within the Green's function. This is an inexact procedure, but established as a good approximation in the high-frequency regime [110, 60]. Following the order swap, the proper time integral is evaluated in two ways. One way involves strictly fixing for the proper time τ at its classical configuration satisfying the equations of motion (10.27). The other way implements a saddle point approximation that considers second-order deviations from the classical value $\sim (\tau - \tau_{cl})^2$. The first approach yields the diffraction integral, demonstrating how our formalism encompasses the standard approach. The second approach leads to the well-known *Feynman-Garrod* direct representation of the propagator.

After assessing the high-frequency behavior, we implemented a perturbative approach in Section 10.5 with respect to the associated potential, considering the limit $\omega\alpha \ll 1$. We demonstrated how the expansion series can be organized in terms of waves propagating freely, except for an always increasing number of scattering events. This is equivalent to organizing the perturbation series such that terms are always of higher orders in the Born approximation. In the ideal limit of $n \rightarrow \infty$, where n represents the perturbation order, the wave scatters with every point of the gravitational potential. Consequently, there is no more free propagation within the region where $V \neq 0$. This observation suggests the possibility of an exact resummation of the perturbation scheme. Indeed, it leads us to the Dyson equation for the wave's propagator, where the gravitational potential clearly plays the role of proper self-energy, as also suggested by the Hamiltonian derivation of Section 10.3.2. Our path integral description also allows for the introduction of a generating functional, Z , implemented through the introduction of a current, Z , coupled to the path $\mathbf{x}(\tau')$.

Subsequently, we explored two explicit applications of our formalism. Firstly, in Section 10.6.1, we present explicit computations for the first-order Green function in the case of a Coulomb potential, providing an estimate of the modifications to the free Green function, or the distance to the source, induced by the presence of the lens in the $\alpha\omega \ll 1$ limit. Our results demonstrate the characteristic divergence in the forward limit of a long-range interaction, alongside the expected attenuation of the effect as we focus in the diffractive limit $\omega MG \rightarrow 0$ or $\lambda \gg \mathcal{R}_s$, where λ represents the wavelength of the wave and \mathcal{R}_s denotes the Schwarzschild radius of the lens. Then, in Section 10.6.2, we apply the formalism to the case of a massive scalar field demonstrating that our method can be straightforwardly implemented by performing a frequency redefinition.

In Section 10.7, we addressed another limitation of the diffraction integral: its neglect of polarization effects, as it is valid only for scalar waves. Previous treatments of the diffractive regime often assume the conservation of polarization associated with gravitational waves during propagation, potentially overlooking valuable insights.

To explore polarization effects, we specialize the lens model to BHs with axisymmetry, and consider a Kerr background. This choice streamlines the derivation, allowing us to leverage results from the established field of BH perturbation theory. For perturbations of various spin ($s = 0, 1/2, 1, 2$), the equations of motion on the Kerr background are well-known in the literature and are described by the Teukolsky equation, in terms of Newman-Penrose scalar. After

separating the radial and angular components, we reformulate the radial part of the Teukolsky equation as a one-dimensional Helmholtz equation, with a potential dependent on frequency, spin, and angular momentum. This formulation enables us to derive the Feynman-Fradkin representation of the propagator for higher spin fields. The role of spin has then been investigated by specializing to the case of non-rotating BH by setting $a = 0$, obtaining two main equations. One which identifies spin-dependent effects within the same regime of validity as the initial Helmholtz equation evaluated for a spherically symmetric lens. Particularly, this result shows that one of the two spin-dependent contributions, $s(s+1)/(\omega^2 r^2)$, contributes to the total angular momentum, with both contributions being suppressed by higher powers of frequency. The other equation demonstrates that spin contributions can become significant in the near-horizon region, where the largeness of ω can be offset the vanishing denominators proportional to $r - 2M$. However, the diffraction integral is valid only for Newtonian lenses, thus it cannot account for near-horizon contributions. These results confirm that neglecting polarization effects in the derivation of the diffraction integral aligns with the underlying assumptions of the eikonal limit (with the frequency lower bound from Eq. (10.6)) and the weak field $\alpha \ll 1$ approximation.

Applying our formalism to specific lensing scenario and assessing its potentiality into investigating the properties of the lenses, will be object of future works.

11.9 Free propagator computations

In this appendix we provide the explicit steps to compute the free propagator $G_\omega(\mathbf{x}_f, \mathbf{x}_i)$ following [51]. We start from its worldline representation, namely Eq. (10.54), where

$$S_0 \equiv \int_0^\tau d\tau' \frac{1}{4} \left(\frac{d\mathbf{x}(\tau')}{d\tau'} \right)^2. \quad (11.99)$$

The first step to perform the path integral, is to redefine the path and consider a perturbation of it from the classical solution, i.e. we define $\mathbf{x}(\tau') = \mathbf{x}_{cl}(\tau') + \mathbf{q}(\tau')$. Since the total and the classical paths satisfy the same boundary conditions, we have that $\mathbf{q}(\tau' = \tau) = \mathbf{q}(\tau' = 0) = 0$. The classical trajectory $\mathbf{x}_{cl}(\tau)$ satisfies the Euler-Lagrange equation of motion, descending from the free action S_0 , namely $\ddot{\mathbf{x}}_{cl} = 0$. This means

$$\dot{\mathbf{x}}_{cl} = \text{const}, \quad \rightarrow \quad \dot{\mathbf{x}}_{cl} = \frac{\mathbf{x}_f - \mathbf{x}_i}{\tau} \quad (11.100)$$

as the motion between \mathbf{x}_f and \mathbf{x}_i occurs in a time period of τ at constant speed. This allows us to write the free action as

$$S_0 = \frac{|\mathbf{x}_f - \mathbf{x}_i|^2}{4\tau^2} + \int_0^\tau d\tau' \frac{\dot{\mathbf{q}}^2}{4}. \quad (11.101)$$

We evaluate the path integral of Eq. (10.54) as

$$\begin{aligned} G_\omega^{(0)}(\mathbf{x}_f, \mathbf{x}_i, \tau) &= \int_{\mathbf{x}(\tau'=0)=\mathbf{x}_i}^{\mathbf{x}(\tau'=\tau)=\mathbf{x}_f} \mathcal{D}\mathbf{x}(\tau') \exp \left[i\omega \int_0^\tau d\tau' \frac{\dot{\mathbf{x}}^2}{4} \right] = \\ &= e^{\frac{i\omega}{4\tau} |\mathbf{x}_f - \mathbf{x}_i|^2} \int_{\mathbf{q}(\tau'=0)=0}^{\mathbf{q}(\tau'=\tau)=0} \mathcal{D}\mathbf{q}(\tau') \exp \left[i\omega \int_0^\tau d\tau' \frac{\dot{\mathbf{q}}^2}{4} \right] = \\ &= \left(\frac{\omega}{4\pi i\tau} \right)^{3/2} e^{\frac{i\omega}{4\tau} |\mathbf{x}_f - \mathbf{x}_i|^2}. \end{aligned} \quad (11.102)$$

Finally, the propagator is obtained by performing the proper time integral

$$G_\omega^{(0)}(\mathbf{x}_f, \mathbf{x}_i) = -\frac{i}{\omega} \int_0^{+\infty} d\tau \left(\frac{\omega}{4\pi i\tau} \right)^{3/2} e^{\frac{i\omega}{4\tau} |\mathbf{x}_f - \mathbf{x}_i|^2 + i\omega\tau} = -\frac{1}{4\pi} \frac{e^{i\omega|\mathbf{x}_f - \mathbf{x}_i|}}{|\mathbf{x}_f - \mathbf{x}_i|}, \quad (11.103)$$

so that, in real time, we have

$$G^{(0)}(\mathbf{x}_f, \mathbf{x}_i, t) = -\frac{1}{4\pi} \int \frac{d\omega}{2\pi} \frac{e^{-i\omega(t-|\mathbf{x}_f-\mathbf{x}_i|)}}{|\mathbf{x}_f-\mathbf{x}_i|} = -\frac{1}{8\pi^2} \frac{\delta(t-|\mathbf{x}_f-\mathbf{x}_i|)}{|\mathbf{x}_f-\mathbf{x}_i|}, \quad (11.104)$$

displaying the correct behavior (up to a normalization factor). To obtain the momentum space Green function, we perform the 3D Fourier transform

$$\begin{aligned} G_\omega^{(0)}(\mathbf{p}) &= -\frac{1}{4\pi} \int d^3x \frac{e^{i\omega|\mathbf{x}_f-\mathbf{x}_i|-i\mathbf{p}\cdot(\mathbf{x}_f-\mathbf{x}_i)}}{|\mathbf{x}_f-\mathbf{x}_i|} = \\ &= -\frac{1}{4\pi} \int_0^{2\pi} d\varphi \int_{-1}^1 d\mu \int_0^{+\infty} dx x^2 \frac{e^{i\omega x - i p x \mu}}{x} = \\ &= \frac{1}{p^2 - \omega^2}. \end{aligned} \quad (11.105)$$

11.10 Coulomb potential computations

An interesting example!!!

In this Appendix we provide the explicit computations of Section 10.6.1. In particular, we illustrate the derivation of Eqs. (10.74) and (10.75). First we point out that the integral in Eq. (10.72) can be easily performed using (see [103])

$$\int_0^T d\tau \frac{e^{-\frac{a}{T-\tau}-\frac{b}{\tau}}}{[\tau(T-\tau)]^{3/2}} = \sqrt{\frac{\pi}{T^3}} \frac{\sqrt{a} + \sqrt{b}}{\sqrt{ab}} e^{-\frac{(\sqrt{a}+\sqrt{b})^2}{T}}. \quad (11.106)$$

Then we show the derivation of Eq. (10.74), namely the worldline form of the first order propagator. As explained in the main text, we consider a scattering situation in which the wave's source and observer are located very far away from the lens. In practice, this means considering the approximation in Eq. (10.73) for $|\mathbf{x}_O - \mathbf{x}^*|$ and $|\mathbf{x}^* - \mathbf{x}_S|$ and approximate the result of Eq. (10.72) considering the lowest order in r^*/r_S and r^*/r_O in the multiplicative parenthesis, while retaining the first order correction in the phase of the exponential. Therefore, we consider the approximation of Eq. (10.72)

$$\begin{aligned} \int_0^\tau d\tau' \tilde{G}_\omega^{(0)}(\mathbf{x}_O, \mathbf{x}^*, \tau - \tau') V(\mathbf{x}^*) \tilde{G}_\omega^{(0)}(\mathbf{x}^*, \mathbf{x}_S, \tau') = \\ \approx \left(\frac{\omega}{4i\pi}\right)^{5/2} \tau^{-3/2} \left(\frac{1}{r_O} + \frac{1}{r_S}\right) e^{\frac{i\omega}{4\tau}(r_S+r_O)^2} e^{\frac{i\omega}{2\tau}\mathbf{x}^*\cdot\mathbf{R}} V(\mathbf{x}^*) \end{aligned} \quad (11.107)$$

where we have introduced the vector $\mathbf{R} \equiv (r_S + r_O)(\boldsymbol{\theta}_S - \boldsymbol{\theta}_O)$. Now we compute the integral

over the interaction point \mathbf{x}^* as prescribed in Eq. (10.60). The steps to obtain Eq. (10.74) are

$$\begin{aligned}
\tilde{G}_\omega^{(1)}(\mathbf{x}_O, \mathbf{x}_S, \tau) &= \left(\frac{\omega}{4i\pi}\right)^{5/2} \tau^{-3/2} \left(\frac{1}{r_O} + \frac{1}{r_S}\right) e^{\frac{i\omega}{4\tau}(r_S+r_O)^2} \int d^3\mathbf{x}^* e^{\frac{i\omega}{2\tau}\mathbf{x}^*\cdot\mathbf{R}} V(\mathbf{x}^*) = \\
&= -4\alpha MG \left(\frac{\omega}{4i\pi}\right)^{5/2} \tau^{-3/2} \left(\frac{1}{r_O} + \frac{1}{r_S}\right) e^{\frac{i\omega}{4\tau}(r_S+r_O)^2} \times \\
&\quad \times \int_0^{2\pi} d\varphi^* \int_{-1}^1 d\mu^* \int_0^{+\infty} dr^* r^{*2} \frac{e^{\frac{i\omega}{2\tau}Rr^*\mu^*-r^*/a}}{r^*} = \\
&= -8\pi\alpha MG \left(\frac{\omega}{4i\pi}\right)^{5/2} \tau^{-3/2} \left(\frac{1}{r_O} + \frac{1}{r_S}\right) e^{\frac{i\omega}{4\tau}(r_S+r_O)^2} \left(\frac{2\tau}{i\omega R}\right) \times \\
&\quad \times \int_0^{+\infty} dr^* \left[e^{r^*\left[\frac{i\omega}{2\tau}R-\frac{1}{a}\right]} - e^{-r^*\left[\frac{i\omega}{2\tau}R+\frac{1}{a}\right]} \right] = \\
&= -8\pi\alpha MG \left(\frac{\omega}{4i\pi}\right)^{5/2} \tau^{-3/2} \left(\frac{1}{r_O} + \frac{1}{r_S}\right) e^{\frac{i\omega}{4\tau}(r_S+r_O)^2} \left(\frac{2\tau}{i\omega R}\right) \times \\
&\quad \times \left[\frac{e^{r^*\left[\frac{i\omega}{2\tau}R-\frac{1}{a}\right]}}{\frac{i\omega}{2\tau}R-\frac{1}{a}} + \frac{e^{-r^*\left[\frac{i\omega}{2\tau}R+\frac{1}{a}\right]}}{\frac{i\omega}{2\tau}R+\frac{1}{a}} \right]_0^{+\infty} = \\
&= 16\pi\alpha MG \left(\frac{\omega}{4i\pi}\right)^{5/2} \left(\frac{1}{r_O} + \frac{1}{r_S}\right) \times \frac{\tau^{-3/2} e^{\frac{i\omega}{4\tau}(r_S+r_O)^2}}{\left(\frac{i\omega R}{2\tau}\right)^2 - \frac{1}{a^2}}, \tag{11.108}
\end{aligned}$$

where μ^* is the cosine of the angle between \mathbf{R} and \mathbf{x}^* . We use this result to perform the integral over the proper time and obtain the first order propagator. That is, we explicitly show the integration that leads to Eq. (10.75), where we have considered a long range interaction and send $a \rightarrow \infty$. These are

$$\begin{aligned}
G_\omega^{(1)}(\mathbf{x}_O, \mathbf{x}_S) &\equiv -\frac{i}{\omega} \int_0^{+\infty} d\tau e^{i\omega\tau} \tilde{G}_\omega^{(1)}(\mathbf{x}, \mathbf{x}_O, \tau) = \\
&= 16\pi\alpha MG \left(-\frac{i}{\omega}\right) \left(\frac{\omega}{4i\pi}\right)^{5/2} \left(\frac{1}{r_O} + \frac{1}{r_S}\right) \left(\frac{-4}{\omega^2 R^2}\right) \int_0^{+\infty} d\tau \sqrt{\tau} e^{\frac{i\omega}{4\tau}(r_S+r_O)^2+i\omega\tau} = \\
&= 16\pi\alpha MG \left(-\frac{i}{\omega}\right) \left(\frac{\omega}{4i\pi}\right)^{5/2} \left(\frac{1}{r_O} + \frac{1}{r_S}\right) \left(\frac{-4}{\omega^2 R^2}\right) \times \\
&\quad \times \left(\frac{-\sqrt{-i\pi\omega}}{2\omega^2}\right) e^{i\omega(r_S+r_O)} (1 - i\omega(r_S+r_O)) = \\
&= \frac{\alpha MG}{\pi\omega^2} \left(\frac{1}{r_O} + \frac{1}{r_S}\right) \frac{e^{i\omega(r_S+r_O)} (1 - i\omega(r_S+r_O))}{R^2}. \tag{11.109}
\end{aligned}$$

Note that, in order to perform the integral, we regularized it by analytically continuing ω in the complex plane, with the prescription $\omega \rightarrow \omega + i\epsilon$, and then took the limit $\epsilon \rightarrow 0$.

Part XII

Preparation to Wave Optics for Gravitational Lensing

Part XIII

Other Experiments

Part XIV

Appendix

.1 Appendix

.1.1 Literature

Bibliography

References

From Regimes in astrophysical lensing: refractive optics, diffractive optics, and the Fresnel scale by Dylan L. Jow, Ue-Li Pen, Job Feldbrugge

References

- [1] Beach T. L., Lovelace R. V. E., 1997, *Radio Science*, **32**, 913
- [2] Brisken W. F., Macquart J. P., Gao J. J., Rickett B. J., Coles W. A., Deller A. T., Tingay S. J., West C. J., 2010, *ApJ*, **708**, 232, [arXiv:0910.5654](#)
- [3] Cherman A., Unsal M., 2014, arXiv e-prints, p. [arXiv:1408.0012](#), [arXiv:1408.0012](#)
- [4] Clegg A. W., Fey A. L., Lazio T. J. W., 1998, *ApJ*, **496**, 253, [arXiv:astro-ph/9709249](#)
- [5] Cognard I., Bourgois G., Lestrade J.-F., Biraud F., Aubry D., Darchy B., Drouhin J.-P., 1993, *NAT*, **366**, 320
- [6] Coles W. A., et al., 2015, *APJ*, **808**, 113, [arXiv:1506.07948](#)
- [7] Cordes J. M., Wolszczan A., 1986, *ApJ*, **307**, L27
- [8] Cordes J. M., Pidwerbetsky A., Lovelace R. V. E., 1986, *ApJ*, **310**, 737
- [9] Cordes J. M., Wasserman I., Hessels J. W. T., Lazio T. J. W., Chatterjee S., Wharton R. S., 2017, *APJ*, **842**, 35, [arXiv:1703.06580](#)
- [10] Dai S., Johnston S., Bell M. E., Coles W. A., Hobbs G., Ekers R. D., Lenc E., 2016, *MNRAS*, **462**, 3115, [arXiv:1607.07740](#)
- [11] Dong L., Petropoulou M., Giannios D., 2018, *MNRAS*, **481**, 2685
- [12] Er X., Wagner J., Mao S., 2022, *MNRAS*, **509**, 5872, [arXiv:2111.09104](#)
- [13] Feldbrugge J., 2020, arXiv e-prints, p. [arXiv:2010.03089](#), [arXiv:2010.03089](#)
- [14] Feldbrugge J., Turok N., 2020, arXiv e-prints, p. [arXiv:2008.01154](#), [arXiv:2008.01154](#)
- [15] Feldbrugge J., Pen U.-L., Turok N., 2019, arXiv e-prints, p. [arXiv:1909.04632](#), [arXiv:1909.04632](#)
- [16] Fiedler R. L., Dennison B., Johnston K. J., Hewish A., 1987, *Nature*, **326**, 675
- [17] Goldreich P., Sridhar S., 2006, *ApJ*, **640**, L159, [arXiv:astro-ph/0602532](#)
- [18] Gozzi E., Pagani C., Reuter M., 2021, *Annals of Physics*, **429**, 168457, [arXiv:2004.08874](#)
- [19] Grillo G., Cordes J., 2018, arXiv e-prints, p. [arXiv:1810.09058](#), [arXiv:1810.09058](#)
- [20] Johnson M. D., et al., 2018, *ApJ*, **865**, 104, [arXiv:1808.08966](#)
- [21] Jow D. L., Foreman S., Pen U.-L., Zhu W., 2020, *MNRAS*, **497**, 4956, [arXiv:2002.01570](#)
- [22] Jow D. L., Lin F. X., Tyhurst E., Pen U.-L., 2021, *MNRAS*, , [arXiv:2103.08687](#)

- [23] Kerr M., Coles W. A., Ward C. A., Johnston S., Tuntsov A. V., Shannon R. M., 2018, *MNRAS*, 474, 4637, [arXiv:1712.00426](#)
- [24] Lee L. C., Jokipii J. R., 1975, *ApJ*, 196, 695
- [25] Lewandowski W., Kijak J., Gupta Y., Krzeszowski K., 2011, *AAP*, 534, A66, [arXiv:1107.3052](#)
- [26] Lin F. X., Main R., Li D., Pen U.-L., van Kerkwijk M. H., 2022, arXiv e-prints, p. [arXiv:2208.13868](#), [arXiv:2208.13868](#)
- [27] Lorimer D. R., Kramer M., 2012, *Handbook of Pulsar Astronomy*
- [28] Main R., et al., 2018, *NAT*, 557, 522, [arXiv:1805.09348](#)
- [29] Nakamura T. T., Deguchi S., 1999, *Progress of Theoretical Physics Supplement*, 133, 137
- [30] Nye J. F., 1999, Natural focusing and fine structure of light: caustics and wave dislocations
- [20] Pen U.-L., King L., 2012, *MNRAS*, 421, L132, [arXiv:1111.6806](#)
- [17] Pen U.-L., Levin Y., 2014, *MNRAS*, 442, 3338, [arXiv:1302.1897](#)
- [33] Rickett B. J., 1977, *ARA&A*, 15, 479
- [52] Rickett B. J., 1990, *ARA&A*, 28, 561
- [35] Rickett B. J., Lyne A. G., Gupta Y., 1997, *MNRAS*, 287, 739
- [36] Salpeter E. E., 1967, *APJ*, 147, 433
- [37] Scheuer P. A. G., 1968, *NAT*, 218, 920
- [1] Schneider P., Ehlers J., Falco E. E., 1992, *Gravitational Lenses*, [doi:10.1007/978-3-662-03758-4](#).
- [39] Shi X., Xu Z., 2021, *MNRAS*, 506, 6039, [arXiv:2104.00940](#)
- [40] Sprenger T., Wucknitz O., Main R., Baker D., Briskin W., 2021, *MNRAS*, 500, 1114, [arXiv:2008.09443](#)
- [41] Stinebring D. R., McLaughlin M. A., Cordes J. M., Becker K. M., Goodman J. E. E., Kramer M. A., Shekard J. L., Smith C. T., 2001, *ApJ*, 549, L97, [arXiv:astro-ph/0010363](#)
- [42] Tanizaki Y., Koike T., 2014, *Annals of Physics*, 351, 250, [arXiv:1406.2386](#)
- [43] Tuntsov A. V., Walker M. A., Koopmans L. V. E., Bannister K. W., Stevens J., Johnston S., Reynolds C., Bignall H. E., 2016, *APJ*, 817, 176, [arXiv:1512.03411](#)
- [44] Turok N., 2014, *New Journal of Physics*, 16, 063006, [arXiv:1312.1772](#)
- [45] Walker M. A., Melrose D. B., Stinebring D. R., Zhang C. M., 2004, *MNRAS*, 354, 43, [arXiv:astro-ph/0403587](#)
- [46] Wang N., Manchester R. N., Johnston S., Rickett B., Zhang J., Yusup A., Chen M., 2005, *MNRAS*, 358, 270, [arXiv:astro-ph/0501218](#)
- [47] Werner M. C., Evans N. W., 2006, *MNRAS*, 368, 1362, [arXiv:astro-ph/0602368](#)

- [48] Witten E., 2010, arXiv e-prints, p. [arXiv:1001.2933](https://arxiv.org/abs/1001.2933), [arXiv:1001.2933](https://arxiv.org/abs/1001.2933)
- [49] Wright F. J., 1980, *Journal of Physics A: Mathematical and General*, 13, 2913

From Wave asymptotics and their applic to astrophys plasma lensing

References

- [1] Bannister, K. W., Stevens, J., Tuntsov, A. V., et al. 2016, *Science*, 351, 354
- [2] Berry, M. V. 1976, *AdPhy*, 25, 1
- [3] Berry, M. V., & Upstill, C. 1980, in *Progress in Optics*, ed. E. Wolf, Vol. 18 (Elsevier), 257–346
- [4] Bleistein, N., & Handelsman, R. A. 1975, *Asymptotic expansions of integrals* (Dover Publications)
- [5] Born, M., & Wolf, E. 1999, *Principles of Optics*, seventh edn. (Cambridge University Press)
- [6] Borovikov, V. A., & Kinber, B. E. 1994, *Geometrical theory of diffraction*, *Electromagnetic Wave Series No. 37* (The Institution of Electrical Engineers)
- [7] Budden, K. G., & Terry, P. D. 1971, *RSPSA*, 321, 275
- [8] Chako, N. 1965, *JApMa*, 1, 372
- [9] Chester, C., Friedman, B., & Ursell, F. 1957, 53, 599
- [10] Clegg, A. W., Fey, A. L., & Fiedler, R. L. 1996, *ApJ*, 457, L23
- [11] Clegg, A. W., Fey, A. L., & Lazio, T. J. W. 1998, *ApJ*, 496, 253
- [12] Cognard, I., Bourgois, G., Lestrade, J.-F., et al. 1993, *Nature*, 366, 320
- [13] Coles, W. A., Filice, J. P., Frehlich, R. G., & Yadlowsky, M. 1995, *ApOpt*, 34, 2089
- [14] Coles, W. A., Kerr, M., Shannon, R. M., et al. 2015, *ApJ*, 808, 113
- [15] Connor, J. 1973, *MolPh*, 25, 181
- [16] Connor, J. N. L. 1973, *MolPh*, 26, 1217
- [17] Cooke, J. C. 1982, *JApMa*, 29, 25
- [18] Cordes, J. M., Pidwerbetsky, A., & Lovelace, R. E. 1986, *ApJ*, 310, 737
- [19] Cordes, J. M., Rickett, B. J., Stinebring, D. R., & Coles, W. A. 2006, *ApJ*, 637, 346
- [20] Cordes, J. M., & Shannon, R. M. 2010. <https://arxiv.org/abs/1010.3785>
- [21] Cordes, J. M., Shannon, R. M., & Stinebring, D. R. 2016, *ApJ*, 817, 16
- [22] Cordes, J. M., & Wasserman, I. 2016, *MNRAS*, 457, 232
- [23] Cordes, J. M., Wasserman, I., Hessels, J. W. T., et al. 2017, *ApJ*, 842, 10

- [24] Cordes, J. M., & Wolszcan, A. 1986, *ApJ*, 307, L27
- [25] Dai, L., & Lu, W. 2017, *ApJ*, 847, 19
- [26] Er, X., & Rogers, A. 2017, *MNRAS*, 475, 867
- [27] Fiedler, R., Dennison, B., Johnston, K. J., et al. 1994, *ApJ*, 430, 581
- [28] Fiedler, R. L., Dennison, B., Johnston, K. J., & Hewish, A. 1987, *Nature*, 326, 675
- [29] Goodman, J. J., Romani, R. W., Blandford, R. D., & Narayan, R. 1987, *MNRAS*, 229, 73
- [30] Gupta, Y., Ghat, N. D. R., & Rao, A. P. 1999, *ApJ*, 520, 173
- [31] Gupta, Y., Rickett, B. J., & Lyne, A. G. 1994, *MNRAS*, 269, 1035
- [32] Howls, C. J. 1997, *RSPSA*, 453, 2271
- [33] Kaminski, D. 1994, *MApAn*, 1, 44
- [34] Katsaounis, T., Kossioris, G. T., & Makrakis, G. N. 2001, *Math. Models Meth. Appl. Sci.*, 11, 199
- [35] Keith, M. J., Coles, W. A., Shannon, R. M., et al. 2013, *MNRAS*, 429, 2161
- [36] Kerr, M., Coles, W. A., Ward, C. A., et al. 2017, *MNRAS*, 474, 4637
- [37] Kravstov, Y. A. 1968, *SPhAc*, 14, 1
- [38] Kravtsov, Y. A., Forbes, G. W., & Asatryan, A. A. 1999, in *Progress in optics*, ed. E. Wolf, Vol. 39 (Elsevier), 1–62
- [39] Kravtsov, Y. A., & Orlov, Y. I. 1999, *Springer Series on Wave Phenomena*, Vol. 15, *Caustics, catastrophes and wave fields*, 2nd edn. (Springer)
- [40] Kryukovskii, A. S., Lukin, D. S., Palkin, E. A., & Rastyagaev, D. S. 2006, *Journal of Communications Technology and Electronics*, 51, 1087
- [41] Lam, M. T., Ellis, J. A., Grillo, G., et al. 2018, *ApJ*, 861, 132
- [42] Ludwig, D. 1966, *Commun. Pure Appl. Math.*, 19, 215
- [43] Main, R., Yang, I. S., Chan, V., et al. 2018, *Nature*, 557, 522
- [44] Melrose, D. B., & Watson, P. G. 2006, *ApJ*, 647, 1131
- [45] Nye, J. F. 1978, *RSPSA*, 361, 21
- [46] Paris, R. B. 1991, *RSPSA*, 432, 391
- [47] Pearcey, T. 1946, *PMag*, 37, 311
- [48] Pen, U. L., & King, L. 2012, *MNRAS*, 421, L132
- [49] Poston, T., & Stewart, I. 1978, *Catastrophe Theory and its Applications* (Dover Publications)
- [50] Pushkarev, A. B., Kovalev, Y. Y., Lister, M. L., et al. 2013, *A&A*, 555, A80

- [51] Qiu, W.-Y., & Wong, R. 2000, RSPSA, 456, 407
- [52] Rickett, B. J. 1990, ARA&A, 28, 561
- [53] Schmidt, J. D. 2010, Numerical simulation of optical wave propagation with examples in MATLAB
- [54] Schneider, P., Ehlers, J., & Falco, E. E. 1992, Gravitational Lenses (Springer)
- [55] Schramm, T., & Kayser, R. 1987, A&A, 174, 361
- [56] Schramm, T., & Kayser, R. 1995, A&A, 299, 1
- [57] Stamnes, J. J. 1986, Waves in Focal Regions: Propagation, Diffraction and Focusing of Light, Sound and Water Waves (Adam Hilger)
- [58] Stinebring, D., Matters, J., & Hemberger, D. 2007, in Astronomical Society of the Pacific Conference Series, Vol. 365, SINS - Small Ionized and Neutral Structures in the Diffuse Interstellar Medium, ed. M. Haverkorn & W. M. Goss, 275
- [59] Thom, R. 1972, Stabilité structurelle et morphogenèse (Benjamin)
- [60] Thorne, K. S., & Blandford, R. D. 2017, Modern Classical Physics (Princeton University Press)
- [61] Tuntsov, A. V., Walker, M. A., Koopmans, L. V. E., et al. 2016, ApJ, 817, 176
- [62] Ursell, F. 1965, PCPS, 61, 113
- [63] Vedantham, H. K., Readhead, A. C. S., Hovatta, T., et al. 2017, ApJ, 845, 90
- [64] Watson, P. G., & Melrose, D. B. 2006, ApJ, 647, 1142
- [65] Wong, R. 2001, Asymptotic approximations of integrals, Vol. 34 (SIAM)

From Proper time path int for GW - an improved wave optics framework

References

- [1] R. Takahashi and T. Nakamura, *Wave effects in gravitational lensing of gravitational waves from chirping binaries*, *Astrophys. J.* **595** (2003) 1039–1051, [[astro-ph/0305055](#)].
- [2] Z. Gao, X. Chen, Y.-M. Hu, J.-D. Zhang, and S.-J. Huang, *A higher probability of detecting lensed supermassive black hole binaries by LISA*, *Mon. Not. Roy. Astron. Soc.* **512** (2022), no. 1 1–10, [[arXiv:2102.10295](#)].
- [3] L. Dai, S.-S. Li, B. Zackay, S. Mao, and Y. Lu, *Detecting Lensing-Induced Diffraction in Astrophysical Gravitational Waves*, *Phys. Rev. D* **98** (2018), no. 10 104029, [[arXiv:1810.00003](#)].
- [4] M. H. Y. Cheung, J. Gais, O. A. Hannuksela, and T. G. F. Li, *Stellar-mass microlensing of gravitational waves*, *Mon. Not. Roy. Astron. Soc.* **503** (2021), no. 3 3326–3336, [[arXiv:2012.07800](#)].

- [5] D. Reitze et al., *Cosmic Explorer: The U.S. Contribution to Gravitational-Wave Astronomy beyond LIGO*, *Bull. Am. Astron. Soc.* **51** (2019), no. 7 035, [[arXiv:1907.04833](#)].
- [6] M. Punturo et al., *The Einstein Telescope: A third-generation gravitational wave observatory*, *Class. Quant. Grav.* **27** (2010) 194002.
- [7] M. Branchesi et al., *Science with the Einstein Telescope: a comparison of different designs*, *JCAP* **07** (2023) 068, [[arXiv:2303.15923](#)].
- [8] **LISA Consortium Waveform Working Group** Collaboration, N. Afshordi et al., *Waveform Modelling for the Laser Interferometer Space Antenna*, [arXiv:2311.01300](#).
- [9] M. Çalıřkan, L. Ji, R. Cotesta, E. Berti, M. Kamionkowski, and S. Marsat, *Observability of lensing of gravitational waves from massive black hole binaries with LISA*, *Phys. Rev. D* **107** (2023), no. 4 043029, [[arXiv:2206.02803](#)].
- [10] G. Tambalo, M. Zumalacárregui, L. Dai, and M. H.-Y. Cheung, *Gravitational wave lensing as a probe of halo properties and dark matter*, *Phys. Rev. D* **108** (2023), no. 10 103529, [[arXiv:2212.11960](#)].
- [11] K.-H. Lai, O. A. Hannuksela, A. Herrera-Martín, J. M. Diego, T. Broadhurst, and T. G. Li, *Discovering intermediate-mass black hole lenses through gravitational wave lensing*, *Physical Review D* **98** (Oct., 2018).
- [12] S. Savastano, G. Tambalo, H. Villarrubia-Rojo, and M. Zumalacarregui, *Weakly lensed gravitational waves: Probing cosmic structures with wave-optics features*, *Phys. Rev. D* **108** (2023), no. 10 103532, [[arXiv:2306.05282](#)].
- [13] S. M. C. Yeung, M. H. Y. Cheung, E. Seo, J. A. J. Gais, O. A. Hannuksela, and T. G. F. Li, *Detectability of microlensed gravitational waves*, *Mon. Not. Roy. Astron. Soc.* **526** (2023), no. 2 2230–2240, [[arXiv:2112.07635](#)].
- [14] J. M. Diego, O. A. Hannuksela, P. L. Kelly, T. Broadhurst, K. Kim, T. G. F. Li, G. F. Smoot, and G. Pagano, *Observational signatures of microlensing in gravitational waves at LIGO/Virgo frequencies*, *Astron. Astrophys.* **627** (2019) A130, [[arXiv:1903.04513](#)].
- [15] X. Guo and Y. Lu, *Probing the nature of dark matter via gravitational waves lensed by small dark matter halos*, *Phys. Rev. D* **106** (2022), no. 2 023018, [[arXiv:2207.00325](#)].
- [16] S. Jung and C. S. Shin, *Gravitational-Wave Fringes at LIGO: Detecting Compact Dark Matter by Gravitational Lensing*, *Phys. Rev. Lett.* **122** (2019), no. 4 041103, [[arXiv:1712.01396](#)].
- [17] M. Fairbairn, J. Urrutia, and V. Vaskonen, *Microlensing of gravitational waves by dark matter structures*, *JCAP* **07** (2023) 007, [[arXiv:2210.13436](#)].
- [18] J. Urrutia and V. Vaskonen, *The dark timbre of gravitational waves*, [arXiv:2402.16849](#).
- [19] J. Urrutia, V. Vaskonen, and H. Veermäe, *Gravitational wave microlensing by dressed primordial black holes*, *Phys. Rev. D* **108** (2023), no. 2 023507, [[arXiv:2303.17601](#)].
- [20] J. Urrutia and V. Vaskonen, *Lensing of gravitational waves as a probe of compact dark matter*, *Mon. Not. Roy. Astron. Soc.* **509** (2021), no. 1 1358–1365, [[arXiv:2109.03213](#)].

- [21] S. Sugiyama, T. Kurita, and M. Takada, *On the wave optics effect on primordial black hole constraints from optical microlensing search*, *Mon. Not. Roy. Astron. Soc.* **493** (2020), no. 3 3632–3641, [[arXiv:1905.06066](#)].
- [22] J. M. Diego, *Constraining the abundance of primordial black holes with gravitational lensing of gravitational waves at LIGO frequencies*, *Phys. Rev. D* **101** (2020), no. 12 123512, [[arXiv:1911.05736](#)].
- [23] M. Oguri and R. Takahashi, *Probing Dark Low-mass Halos and Primordial Black Holes with Frequency-dependent Gravitational Lensing Dispersions of Gravitational Waves*, *Astrophys. J.* **901** (2020), no. 1 58, [[arXiv:2007.01936](#)].
- [24] H. Gil Choi, S. Jung, P. Lu, and V. Takhistov, *Co-Existence Test of Primordial Black Holes and Particle Dark Matter*, [arXiv:2311.17829](#).
- [25] S. Basak, A. Ganguly, K. Haris, S. Kapadia, A. K. Mehta, and P. Ajith, *Constraints on Compact Dark Matter from Gravitational Wave Microlensing*, *Astrophys. J.* **926** (2022), no. 2 L28, [[arXiv:2109.06456](#)].
- [26] T. T. Nakamura and S. Deguchi, *Wave optics in gravitational lensing*, *Progress of Theoretical Physics Supplement* **133** (1999) 137–153.
- [27] T. T. Nakamura, *Gravitational lensing of gravitational waves from inspiraling binaries by a point mass lens*, *Phys. Rev. Lett.* **80** (1998) 1138–1141.
- [28] R. Takahashi, *Quasigeometrical optics approximation in gravitational lensing*, *Astron. Astrophys.* **423** (2004) 787–792, [[astro-ph/0402165](#)].
- [29] R. Takahashi, T. Suyama, and S. Michikoshi, *Scattering of gravitational waves by the weak gravitational fields of lens objects*, *Astron. Astrophys.* **438** (2005) L5, [[astro-ph/0503343](#)].
- [30] R. Takahashi, *Amplitude and phase fluctuations for gravitational waves propagating through inhomogeneous mass distribution in the universe*, *Astrophys. J.* **644** (2006) 80–85, [[astro-ph/0511517](#)].
- [31] P. Schneider, J. Ehlers, and E. E. Falco, *Gravitational Lenses*. Springer Berlin Heidelberg, 1992.
- [32] M. Born, E. Wolf, A. B. Bhatia, P. C. Clemmow, D. Gabor, A. R. Stokes, A. M. Taylor, P. A. Wayman, and W. L. Wilcock, *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*. Cambridge University Press, 7 ed., 1999.
- [33] C. Leung, D. Jow, P. Saha, L. Dai, M. Oguri, and L. V. E. Koopmans, *Wave mechanics, interference, and decoherence in strong gravitational lensing*, 2023.
- [34] H. Goldstein, C. Poole, and J. Safko, *Classical Mechanics*. Addison Wesley, 2002.
- [35] J. Masoliver and A. Ros, *From classical to quantum mechanics through optics*, *European Journal of Physics* **31** (Nov., 2009) 171–192. **good title, but article looks like a rewritten textbook with not so much explanations.**
- [36] J. Feldbrugge, U.-L. Pen, and N. Turok, *Oscillatory path integrals for radio astronomy*, *Annals Phys.* **451** (2023) 169255, [[arXiv:1909.04632](#)].

- [37] L. Fishman and J. J. McCoy, *Derivation and application of extended parabolic wave theories. II. Path integral representations*, *Journal of Mathematical Physics* **25** (02, 1984) 297–308. [article unavailable](#).
- [38] L. Fishman, *Helmholtz path integrals*, *AIP Conference Proceedings* **834** (05, 2006).
- [39] L. S. Schulman and C. DeWitt-Morette, *Techniques and applications of path integration*, 1981.
- [40] M. Eve, *The use of path integrals in guided wave theory*, .
- [41] E. T. G. Alvarez and F. H. Gaioli *Foundations of Physics* **28** (1998), no. 10 1529–1538.
- [42] R. P. Feynman, *Space-time approach to quantum electrodynamics*, *Phys. Rev.* **76** (Sep, 1949) 769–789.
- [43] Y. Nambu, *The use of the Proper Time in Quantum Electrodynamics*, *Prog. Theor. Phys.* **5** (1950) 82–94.
- [44] S. S. Schweber, *Feynman and the visualization of space-time processes*, *Rev. Mod. Phys.* **58** (Apr, 1986) 449–508.
- [45] J. Schwinger, *On gauge invariance and vacuum polarization*, *Phys. Rev.* **82** (Jun, 1951) 664–679.
- [46] V. Fock, *Proper time in classical and quantum mechanics*, *Phys. Z. Sowjetunion* **12** (1937) 404–425.
- [47] M. J. Strassler, *Field theory without feynman diagrams: One-loop effective actions*, *Nuclear Physics B* **385** (Oct., 1992) 145–184.
- [48] F. Bastianelli and A. Zirotti, *Worldline formalism in a gravitational background*, *Nuclear Physics B* **642** (Oct., 2002) 372–388.
- [49] F. Bastianelli, F. Comberiati, and L. de la Cruz, *Light bending from eikonal in worldline quantum field theory*, *JHEP* **02** (2022) 209, [[arXiv:2112.05013](#)].
- [50] F. Bastianelli and M. D. Paciarini, *Worldline path integrals for the graviton*, [arXiv:2305.06650](#).
- [51] O. Corradini, C. Schubert, J. P. Edwards, and N. Ahmadinia, *Spinning particles in quantum mechanics and quantum field theory*, 2021.
- [52] L. Bonora, M. Cvitan, P. Dominis Prester, S. Giaccari, M. Paulišić, and T. Štemberga, *Worldline quantization of field theory, effective actions and L_∞ structure*, *JHEP* **04** (2018) 095, [[arXiv:1802.02968](#)].
- [53] J. P. Edwards and C. Schubert, *Quantum mechanical path integrals in the first quantised approach to quantum field theory*, 12, 2019. [arXiv:1912.10004](#).
- [54] J. P. Edwards, C. M. Mata, and C. Schubert, *One-loop amplitudes in the worldline formalism*, *Phys. Scripta* **97** (2022), no. 6 064002, [[arXiv:2201.12457](#)].
- [55] J. Feldbrugge, J.-L. Lehners, and N. Turok, *Lorentzian Quantum Cosmology*, *Phys. Rev. D* **95** (2017), no. 10 103508, [[arXiv:1703.02076](#)].

- [56] C. Teitelboim, *Quantum Mechanics of the Gravitational Field*, *Phys. Rev. D* **25** (1982) 3159.
- [57] C. Teitelboim, *Causality Versus Gauge Invariance in Quantum Gravity and Supergravity*, *Phys. Rev. Lett.* **50** (1983) 705.
- [58] C. Teitelboim, *The Proper Time Gauge in Quantum Theory of Gravitation*, *Phys. Rev. D* **28** (1983) 297.
- [59] M. C. Gutzwiller, *Phase-Integral Approximation in Momentum Space and the Bound States of an Atom*, *Journal of Mathematical Physics* **8** (12, 2004) 1979–2000, [https://pubs.aip.org/aip/jmp/article-pdf/8/10/1979/8182197/1979_1_online.pdf].
- [60] M. C. Gutzwiller, *Phase-Integral Approximation in Momentum Space and the Bound States of an Atom*, *Journal of Mathematical Physics* **8** (12, 2004) 1979–2000, [https://pubs.aip.org/aip/jmp/article-pdf/8/10/1979/8182197/1979_1_online.pdf].
- [61] J. Babington, *Ray-wave duality in classical optics: Crossing the feynman bridge*, *Optics Letters* **43** (11, 2018) 5591.
- [62] J. Babington, *Ray-wave duality of electromagnetic fields: a feynman path integral approach to classical vectorial imaging*, *Journal of the Optical Society of America A* **38** (May, 2021) 817.
- [63] J. H. Hannay, *Path-linking interpretation of kirchhoff diffraction*, *Proceedings: Mathematical and Physical Sciences* **450** (1995), no. 1938 51–65.
- [64] R. B. Schlottmann, *A path integral formulation of acoustic wave propagation*, *Geophysical Journal International* **137** (May, 1999) 353–363.
- [65] D. R. Palmer, *A path-integral approach to the parabolic approximation. I*, *The Journal of the Acoustical Society of America* **66** (09, 1979) 862–871, [https://pubs.aip.org/asa/jasa/article-pdf/66/3/862/11940861/862_1_online.pdf].
- [66] **NANOGrav** Collaboration, G. Agazie et al., *The NANOGrav 15 yr Data Set: Evidence for a Gravitational-wave Background*, *Astrophys. J. Lett.* **951** (2023), no. 1 L8, [[arXiv:2306.16213](https://arxiv.org/abs/2306.16213)].
- [67] **EPTA, InPTA:** Collaboration, J. Antoniadis et al., *The second data release from the European Pulsar Timing Array - III. Search for gravitational wave signals*, *Astron. Astrophys.* **678** (2023) A50, [[arXiv:2306.16214](https://arxiv.org/abs/2306.16214)].
- [68] D. J. Reardon et al., *Search for an Isotropic Gravitational-wave Background with the Parkes Pulsar Timing Array*, *Astrophys. J. Lett.* **951** (2023), no. 1 L6, [[arXiv:2306.16215](https://arxiv.org/abs/2306.16215)].
- [69] H. Xu et al., *Searching for the Nano-Hertz Stochastic Gravitational Wave Background with the Chinese Pulsar Timing Array Data Release I*, *Res. Astron. Astrophys.* **23** (2023), no. 7 075024, [[arXiv:2306.16216](https://arxiv.org/abs/2306.16216)].
- [70] D. G. Figueroa, M. Pieroni, A. Ricciardone, and P. Simakachorn, *Cosmological Background Interpretation of Pulsar Timing Array Data*, *Phys. Rev. Lett.* **132** (2024), no. 17 171002, [[arXiv:2307.02399](https://arxiv.org/abs/2307.02399)].
- [71] C. R. Contaldi, *Anisotropies of Gravitational Wave Backgrounds: A Line Of Sight Approach*, *Phys. Lett. B* **771** (2017) 9–12, [[arXiv:1609.08168](https://arxiv.org/abs/1609.08168)].

- [72] **LISA Cosmology Working Group** Collaboration, N. Bartolo et al., *Probing anisotropies of the Stochastic Gravitational Wave Background with LISA*, *JCAP* **11** (2022) 009, [[arXiv:2201.08782](#)].
- [73] N. Bartolo, D. Bertacca, S. Matarrese, M. Peloso, A. Ricciardone, A. Riotto, and G. Tasinato, *Anisotropies and non-Gaussianity of the Cosmological Gravitational Wave Background*, *Phys. Rev. D* **100** (2019), no. 12 121501, [[arXiv:1908.00527](#)].
- [74] N. Bartolo, D. Bertacca, S. Matarrese, M. Peloso, A. Ricciardone, A. Riotto, and G. Tasinato, *Characterizing the cosmological gravitational wave background: Anisotropies and non-Gaussianity*, *Phys. Rev. D* **102** (2020), no. 2 023527, [[arXiv:1912.09433](#)].
- [75] A. Ricciardone, L. V. Dall’Armi, N. Bartolo, D. Bertacca, M. Liguori, and S. Matarrese, *Cross-Correlating Astrophysical and Cosmological Gravitational Wave Backgrounds with the Cosmic Microwave Background*, *Phys. Rev. Lett.* **127** (2021), no. 27 271301, [[arXiv:2106.02591](#)].
- [76] **LISA Cosmology Working Group** Collaboration, P. Auclair et al., *Cosmology with the Laser Interferometer Space Antenna*, *Living Rev. Rel.* **26** (2023), no. 1 5, [[arXiv:2204.05434](#)].
- [77] D. Bertacca, A. Ricciardone, N. Bellomo, A. C. Jenkins, S. Matarrese, A. Raccanelli, T. Regimbau, and M. Sakellariadou, *Projection effects on the observed angular spectrum of the astrophysical stochastic gravitational wave background*, *Phys. Rev. D* **101** (2020), no. 10 103513, [[arXiv:1909.11627](#)].
- [78] F. Schulze, L. Valbusa Dall’Armi, J. Lesgourgues, A. Ricciardone, N. Bartolo, D. Bertacca, C. Fidler, and S. Matarrese, *GW_CLASS: Cosmological Gravitational Wave Background in the cosmic linear anisotropy solving system*, *JCAP* **10** (2023) 025, [[arXiv:2305.01602](#)].
- [79] A. Malhotra, E. Dimastrogiovanni, G. Domènech, M. Fasiello, and G. Tasinato, *New universal property of cosmological gravitational wave anisotropies*, *Phys. Rev. D* **107** (2023), no. 10 103502, [[arXiv:2212.10316](#)].
- [80] L. Valbusa Dall’Armi, A. Mierna, S. Matarrese, and A. Ricciardone, *Adiabatic or Non-Adiabatic? Unraveling the Nature of Initial Conditions in the Cosmological Gravitational Wave Background*, [arXiv:2307.11043](#).
- [81] R. A. Isaacson, *Gravitational radiation in the limit of high frequency. i. the linear approximation and geometrical optics*, *Phys. Rev.* **166** (Feb, 1968) 1263–1271.
- [82] R. A. Isaacson, *Gravitational radiation in the limit of high frequency. ii. nonlinear terms and the effective stress tensor*, *Phys. Rev.* **166** (Feb, 1968) 1272–1280.
- [83] L. Pizzuti, A. Tomella, C. Carbone, M. Calabrese, and C. Baccigalupi, *Boltzmann equations for astrophysical Stochastic Gravitational Wave Backgrounds scattering off of massive objects*, *JCAP* **02** (2023) 054, [[arXiv:2208.02800](#)].
- [84] G. Cusin, R. Durrer, and P. G. Ferreira, *Polarization of a stochastic gravitational wave background through diffusion by massive structures*, *Phys. Rev. D* **99** (2019), no. 2 023534, [[arXiv:1807.10620](#)].

- [85] A. Garoffolo, *Wave-optics limit of the stochastic gravitational wave background*, *Phys. Dark Univ.* **44** (2024) 101475, [[arXiv:2210.05718](#)].
- [86] C. Dalang, G. Cusin, and M. Lagos, *Polarization distortions of lensed gravitational waves*, *Phys. Rev. D* **105** (2022), no. 2 024005, [[arXiv:2104.10119](#)].
- [87] G. Cusin and M. Lagos, *Gravitational wave propagation beyond geometric optics*, *Phys. Rev. D* **101** (2020), no. 4 044041, [[arXiv:1910.13326](#)].
- [88] M. Pijnenburg, G. Cusin, C. Pitrou, and J.-P. Uzan, *Wave optics lensing of gravitational waves: theory and phenomenology of triple systems in the LISA band*, [arXiv:2404.07186](#).
- [89] M. A. Oancea, R. Stiskalek, and M. Zumalacárregui, *Frequency- and polarization-dependent lensing of gravitational waves in strong gravitational fields*, [arXiv:2209.06459](#).
- [90] M. A. Oancea, R. Stiskalek, and M. Zumalacárregui, *Probing general relativistic spin-orbit coupling with gravitational waves from hierarchical triple systems*, [arXiv:2307.01903](#).
- [91] S. Chandrasekhar, *The mathematical theory of black holes*. 1985.
- [92] E. Berti, V. Cardoso, and A. O. Starinets, *Quasinormal modes of black holes and black branes*, *Class. Quant. Grav.* **26** (2009) 163001, [[arXiv:0905.2975](#)].
- [93] R. A. Konoplya and A. Zhidenko, *Quasinormal modes of black holes: From astrophysics to string theory*, *Rev. Mod. Phys.* **83** (2011) 793–836, [[arXiv:1102.4014](#)].
- [94] T. Hinderer, *Tidal Love numbers of neutron stars*, *Astrophys. J.* **677** (2008) 1216–1220, [[arXiv:0711.2420](#)]. [Erratum: *Astrophys. J.* 697, 964 (2009)].
- [95] T. Binnington and E. Poisson, *Relativistic theory of tidal Love numbers*, *Phys. Rev. D* **80** (2009) 084018, [[arXiv:0906.1366](#)].
- [96] T. Damour and A. Nagar, *Relativistic tidal properties of neutron stars*, *Phys. Rev. D* **80** (2009) 084035, [[arXiv:0906.0096](#)].
- [97] A. Pound and B. Wardell, *Black hole perturbation theory and gravitational self-force*, [arXiv:2101.04592](#).
- [98] S. A. Teukolsky, *Perturbations of a rotating black hole. 1. Fundamental equations for gravitational electromagnetic and neutrino field perturbations*, *Astrophys. J.* **185** (1973) 635–647.
- [99] E. Newman and R. Penrose, *An Approach to gravitational radiation by a method of spin coefficients*, *J. Math. Phys.* **3** (1962) 566–578.
- [100] D. C. Gloge and D. Marcuse, *Formal quantum theory of light rays*, *Journal of the Optical Society of America* **59** (1969) 1629–1631.
- [101] D. Stoler, *Operator methods in physical optics*, *Journal of the Optical Society of America (1917-1983)* **71** (Mar., 1981) 334.
- [102] Z. Guralnik, *A new look at the helmholtz equation: Lefschetz thimbles and the einbein action*, 2019.

- [103] R. P. Feynman and A. R. Hibbs, *Quantum mechanics and path integrals*. International series in pure and applied physics. McGraw-Hill, New York, NY, 1965.
- [104] J. Feldbrugge and N. Turok, *Gravitational lensing of binary systems in wave optics*, 2020.
- [105] G. Tambalo, M. Zumalacárregui, L. Dai, and M. H.-Y. Cheung, *Lensing of gravitational waves: Efficient wave-optics methods and validation with symmetric lenses*, *Physical Review D* **108** (Aug., 2023).
- [106] M. H.-Y. Cheung, K. K. Y. Ng, M. Zumalacárregui, and E. Berti, *Probing minihalo lenses with diffracted gravitational waves*, 2024.
- [107] M. Çalışkan, L. Ji, R. Cotesta, E. Berti, M. Kamionkowski, and S. Marsat, *Observability of lensing of gravitational waves from massive black hole binaries with lisa*, *Phys. Rev. D* **107** (Feb, 2023) 043029.
- [108] D. L. Jow, U.-L. Pen, and J. Feldbrugge, *Regimes in astrophysical lensing: refractive optics, diffractive optics, and the Fresnel scale*, *Mon. Not. Roy. Astron. Soc.* **525** (2023), no. 2 2107–2124, [[arXiv:2204.12004](#)].
- [109] K. Thorne and R. Blandford, *Optics: Volume 2 of Modern Classical Physics*. Princeton University Press, 2021.
- [110] C. GARROD, *Hamiltonian path-integral methods*, *Rev. Mod. Phys.* **38** (Jul, 1966) 483–494.
- [111] J. Feldbrugge, D. L. Jow, and U.-L. Pen, *Complex classical paths in quantum reflections and tunneling*, [arXiv:2309.12420](#).
- [112] C. DeWitt-Morette, A. N. Maheshwari, and B. Nelson, *Path integration in non-relativistic quantum mechanics*, *Physics Reports* **50** (1979) 255–372.
- [113] K. S. Thorne and R. D. Blandford, *Modern Classical Physics: Optics, Fluids, Plasmas, Elasticity, Relativity, and Statistical Physics*. 2017.
- [114] M. J. Pfenning and E. Poisson, *Scalar, electromagnetic, and gravitational selfforces in weakly curved space-times*, *Phys. Rev. D* **65** (2002) 084001, [[gr-qc/0012057](#)].
- [115] Y.-Z. Chu, K. Pasmatsiou, and G. D. Starkman, *Finite-size effects on the self-force*, *Phys. Rev. D* **101** (2020), no. 10 104020, [[arXiv:1910.02924](#)].
- [116] C. Copi and G. D. Starkman, *Gravitational Glint: Detectable Gravitational Wave Tails from Stars and Compact Objects*, *Phys. Rev. Lett.* **128** (2022), no. 25 251101, [[arXiv:2201.03684](#)].
- [117] Y.-Z. Chu and G. D. Starkman, *Retarded Green's Functions In Perturbed Spacetimes For Cosmology and Gravitational Physics*, *Phys. Rev. D* **84** (2011) 124020, [[arXiv:1108.1825](#)].
- [118] M. E. Peskin and D. V. Schroeder, *An Introduction to quantum field theory*. Addison-Wesley, Reading, USA, 1995.
- [119] A. L. Fetter and J. D. Walecka, *Quantum Theory of Many-Particle Systems*. McGraw-Hill, Boston, 1971.

- [120] M. D. Schwartz, *Quantum Field Theory and the Standard Model*. Cambridge University Press, 3, 2014.
- [121] M. Srednicki, *Quantum field theory*. Cambridge University Press, 1, 2007.
- [122] D. Bertacca, A. Raccanelli, N. Bartolo, and S. Matarrese, *Cosmological perturbation effects on gravitational-wave luminosity distance estimates*, *Phys. Dark Univ.* **20** (2018) 32–40, [[arXiv:1702.01750](#)].
- [123] G. Tambalo, M. Zumalacárregui, L. Dai, and M. H.-Y. Cheung, *Lensing of gravitational waves: Efficient wave-optics methods and validation with symmetric lenses*, *Phys. Rev. D* **108** (2023), no. 4 043527, [[arXiv:2210.05658](#)].
- [124] T. Morita and J. Soda, *Arrival Time Differences of Lensed Massive Gravitational Waves*, [arXiv:1911.07435](#).
- [125] T. Regge and J. A. Wheeler, *Stability of a Schwarzschild singularity*, *Phys. Rev.* **108** (1957) 1063–1069.
- [126] F. J. Zerilli, *Effective potential for even parity Regge-Wheeler gravitational perturbation equations*, *Phys. Rev. Lett.* **24** (1970) 737–738.
- [127] M. Maggiore, *Gravitational Waves. Vol. 2: Astrophysics and Cosmology*. Oxford University Press, 3, 2018.
- [128] W. H. Press and S. A. Teukolsky, *Perturbations of a Rotating Black Hole. II. Dynamical Stability of the Kerr Metric*, *Astrophys. J.* **185** (1973) 649–674.
- [129] L. Andersson, J. Joudioux, M. A. Oancea, and A. Raj, *Propagation of polarized gravitational waves*, *Phys. Rev. D* **103** (2021), no. 4 044053, [[arXiv:2012.08363](#)].

From Wave optics for rotating stars

References

- [1] C. Baraldo, A. Hosoya, and T. T. Nakamura, *Phys. Rev. D* **59**, 083001 (1999).
- [2] E. Petroff, J. W. T. Hessels, and D. R. Lorimer, *AAPP* **27**, 4 (2019), [[arXiv:1904.07947 astro-ph.HE](#)] .
- [3] e. a. Abbott, B. P., *Phys. Rev. Lett.* **116**, 061102 (2016), [[arXiv:1602.03837 gr-qc](#)] .
- [2] T. T. Nakamura and S. Deguchi, *Progress of Theoretical Physics Supplement* **133**, 137 (1999).
- [25] G. Grillo and J. Cordes, *arXiv e-prints* , [[arXiv:1810.09058](#)] (2018), [[arXiv:1810.09058 astro-ph.CO](#)] .
- [6] J. Feldbrugge and N. Turok, *arXiv e-prints* , [[arXiv:2008.01154](#)] (2020), [[arXiv:2008.01154 gr-qc](#)] .
- [7] G. Tambalo, M. Zumalacárregui, L. Dai, and M. H.-Y. Cheung, *Phys. Rev. D* **108**, 043527 (2023), [[arXiv:2210.05658 gr-qc](#)] .
- [8] J. Feldbrugge, U.-L. Pen, and N. Turok, *Annals of Physics* **451**, 169255 (2023).

- [9] J. Feldbrugge, *MNRAS* **520**, 2995 (2023), arXiv:2010.03089 [astro-ph.CO] .
- [10] D. L. Jow, U.-L. Pen, and J. Feldbrugge, *MNRAS* **525**, 2107 (2023), arXiv:2204.12004 [astro-ph.HE] .
- [11] G. Braga, A. Garoffolo, A. Ricciardone, N. Bartolo, and S. Matarrese, arXiv e-prints , arXiv:2405.20208 (2024), arXiv:2405.20208 [astro-ph.CO] .
- [12] H. Thirring, *Physikalische Zeitschrift* **19**, 33 (1918).
- [13] J. Lense and H. Thirring, *Physikalische Zeitschrift* **19**, 156 (1918).
- [14] J. C. A. Miller-Jones, A. J. Tetarenko, G. R. Sivakoff, M. J. Middleton, D. Altamirano, G. E. Anderson, T. M. Belloni, R. P. Fender, P. G. Jonker, E. G. Körding, H. A. Krimm, D. Maitra, S. Markoff, S. Migliari, K. P. Mooley, M. P. Rupen, D. M. Russell, T. D. Russell, C. L. Sarazin, R. Soria, and V. Tudose, *Nature* **569**, 374 (2019), arXiv:1906.05400 [astro-ph.HE] .
- [15] V. Venkatraman Krishnan, M. Bailes, W. van Straten, N. Wex, P. C. C. Freire, E. F. Keane, T. M. Tauris, P. A. Rosado, N. D. R. Bhat, C. Flynn, A. Jameson, and S. Osłowski, *Science* **367**, 577 (2020), arXiv:2001.11405 [astro-ph.HE] .
- [16] M. Grould, F. H. Vincent, T. Paumard, and G. Perrin, in *The Multi-Messenger Astrophysics of the Galactic Centre*, IAU Symposium, Vol. 322, edited by R. M. Crocker, S. N. Longmore, and G. V. Bicknell (2017) pp. 25–30.
- [17] A. Ulmer and J. Goodman, *Astrophys. J.* **442**, 67 (1995), arXiv:astro-ph/9406042 .
- [18] R. Takahashi and T. Nakamura, *Astrophys. J.* **595**, 1039 (2003), arXiv:astro-ph/0305055 .
- [19] N. Matsunaga and K. Yamamoto, *JCAP* **01**, 023 (2006), arXiv:astro-ph/0601701 .
- [20] H. Asada and M. Kasai, *Prog. Theor. Phys.* **104**, 95 (2000), arXiv:astro-ph/0006157 .
- [21] M. Sereno, *Phys. Lett. A* **305**, 7 (2002), arXiv:astro-ph/0209148 .
- [22] H. Ebrahimnejad Rahbari, M. Nouri-Zonoz, and S. Rahvar, arXiv e-prints , astro-ph/0508477 (2005), arXiv:astro-ph/0508477 [astro-ph] .
- [23] R. Feynman and A. Zee, *QED: The Strange Theory of Light and Matter*, Alix G. Mautner memorial lectures (Princeton University Press, 2006).
- [24] P. Schneider, J. Ehlers, and E. E. Falco, *Gravitational Lenses* (1992).
- [25] J. Ibanez, *A&A* **124**, 175 (1983).
- [26] I. G. Dymnikova, in *Relativity in Celestial Mechanics and Astrometry. High Precision Dynamical Theories and Observational Verifications*, IAU Symposium, Vol. 114, edited by J. Kovalevsky and V. A. Brumberg (1986) p. 411.
- [27] J. F. Glicenstein, *A&A* **343**, 1025 (1999).
- [28] J. F. Nye, *Proceedings of the Royal Society of London Series A* **361**, 21 (1978).
- [29] J. F. Nye, *Proceedings of the Royal Society of London Series A* **403**, 1 (1986).

- [30] M. Sereno, *MNRAS* **344**, 942 (2003), [arXiv:astro-ph/0307243 \[astro-ph\]](#) .
- [31] J. Feldbrugge and N. Turok, *Annals of Physics* **454**, 169315 (2023), [arXiv:2207.12798 \[hep-th\]](#) .
- [32] M. Çalıřkan, L. Ji, R. Cotesta, E. Berti, M. Kamionkowski, and S. Marsat, *Phys. Rev. D* **107**, 043029 (2023), [arXiv:2206.02803 \[astro-ph.CO\]](#) .
- [33] K. S. Thorne, *Rev. Mod. Phys.* **52**, 299 (1980).
- [34] K. Thorne and R. Blandford, *Modern Classical Physics: Optics, Fluids, Plasmas, Elasticity, Relativity, and Statistical Physics* (Princeton University Press, 2017).
- [35] D. C. Robinson, *Phys. Rev. Lett.* **34**, 905 (1975).

From Gravit len of binary systems in wave optics

References

- [1] P. Schneider, J. Ehlers and E.E. Falco, *Gravitational Lenses*. Springer, Berlin, 1992.
- [2] T. T. Nakamura and S. Deguchi, “Wave Optics in Gravitational Lensing,” *Progress of Theoretical Physics Supplement* **133** (Jan, 1999) 137–153.
- [3] S. Sugiyama, T. Kurita, and M. Takada, “On the wave optics effect on primordial black hole constraints from optical microlensing search,” *Mon. Not. Roy. Astron. Soc.* **493** (2020) no. 3, 3632–3641, [arXiv:1905.06066 \[astro-ph.CO\]](#) .
- [4] A. Katz, J. Kopp, S. Sibiryakov, and W. Xue, “Looking for MACHOs in the Spectra of Fast Radio Bursts,” [arXiv:1912.07620 \[astro-ph.CO\]](#) .
- [5] D. L. Jow, S. Foreman, U.-L. Pen, and W. Zhu, “Wave effects in the microlensing of pulsars and FRBs by point masses,” [arXiv:2002.01570 \[astro-ph.HE\]](#) .
- [6] M. W. Sammons, J.-P. Macquart, R. D. Ekers, R. M. Shannon, H. Cho, J. X. Prochaska, A. T. Deller, and C. K. Day, “First Constraints on Compact Dark Matter from Fast Radio Burst Microstructure,” [arXiv:2002.12533 \[astro-ph.CO\]](#) .
- [7] O. Wucknitz, L. G. Spitler, and U. L. Pen, “Cosmology with gravitationally lensed repeating Fast Radio Bursts,” [arXiv:2004.11643 \[astro-ph.CO\]](#) .
- [8] S. Mao, “Astrophysical Applications of Gravitational Microlensing,” *Res. Astron. Astrophys.* **12** (2012) 947–972, [arXiv:1207.3720 \[astro-ph.GA\]](#) .
- [9] A. Herrera-Martín, M. D. Albrow, A. Udalski, A. Gould, Y.-H. Ryu, J. C. Yee, S.-J. Chung, C. Han, K.-H. Hwang, Y. K. Jung, and et al., “Ogle-2018-blg-06771b: A super-earth near the galactic bulge,” *The Astronomical Journal* **159** (May, 2020) 256. <http://dx.doi.org/10.3847/1538-3881/ab893e>.
- [10] E. Petroff, J. W. T. Hessels, and D. R. Lorimer, “Fast Radio Bursts,” *Astron. Astrophys. Rev.* **27** (2019) no. 1, 4, [arXiv:1904.07947 \[astro-ph.HE\]](#) .
- [11] A. Einstein, “Lens-Like Action of a Star by the Deviation of Light in the Gravitational Field,” *Science* **84** (1936) no. 2188, 506–7.

- [12] J. Feldbrugge, U.-L. Pen, and N. Turok, “Oscillatory path integrals for radio astronomy,” [arXiv:1909.04632 \[astro-ph.HE\]](#).
- [13] <https://p-lpi.github.io/glwo.html>.
- [14] J. Feldbrugge and N. Turok, “in preparation.”
- [15] V. Gajjar *et al.*, “Highest Frequency Detection of FRB 121102 at 4–8 GHz Using the Breakthrough Listen Digital Backend at the Green Bank Telescope,” *Astrophys. J.* **863** (2018) no. 1, 2, [arXiv:1804.04101 \[astro-ph.HE\]](#).

From On the wave optics ef on PBH constr fr opt microl search

References

- [1] Alcock C., et al., 1996, *ApJ*, 461, 84
- [2] Alcock C., et al., 2000, *ApJ*, 542, 281
- [3] Ali-Haimoud Y., et al., 2019, arXiv e-prints, p. [arXiv:1903.04424](#)
- [4] Arina C., 2018,] 10.3389/fspas.2018.00030
- [5] Bai Y., Orlofsky N., 2018, arXiv e-prints, p. [arXiv:1812.01427](#)
- [6] Bechtol K., et al., 2019, arXiv e-prints, p. [arXiv:1903.04425](#)
- [7] Bird S., Cholis I., Muñoz J. B., Ali-Haimoud Y., Kamionkowski M., Kovetz E. D., Raccanelli A., Riess A. G., 2016, *Physical Review Letters*, 116, 201301
- [8] Carr B. J., 1975, *ApJ*, 201, 1
- [9] Carr B., 2019, in Symposium on Illuminating Dark Matter Kruen, Germany, May 13-19, 2018. ([arXiv:1901.07803](#))
- [10] Carr B. J., Kohri K., Sendouda Y., Yokoyama J., 2010, *Phys. Rev. D*, 81, 104019
- [11] Carr B., Kühnel F., Sandstad M., 2016, *Phys. Rev. D*, 94, 083504
- [12] Carr B., Raidal M., Tenkanen T., Vaskonen V., Veermäe H., 2017, *Phys. Rev. D*, 96, 023514
- [13] Cieplak A. M., Griest K., 2013, *ApJ*, 767, 145
- [14] Gould A., 1992, *ApJ*, 386, L5
- [15] Griest K., et al., 1991, *ApJ*, 372, L79
- [16] Guenther B. D., 2015, *Modern Optics*. Oxford University Press
- [17] Hawking S. W., 1974, *Nature*, 248, 30
- [18] Hooper D., 2018, arXiv e-prints, p. [arXiv:1812.02029](#)
- [19] Inomata K., Kawasaki M., Mukaida K., Tada Y., Yanagida T. T., 2017, *Phys. Rev. D*, 96, 043504

- [20] Jungman G., Kamionkowski M., Griest K., 1996, *Phys. Rep.*, 267, 195
- [21] Katz A., Kopp J., Sibiryakov S., Xue W., 2018, *Journal of Cosmology and Astro-Particle Physics*, 2018, 005
- [22] Klypin A., Zhao H., Somerville R. S., 2002, *ApJ*, 573, 597
- [23] Kochanek C. S., 2006, in Meylan G., Jetzer P., North P., Schneider P., Kochanek C. S., Wambsganss J., eds, *Saas-Fee Advanced Course 33: Gravitational Lensing: Strong, Weak and Micro*. pp 91–268
- [24] Matsunaga N., Yamamoto K., 2006, *J. Cosmology Astropart. Phys.*, 1, 023
- [25] Mróz P., et al., 2017, *Nature*, 548, 183
- [26] Naderi T., Mehrabi A., Rahvar S., 2018, *Phys. Rev. D*, 97, 103507
- [27] Nakamura T. T., 1998, *Physical Review Letters*, 80, 1138
- [28] Nakamura T. T., Deguchi S., 1999, *Progress of Theoretical Physics Supplement*, 133, 137
- [29] Navarro J. F., Frenk C. S., White S. D. M., 1997, *ApJ*, 490, 493
- [30] Niikura H., et al., 2019a, *Nature Astronomy*, p. 238
- [31] Niikura H., Takada M., Yokoyama S., Sumi T., Masaki S., 2019b, *Phys. Rev. D*, 99, 083503
- [32] Paczynski B., 1986, *ApJ*, 304, 1
- [33] Sasaki M., Suyama T., Tanaka T., Yokoyama S., 2016, *Physical Review Letters*, 117, 061101
- [34] Sasaki M., Suyama T., Tanaka T., Yokoyama S., 2018, *Class. Quant. Grav.*, 35, 063001
- [35] Sato-Polito G., Kovetz E. D., Kamionkowski M., 2019, arXiv e-prints, p. arXiv:1904.10971
- [36] Schneider P., Ehlers J., Falco E. E., 1992, *Gravitational Lenses*, doi:10.1007/978-3-662-03758-4.
- [37] Sumi T., et al., 2003, *ApJ*, 591, 204
- [38] Takahashi R., Nakamura T., 2003, *ApJ*, 595, 1039
- [39] Tisserand P., et al., 2007, *A&A*, 469, 387
- [40] Witt H. J., Mao S., 1994, *ApJ*, 430, 505

From “Oscillatory path integrals for radio astronomy” by Feldbrugge, Pen, Turok

References

- [1] R. Feynman and A. Hibbs, *Quantum mechanics and path integrals*. International series in pure and applied physics. McGraw-Hill, 1965.
- [2] J. Klauder, *A Modern Approach to Functional Integration*. Applied and Numerical Harmonic Analysis. Birkhäuser Boston, 2010.
https://books.google.ca/books?id=QYA_AAAAQBAJ.
- [3] M. V. Berry and J. F. Nye, “Fine structure in caustic junctions,” *Nature* **267** (May, 1977) 34–36.
- [4] M. V. Berry, “Focusing and twinkling: critical exponents from catastrophes in non-Gaussian random short waves,” *Journal of Physics A Mathematical General* **10** (Dec, 1977) 2061–2081.
- [5] M. V. Berry, J. F. Nye, and F. J. Wright, “The Elliptic Umbilic Diffraction Catastrophe,” *Philosophical Transactions of the Royal Society of London Series A* **291** (Apr, 1979) 453–484.
- [6] M. V. Berry and C. Upstill, “IV Catastrophe Optics: Morphologies of Caustics and Their Diffraction Patterns,” *Progress in Optics* **18** (Jan, 1980) 257–346.
- [7] M. V. Berry, “Focused tsunami waves,” *Proceedings of the Royal Society of London Series A* **463** (Nov, 2007) 3055–3071.
- [8] L. G. Spitler, J. M. Cordes, J. W. T. Hessels, D. R. Lorimer, M. A. McLaughlin, S. Chatterjee, F. Crawford, J. S. Deneva, V. M. Kaspi, R. S. Wharton, B. Allen, S. Bogdanov, A. Brazier, F. Camilo, P. C. C. Freire, F. A. Jenet, C. Karako-Argaman, B. Knispel, P. Lazarus, K. J. Lee, J. van Leeuwen, R. Lynch, S. M. Ransom, P. Scholz, X. Siemens, I. H. Stairs, K. Stovall, J. K. Swiggum, A. Venkataraman, W. W. Zhu, C. Aulbert, and H. Fehrmann, “Fast Radio Burst Discovered in the Arecibo Pulsar ALFA Survey,” *ApJ* **790** (Aug, 2014) 101, [arXiv:1404.2934](https://arxiv.org/abs/1404.2934) [astro-ph.HE].
- [9] CHIME/FRB Collaboration, M. Amiri *et al.*, “Observations of fast radio bursts at frequencies down to 400 megahertz,” *Nature* **566** (2019) no. 7743, 230–234, [arXiv:1901.04524](https://arxiv.org/abs/1901.04524) [astro-ph.HE].
- [10] A. Josephy *et al.*, “CHIME/FRB Detection of the Original Repeating Fast Radio Burst Source FRB 121102,” [arXiv:1906.11305](https://arxiv.org/abs/1906.11305) [astro-ph.HE].
- [11] D. B. Melrose and P. G. Watson, “Scintillation of Radio Sources: The Role of Caustics,” *ApJ* **647** (Aug, 2006) 1131–1141.
- [12] J. M. Cordes, I. Wasserman, J. W. T. Hessels, T. J. W. Lazio, S. Chatterjee, and R. S. Wharton, “Lensing of Fast Radio Bursts by Plasma Structures in Host Galaxies,” *ApJ* **842** (Jun, 2017) 35, [arXiv:1703.06580](https://arxiv.org/abs/1703.06580) [astro-ph.HE].
- [13] R. Main, I. S. Yang, V. Chan, D. Li, F. X. Lin, N. Mahajan, U.-L. Pen, K. Vanderlinde, and M. H. van Kerkwijk, “Pulsar emission amplified and resolved by plasma lensing in an eclipsing binary,” *Nature* **557** (May, 2018) 522–525, [arXiv:1805.09348](https://arxiv.org/abs/1805.09348) [astro-ph.HE].

- [14] U.-L. Pen and L. King, “Refractive convergent plasma lenses explain extreme scattering events and pulsar scintillation,” *MNRAS* **421** (Mar, 2012) L132–L136, [arXiv:1111.6806 \[astro-ph.GA\]](#).
- [15] B. J. Rickett, “Interstellar scattering and scintillation of radio waves,” *Annual review of astronomy and astrophysics* **15** (Jan, 1977) 479–504.
- [16] A. Ishimaru, “Optical scattering and diffusion in turbulence and scatterers (A),” *Journal of the Optical Society of America (1917-1983)* **68** (Oct, 1978) 1368.
- [17] U.-L. Pen and Y. Levin, “Pulsar scintillations from corrugated reconnection sheets in the interstellar medium,” *MNRAS* **442** (Aug, 2014) 3338–3346, [arXiv:1302.1897 \[astro-ph.GA\]](#).
- [18] A. V. Bilous, S. M. Ransom, and P. Demorest, “Unusually Bright Single Pulses from the Binary Pulsar B1744-24A: A Case of Strong Lensing?,” *ApJ* **877** (Jun, 2019) 125, [arXiv:1811.05766 \[astro-ph.HE\]](#).
- [19] R. L. Fiedler, B. Dennison, K. J. Johnston, and A. Hewish, “Extreme scattering events caused by compact structures in the interstellar medium,” *Nature* **326** (Apr, 1987) 675–678.
- [20] U.-L. Pen and L. King, “Refractive convergent plasma lenses explain extreme scattering events and pulsar scintillation,” *MNRAS* **421** (Mar, 2012) L132–L136, [arXiv:1111.6806 \[astro-ph.GA\]](#).
- [21] L. Dai and T. Venumadhav, “On the waveforms of gravitationally lensed gravitational waves,” *arXiv e-prints* (Feb, 2017) [arXiv:1702.04724](#), [arXiv:1702.04724 \[gr-qc\]](#).
- [22] X. Er and A. Rogers, “Two families of astrophysical diverging lens models,” *MNRAS* **475** (Mar, 2018) 867–878, [arXiv:1712.06900 \[astro-ph.GA\]](#).
- [23] W. A. Coles, J. P. Filice, R. G. Frehlich, and M. Yadlowsky, “Simulation of wave propagation in three-dimensional random media,” *Applied Optics* **34** (Apr, 1995) 2089.
- [24] W. A. Coles, B. J. Rickett, J. J. Gao, G. Hobbs, and J. P. W. Verbiest, “Scattering of Pulsar Radio Emission by the Interstellar Plasma,” *The Astrophysical Journal* **717** (Jul, 2010) 1206–1221, [arXiv:1005.4914 \[astro-ph.GA\]](#).
- [25] G. Grillo and J. Cordes, “Wave asymptotics and their application to astrophysical plasma lensing,” *arXiv e-prints* (Oct, 2018) [arXiv:1810.09058](#), [arXiv:1810.09058 \[astro-ph.CO\]](#).
- [26] E. Witten, “A New Look At The Path Integral Of Quantum Mechanics,” *arXiv e-prints* (Sep, 2010) [arXiv:1009.6032](#), [arXiv:1009.6032 \[hep-th\]](#).
- [27] J. Feldbrugge, J.-L. Lehners, and N. Turok, “Lorentzian Quantum Cosmology,” *Phys. Rev.* **D95** (2017) no. 10, 103508, [arXiv:1703.02076 \[hep-th\]](#).
- [28] J. Feldbrugge, J.-L. Lehners, and N. Turok, “No Smooth Beginning for Spacetime,” *Physical Review Letters* **119** (Oct, 2017) 171301, [arXiv:1705.00192 \[hep-th\]](#).
- [29] J. Feldbrugge, J.-L. Lehners, and N. Turok, “No rescue for the no boundary proposal: Pointers to the future of quantum cosmology,” *Phys. Rev.* **D97** (2018) no. 2, 023509, [arXiv:1708.05104 \[hep-th\]](#).

- [30] J. Feldbrugge, J.-L. Lehners, and N. Turok, “Inconsistencies of the New No-Boundary Proposal,” *Universe* **4** (2018) no. 10, 100, [arXiv:1805.01609 \[hep-th\]](#).
- [31] A. Di Tucci, J. Feldbrugge, J.-L. Lehners, and N. Turok, “Quantum Incompleteness of Inflation,” *arXiv e-prints* (June, 2019) [arXiv:1906.09007](#), [arXiv:1906.09007 \[hep-th\]](#).
- [32] J. Feldbrugge, A. Fertig, L. Sberna, , and N. Turok, “In preparation: The Spacetime Amplitude Approach to Relativistic Quantum Mechanics: the Schwinger Effect,”.
- [33] L. Dai and W. Lu, “Probing Motion of Fast Radio Burst Sources by Timing Strongly Lensed Repeaters,” *ApJ* **847** (Sep, 2017) 19, [arXiv:1706.06103 \[astro-ph.HE\]](#).
- [34] M. Born and E. Wolf, *Principles of Optics*. Cambridge university press, 1999.
- [35] J. Jackson, *Classical Electrodynamics*. John Wiley, New York, 1975.
- [36] J. Feldbrugge and N. Turok, “In preparation,”.
- [37] G. Basar, G. V. Dunne, and M. Ünsal, “Resurgence theory, ghost-instantons, and analytic continuation of path integrals,” *Journal of High Energy Physics* **2013** (Oct, 2013) 41, [arXiv:1308.1108 \[hep-th\]](#).
- [38] A. Behtash, G. V. Dunne, T. Schaefer, T. Sulejmanpasic, and M. Unsal, “Toward Picard-Lefschetz Theory of Path Integrals, Complex Saddles and Resurgence,” *arXiv e-prints* (Oct, 2015) [arXiv:1510.03435](#), [arXiv:1510.03435 \[hep-th\]](#).
- [39] G. V. Dunne and M. Unsal, “What is QFT? Resurgent trans-series, Lefschetz thimbles, and new exact saddles,” *arXiv e-prints* (Nov, 2015) [arXiv:1511.05977](#), [arXiv:1511.05977 \[hep-lat\]](#).
- [40] A. Behtash, G. V. Dunne, T. Schäfer, T. Sulejmanpasic, and M. Ünsal, “Complexified Path Integrals, Exact Saddles, and Supersymmetry,” *Phys. Rev. Lett.* **116** (Jan, 2016) 011601, [arXiv:1510.00978 \[hep-th\]](#).
- [41] A. Behtash, G. V. Dunne, T. Schäfer, T. Sulejmanpasic, and M. Ünsal, “Critical points at infinity, non-Gaussian saddles, and bions,” *Journal of High Energy Physics* **2018** (Jun, 2018) 68, [arXiv:1803.11533 \[hep-th\]](#).
- [42] M. Serone, G. Spada, and G. Villadoro, “The power of perturbation theory,” *Journal of High Energy Physics* **2017** (May, 2017) 56, [arXiv:1702.04148 \[hep-th\]](#).
- [43] Y. Tanizaki and T. Koike, “Real-time Feynman path integral with Picard-Lefschetz theory and its applications to quantum tunneling,” *Annals of Physics* **351** (Dec, 2014) 250–274, [arXiv:1406.2386 \[math-ph\]](#).
- [44] G. Brooker, *Modern Classical Optics*. Oxford Master Series in Physics. OUP Oxford, 2003. <https://books.google.ca/books?id=Z1bY1eW4znIC>.
- [45] W. A. Coles, “Interplanetary Scintillation,” *Space Sci. Rev.* **21** (Feb, 1978) 411–425.
- [46] M. v. d. Kamp, P. S. Cannon, and M. Terkildsen, “Effect of the ionosphere on defocusing of space-based radars,” *Radio Science* **44** (Feb, 2009) 1–15.
- [47] V. I. Arnol’d, “Normal forms for functions near degenerate critical points, the Weyl groups of A_k, D_k, E_k and Lagrangian singularities,” *Functional Anal. Appl* **6** (1972) 1972.

- [48] V. I. Arnol'd, "Wave front evolution and equivalent Morse lemma," *Communications in Pure Applied Mathematics* **29** (Nov., 1976) 557–582.
- [49] R. Thom, *Structural Stability And Morphogenesis*. CRC Press, 2018.
<https://books.google.ca/books?id=YVNPdWAAQBAJ>.
- [50] V. I. Arnol'd, "Remarks on the Stationary Phase Method and Coxeter Numbers," *Russian Mathematical Surveys* **28** (Oct, 1973) 19–48.
- [51] V. I. Arnol'd, "Critical Points of Smooth Functions and Their Normal Forms," *Russian Mathematical Surveys* **30** (Oct, 1975) 1–75.
- [52] J. Feldbrugge, R. van de Weygaert, J. Hidding, and J. Feldbrugge, "Caustic Skeleton & Cosmic Web," *Journal of Cosmology and Astro-Particle Physics* **2018** (May, 2018) 027, [arXiv:1703.09598](https://arxiv.org/abs/1703.09598) [astro-ph.CO].
- [53] I. Thompson, "NIST Handbook of Mathematical Functions, edited by Frank W.J. Olver, Daniel W. Lozier, Ronald F. Boisvert, Charles W. Clark," *Contemporary Physics* **52** (Sep, 2011) 497–498.
- [54] A. N. Varchenko, "Zeta-Function of Monodromy and Newton's Diagram," *Inventiones Mathematicae* **37** (Jan, 1976) 253.
- [55] M. Morse, "Relations between the critical points of a real function on n independent variables," *Transaction of the American Mathematical Society* **27** (1925) 345–396.
- [56] E. Stein, J. Milnor, M. Spivak, R. Wells, R. Wells, and J. Mather, *Morse Theory*. Annals of mathematics studies. Princeton University Press, 1963.
<https://books.google.ca/books?id=uj44vQEACAAJ>.
- [57] R. M. Range, "Complex analysis: A brief tour into higher dimensions," *The American Mathematical Monthly* **110** (2003) 89–108.
- [58] D. R. Lorimer, M. Bailes, M. A. McLaughlin, D. J. Narkevic, and F. Crawford, "A Bright Millisecond Radio Burst of Extragalactic Origin," *Science* **318** (Nov, 2007) 777, [arXiv:0709.4301](https://arxiv.org/abs/0709.4301) [astro-ph].
- [59] CHIME/FRB Collaboration, "A second source of repeating fast radio bursts," *Nature* **566** (Jan, 2019) 235–238, [arXiv:1901.04525](https://arxiv.org/abs/1901.04525) [astro-ph.HE].
- [60] U.-L. Pen, "The nature of fast radio bursts," *Nature Astronomy* **2** (Oct, 2018) 842–844, [arXiv:1811.00605](https://arxiv.org/abs/1811.00605) [astro-ph.HE].
- [61] E. Platts, A. Weltman, A. Walters, S. P. Tendulkar, J. E. B. Gordin, and S. Kandhai, "A Living Theory Catalogue for Fast Radio Bursts," *arXiv e-prints* (Oct, 2018) [arXiv:1810.05836](https://arxiv.org/abs/1810.05836), [arXiv:1810.05836](https://arxiv.org/abs/1810.05836) [astro-ph.HE].
- [62] S. Chatterjee, C. J. Law, R. S. Wharton, S. Burke-Spolaor, J. W. T. Hessels, G. C. Bower, J. M. Cordes, S. P. Tendulkar, C. G. Bassa, P. Demorest, B. J. Butler, A. Seymour, P. Scholz, M. W. Abruzzo, S. Bogdanov, V. M. Kaspi, A. Keimpema, T. J. W. Lazio, B. Marcote, M. A. McLaughlin, Z. Paragi, S. M. Ransom, M. Rupen, L. G. Spitler, and H. J. van Langevelde, "A direct localization of a fast radio burst and its host," *Nature* **541** (Jan, 2017) 58–61, [arXiv:1701.01098](https://arxiv.org/abs/1701.01098) [astro-ph.HE].

- [63] S. Chatterjee, R. Wharton, C. J. Law, J. Hessels, S. Burke-Spolaor, G. C. Bower, M. W. AbruZZo, C. Bassa, B. J. Butler, J. M. Cordes, D. Paul, V. M. Kaspi, M. McLaughlin, S. M. Ransom, P. Scholz, A. Seymour, L. Spitler, S. P. Tendulkar, PALFA Survey Team, VLA+AO FRB121102 Simultaneous Campaign Team, and EVN FRB121102 Campaign Team, “Localizing the Fast Radio Burst 121102,” in *American Astronomical Society Meeting Abstracts #229*, vol. 229 of *American Astronomical Society Meeting Abstracts*, p. 330.01. Jan, 2017.
- [64] D. Michilli, A. Seymour, J. W. T. Hessels, L. G. Spitler, V. Gajjar, A. M. Archibald, G. C. Bower, S. Chatterjee, J. M. Cordes, K. Gourdji, G. H. Heald, V. M. Kaspi, C. J. Law, C. Sobey, E. A. K. Adams, C. G. Bassa, S. Bogdanov, C. Brinkman, P. Demorest, F. Fernandez, G. Hellbourg, T. J. W. Lazio, R. S. Lynch, N. Maddox, B. Marcote, M. A. McLaughlin, Z. Paragi, S. M. Ransom, P. Scholz, A. P. V. Siemion, S. P. Tendulkar, P. van Rooy, R. S. Wharton, and D. Whitlow, “An extreme magneto-ionic environment associated with the fast radio burst source FRB 121102,” *Nature* **553** (Jan, 2018) 182–185, [arXiv:1801.03965](https://arxiv.org/abs/1801.03965) [astro-ph.HE].
- [65] J. Doob, *Measure Theory*. Graduate Texts in Mathematics. Springer New York, 2012. <https://books.google.ca/books?id=HOPhBwAAQBAJ>.