

WAVE ASYMPTOTICS AND THEIR APPLICATION TO ASTROPHYSICAL PLASMA LENSING

A Thesis

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

M. Sc.

by

Gianfranco Grillo Csaszar

May 2018

© 2018 Gianfranco Grillo Csaszar
ALL RIGHTS RESERVED

ABSTRACT

Plasma lensing events can have significant observational consequences, including flux modulations and perturbations in pulse arrival times. In this paper we develop and apply a formalism based on an extension of geometric optics that can be used 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 TOA perturbations due to lensing. We find that the enhanced geometric optics approach successfully reproduces the predictions from wave optics, and that it can be efficiently used to simulate multifrequency TOA residuals during lensing events. Lensing will introduce perturbations in these residuals that will manifest as an increased spreading in the data points at lower frequencies, and will deviate from the expected dispersive ν^{-2} scaling most significantly when including observations at low frequencies, $\nu < 0.7$ GHz.

BIOGRAPHICAL SKETCH

Originally from Chile, Gianfranco came to Cornell in August 2012 to study physics and philosophy, having been inspired by the writings and TV appearances of his hero Carl Sagan. After graduating in 2016, he had the pleasure of remaining in Ithaca for two more years in order to pursue a Master's degree. An avid tennis and soccer player, he's had the honor of representing Cornell as part of the Club Soccer team Academy FC, and as the longest running member in the history of Cornell Club Tennis. You can find him yelling at the TV whenever his favorite soccer team, Audax Italiano, is losing yet another game.

This work is dedicated to the memory of my maternal grandfather, Gabriel Cszasz.

ACKNOWLEDGEMENTS

I am extremely grateful to my advisor, Prof. James Cordes, for his patient guidance throughout all the stages of this investigation. I don't think it is possible to overstate how much I have learned from him, and I was very lucky to have taken his Experimental Astronomy class the first semester of my senior year. I would also like to thank Dr. Shami Chatterjee, who supervised me the first summer I engaged in research at Cornell, and from whom I also learned a lot over the past two years. Finally, I would like to thank my good friend Manuel Fernandez, for always being there for me whenever Ithaca and Cornell succeeded in overwhelming me, and my family, for always believing in me, even (or perhaps especially) during those times I had trouble believing in myself.

TABLE OF CONTENTS

| | |
|--|-----------|
| Biographical Sketch | iii |
| Dedication | iv |
| Acknowledgements | v |
| Table of Contents | vi |
| 1 Introduction | 1 |
| 2 Zeroth and first order geometrical optics | 5 |
| 2.1 Geometrical picture | 5 |
| 2.2 Zeroth order gain | 10 |
| 2.3 The 2D Kirchhoff diffraction integral | 11 |
| 2.4 Accuracy and regions of applicability | 14 |
| 3 Second order geometric optics | 17 |
| 3.1 Complex rays | 17 |
| 3.2 Caustic location and extent of the caustic zone | 18 |
| 3.3 Gain inside the caustic zone: catastrophe theory and uniform asymptotics | 22 |
| 4 TOA perturbations | 33 |
| 4.1 Geometry and dispersion | 34 |
| 4.2 Image interference and pulse addition | 38 |
| 4.3 Dispersion measure perturbation | 42 |
| 5 Summary and conclusions | 47 |
| A Solving the KDI using the FFT | 50 |
| B Estimation of the value of the gain at a caustic | 54 |
| C Numerics | 55 |
| D More numerical examples and lens colormaps | 58 |

CHAPTER 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. 1987) and early 1990s (Cognard et al. 1993), 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. (1994) and Clegg et al. (1996) mention the idea, introduced in Cognard et al. (1993), 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. (1998) 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. 2013; Tuntsov et al. 2016; Vedantham et al. 2017; Kerr et al. 2017).

More recently, plasma lensing has also been suggested as a possible mechanism to explain certain properties of FRBs (Cordes et al. 2017; Dai and Lu 2017), and other works have examined different kinds of lens models, as well as their possible observational signatures (Pen and King 2012; Er and Rogers 2017), although most, if not all, 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 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 minute deviations in pulse arrival times. In fact, some plasma lensing events have been inferred by their effects on observed pulsar TOAs (Lam et al. 2017) and dispersion measures (Coles et al. 2015), instead of their effects on measured flux, 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 thought to be produced by larger scales density inhomogeneities in the ISM. This motivates a different kind of analysis than the common statistical one used to describe scintillations based on turbulent phase screens with a Kolmogorov wavenumber spectrum. 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 and Melrose 2006; Melrose and Watson 2006). 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, from the observer's point of view, as the lens crosses the line of sight. Such multiple imaging has been directly observed in cases in which the angular separation of each of the images is large enough (Gupta et al. 1999; Pushkarev

et al. 2013), and can be inferred from the existence of fringes in the dynamic spectra of pulsars during certain epochs of observation (Cordes et al. 2006). 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 what happens to the intensity close to these regions is of importance both in the context of plasma lensing and scintillation (Goodman et al. 1987; Melrose and Watson 2006; Cordes et al. 2017). The geometric optics framework, however, is useful because it provides information about the different images, including their amplitudes, phases, and locations, which can be used to calculate TOA perturbations, 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 be able to describe what happens to the flux 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 and Kinber 1994) and catastrophe optics (Poston and Stewart 1978; Berry and Upstill 1980; Stamnes 1986; Kravtsov and Orlov 1999; Katsaounis et al. 2001; Kryukovskii et al. 2006), and the solution applied 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 will restrict ourselves to cases in which the source of radiation can be accurately regarded as a point source, and will focus on the potential effects of plasma lensing with regards to pulsar timing. The paper is divided as follows. In chapter 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 chapter 3 we use wave asymptotic methods to construct a second order geometric optics description. In chapter 4, we use the concepts developed in chapters 2 and 3 to examine the TOA perturbations due to a specific plasma lens realization, and we summarize conclusions in chapter 5. We expect to apply the methodology presented here to specific events in subsequent work.

CHAPTER 2

ZEROTH AND FIRST ORDER GEOMETRICAL OPTICS

2.1 Geometrical picture

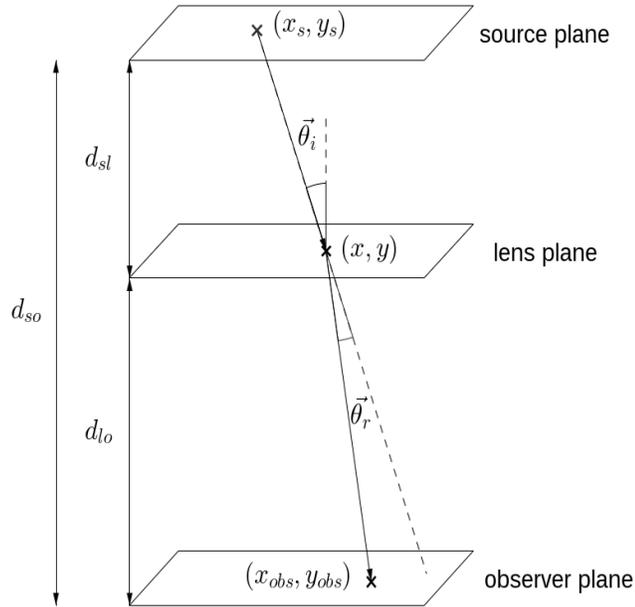


Figure 1: Lensing geometry.

We follow the basics of the treatment given in Clegg et al. (1998) and Cordes et al. (2017) but extend their results to two dimensions. We start by defining planes for the source, the lens, and the observer with coordinates \vec{x}_s , \vec{x} , and \vec{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 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 $\vec{\theta}_i$ and the its deviation angle are given (in the

paraxial approximation) by

$$\vec{\theta}_i = \frac{\vec{x}_s - \vec{x}}{d_{sl}} \quad (1)$$

$$\vec{\theta}_r = \frac{\vec{x}_{obs} - \vec{x}}{d_{lo}} - \vec{\theta}_i \quad (2)$$

Combining into a single equation in terms of $\vec{\theta}_r$ gives the lens equation,

$$\vec{x}_s \left(\frac{d_{lo}}{d_{so}} \right) + \vec{x}_{obs} \left(\frac{d_{sl}}{d_{so}} \right) = \vec{x} + \vec{\theta}_r \left(\frac{d_{sl}d_{lo}}{d_{so}} \right) \quad (3)$$

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

$$\vec{x}' \equiv \vec{x}_s \left(\frac{d_{lo}}{d_{so}} \right) + \vec{x}_{obs} \left(\frac{d_{sl}}{d_{so}} \right) \quad (4)$$

and write the lens equation in the simpler form

$$\vec{x}' = \vec{x} + \vec{\theta}_r \left(\frac{d_{sl}d_{lo}}{d_{so}} \right) \quad (5)$$

This expression is perfectly general and not only applies to plasma lensing, but to gravitational lensing as well (Schneider et al. 1992). The nature of the lensing is what determines the formula for the deviation angle $\vec{\theta}_r$. 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 is greater or smaller than unity, because the phase velocity will be greater or smaller than the speed of light. More precisely, we can write this phase difference $\delta\phi_{lens}$ as

$$\delta\phi_{lens} = \omega\tau = kc\tau \quad (6)$$

where τ is the propagation time difference between an unlensed ray and a lensed ray, $k = 2\pi/\lambda$ is the wavenumber and $\omega = 2\pi c/\lambda$ is the radiation's angular frequency. For a rectangular lens of length l parallel to the direction of propagation, this is

$$\tau = \frac{l}{c} (1 - n) \quad (7)$$

In the case of plasma lensing, the frequency dependent index of refraction corresponds to that experienced by a radio wave in an unmagnetized plasma, which is given by

$$\begin{aligned} n_r &= \sqrt{1 - \left(\frac{\omega_e}{\omega}\right)^2} \\ &= \sqrt{1 - \frac{\lambda^2 r_e n_e}{\pi}} \\ &\approx 1 - \frac{\lambda^2 r_e n_e}{2\pi} \end{aligned} \quad (8)$$

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 last 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 and Wolf 1999, Ch. 3), so the refractive angle $\vec{\theta}_r$ is given by

$$\vec{\theta}_r = -\frac{1}{k} \nabla \delta\phi_{lens} \quad (9)$$

This means that the condition for $\vec{\theta}_r \neq 0$ is that $\delta\phi_{lens}$ not be a constant, which happens when the electron column density or dispersion measure perturbation $DM = n_e l$ at the lens plane varies as a function of position, $DM \rightarrow DM(\vec{x})$. Putting all this together, the phase perturbation becomes

$$\delta\phi_{lens}(\vec{x}) = -\lambda r_e DM(\vec{x}) \quad (10)$$

which implies that the refractive angle is

$$\begin{aligned}\vec{\theta}_r &= -\frac{\lambda^2 r_e}{2\pi} \nabla DM(\vec{x}) \\ &= -\frac{c^2 r_e}{2\pi v^2} \nabla DM(\vec{x})\end{aligned}\tag{11}$$

For convenience, we write $DM(\vec{x})$ as the product of a maximum perturbation DM_l and a function with unit maximum $\psi(\vec{x})$, and take the origin of the lens plane's coordinate system to coincide with the lens's center. Thus (11) takes the form

$$\vec{\theta}_r = -\frac{c^2 r_e DM_l}{2\pi v^2} \nabla \psi(\vec{x})\tag{12}$$

At this point, it becomes important to emphasize two things. First, the function $\psi(\vec{x})$ is what determines the lens's shape, and as such it is a central component of our analysis. It is not clear whether it is possible to use physical arguments to determine the most likely form that this function would take in the case of large scales inhomogeneities in the ISM, but past work has focused mostly on realizations of ψ that result in a maximum column density located at the lens center, like the Gaussian lens. However, even if we assume that lenses tend to adopt a form with this characteristic, changing the rate at which the dispersion measure drops as one moves away from the center can have dramatic effects, as will be illustrated later.

Second, DM_l can be either positive or negative. Physically, a positive DM_l describes a lens that is the result of an overdensity in the ISM, and is in general divergent because it bends the incident rays away from its center. In contrast, a lens with a negative DM_l is the product of an underdensity in the ISM, and is convergent. We note that the vast majority of the work done so far on this topic has focused on the case $DM_l > 0$, but there is evidence that at least some of the lensing events observed historically are best described as the result of a lens with $DM_l < 0$, in particular the two timing events in the

direction of the pulsar J1713+0747 (Coles et al. 2015; Lam et al. 2017). The observable effects of both types of lenses are in general very different, as will be shown in later sections.

We now define the Fresnel scale as $r_F = \sqrt{cd_{sl}d_{lo}/2\pi\nu}$, the lens's frequency dependent strength parameter as $\phi_0 = -cr_e DM_l/\nu$, and a new parameter $\alpha = r_F^2 \phi_0$, and substitute (12) in terms of these new quantities into the lens equation, which yields a more compact form that is specific to plasma lensing,

$$\vec{x}' = \vec{x} + \alpha \nabla \psi(\vec{x}) \quad (13)$$

Finally, we adimensionalize in terms of the characteristic lens scales a_x and a_y , defining $u'_x = x/a_x$ and $u'_y = y/a_y$, and explicitly write (13) in its adimensionalized component form. Using the notation $\psi_{ij} = \partial^{i+j} \psi / \partial u_x^i \partial u_y^j$, we get

$$\begin{aligned} \begin{bmatrix} u'_x \\ u'_y \end{bmatrix} &= \begin{bmatrix} u_x + \frac{\alpha}{a_x^2} \psi_{10}(u_x, u_y) \\ u_y + \frac{\alpha}{a_y^2} \psi_{01}(u_x, u_y) \end{bmatrix} \\ &= \begin{bmatrix} u_x + \alpha_x \psi_{10} \\ u_y + \alpha_y \psi_{01} \end{bmatrix} \end{aligned} \quad (14)$$

Depending on the nature of ψ , the mapping between the u' plane and the u plane can be solved for explicitly; in most cases, however, this is not possible, and (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 C. In this context, solving the lens equation implies finding the u_x and u_y that satisfy (14) given a set of coordinates in the u' plane and parameters α_i . These coordinates will change as a function of time as the Earth, the lens, and the source move with different velocities, and the way the coordinates 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 also vary as a function of \vec{u}' and the parameters α_i .

2.2 Zeroth order gain

A large majority of the existing literature on plasma lensing (Clegg et al. 1998; Pen and King 2012; Tuntsov et al. 2016; Cordes et al. 2017; Er and Rogers 2017; Vedantham et al. 2017) derives the gain (or magnification) for an individual image G_j directly from some version of (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 $\vec{u} = \vec{u}_j^0$,

$$G_j = |\mathcal{J}|^{-1} = \left| (1 + \alpha_x \psi_{20})(1 + \alpha_y \psi_{02}) - \alpha_x \alpha_y \psi_{11}^2 \right|^{-1} \quad (15)$$

and the total gain is

$$G = \sum_{j=1}^n G_j \quad (16)$$

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 image interference. Instead, it gives a moving average of the gain that can be useful to calculate when $|\phi_o| \gg 1$. When this condition is not met, the oscillations tend to dominate, as noted in the analysis by Melrose and Watson (2006) (see also Figure 2). Large numbers of images tend to decrease the frequency of the oscillations, which also increases the oscillations’ visibility. An accurate description of the interference pattern can be obtained by solving the Kirchhoff diffraction integral (KDI), which we introduce below.

2.3 The 2D Kirchoff 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 and Wolf 1999, Ch. 8; Thorne and Blandford 2017, Ch. 8). In the paraxial approximation and for the near field, as is the case for AU sized lenses and astronomical distances, the adimensionalized integral can be written as (Goodman et al. 1987; Melrose and Watson 2006; Cordes et al. 2017)

$$\varepsilon(\vec{u}') = \frac{a_x a_y}{2\pi r_F^2} \iint d^2 u \exp(i\Phi) \quad (17)$$

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

$$\Phi(\vec{u}', \vec{u}) = \frac{1}{2r_F^2} \left[a_x^2 (u_x - u'_x)^2 + a_y^2 (u_y - u'_y)^2 \right] + \phi_0\psi(\vec{u}) \quad (18)$$

The integral is normalized such that in the absence of a lens (ie. $\delta\phi_{lens} = 0$), $\varepsilon(\vec{u}') = 1$ for all \vec{u}' . Analytic solutions to the integral are only available for a few specific forms of ψ (Watson and Melrose 2006). As detailed in Appendix A, 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.

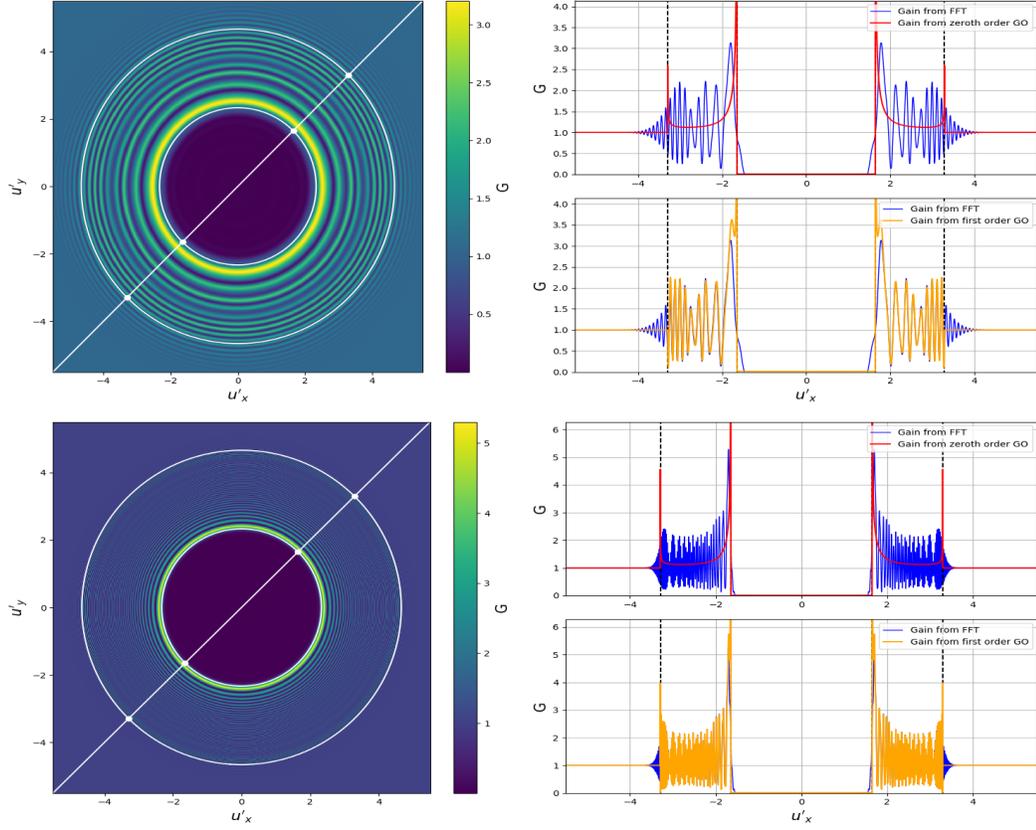


Figure 2: Comparison of the gains obtained from a full numerical solution of the KDI, zeroth order geometrical optics, and first order geometrical optics. The top panel corresponds to a lens with $\phi_0 = -50$ and the bottom panel corresponds to one with $\phi_0 = -250$ (thus $DM_l > 0$ in both cases, 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, $\psi(\vec{u}) = \exp(-u_x^2 - u_y^2)$. 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.

An approximate solution that grows more accurate as the strength of the lens increases follows by applying the method of stationary phase to (17). For a rapidly oscillating two dimensional integral of the form $I(\vec{x}) = \iint d^2\vec{x}g(\vec{x}) \exp[if(\vec{x})]$, the stationary phase lemma (Bleistein and Handelsman 1975) 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, $f_{10} = f_{01} = 0$. Each of these points $\vec{x} = \vec{x}_j^0$ provides a contribution to the integral I_j given by (Dingle 1973; Bleistein and Handelsman 1975; Cooke 1982)

$$I_j = \frac{2\pi g_j \exp\left[if_j + \frac{i\pi}{4}(\delta_j + 1)\sigma_j\right]}{|f_{20}f_{02} - f_{11}^2|^{1/2}} \quad (19)$$

where $\sigma_j = \text{sgn}(f_{20}f_{02} - f_{11}^2)$, $\delta_j = \text{sgn}(f_{02})$, $f_j = f(\vec{x} = \vec{x}_j^0)$, $g_j = g(\vec{x} = \vec{x}_j^0)$, and all the second derivatives are also evaluated at \vec{x}_j^0 . A good approximation to the integral is then obtained by summing over the n stationary points present.

In the case of the KDI as given in the form (17), the points of stationary phase correspond to the points that satisfy the two dimensional equation

$$\begin{bmatrix} \Phi_{10} \\ \Phi_{01} \end{bmatrix} = \begin{bmatrix} \frac{\alpha_x^2(u_x - u'_x)}{r_F^2} + \phi_0\psi_{10} \\ \frac{\alpha_y^2(u_y - u'_y)}{r_F^2} + \phi_0\psi_{01} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \quad (20)$$

A quick examination reveals that this is precisely equivalent to the lens equation (14), given our definitions of the parameters α_x and α_y , which therefore implies that solving the KDI by the method of stationary phase leads to geometric optics¹. Performing the appropriate substitutions in (19), we have that the scalar field ε_j due to each of the

¹It is also possible to derive the geometric optics quantities by directly solving the Helmholtz equation via WKB methods (see, e.g. Born and Wolf 1999, Ch. 3; Katsounis et al. 2001; Poston and Stewart 1978, Ch. 12).

stationary points (or, equivalently, due to each of the solutions of the lens equation) is

$$\begin{aligned}\varepsilon_j(\vec{u}') &= \frac{\alpha_x \alpha_y \exp\left[i\Phi_j + \frac{i\pi}{4}(\delta_j + 1)\sigma_j\right]}{r_F^2 \left|\Phi_{20}\Phi_{02} - \Phi_{11}^2\right|^{1/2}} \\ &= \frac{\exp\left[i\Phi_j + \frac{i\pi}{4}(\delta_j + 1)\sigma_j\right]}{\left|(1 + \alpha_x \psi_{20})(1 + \alpha_y \psi_{02}) - \alpha_x \alpha_y \psi_{11}^2\right|^{1/2}}\end{aligned}\quad (21)$$

It is instructive to think about (21) as the oscillating electric field of each of the source's images, with maximum amplitude

$$A_j = |\mathcal{J}|^{-1/2} = \left|(1 + \alpha_x \psi_{20})(1 + \alpha_y \psi_{02}) - \alpha_x \alpha_y \psi_{11}^2\right|^{-1/2}\quad (22)$$

and an oscillating component with phase

$$\beta_j = \Phi_j + \frac{\pi}{4}(\delta_j + 1)\sigma_j\quad (23)$$

whose order of magnitude depends on the Fresnel scale r_F and the lens parameters. The total scalar field is simply the sum of the contributions from the n stationary points,

$$\varepsilon(\vec{u}') = \sum_{j=1}^n \varepsilon_j = \sum_{j=1}^n A_j e^{i\beta_j}\quad (24)$$

The gain can then be obtained by taking the squared modulus of this last expression, $G = |\varepsilon(\vec{u}')|^2$. 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 equation (16), which can be recovered by taking the square modulus of (21) before summing over all of the present images, ie. by summing the individual intensities directly.

2.4 Accuracy and regions of applicability

A curious feature of the phasors that emerge from solving the KDI by the method of stationary phase 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 point. 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 and Upstill 1980; Kravtsov and Orlov 1999). 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, which do not allow the amplitude to grow without bounds. Caustics also correspond to boundaries that separate regions in parameter space that contain different numbers of solutions to the lens equation.

From Figure 2, 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 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 example, the two dimensional Gaussian lens described by $\psi(\vec{u}) = \exp(-u_x^2 - u_y^2)$ with $DM_l > 0$ and $a_x = a_y$ can show up to two sets of circular caustics in the u' plane, as shown in the figure. 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, which translates into a single image of the source. Crossing into the region between the two caustics, the oscillations's amplitude reach a peak shortly after the boundary, and

the observer sees the interference of 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. (1998), Melrose and Watson (2006), and Cordes et al. (2017) studied an analogous lens shape in one dimension. The short paper by Stinebring et al. (2007) presents similar plots to the ones above 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).

CHAPTER 3
SECOND ORDER GEOMETRIC OPTICS

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 whose imaginary part grows as we move farther into the caustic's shadow side.

Equation (21) implies that the field's magnitude of the solution whose imaginary part is positive will decrease exponentially, and the contribution due to its conjugate pair will increase exponentially. The exponentially increasing solution can be disregarded as unphysical (Kravtsov et al. (1999)), but the exponentially decaying contribution from the other complex ray can be taken at face value and included as part of the asymptotic expansion of 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 their amplitude blows up, at least according to (21).

The idea of looking for complex solutions to the lens equation has surfaced in a variety of contexts. Schramm and Kayser (1995) apply the concept to gravitational lensing, and Budden and Terry (1971) apply it to atmospheric radio ray tracing. There is also a direct connection between complex stationary points, the method of steepest descent, and hyperasymptotics of oscillatory integrals (Kaminski 1994; Howls 1997).

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

$$(1 + \alpha_x \psi_{20})(1 + \alpha_y \psi_{02}) - \alpha_x \alpha_y \psi_{11}^2 = 0 \quad (25)$$

for at least some value of \vec{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 columns of Figure 2 and 4 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

$$\begin{aligned} \frac{\psi_{10}\psi_{01}}{\Delta u_x \Delta u_y} + \frac{\psi_{20}\psi_{01}}{\Delta u_y} + \frac{\psi_{02}\psi_{10}}{\Delta u_x} + \psi_{20}\psi_{02} - \psi_{11}^2 &= 0 \\ \left(\frac{a_x}{a_y}\right)^2 \frac{\Delta u_x}{\psi_{10}} - \frac{\Delta u_y}{\psi_{01}} &= 0 \end{aligned} \quad (26)$$

for \vec{u} , where $\Delta u_i = u'_i - u_i$. The caustics will be located at frequencies ν_{caus} , given by

$$\begin{aligned} \nu_{caus} &= \frac{c}{a_x} \left[\frac{d_{sl} d_{lo} r_e DM_l}{2\pi d_{so} \Delta u_x} \psi_{10} \right]^{1/2} \\ &= \frac{c}{a_y} \left[\frac{d_{sl} d_{lo} r_e DM_l}{2\pi d_{so} \Delta u_y} \psi_{01} \right]^{1/2} \end{aligned} \quad (27)$$

evaluated at the solutions of (26) 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 $\alpha_i < -1.2$ (for the positive DM_l case)

and $\alpha_i > 0.5$ (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 $DM_l = \pm 10^{-6}$ pc cm⁻³, which corresponds to a strength parameter of $\phi_0 \approx \mp 33$ at 0.8 GHz, yields a maximum value of $a_i \approx 2.4 \times 10^{-2}$ AU for the overdense case and $a_i \approx 3.6 \times 10^{-2}$ AU for the underdense case. Ray crossings for lenses with $a_i \approx 1$ AU would require a minimum value of $|DM_l| \approx 2 \times 10^{-3}$ pc cm⁻³ for the diverging lens and $|DM_l| \approx 7 \times 10^{-4}$ pc cm⁻³ for the converging lens, which correspond to strength parameters at 0.8 GHz of $\phi_o \approx -5.8 \times 10^4$ and $\phi_o \approx 2.4 \times 10^4$, respectively. Changing the lens-observer distance d_{so} and source-lens distance d_{sl} also leads to changes in α_i , although not in a very simple way because the value of $d_{lo} = d_{so} - d_{sl}$ also factors into the expression. In general, however, if we keep d_{sl} fixed at $d_{so}/2$, increasing d_{so} also increases α_i and makes caustic formation more likely. The radius of the caustic curves tends to increase linearly with $|\alpha_i|^1$.

Caustics will show up as a function of ν at a fixed value of \vec{u}' if we search within a range of frequencies that contains a value of ν that leads to at least one of the α_i parameters have a magnitude larger than the required minimum. Since $|\alpha_i| \propto \nu^{-2}$, caustic curves in dynamic spectra, such as the ones depicted in Figures 3 and 5, will be more likely to show up only at low frequencies.

In practical terms, it is useful to be able to locate caustics as functions of both \vec{u}' and ν . Telescope observations made during an observing epoch correspond roughly

¹An important exception is the underdense ($DM_l < 0$) 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.

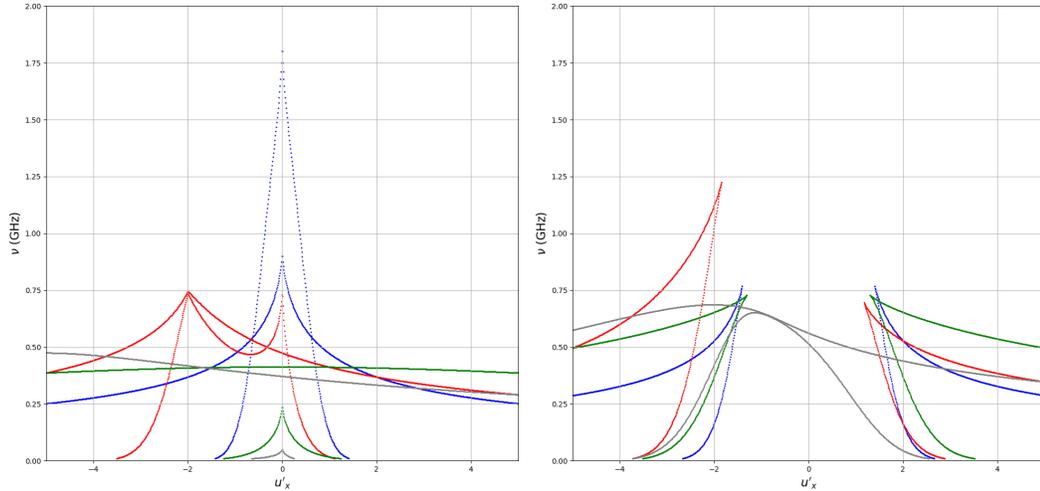


Figure 3: Caustic curves in the dynamic spectra of underdense (left) and overdense (right) two dimensional Gaussian lenses for different paths across the u' plane. The blue caustics derive from a path with slope $m = 1$ and y-intercept $n = 0$, the red caustics have $m = 0.5$ and $n = 1$, the green caustics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

to observations made at a fixed \vec{u}' , as long as the epoch duration is not too long and the effective velocity of the lens is not too large. Given that we regularly observe at more than one frequency band, we would expect to be able 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 at a single frequency band over a range of epochs. As shown in Figure 3, repeated application of equations (26) and (27) for different u' coordinates can be used to construct the caustic curves for the dynamic spectrum due to a lensing event.

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 indistinguishable from each other. As noted by Kravtsov and Orlov (1999), 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 ,

$$|\Delta\Phi_{ij}| \lesssim \pi \quad (28)$$

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 1965; Bleistein and Handelsman 1975; Cooke 1982; Wong 2001; Cordes et al. 2017) 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 B) 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, ie. close to the boundary defined by (28). Large magnitudes of ϕ_0 decrease the size of the caustic region, but in 2D it is possible for the observer to describe paths in the u' plane that cross the caustic curves at shallow angles. When this happens, it is possible for the observations to remain within the caustic zone for an extended set of observer positions. This suggests that a more general approach is needed.

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

Catastrophe theory, first developed by the mathematician Rene Thom (Thom 1972) and subsequently applied to optics by Sir Michael Berry and others (Berry 1976; Nye 1978; Berry and Upstill 1980), 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 rigourously construct a mapping that takes the global form of the phase to the standard form (Kravtsov and Orlov 1999). Instead, the mapping is performed by expanding the phase in a Taylor series at the point that satisfies both the lens equation and equation (25), 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 and Orlov (1999), and performed specifically for the case of two dimensional scattering screens in the context of scintillation by Goodman et al. (1987).

Watson and Melrose (2006) 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 maps

into the Airy function, defined by

$$\text{Ai}(\xi) = \frac{1}{2\pi} \int_{-\infty}^{\infty} dt \exp [i(t^3/3 + \xi t)] \quad (29)$$

where ξ denotes a control variable and t denotes a state variable. 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 (after determining the proper relationship between ξ , t , and the experimental variables from the mapping procedure) within the caustic region, and revert back to the regular geometrical optics description far away from the caustic, which is what Watson and Melrose (2006) do. However, this suffers from the primary problem that, although the caustic boundary can be operationally defined by (28), some further work is still required to be able to match the solutions close to and far away from the caustic, which is why the authors also allocate a portion of their paper to describe the criterion they use to go from one type of approximation to the other.

A better solution is to employ the method of uniform asymptotics, as initially developed by Chester et al. (1957), Ursell (1965), and Ludwig (1966) for oscillatory integrals, and later explicitly applied to optics and related to catastrophe theory by Kravstov (1968); Kravtsov and Orlov (1999). 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. (1957)², and in general there are slight variations between each of the presented expressions. We derive it here in the simplest way, using a procedure similar to the one employed by Kravtsov and Orlov (1999).

The general scheme consists in starting with a guess solution 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 that can be a function of \vec{u}' , 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 guess solution 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 argument. Thus, we start with a guess solution of the form,

$$\varepsilon(\vec{u}') = e^{i\chi} [g_1 \text{Ai}(\xi) + g_2 \text{Ai}'(\xi)] \quad (30)$$

where g_j , χ , and ξ are all potentially functions of \vec{u}' . From (24), we have that the first order geometrical optics solution in the case of two rays can be written as

$$\varepsilon(\vec{u}') = A_1 e^{i\beta_1} + A_2 e^{i\beta_2} \quad (31)$$

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

$$\text{Ai}(\xi) \approx \frac{1}{\sqrt{\pi}} (-\xi)^{-1/4} \cos \left[\frac{2}{3} (-\xi)^{3/2} - \frac{\pi}{4} \right] \quad (32)$$

$$\text{Ai}'(\xi) \approx \frac{1}{\sqrt{\pi}} (-\xi)^{1/4} \sin \left[\frac{2}{3} (-\xi)^{3/2} - \frac{\pi}{4} \right] \quad (33)$$

²Also see Ludwig (1966); Stamnes (1986); Borovikov and Kinber (1994); Kravtsov and Orlov (1999); Qiu and Wong (2000); Katsaounis et al. (2001)

Let $\gamma = 2(-\xi)^{3/2}/3 - \pi/4$. Using Euler's identity and substituting into (30), we get

$$\begin{aligned}\varepsilon(\vec{u}') &= \frac{e^{i\chi}}{2\sqrt{\pi}} \left[g_1(-\xi)^{-1/4} (e^{i\gamma} + e^{-i\gamma}) + ig_2(-\xi)^{1/4} (e^{i\gamma} - e^{-i\gamma}) \right] \\ &= \frac{e^{i\chi}}{2\sqrt{\pi}} \left\{ e^{i\gamma} [g_1(-\xi)^{-1/4} + ig_2(-\xi)^{1/4}] + e^{-i\gamma} [g_1(-\xi)^{-1/4} - ig_2(-\xi)^{1/4}] \right\}\end{aligned}\quad (34)$$

Matching this to (31), 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 . The first set is

$$\begin{aligned}A_1 &= \frac{1}{2\sqrt{\pi}} [g_1(-\xi)^{-1/4} + ig_2(-\xi)^{1/4}] \\ A_2 &= \frac{1}{2\sqrt{\pi}} [g_1(-\xi)^{-1/4} - ig_2(-\xi)^{1/4}]\end{aligned}\quad (35)$$

Solving for g_1 and g_2 gives

$$\begin{aligned}g_1 &= \sqrt{\pi}(A_1 + A_2)(-\xi)^{1/4} \\ g_2 &= i\sqrt{\pi}(A_1 - A_2)(-\xi)^{-1/4}\end{aligned}\quad (36)$$

The second set of equations is

$$\begin{aligned}\chi + \gamma &= \beta_1 \\ \chi - \gamma &= \beta_2\end{aligned}\quad (37)$$

which leads to

$$\begin{aligned}\chi &= \frac{1}{2}(\beta_1 + \beta_2) \\ \xi &= -\left[\frac{3}{4}(\beta_1 - \beta_2 + \frac{\pi}{2}) \right]^{2/3}\end{aligned}\quad (38)$$

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

$$\varepsilon_{fold}(\vec{u}') = \sqrt{\pi} e^{i\chi} \left[(A_1 + A_2)(-\xi)^{1/4} \text{Ai}(\xi) + i(A_1 - A_2)(-\xi)^{-1/4} \text{Ai}'(\xi) \right] \quad (39)$$

The ambiguity in the labeling is resolved by the condition $\beta_1 + \pi/2 > \beta_2$, which is equivalent to $\Phi_1 - \Phi_2 > 0$, since the two interfering images have opposite parity (Stamnes 1986; Qiu and Wong 2000). For the dark side, the complex conjugate nature of the solutions imply that $A_1 = A_2$, $Re(\beta_1) = Re(\beta_2)$, and $Im(\beta_1) = -Im(\beta_2)$, and thus we can drop the indices and write the field as

$$\varepsilon_{fold}(\vec{u}') = 2\sqrt{\pi} A \xi^{1/4} e^{i\chi} \quad (40)$$

where $\chi = Re(\beta)$, $\xi = \left[\frac{3}{2} |Im(\beta)| \right]^{2/3}$, and A and β are respectively the first order geometrical optics amplitude and phase of either of the solutions. Note that even though the A_1 and A_2 diverge as they approach the singularity, the quantity $(A_1 + A_2)(-\xi)^{1/4}$ goes to a finite limit, because $\xi \rightarrow 0$ at the caustic. By the same token, although $(-\xi)^{-1/4}$ goes to infinity at the singularity, the quantity $(A_1 - A_2)(-\xi)^{-1/4}$ does not, because $A_1 - A_2 \rightarrow 0$.

For the present case of plasma lenses and astrophysical distances, the idealized situation presented above involving two images on the caustic's bright side and no 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, expressions (39) and (40) 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 corresponds to the Pearcey function,

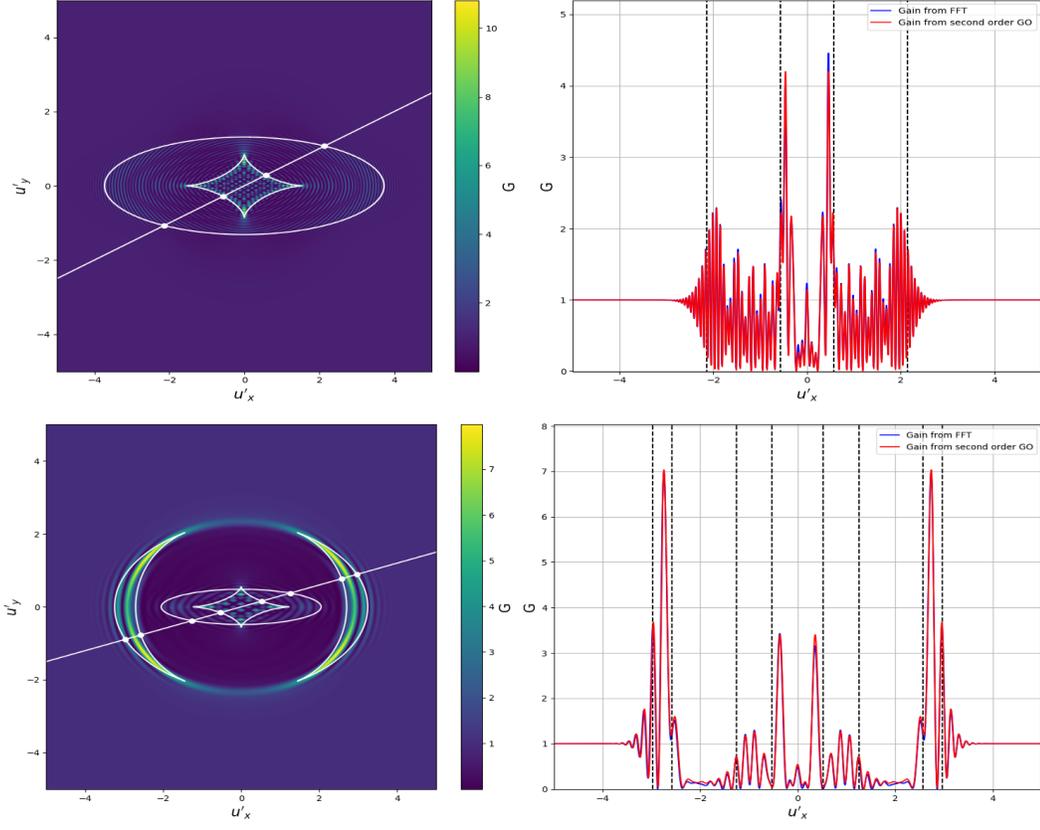


Figure 4: Comparison of the gains obtained from a full numerical solution of the KDI and second order geometric optics. The left column shows color maps of the gain obtained by solving the KDI via the FFT. The white circles correspond to caustic curves, and the straight white line shows the path of the observer through the u' plane. The right column shows the gain along this path as calculated via the FFT method 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 elliptical Gaussian lens with $\psi(\vec{u}) = \exp(-u_x^2 - u_y^2)$, strength parameter $\phi_0 = 100$, 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 $\psi(\vec{u}) = 2.72(u_x^2 + u_y^2) \exp(-u_x^2 - u_y^2)$, strength parameter $\phi_0 = -30$, and lens scales $a_x = 1.5 \times 10^{-2}$ AU and $a_y = 2.5 \times 10^{-2}$ AU. The frequency of observation is $\nu = 0.8$ GHz, and the distances are $d_{so} = 1$ kpc, and $d_{sl} = 0.5$ 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 and Upstill 1980).

$$P(\xi_1, \xi_2) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} dt \exp \left[i \left(\xi_1 t + \frac{1}{2} \xi_2 t^2 - \frac{1}{4} t^4 \right) \right] \quad (41)$$

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 2. The corresponding guess solution for the bright side of the caustic would then be

$$\varepsilon(\vec{r}') = e^{i\chi} \left[g_1 P(\xi_1, \xi_2) + g_2 \frac{\partial P}{\partial \xi_1} + g_3 \frac{\partial P}{\partial \xi_2} \right] \quad (42)$$

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 Pearcey integral does not have a simple asymptotic expansion for large parameter like the Airy integral. Instead, the correct strategy involves mapping the stationary points of the canonical integral to those obtained from the lens equation, as described in detail by Connor (1990) and Katsaounis et al. (2001). 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 closed form, and the systems of equations that result from the mapping must be solved numerically.

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 lines which correspond to fold catastrophes, where only two images merge. As noted by Ludwig (1966), far from these cusps points (and far from fold curve intersections), (42) 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 or to fold line intersections, the total field can be written as

$$\varepsilon_{tot}(\vec{u}') = \varepsilon_{fold} + \sum_{j=1}^n A_j e^{i\beta_j} \quad (43)$$

where ε_{fold} is given by (39) or (40) depending on whether we are at the caustic's dark side or bright side.

Figure 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 strength parameters ϕ_0 . Unlike the case of the circular Gaussian lens with positive DM_l depicted in Figure 2, the lenses in these figures show cusps as well as folds. In Figure 4, both the elliptical Gaussian with $DM_l < 0$ and the ring-like lens with $DM_l > 0$ 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 panel, 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 included as part of Appendices A and C.

In general, because folds manifest as lines in the u' plane, cusps manifest as points, and higher order catastrophes only occur for more exotic lens shapes, it is overwhelmingly likely that in the course of observing at a single frequency, we will encounter fold caustics as opposed to caustics corresponding to higher order catastrophes during a lensing event. Nevertheless, the implementation of an algorithm that can effectively model the field near cusps and higher order catastrophes is something worthy of exploration

as we do expect to encounter these catastrophes in dynamic spectra, even in the case of relatively simple lenses. As illustrated in Figure 3, dynamic spectra of both over and underdense Gaussian lenses for which $a_x \neq a_y$ will show cusps for paths traversed by the observer along the u' plane that pass close to the caustic structure's center.

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 and a_y 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 (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. As explained in section 3.2, these plots can be compared to observations made at different frequency bands during a single epoch.

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 5. From the figure, it is also apparent that larger magnitudes of the maximum column density $|DM_l|$ induce faster oscillations in the gain, and the contributions from complex rays in the shadow sides of caustics become less important. In such instances, the second order approach becomes less useful, and reverting to the zeroth order approximation to obtain the average value of the oscillations is justifiable in cases

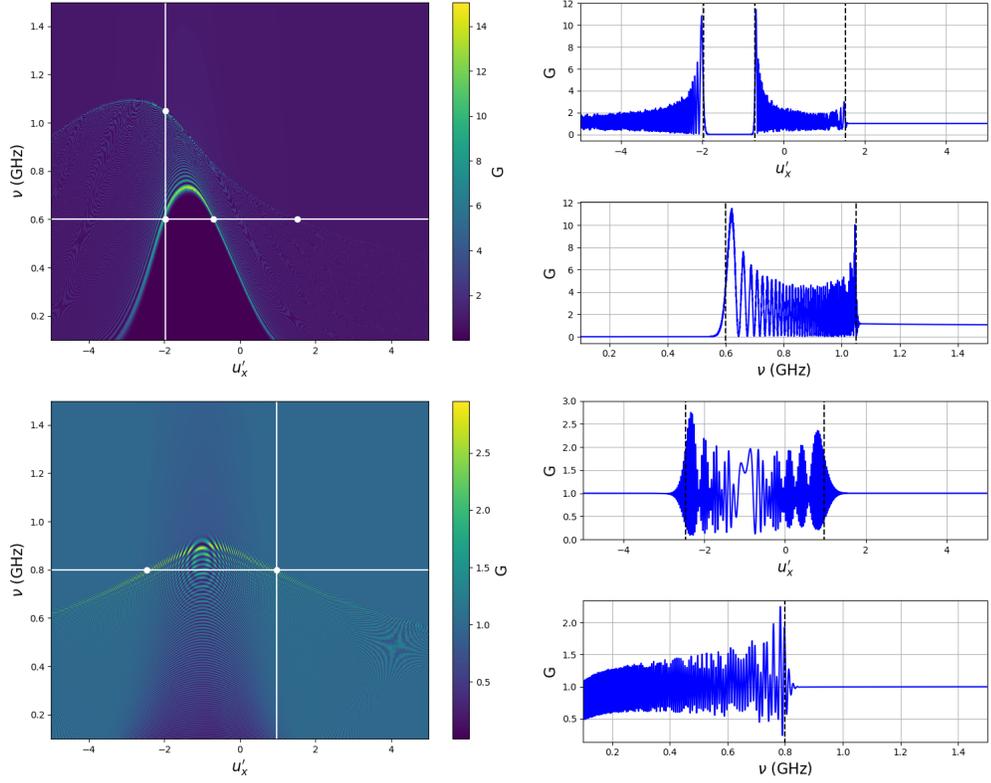


Figure 5: Sections of dynamic spectra and slices across them for overdense and underdense perturbed Gaussian lenses with $\psi(\vec{u}) = \exp(-u_x^2 - u_y^2) \left\{ 1 - A \left[\sin(Bu_x) + \sin(Bu_y) \right] \right\}$ and different DM_l magnitudes. The left column shows the two dimensional spectrum 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 $DM_l = 10^{-4} \text{ pc cm}^{-3}$ and lens scales of $a_x = 0.1 \text{ AU}$ and $a_y = 0.2 \text{ AU}$, whereas the underdense lens has $DM_l = -10^{-5} \text{ pc cm}^{-3}$, and $a_x = a_y = 0.04 \text{ AU}$. Both lenses have perturbation parameters $A = 1.5 \times 10^{-2}$ and $B = 5$, source-observer distance $d_{so} = 1 \text{ kpc}$, and source-lens distance $d_{sl} = 0.5 \text{ kpc}$. The path through the u' plane in both cases is a straight line with slope $m = 0.5$ and y-intercept $n = 2.5$.

in which the only quantity of interest is the source's light curve. We thus expect the concepts developed in this section to be most useful for relatively weak lensing events. However, we will see in the next section that even for the case of large DM_l magnitudes, accounting for image interference can be very important when we model the effects of lensing on multifrequency TOA residuals.

Second order geometric optics also preserves the advantages of the zeroth and first order approximations over the FFT method in that it gives us information about the number of images that can be seen by the observer, the delay or advance in the time of arrival of each of these images, and the individual image amplification, as determined by their location in the u plane obtained from solving the lens equation. We show how we can use this information to determine overall TOA perturbations from lensing events in the next chapter.

CHAPTER 4

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 and Wolszcan 1986; Cordes et al. 1986) and has resurfaced more recently given the potential of PTAs to detect low frequency gravitational waves (Cordes and Shannon 2010; Cordes et al. 2016) and in the context of FRBs (Cordes et al. 2017; Dai and Lu 2017). Our analysis will rely on examples that use parameters that are more likely to be relevant for pulsar timing, that is, 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. We focus on the underdense case, but the formalisms and procedures are identical for overdense lenses.

Note that in the following analysis we will apply the concepts of image interference and make use of the uniform asymptotic formulas developed in the previous section in order to describe image amplification within caustic regions. However, we will ignore contributions from complex rays, due to both practical and conceptual considerations. As noted previously, complex ray contributions to the field become negligible for large magnitudes of DM_l , and numerical results indicate that TOA perturbations in the order of microseconds are only produced by lenses with $|DM_l| \gtrsim 10^{-4} \text{ pc cm}^{-3}$, which corresponds to a relatively large magnitude of dispersion measure perturbation. On the other hand, it is not clear how to properly determine the arrival time of radiation derived from a complex solution to the lens equation. Real solutions are physically identified with

images of the source, and pulse arrival times for each of these images can be determined from the dispersive and geometrical characteristics of the pulse's path of propagation. Complex solutions, on the other hand, have no straightforward physical analog, and direct application of the soon to be derived expressions for the geometrical and dispersive TOA perturbations yield contributions from the complex solutions that are themselves complex. It is not clear to us how this makes sense from a physical point of view, and proper examination of the issue is left to future work.

4.1 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, which means that it takes longer for the ray to reach the observer plane, and the TOA perturbation is said to be positive. Referring to the geometry of Figure 1, we can write the magnitude of the geometric delay Δt_{geo} as

$$\Delta t_{geo} = t_{g_x}(u_x - u'_x)^2 + t_{g_y}(u_y - u'_y)^2 \quad (44)$$

where the $t_{g_i} = a_i^2 d_{so}/2cd_{sl}d_{lo}$ are the geometrical delay coefficients along the u_i 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 $\vec{u} = \vec{u}'_j$, we can express the geometric delay as

$$\Delta t_{geo} = t_{g_x}\alpha_x^2\psi_{10}^2 + t_{g_y}\alpha_y^2\psi_{01}^2 \quad (45)$$

Independently of the geometric delay, the lens will also introduce a dispersive perturbation in pulse arrival time. This perturbation follows from the well known frequency

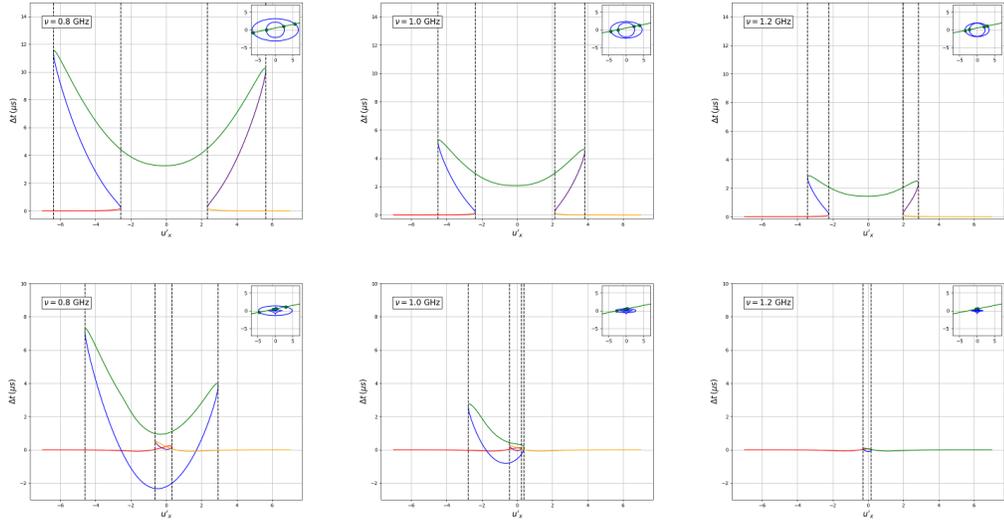


Figure 6: Individual image perturbation as a function of position and frequency in the u' plane for overdense (top panel) and underdense (bottom panel) lenses with $DM_l = \pm 5 \times 10^{-4} \text{ pc cm}^{-3}$. Both lenses have a Lorentzian shape with $\psi(\vec{u}) = 1/[(u_x^2 + u_y^2) + 1]$ and 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 path of the observer through the u' plane and the (green) path of the observer through the u' plane. The different colors in the Δt vs u'_x plots correspond to different images of the source.

dependent dispersive effect of plasma, and is given by

$$\begin{aligned} \Delta t_{DM} &= \frac{c r_e DM_l}{2\pi\nu^2} \psi(\vec{u}) \\ &= 4.149 \text{ ms} \times \frac{DM_l}{\nu^2} \psi(\vec{u}) \end{aligned} \quad (46)$$

where the second equality applies for a DM_l in units of pc cm^{-3} and ν in GHz. As we have already mentioned, it is possible for DM_l to be positive or negative, depending on whether the lens is more or less dense than the surrounding interstellar medium. If $DM_l > 0$, the dispersive perturbation will introduce a TOA delay, as the column density of electrons along the line of sight will increase¹. On the other hand, if $DM_l < 0$, the

¹Of course, it is possible to have a lens function ψ that is both positive and negative depending on \vec{u} , such as $\psi(\vec{u}) = \sin(u_x) + \sin(u_y)$. Thus, this statement is correct only for lens realizations that have $\psi > 0$ for all \vec{u} , which is the case for all the examples shown throughout this work.

lens will constitute a sort of “pinhole” in the interstellar medium, and radiation passing through the lens will experience less of a dispersive delay than radiation traveling outside of it, which means that the dispersive TOA perturbation will be negative.

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

$$\Delta t_j = \Delta t_{geo}^j + \Delta t_{DM}^j \quad (47)$$

When $DM_l > 0$, 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 $DM_l < 0$, Δt_j can be either positive or negative depending on whether the magnitude of the geometric contribution is larger than the magnitude of the dispersive contribution. 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} , 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 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 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 $DM_l = \pm 5 \times 10^{-4}$ pc cm⁻³, which gives a strength parameter ϕ_o for each of

the frequencies of $\sim \mp 1.6 \times 10^4$, $\mp 1.3 \times 10^4$, and $\mp 1.1 \times 10^4$, respectively. For the overdense lens sequence in the top panel, 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 that 1. 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 2. 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 2.

The bottom panel, corresponding to the lens with $DM_l < 0$, shows a different pattern. This time the maximum number of images (five)² 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

²This type of lens can actually lead to the formation of up to nine images of the source in the underdense case if the line of sight crosses through the center of the u' plane. We chose to use a path that did not take us exactly through the center, in order to make the plot simpler to understand. Similar plots with lenses and paths that lead to numbers of images greater than five can be found in Appendices A and D.

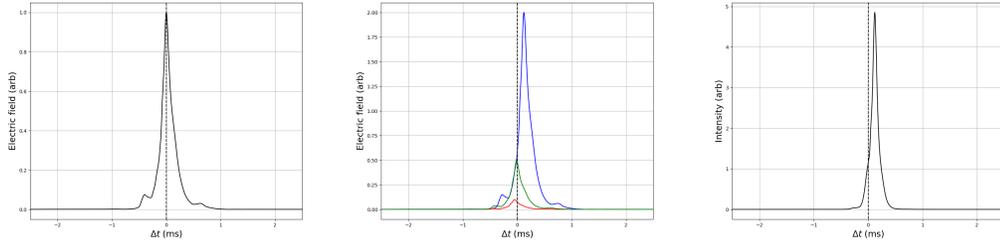


Figure 7: Illustration of the process that leads to TOA perturbation during epochs of multiple imaging. The left panel corresponds to the idealized shape and TOA of a single electric field pulse due to a five millisecond pulsar. The middle panel shows, in different colors, the pulses due to three different images arriving at the telescope, each with a different amplification A_j , shift Δt_j , and phase β_j . The rightmost plot shows the resulting intensity measured by the telescope due to the superposition of the three electric field pulses from the middle panel. The template used in this example is used by the NANOGrav collaboration to determine TOAs from the pulsar J1713+0747.

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.

4.2 Image interference and pulse addition

Within regions of parameter space that contain more than one solution to the lens equation, it is in principle possible for the observer to see more than one image of the source, and there have been reported ESEs during which more than one image of the source showed up in the data (Gupta et al. 1999; Pushkarev et al. 2013). In practice, the distances involved are such that the angular separation between the images is too small for each of them to be resolved individually. Instead, what we see is an interference pattern, whose fluctuations are dominated by those images of the source that have a larger magnification.

During a lensing event, each of the images arrives at the telescope in the form of an electric field pulse with an amplitude proportional to A_j as given by equation (22) and a phase β_j as given by equation (23). Additionally, each of the images will arrive with a different TOA perturbation Δt_j , whose value can be calculated from equations (45)-(47). The amount of overlap between the pulses will depend on Δt_j , the shape of the pulse, and its width. The telescope detects the intensity that results from adding the pulses and taking the squared magnitude. The result is that the pulse shape is distorted with respect to the unlensed shape, and the TOA resulting from the matching procedure is altered.

Figure 7 graphically describes an idealized version of the process. In the absence of lensing and all other kinds of noise, the electric field pulse arrives with a characteristic shape that depends on the specific pulsar being observed, and the telescope detects a pulse intensity that matches the shape of a previously constructed template. The TOA in this idealized case matches the expected TOA based on our timing model, and the timing residual equals zero. This is shown in the leftmost panel of the figure. During a lensing event that leads to the formation of multiple images, the total electric field corresponds to the sum of a number of electric field pulses that arrive at the telescope with different TOAs, amplitudes, and phases, as shown in the middle panel. The resulting intensity, depicted in the rightmost panel, has a different shape than expected, and the center of the pulse is shifted, depending on the Δt_j and A_j of each of the pulses corresponding to the different images, as well as their respective phase differences. This means that the pulse TOA does no longer match the expected TOA based on the timing model.

Note that for this kind of analysis to work, we do not employ the method of uniform asymptotics to calculate the field due to two images merging at a caustic for regions that are far away from the caustic. Equation (39) is useful for determining the total field due to these two images, but it does not tell us anything about the shape and arrival

time shift of the resulting intensity pulse. We know that non-uniform geometric optics matches the uniform asymptotic version far away from the caustics, when the geometric phase difference between the two roots that merge at the singularity satisfies $|\Delta\Phi| \gtrsim \pi$. Consequently, we can directly apply the procedure depicted in Figure 7 at regions outside of the caustic zone. Inside the caustic zone, the two images are very close to each other, and we cannot calculate their individual amplitudes. What we do instead in these regions is to calculate the combined amplitude and phase using (39), and find the TOA shift of the pulse by averaging the Δt_j of the two images, which in fact converges to the same value at the caustic when both images merge. In practical terms, this means that we treat the two merging images as a single image in the caustic zone, which means that our pulse addition procedure involves one less pulse inside the caustic zone than it does in the bright side of the caustic, far away from the singularity.

For large values of DM_l , the phases of each of the images oscillate very fast, and the interference varies rapidly from constructive to destructive and viceversa. Thus, we expect to see a greater spread in the TOAs in regions that contain multiple images. At caustics, images appear or disappear, which can lead to a sharp increase or decrease in the average residuals, which is accentuated because the amplitude of the merging images achieves a maximum close to the caustics. The right panel of Figure 8 displays a simulation of the residuals as a function of frequency close to the epoch of maximum visibility of a lensing event due to an underdense lens. To simplify the analysis, we assume that the observer crosses the u' plane through its center, and that the lens is a Gaussian with maximum dispersion measure perturbation located at the origin of the u plane. Combining the gains and TOAs as described above, and performing the template matching procedure using PyPulse³ (after adding a small amount of white noise), we obtain a scatter plot of the residuals like the one shown in the right panel, which are

³Lam, M. T., 2017, PyPulse, Astrophysics Source Code Library, record ascl:1706.011

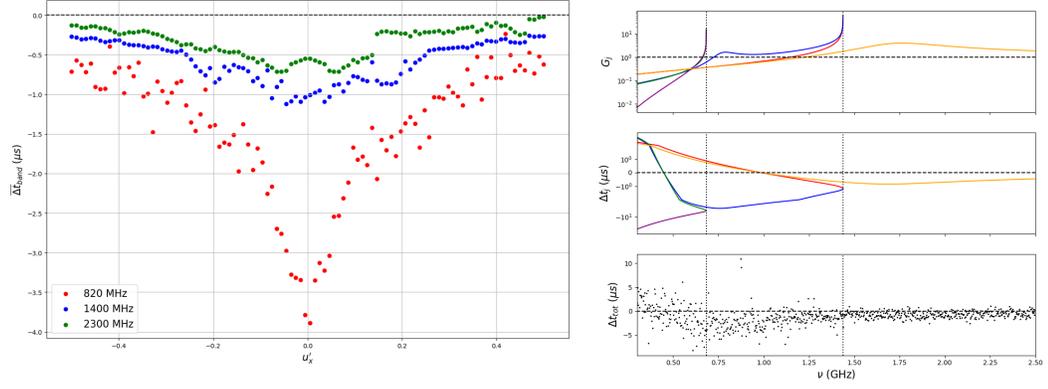


Figure 8: *Left*: Simulation of averaged TOA residuals per frequency band as a function of position in the u' plane for an underdense Gaussian lens with $DM_l = -7 \times 10^{-4}$ pc cm $^{-3}$, $d_{so} = 1.1$ kpc, $d_{sl} = 0.55$ kpc, $a_x = 0.5$ AU, and $a_y = 1.1$ AU. The observer moving through the plane following a path that passes through the origin with slope $m = 0.5$. *Right*: Simulated quantities at $u'_x = 0.05$, close to the epoch of maximum TOA perturbation. The top and middle panels show the individual geometric optic image gains G_j and TOA perturbations Δt_j , with each color representing a different image of the source. For the bottom panel, we perform the procedure described in the main text to obtain the total TOA residuals from the individual G_j and Δt_j . Caustic locations are denoted by the dotted vertical lines, and the horizontal dashed lines correspond to $G = 1$ in the top panel, and $\Delta t = 0$ in the middle and bottom panels.

averaged over frequency subbands with width 3 MHz. We can see that the presence of multiple images results in a larger spread in the data points, and that the maximum averaged perturbations occur at lower frequencies. Different variations of this plot can be obtained by altering the lens parameters, whereupon it is possible to move the caustic boundaries to larger or lower frequencies, and different lens shapes can yield larger numbers of images in certain regions of parameter space.

We can also plot the averaged TOA perturbation per band as a function of the coordinates in the u' plane, as shown in the left panel of Figure 8, with the observer moving along a path that passes through the origin of the u' plane with slope $m = 0.5$. The magnitude of the perturbation increases gradually as a function of u' , reaches a peak at the origin, and then goes back to zero in a similar fashion. The right panel corresponds to the observations made for $u'_x = 0.05$, close to where the magnitude of the averaged

TOA perturbations in the left panel reach a maximum.

We note that the qualitative characteristics of these examples show that the lensing hypothesis constitutes a very plausible explanation for the chromatic timing event observed in the direction of the pulsar J1713+0747, as reported by Lam et al. (2017). Even though the light curve of the pulsar did not show obvious evidence of lensing, the relatively small value of the maximum TOA perturbation (about $3 \mu s$ at the 820 MHz band) suggests that if the event was indeed due to lensing, the refractive strength of the lens was relatively weak, and thus the effects of the lens in the pulsar’s light curve could be masked by scintillation. Unlike the scatter plot depicted in the left panel of Figure 8, the observed event was asymmetric in time, in the sense that the residuals dropped very rapidly at the event’s onset, and then returned gradually to an average value close to zero. Another complication is that a similar event was observed in the direction of that same pulsar some eight years earlier, and it would be desirable to be able to explain both events as the consequence of a single lens crossing our line of sight, which is not possible using a Gaussian-like model. We expect to analyze these issues in further work by specifically applying the methods introduced in this paper in the context of those observations.

4.3 Dispersion measure perturbation

Under ordinary circumstances, when lensing effects due to the interstellar medium are negligible, the TOA perturbation delay is purely dispersive. In order to get accurate TOAs, we need to include this perturbation in our timing model.

The frequency dependent dispersive delay Δt_{DM} is equivalent to the one given in equation (46) with $\psi(\vec{u}) = 1$ and $DM_l = DM$ corresponding to the total integrated

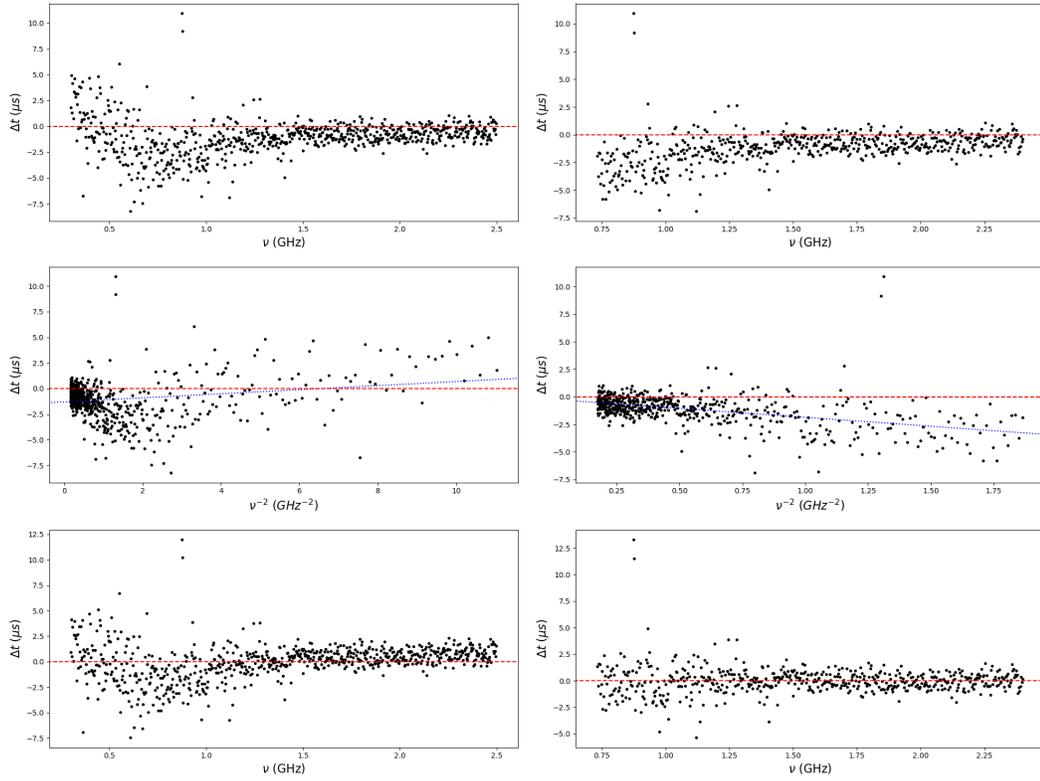


Figure 9: Dedispersion of lensing residuals. The data used is the same as in Figure 8. The left column shows the steps and results for the dedispersion procedure using all the data from the simulation, spanning a frequency range of 0.3 GHz $\leq \nu \leq$ 2.5 GHz. The right column only performs the procedure using the data for 0.74 GHz $\leq \nu \leq$ 2.5 GHz. The top panel shows the residuals as a function of frequency for both cases. The middle panel shows a plot of the residuals as a function of ν^{-2} , with the blue dotted lines corresponding to the lines of best fit. The bottom panel shows the dedispersed residuals.

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, but also introduces a geometric perturbation in the TOAs. Furthermore, these perturbations are 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 does not necessarily hold.

Nevertheless, our results show that it is possible for the deviation to be very small,

and not easy to detect. We illustrate this with a graphical argument. The top panel of Figure 9 shows the lensing residuals from the bottom right panel of Figure 8, for different frequency ranges. The left column shows the residuals for $0.3 \text{ GHz} \leq \nu \leq 2.5 \text{ GHz}$, whereas the right column shows the residuals for $0.74 \text{ GHz} \leq \nu \leq 2.5 \text{ GHz}$. Common PTA observations are made in the latter, more restricted range of frequencies, so it is important to analyze it independently. The next row plots these residuals as a function of ν^{-2} , and calculates the best fit line in order to determine the value of the DM parameter. We can immediately see how the inclusion of smaller frequencies drastically affects the slope of the best fit line, plotted in blue. If we include the lower frequencies, the slope is positive, which leads to a positive estimate of the DM perturbation. On the other hand, if we include only higher frequencies, the resulting DM perturbation is negative. Blindly dedispersing using the derived DM s, we see that the dedispersed residuals look much better in the right column than in the left column.

By eye, the dedispersion procedure appears to work well for the higher frequency case, implying that the deviation from ν^{-2} dependence is not very large. Nevertheless, the value of the dispersion measure obtained from this fitting procedure $DM_{fit} = -3.87 \times 10^{-4} \text{ pc cm}^{-3}$ is about half of the value of the real DM_l perturbation from the lens center, $DM_l = -7 \times 10^{-4} \text{ pc cm}^{-3}$, an effect that can be associated to the geometric delay from refraction. Since the geometric delay decreases as ν^{-4} , including the lower frequencies in the analysis as shown in the left column leads to a highly deceptive value for the DM , and blindly attempting to dedisperse the TOAs using this value will lead to a very misleading form for the residuals. Results like this one can potentially be obtained even in the case where we include only the range of frequencies commonly used for pulsar timing, for cases in which the lens's refractive power is stronger.

It is not clear, however, that the best way to search for refractive contributions to

the TOAs is by looking for delays that scale with ν^{-4} , at least for the case in which the lensing is strong enough to induce caustic formation. The reason is that this scaling applies only for the TOA delay of the individual images, and in practice we do not detect each of these images individually, as explained in the previous section. Since each of the individual images will have its own magnification, some of them will affect the TOA of the pulse as measured by the telescope more than others. These magnifications will depend on the solutions to the lens equation, which in turn depend on a variety of other factors, like the lens shape. The appearance or disappearance of images will also alter the way the residuals change as a function of frequency. The upshot is that the delay resulting from geometry, as measured by the telescope, will depend on more than just ν , and as such it is plausible that these other factors mask the expected ν^{-4} dependence.

Perhaps a better way of looking for the signature of lensing events in timing observations is by looking at the scatter of the points in the frequency dependent residuals. If multiple imaging occurs, the interference between the images will increase the scatter of the TOAs because the phase for each of the images will necessarily vary very rapidly if the lensing is to be strong enough to significantly alter the average residuals per band. In the case of the underdense lens, the number of images is larger at lower frequencies, and thus we expect the scatter to be more pronounced in these regions. Furthermore, we expect the amount of scatter to decrease dramatically when crossing caustics, although this can be hard to detect in practice because we observe at separated frequency bands, and the caustics can potentially fall within the regions of separation.

We expect this section's analysis to be the starting point for more sophisticated simulations that yield more quantitative results, and that work with a larger variety of two dimensional lenses and paths through the u' plane. Some other interesting research directions along these lines include investigating lensing effects in dimensional coordi-

nates, analyzing their effects on TOA uncertainties, and quantitatively determining how much of a ν^{-4} dependence in the residuals we should expect to observe during a lensing event. An important requirement for a more complete analysis will be the implementation of asymptotic methods that work for higher order caustics and caustic intersections, as we invariably expect to encounter these in most realizations of dynamic spectra and for more complicated lens shapes.

CHAPTER 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 only somewhat useful in the limit of very strong lensing, and that the effects of weaker lenses on flux time series involve oscillations that are not well captured by said approach.

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 applied 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 fluctuation acquires the form of an overdensity or underdensity, and a series of other parameters such as the lens size, distances, and the frequencies at which observations are made. 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 in simulations.

Consistent with the results of previous works (Goodman et al., 1987; Melrose and Watson, 2006; Watson and Melrose, 2006; Stinebring et al., 2007), we find that lensing effects tend to be stronger at lower frequencies, primarily because ray crossings are more likely to occur at lower frequencies than at larger 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 (Clegg et al., 1998; Stinebring et al., 2007; Cordes et al., 2017; Er and Rogers, 2017). 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. (2015) and Lam et al. (2017) 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 conceivable 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.

APPENDIX A
SOLVING THE KDI USING THE FFT

The two dimensional Kirchhoff diffraction integral (KDI) introduced in section 2.3, gives the normalized wave optics field ε as a function of the observer coordinates \vec{u}' by integrating over an angular spectrum of plane waves,

$$\varepsilon(\vec{u}') = \frac{a_x a_y}{2\pi r_F^2} \iint d^2 u \exp(i\Phi) \quad (\text{A1})$$

where Φ is the geometric phase,

$$\Phi(\vec{u}', \vec{u}) = \frac{1}{2r_F^2} \left[a_x^2 (u_x - u'_x)^2 + a_y^2 (u_y - u'_y)^2 \right] + \phi_0 \psi(\vec{u}) \quad (\text{A2})$$

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,

$$\varepsilon(\vec{u}') = \iint d^2 u G(\vec{u} - \vec{u}') H(\vec{u}) \quad (\text{A3})$$

where

$$G(\vec{u}) = \frac{a_x a_y}{2\pi r_F^2} \exp \left[\frac{i}{2r_F^2} (a_x^2 u_x^2 + a_y^2 u_y^2) \right] \quad (\text{A4})$$

$$H(\vec{u}) = \exp [i\phi_0 \psi(\vec{u})] \quad (\text{A5})$$

From the discrete version of the convolution theorem (Schmidt 2010), we have that

$$\varepsilon(\vec{u}') = \mathcal{F}^{-1} \{ \mathcal{F} [G(\vec{u})] \cdot \mathcal{F} [H(\vec{u})] \} \quad (\text{A6})$$

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 and Melrose (2006) and Melrose and Watson (2006), and in two dimensions by Stinebring et al. (2007), 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.

Although useful, this approach suffers from two serious limitations. First, it does not give information about the number of images of the source that can be seen by the observer or the respective amplifications, phases, and TOAs of each of these images, so it is not clear how we could use it to find TOA perturbations due to lensing. Second, in practice the method can only be applied for a restricted range of lens scales a_i and relatively small values of ϕ_0 . The issue is the grid size necessary to properly sample the oscillations of the functions $G(\vec{u})$ and $H(\vec{u})$. 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 2010)

$$n_{max} = \frac{\pi r_F^2}{(\Delta x)^2} \quad (\text{A7})$$

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

$$\Delta x = \frac{2au'_{max}}{N} \quad (\text{A8})$$

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

$$N = \frac{4a^2(u'_{max})^2}{\pi r_F^2} \quad (\text{A9})$$

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 (2010).

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 A1 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.

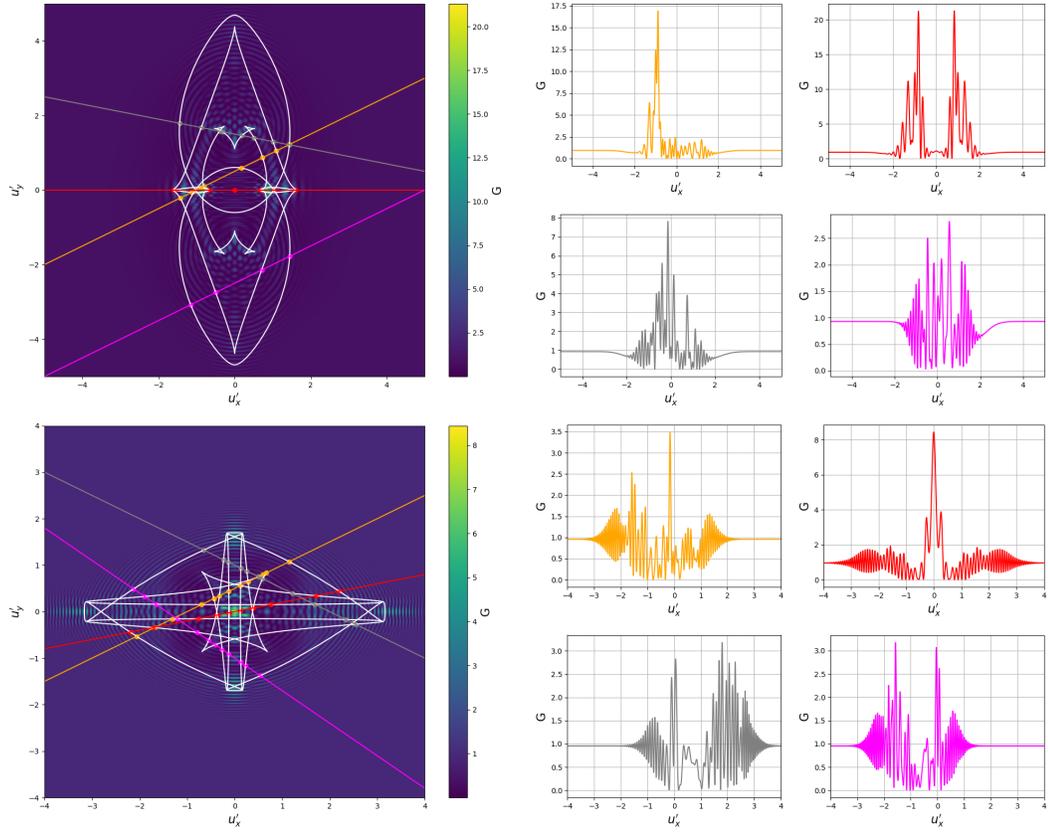


Figure A1: *Left*: Colormap of the gain in the u' plane overlaid with the caustic curves in white and slices along the plane in different colors for two different lens shapes. The points of intersection between the slices and the caustic lines are marked by points. The top panel corresponds to a lens with shape $\psi(\vec{u}) = 0.74(u_x^2 + u_y^6) \exp(-u_x^2 - u_y^2)$, and parameters $a_x = a_y = 0.02$ AU, $DM_l = -1.5 \times 10^{-6}$ pc cm $^{-3}$, $\nu = 0.8$ GHz, $d_{so} = 1$ kpc, and $d_{sl} = 0.5$ kpc. The bottom lens has $\psi(\vec{u}) = \exp(-u_x^4 - u_y^4)$, $a_x = 0.04$ AU, $a_y = 0.05$ AU, $DM_l = -2 \times 10^{-6}$ pc cm $^{-3}$, $\nu = 1.0$ GHz, $d_{so} = 5$ kpc, $d_{sl} = 2.5$ kpc. In both cases, $\phi_0 \approx 50$. *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.

APPENDIX B

ESTIMATION OF THE VALUE OF THE GAIN AT A 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 (equation (16)), since the oscillations due to multiple imaging will give a value of the flux consistent with the prediction from that equation once we average over the size of a frequency subband. 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 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 (1965) and Wong (2001), the gain at the singularity is

$$G_{max} = \frac{a_x^2 a_y^2}{12\pi r_F^4} \frac{\Gamma^2(1/3)}{|\Phi_{20}| |\Phi_{03}|^{2/3}} \quad (\text{B10})$$

Bleistein and Handelsman (1975) and Cooke (1982) give a more complicated expression,

$$G_{max} = \frac{a_x^2 a_y^2}{4\pi^2 r_F^4} |\Phi_{20}| \Gamma^2(1/3) \left(\frac{32\pi^2}{3|B|^2} \right)^{1/3} \quad (\text{B11})$$

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 (B1) and (B2) 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. Establishing which one is more accurate is outside the scope of this work.

APPENDIX C

NUMERICS

The key behind successful application of geometric optics as presented in the main text is the ability to numerically solve the lens equation (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 (that is, the equation can be multivalued), 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

$$\begin{bmatrix} f(x, y) \\ g(x, y) \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \quad (\text{C12})$$

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 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 two 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 and Kayser (1987) 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 it 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 SciPy¹ 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 they don't provide a practical method 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 \vec{u}' , we apply the contour plotting algorithm to find the intersections between the curves in the u' plane that satisfy (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 (26).

This step allows us to separate the regions in parameter space that contain different numbers of solutions to the lens equation. Now, we 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

¹Jones E, Oliphant E, Peterson P, et al. SciPy: Open Source Scientific Tools for Python, 2001-, <http://www.scipy.org/>

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 at many points in four dimensional space is not practically 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.

APPENDIX D

MORE NUMERICAL EXAMPLES AND LENS COLORMAPS

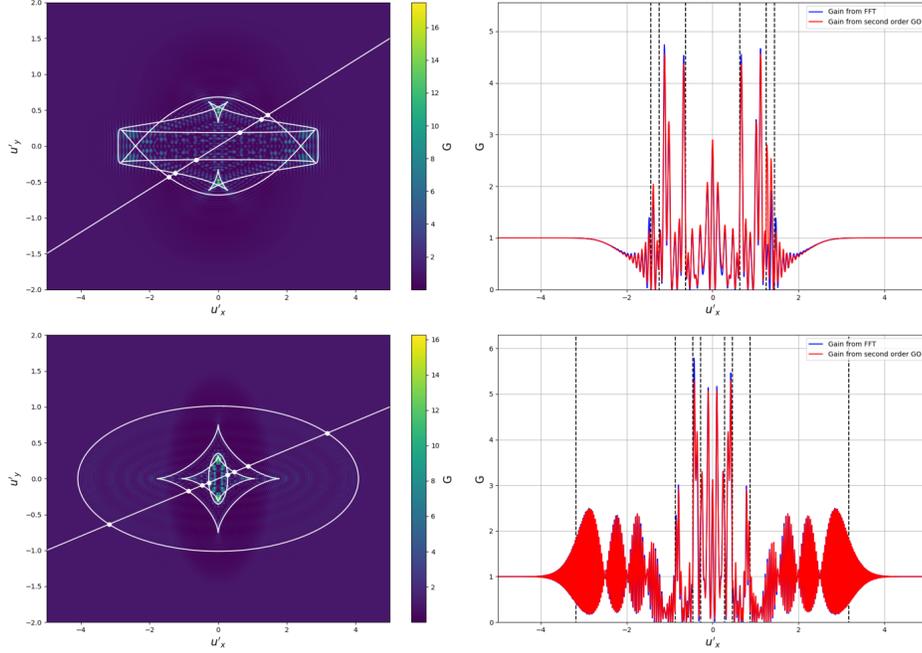


Figure D2: Comparison of the gains obtained from a full numerical solution of the KDI and second order geometric optics. The left column shows color maps of the gain obtained by solving the KDI via the FFT. The white circles correspond to caustic curves, and the straight white line shows the path of 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 $\psi(\vec{u}) = \exp(-u_x^2 - u_y^4)$, strength parameter $\phi_0 = 80$, 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 $\psi(\vec{u}) = \exp\left[-(u_x^2 + u_y^2)^3\right]$, strength parameter $\phi_0 = 120$, 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.

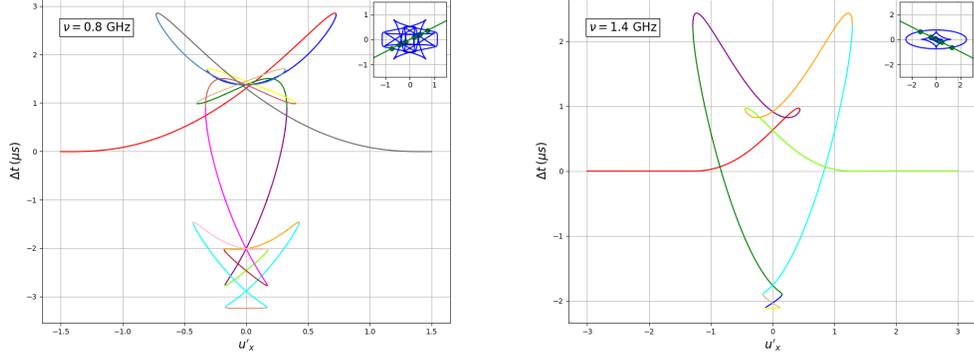


Figure D3: Individual image TOAs for two different lenses and different paths through the u' plane. The left panel corresponds to a square Gaussian lens with $\psi(\vec{u}) = \exp(-u_x^4 - u_y^4)$, $DM_l = -5 \times 10^{-4} \text{ pc cm}^{-3}$, $d_{so} = 1 \text{ kpc}$, $d_{sl} = 0.5 \text{ kpc}$, $a_x = 0.5 \text{ AU}$ and $a_y = 0.6 \text{ AU}$. The frequency of observation is $\nu = 0.8 \text{ GHz}$, which gives a strength parameter of $\phi_o \approx 1.63 \times 10^4$. The right panel corresponds to a super-Gaussian lens with $\psi(\vec{u}) = \exp\left[-(u_x^2 + u_y^2)^2\right]$, $DM_l = -1 \times 10^{-2} \text{ pc cm}^{-3}$, $d_{so} = 5 \text{ kpc}$, $d_{sl} = 2.5 \text{ kpc}$, $a_x = 0.7 \text{ AU}$ and $a_y = 1 \text{ AU}$, with $\nu = 1.4 \text{ GHz}$, and thus $\phi_o \approx 1.86 \times 10^4$. 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.

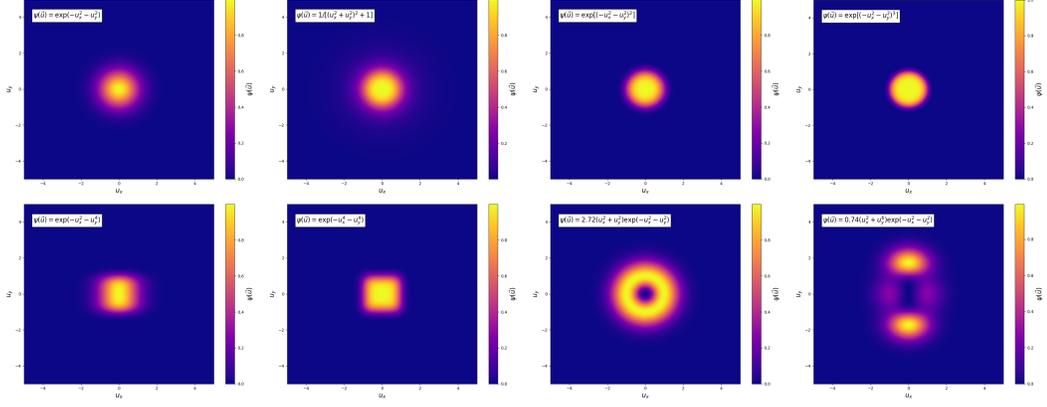


Figure D4: Colormaps of the different types of lensing structures used in the text and appendices. *Top row, from left to right:* Gaussian lens, $\psi(\vec{u}) = \exp(-u_x^2 - u_y^2)$, Lorentzian lens, $\psi(\vec{u}) = 1/[(u_x^2 + u_y^2) + 1]$, and super-Gaussian lenses $\psi(\vec{u}) = \exp[-(u_x^2 + u_y^2)^2]$ and $\psi(\vec{u}) = \exp[-(u_x^2 + u_y^2)^3]$. *Bottom row, from left to right:* Rectangular Gaussian lens, $\psi(\vec{u}) = \exp(-u_x^2 - u_y^4)$, square Gaussian lens, $\psi(\vec{u}) = \exp(-u_x^4 - u_y^4)$, ring-like lens $\psi(\vec{u}) = 2.72(u_x^2 + u_y^2) \exp(-u_x^2 - u_y^2)$, and double lens $\psi(\vec{u}) = 0.74(u_x^2 + u_y^6) \exp(-u_x^2 - u_y^2)$.

BIBLIOGRAPHY

- M. V. Berry. Waves and Thom's Theorem. *Advances in Physics*, 25(1):1–26, 1976.
- M. V. Berry and C. Upstill. Catastrophe optics: morphologies of caustics and their diffraction patterns. In Emil Wolf, editor, *Progress in Optics*, volume 18, pages 257–346. Elsevier, 1980.
- N. Bleistein and R. A. Handelsman. *Asymptotic expansions of integrals*. Courier Corporation, 1975.
- M. Born and E. Wolf. *Principles of Optics*. Cambridge University Press, seventh edition, October 1999.
- V. A. Borovikov and B. E. Kinber. *Geometrical theory of diffraction*. Iet, 1994.
- K. G. Budden and P. D. Terry. Radio Ray Tracing in Complex Space. *Proceedings of the Royal Society of London Series A*, 321:275–301, February 1971. doi: 10.1098/rspa.1971.0033.
- N. Chako. Asymptotic expansions of double and multiple integrals occurring in diffraction theory. *IMA Journal of Applied Mathematics*, 1(4):372–422, 1965.
- C. Chester, B. Friedman, and F. Ursell. An extension of the method of steepest descents. *Mathematical Proceedings of the Cambridge Philosophical Society*, 53(3):599–611, 1957.
- A. W. Clegg, A. L. Fey, and R. L. Fiedler. VLA polarization observations of the extragalactic source 1741-038 during and extreme scattering event. *The Astrophysical Journal*, 457:L23–L26, 1996.

- A. W. Clegg, A. L. Fey, and T. J. W. Lazio. The Gaussian Plasma Lens in Astrophysics: Refraction. *The Astrophysical Journal*, 496:253–266, March 1998. doi: 10.1086/305344.
- I. Cognard, G. Bourgois, J-F. Lestrade, et al. An extreme scattering event in the direction of the millisecond pulsar 1937+21. *Nature*, 366:320–322, 1993.
- W. A. Coles, M. Kerr, R. M. Shannon, et al. Pulsar observations of extreme scattering events. *The Astrophysical Journal*, 808(2):113–119, 2015.
- J. N. L. Connor. Practical methods for the uniform asymptotic evaluation of oscillating integrals with several coalescing saddle points. In R. Wong, editor, *Asymptotic and Computational Analysis*, chapter 6, pages 137–173. Marcel Dekker, Inc, 1990.
- J. C. Cooke. Stationary phase in two dimensions. *IMA Journal of Applied Mathematics*, 29(1):25–37, 1982.
- J. M. Cordes and R. M. Shannon. A measurement model for precision pulsar timing. 2010.
- J. M. Cordes and A. Wolszcan. Multiple imaging of pulsars by refraction in the interstellar medium. *The Astrophysical Journal*, 307:L27–L31, 1986.
- J. M. Cordes, A. Pidwerbetsky, and R. E. Lovelace. Refractive and diffractive scattering in the interstellar medium. *The Astrophysical Journal*, 310:737–767, 1986.
- J. M. Cordes, B. J. Rickett, D. R. Stinebring, and W. A. Coles. Theory of parabolic arcs in interstellar scintillation spectra. *The Astrophysical Journal*, 637(1):346–365, 2006.
- J. M. Cordes, R. M. Shannon, and D. R. Stinebring. Frequency-dependent dispersion measures and implications for pulsar timing. *The Astrophysical Journal*, 817:16–34, 2016.

- 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. *The Astrophysical Journal*, 842:35, June 2017. doi: 10.3847/1538-4357/aa74da.
- L. Dai and W. Lu. Probing motion of fast radio burst sources by timing strongly lensed repeaters. *The Astrophysical Journal*, 847:19–35, 2017.
- R. B. Dingle. *Asymptotic expansions: their derivation and interpretation*, volume 48. Academic Press London, 1973.
- X. Er and A. Rogers. Two families of astrophysical diverging lens models. *Monthly Notices of the Royal Astronomical Society*, 475:867–878, 2017.
- R. Fiedler, B. Dennison, K. J. Johnston, et al. A summary of extreme scattering events and a descriptive model. *The Astrophysical Journal*, 430:581–594, 1994.
- R. L. Fiedler, B. Dennison, K. J. Johnston, and A. Hewish. Extreme scattering events caused by compact structures in the interstellar medium. *Nature*, 326(16):675–678, 1987.
- J. J. Goodman, R. W. Romani, R. D. Blandford, and R. Narayan. The effects of caustics on scintillating radio sources. *Monthly Notices of the Royal Astronomical Society*, 229:73–102, November 1987. doi: 10.1093/mnras/229.1.73.
- Y. Gupta, N. D. R. Ghat, and A. P. Rao. Multiple imaging of PSR B1133+16 by the interstellar medium. *The Astrophysical Journal*, 520(1):173–181, 1999.
- C. J. Howls. Hyperasymptotics for multidimensional integrals, exact remainder terms and the global connection problem. *Proceedings of the Royal Society of London Series A*, 453:2271–2294, 1997.

- D. Kaminski. Exponentially improved stationary phase approximations for double integrals. *Methods and Applications of Analysis*, 1(1):44–56, 1994.
- T. Katsaounis, G. T. Kossioris, and G. N. Makrakis. Computation of high frequency fields near caustics. *Mathematical Models and Methods in Applied Sciences*, 11(2): 199–228, 2001.
- M. Kerr, W. A. Coles, C. A. Ward, et al. Extreme scattering events towards two young pulsars. *Monthly Notices of the Royal Astronomical Society*, 474(4):4637–4647, 2017.
- Yu. A. Kravstov. Two new asymptotic methods in the theory of wave propagation in inhomogeneous media (review). *Soviet Physics: Acoustics*, 14:1–14, 1968.
- Yu A Kravtsov and Yu I Orlov. *Caustics, catastrophes and wave fields*, volume 15. Springer Science & Business Media, 1999.
- Yu A Kravtsov, G. W. Forbes, and A. A. Asatryan. I theory and applications of complex rays. In Emil Wolf, editor, *Progress in optics*, volume 39, pages 1–62. Elsevier, 1999.
- A. S. Kryukovskii, D. S. Lukin, E. A. Palkin, and D. S. Rastyagaev. Wave catastrophes: types of focusing in diffraction and propagation of electromagnetic waves. *Journal of Communications Technology and Electronics*, 51(10):1087, 2006.
- M. T. Lam, J. A. Ellis, G. Grillo, et al. A second chromatic timing event of interstellar origin toward PSR J1713+0747. 2017.
- D. Ludwig. Uniform asymptotic expansions at a caustic. *Communications on Pure and Applied Mathematics*, 19(2):215–250, 1966.
- D. B. Melrose and P. G. Watson. Scintillation of Radio Sources: The Role of Caustics. *The Astrophysical Journal*, 647:1131–1141, August 2006. doi: 10.1086/505589.

- J. F. Nye. Optical caustics in the near field from liquid drops. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical*, 361(1704):21–41, 1978.
- U. L. Pen and Lindsay King. Refractive convergent plasma lenses explain extreme scattering events and pulsar scintillation. *Monthly Notices of the Royal Astronomical Society: Letters*, 421:L132–L136, 2012.
- T. Poston and I. Stewart. *Catastrophe Theory and its Applications*. Dover Publications, 1978.
- A. B. Pushkarev, Y. Y. Kovalev, M. L. Lister, et al. VLBA observations of a rare multiple quasar imaging event caused by refraction in the interstellar medium. *Astronomy and Astrophysics*, 555:80–92, 2013.
- W.-Y. Qiu and R. Wong. Uniform asymptotic expansions of a double integral: coalescence of two stationary points. *Proceedings of the Royal Society of London Series A*, 456:407–431, February 2000. doi: 10.1098/rspa.2000.0523.
- J. D. Schmidt. *Numerical Simulation of Optical Wave Propagation with examples in MATLAB*. SPIE, 2010.
- P. Schneider, J. Ehlers, and E. E. Falco. *Gravitational Lenses*. Springer, 1992.
- T. Schramm and R. Kayser. A simple imaging procedure for gravitational lenses. *Astronomy & Astrophysics*, 174:361–364, March 1987.
- T. Schramm and R. Kayser. The complex theory of gravitational lensing. *Astronomy and Astrophysics*, 299:1–10, 1995.
- J. J. Stamnes. *Waves in Focal Regions: Propagation, Diffraction and Focusing of Light, Sound and Water Waves*. Adam Hilger, 1986.

- D. Stinebring, J. Matters, and D. Hemberger. Diffraction from 2d Lenses in the ISM. In M. Haverkorn and W. M. Goss, editors, *SINS - Small Ionized and Neutral Structures in the Diffuse Interstellar Medium*, volume 365 of *Astronomical Society of the Pacific Conference Series*, page 275, July 2007.
- R. Thom. *Stabilite structurelle et morphogenese*. Benjamin, 1972.
- K. S. Thorne and R. D. Blandford. *Modern Classical Physics*. Princeton University Press, 2017.
- A. V. Tuntsov, M. A. Walker, L. V. E. Koopmans, K. W. Bannister, J. Stevens, S. Johnston, C. Reynolds, and H. E. Bignall. Dynamic spectral mapping of interstellar plasma lenses. *The Astrophysical Journal*, 817(2):176, 2016.
- F. Ursell. Integrals with a large parameter. the continuation of uniformly asymptotic expansions. *Proceedings of the Cambridge Philosophical Society*, 61:113–128, 1965.
- H. K. Vedantham, A. C. S. Readhead, T. Hovatta, et al. The peculiar light curve of J1415+1320: A case study in extreme scattering events. *The Astrophysical Journal*, 845:90–98, 2017.
- P. G. Watson and D. B. Melrose. Scintillation of Radio Sources: The Signature of a Caustic. *The Astrophysical Journal*, 647:1142–1150, August 2006. doi: 10.1086/505591.
- Roderick Wong. *Asymptotic approximations of integrals*, volume 34. SIAM, 2001.