

Oscillatory path integrals for radio astronomy

Job Feldbrugge,^{1,*} Ue-Li Pen,^{2,†} and Neil Turok^{1,‡}

¹*Perimeter Institute, 31 Caroline St N, Ontario, N2L 2Y5, Canada*

²*Canadian Institute for Theoretical Astrophysics,
University of Toronto, M5S 3H8, ON, Canada*

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

*Electronic address: jfeldbrugge@perimeterinstitute.ca

†Electronic address: pen@cita.utoronto.ca

‡Electronic address: nturok@perimeterinstitute.ca

Contents

I. Introduction	3
II. From Feynman to Fermat to Fresnel-Kirchhoff	9
III. Evaluating the Fresnel-Kirchhoff integral	12
A. The geometric optics limit	15
B. Catastrophe theory	19
IV. Example: a one-dimensional lens	23
A. Geometric optics approximation	25
B. Finding the thimbles	27
1. Method 1: following steepest ascent contours	28
2. Method 2: flowing the integration domain	29
C. Integrating along the thimbles	34
V. The elementary catastrophes	37
A. The fold A_2	37
1. Asymptotics	38
B. The cusp A_3	39
1. Numerics	41
2. Asymptotics	42
C. The swallowtail A_4	43
1. Numerics	48
D. The elliptic umbilic D_4^-	49
1. Numerics	54
E. The hyperbolic umbilic D_4^+	54
1. Numerics	58
VI. Two dimensional localized lenses	59
A. A generic peak	60
B. A degenerate peak	64
C. The swallowtail caustic	65

D. The elliptic umbilic caustic	68
VII. Signatures of caustics in fast radio bursts	71
VIII. Conclusions	73
Acknowledgments	74
References	75
A. Defining oscillatory integrals	80
1. One-dimensional integral	80
2. Multi-dimensional integrals	82
B. Young’s double-slit experiment	83

I. INTRODUCTION

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

In quantum mechanics, these difficulties run deep. In fact, so far they have thwarted all efforts to rigorously define nontrivial real-time Feynman path integrals, even in non-relativistic quantum mechanics [1]. The only available existence proofs involve a Wick rotation from real, Lorentzian time to imaginary, Euclidean time, which maps the phase factor to a real Boltzmann weight (for a recent review see, *e.g.*, [2]). Unfortunately, securing mathemati-

cal rigour this way comes at a high price: the system’s dynamics can only be described in imaginary time instead of real time where experiments and observations actually take place. Analytic continuation back to real time is often only possible for certain quantities, such as perturbative S-matrix elements and, even then, is often hard. Furthermore, for some theories, including general relativity and quantum condensed matter models with a “sign problem,” *e.g.* the Hubbard model, the Wick rotation trick does not work.

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

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

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

Spatial variations in the refractive index of astrophysical plasmas can arise due to turbu-

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

There is already an extensive astrophysical literature on the computation of interference patterns in wave optics [23, 24], but published methods tend to converge slowly [25]. They are expensive to implement and the results are sometimes inconclusive. In this paper, we shall present faster and more reliable methods. Our approach builds on Picard-Lefschetz theory, a general, exact approach to multidimensional oscillatory integrals based upon saddle point and steepest descent techniques (for an introduction [26]; for applications to quantum cosmology, see [27–31] and to relativistic quantum mechanics, see [32]). As we shall show, our methods allow for the fast and reliable computation of even very intricate “diffraction catastrophe” patterns. The calculations of these patterns by Berry, Nye and collaborators were an analytical *tour de force*, but relied heavily on the particularities of Thom’s canonical “normal forms” of catastrophes, and the mathematical properties of the related special functions, with each case treated separately. Unfortunately, while the normal forms represent the correct universal forms locally, they diverge at large distances. Hence, they are unrealistic as models for natural lenses. Realistic modeling requires a more versatile method which can be efficiently and straightforwardly implemented numerically. We present just such a method here.

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

Instead of using specific properties of special functions and symmetries, our method exploits Cauchy’s theorem to exactly transform an integral of an oscillatory phase factor into a sum of absolutely convergent integrals taken over “Lefschetz thimbles” in the space of complexified coordinates. These “thimble” integrals are fast to compute numerically, requiring only polynomial time. They are insensitive to numerical cutoffs and may be performed iteratively in any order with no change to the result. In this paper, we demonstrate the efficacy of our method by computing the interference patterns for one- and two-dimensional thin lenses. We study the most observationally accessible catastrophes, both in their “normal forms” and in a set of more realistic, localized lens models where these catastrophes appear. Our one-dimensional numerical code, capable of handling generic one-dimensional lenses, is now publicly available [online](#)¹.

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

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

Here, $\Psi(\mu, \alpha, \nu)$ is the amplitude whose square $|\Psi(\mu, \alpha, \nu)|^2$ gives the intensity of light observed at a position, frequency and lens strength controlled by the parameters μ , ν and α . The control parameter μ is determined by the transverse positions of the observer and the source relative to the lens (See Fig. 1 and Eq. (13) below). The frequency of the waves

¹ See https://github.com/jfeldbrugge/Picard_Lefschetz_Integrator.

is proportional to ν so the spacing of interference fringes shrinks as ν is increased. The eikonal limit is $\nu \rightarrow \infty$. Finally, α controls the strength of the lens which, in this example, is taken to have a Lorentzian profile. The integral (1) is analytically intractable. However, it is simple to compute numerically, for reasonable values of ν , μ and α , using the methods we shall describe below.

In the eikonal limit of large ν , only real saddle point solutions – real stationary points of the phase $\phi(x)$ – contribute significantly to the amplitude. Each one corresponds to a particular ray. For $\alpha < 1$ the lens is “weak” and there is only one real solution of $\partial_x \phi(x) = 0$. Hence there is only one contributing ray at each value of μ . For $\alpha > 1$, the lens is “strong:” for a finite range of μ values centred on zero, there are three real solutions of $\partial_x \phi(x) = 0$ hence three contributing rays. Correspondingly, one finds three images of the source in this range of μ . The values of μ bounding this range mark a transition from three contributing rays (*i.e.* three real saddles in the phase) to one. At these values of μ , a maximum and a minimum of $\phi(x)$ merge into a *cubic* stationary point (*i.e.* a point of inflexion), creating the simplest “fold” catastrophe. If we now decrease the strength of the lens α towards unity, the two “fold” catastrophes approach the point $\mu = 0$ where they merge to form a “cusp” catastrophe, in which there is a *quartic* stationary point in the phase $\phi(x)$. Since the phase (viewed as a function of x) is flatter in the vicinity of higher order stationary points, there is less destructive interference. The intensity of light grows more rapidly as ν is increased as compared to the intensity from a quadratic saddle, so that “folds” become increasingly bright compared to the unlensed image and “cusps” become even brighter. While higher order catastrophes are rarer, their brightness makes them easier to detect. This has encouraged the conjecture, yet to be verified [8, 12, 14, 25, 33], that the brightest sources seen may be those which happen to be lensed into high order catastrophes.

In order to emphasize the foundational character of the physics at play and by way of a pedagogical introduction, we show how the standard Fresnel-Kirchhoff integral (see, *e.g.*, Ref. [34], Chapter 8), central to the description of lensing in radioastronomy and in optics²,

² The integral formula has a fascinating history of successive approximate derivations and subsequent critiques, reviewed in Ref. [34], Chapter 8. Exact solutions of Maxwell’s equations (or their scalar version) representing quasi-realistic interference patterns created by diffraction around physical obstacles of various types are still few in number, and are reviewed in Chapter 11 of the same work. It would be interesting to revisit these solutions and, perhaps to find others, using the ideas we develop here.

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

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

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

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

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

II. FROM FEYNMAN TO FERMAT TO FRESNEL-KIRCHHOFF

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

the directness and economy of this approach compared to more standard (and cumbersome) derivations based on Maxwell's equations, or their scalar counterpart.

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

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

Here, ω and k are the angular frequency and wavenumber of the waves, c is the speed of light and $\omega_p(\mathbf{x})$ is the plasma frequency at position \mathbf{x} , determined by the local density of electrons, assumed to vary across space on scales much larger than the wavelength of the electromagnetic waves. Notice that (2) takes exactly the same form as the dispersion relation for a relativistic particle whose mass varies with spatial position.

The dispersion relation (2) yields a phase propagation speed

$$v_p(\mathbf{x}) \equiv \frac{\omega}{k} = c \sqrt{1 + \frac{\omega_p^2(\mathbf{x})}{k^2 c^2}}, \quad (3)$$

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

The quanta of these waves may be described as relativistic particles, following parameterized worldlines in spacetime: $x^\mu(\lambda) = (ct(\lambda), \mathbf{x}(\lambda))$. Reparameterizations $\lambda \rightarrow \tilde{\lambda}(\lambda)$ are generated by a Hamiltonian, and reparameterization invariance corresponds to the constraint that the Hamiltonian vanishes, $\mathcal{H} = 0$. The correct expression for the Hamiltonian \mathcal{H} may be read off from the dispersion relation (2), using the correspondence $p_\mu = (p_0, \mathbf{p}) \leftrightarrow \hat{p}_\mu = -i\hbar\partial_\mu = \hbar(-\omega/c, \mathbf{k})$:

$$\mathcal{H} = -p_0^2 c^2 + \mathbf{p}^2 c^2 + \hbar^2 \omega_p^2(\mathbf{x}). \quad (4)$$

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

$$S[x; x^\mu(0), x^\mu(1)] = \int_0^1 d\lambda (p_0 \dot{x}^0 + \mathbf{p} \cdot \dot{\mathbf{x}} - \tau(\lambda) \mathcal{H}). \quad (5)$$

where dots denote derivatives with respect to λ , taken to run from 0 to 1 as the the par-

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

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

$$S[\mathbf{x}; E, \mathbf{x}_i, t_f, \mathbf{x}_f] = p_0 x^0(0) + \int_0^1 d\lambda (p_0 \dot{x}^0 + \mathbf{p} \dot{\mathbf{x}} - \tau(\lambda) \mathcal{H}), \quad (6)$$

with \mathcal{H} given in (4).

Since the action (6) is quadratic in the momenta and linear in τ , we can integrate out those variables. At the relevant saddle, we may use Hamilton’s equations for the momenta, and the constraint, to obtain a reduced action expressed purely in terms of reparameterization-invariant quantities:

$$S_r[\mathbf{x}] = -Et_i - \int_{t_i}^{t_f} dt \frac{\hbar^2 \omega_p^2(\mathbf{x}(t))}{E}. \quad (7)$$

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

$$S_r[\mathbf{x}] = E \int_{t_i}^{t_f} dt \left(1 - \frac{\hbar^2 \omega_p^2(\mathbf{x})}{E^2} \right) = E \int_{\mathbf{x}_i}^{\mathbf{x}_f} \frac{|d\mathbf{x}|}{c} \left(1 - \frac{\hbar^2 \omega_p^2(\mathbf{x})}{E^2} \right)^{\frac{1}{2}} = \int_{\mathbf{x}_i}^{\mathbf{x}_f} |d\mathbf{x}| |\mathbf{p}|, \quad (8)$$

where, again, we used Hamilton’s equations for the momenta and the Hamiltonian constraint.

Finally, we express the result in terms of the phase velocity (3), obtaining

$$S_r[\mathbf{x}] = E \int_{\mathbf{x}_i}^{\mathbf{x}_f} \frac{|\mathbf{d}\mathbf{x}|}{c_p(\mathbf{x})}. \quad (9)$$

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

The reduced action (9) embodies Fermat's principle of least time or, more correctly, the principle that the time taken is stationary on dominant classical trajectories. The path integral over all paths, weighted by $e^{iS_r[\mathbf{x}]/\hbar}$, is the Fresnel-Kirchhoff integral we seek.

As an aside, note that one may, equally well, obtain the result (9) starting from the square root (Nambu-type) action for a particle with a spatially dependent mass $m(\mathbf{x})$, by making use of the correspondence $m(\mathbf{x}) c^2 \leftrightarrow \hbar\omega_p(\mathbf{x})$, namely

$$\mathcal{S}[\mathbf{x}] = - \int_{t_i}^{t_f} dt \hbar\omega_p(\mathbf{x}(t)) \left(1 - \frac{\dot{\mathbf{x}}(t)^2}{c^2}\right)^{\frac{1}{2}}, \quad (10)$$

where the dot now denotes a t derivative. This action is explicitly reparameterization invariant from the start. However, it is the action appropriate to fixing the initial time t_i whereas we need to fix the initial energy E . As before, we must supplement the action (10) by a boundary term, which turns out to be $+E(t_f - t_i)$. One can easily check that the identity $\partial S_{cl}/\partial t_i = E$ for Hamilton's principal function S_{cl} implies the total action is stationary, provided the desired boundary conditions and the equations of motion are fulfilled. Using $E = \hbar\omega_p/(1 - \dot{\mathbf{x}}^2/c^2)^{\frac{1}{2}}$, the total action reduces (again, up to a constant phase) to (8) as before.

III. EVALUATING THE FRESNEL-KIRCHHOFF INTEGRAL

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




Figure 1: The geometry of interfering paths passing through a thin lens.

a possible bend at the lens, and the integral over these paths reduces to an ordinary integral over the lens plane [1]. The path integral amplitude for a (spinless) photon is obtained by integrating over all paths weighted by the phase factor $e^{iS_r[\mathbf{x}]/\hbar}$ obtained from (9):

$$\Psi(\mathbf{x}_{obs}, \mathbf{x}_s) = \int d\mathbf{x} \exp \left[i\omega \int_{\mathbf{x}_s}^{\mathbf{x}_{obs}} |d\mathbf{x}| \frac{n(\mathbf{x})}{c} \right]. \quad (11)$$

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

In the thin lens approximation, variations in the phase arise in part geometrically, from variations in the length of the straight line segments on either side of the lens, and in part from the passage through the lens. The former are straightforwardly computed using the Pythagorean theorem in the approximation that the relative horizontal displacements in Fig. 1, $|\mathbf{x}_{obs} - \mathbf{x}|$ and $|\mathbf{x}_s - \mathbf{x}|$ are much smaller than the vertical distances d_{lo} and d_{sl} . The latter are likewise computed approximately, noting that, to lowest order, the paths pass vertically through the lens so we may replace $\int n_e(\mathbf{x}, z) dz$ with $\Sigma_e(\mathbf{x})$, the electron surface density.

The path integral amplitude then becomes

$$\Psi(\mathbf{x}_{obs}, \mathbf{x}_s; \nu) \propto \int d\mathbf{x} \exp \left[i \frac{\omega}{2c} \left(\frac{(\mathbf{x} - \boldsymbol{\mu})^2}{\bar{d}} - \frac{\Sigma_e(\mathbf{x})e^2}{m_e \epsilon_0 \omega^2} \right) \right], \quad (12)$$

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

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

$$\Psi(\boldsymbol{\mu}; \nu) = \left(\frac{\nu}{\pi} \right)^{N/2} \int_{\mathbb{R}^N} d\mathbf{x} \exp [i\phi(\mathbf{x}; \boldsymbol{\mu})\nu], \quad \text{with} \quad \phi(\mathbf{x}; \boldsymbol{\mu}) = (\mathbf{x} - \boldsymbol{\mu})^2 + \varphi(\mathbf{x}), \quad (13)$$

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

The intensity corresponding to the amplitude (13) is proportional to the probability for a photon to be detected at $\boldsymbol{\mu}$:

$$I(\boldsymbol{\mu}; \nu) \propto |\Psi(\boldsymbol{\mu}; \nu)|^2. \quad (14)$$

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

Except in special cases, the Fresnel-Kirchhoff integral (13) is not possible to evaluate analytically. At large ν (and with the dimensionless form of the lens $\varphi(\mathbf{x})$ held fixed) and in the geometric optics limit, one can easily model the intensity, as we shall explain. However, the most interesting regime for astrophysical plasma lenses occurs in the intermediate regime, where focusing and caustic catastrophes generate bright features whose peak intensity is controlled by diffraction. In this regime, there are characteristic patterns in the intensity, controlled by the topological character of the lens. In this intermediate- ν regime, conventional integration techniques typically fail, and it is hard to capture the complex, oscillatory interference pattern numerically. For example, G. Grillo and J. M. Cordes [25] implemented a procedure based on Fourier methods but found this technique to generate numerical artifacts. Here, motivated by our earlier work on Picard-Lefschetz theory, we instead employ analytic continuation and Cauchy's theorem to unambiguously define and to evaluate the relevant oscillatory integrals. We have developed a custom numerical scheme (made available [online](#)³.) which is fast and accurate, and applicable to a generic one dimensional oscillatory integral. A two dimensional version will be made available shortly. A nice feature of our method is that it typically becomes more efficient, *i.e.*, its convergence is improved, as the integrand becomes more oscillatory and difficult to handle via conventional techniques.

A. The geometric optics limit

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

³ See https://github.com/jfeldbrugge/Picard_Lefschetz_Integrator.

useful to begin by studying the geometric optics limit.

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

$$\nabla_{\mathbf{x}}\phi(\mathbf{x}; \boldsymbol{\mu}) = 0, \quad (15)$$

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

$$\nabla_{\mathbf{x}}\phi(\mathbf{x}; \boldsymbol{\mu})|_{\boldsymbol{\mu}=\boldsymbol{\xi}(\mathbf{x})} = 0. \quad (16)$$

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

$$\nabla_{\mathbf{x}}\phi(\mathbf{x}; \boldsymbol{\mu}) = 2(\mathbf{x} - \boldsymbol{\mu}) + \nabla\varphi(\mathbf{x}) \quad \implies \quad \boldsymbol{\mu} = \boldsymbol{\xi}(\mathbf{x}) = \mathbf{x} + \frac{1}{2}\nabla\varphi(\mathbf{x}). \quad (17)$$

The Lagrangian map $\boldsymbol{\xi}$ determines the optical rays, giving a purely geometric description of the lens. Every point \mathbf{x} is mapped to a point $\boldsymbol{\mu}$ in the space of observational parameters. In general, a point $\boldsymbol{\mu} \in M$ might be obtained from several points in $\mathbf{x} \in X$, *i.e.*, the Lagrangian map can be many-to-one. The regions in $\boldsymbol{\mu}$ where each point is obtained from n points in X are known as n -image regions. In multi-image regions, one adds the intensities due to each of the contributing paths: performing the relevant Gaussian integrals one finds for the normalized intensity

$$I(\boldsymbol{\mu}; \infty) = \sum_{\mathbf{x} \in \boldsymbol{\xi}^{-1}(\boldsymbol{\mu})} \frac{2}{|\lambda_1(\mathbf{x})| \dots |\lambda_N(\mathbf{x})|}, \quad (18)$$

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




Figure 2: The Lagrangian map in geometrical optics. The image consists of two single and one triple image regions separated by a fold caustic at which the normalized intensity spikes.

mation tensor,

$$\mathcal{M}_{ij}(\mathbf{x}) = \frac{\partial^2 \phi(\mathbf{x}; \boldsymbol{\mu})}{\partial x_i \partial x_j} = 2 \frac{\partial \xi_i(\mathbf{x})}{\partial x_j}, \quad (19)$$

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

At the boundaries between regions with a different number of images, at least one of the eigenvalue fields λ_i must vanish. At infinite ν this leads to an infinite spike in the normalized intensity map, signalling a caustic. See Fig. 2 for an illustration of the Lagrangian map corresponding to a one-dimensional lens with a single- and a triple-image region. The triple-image region is separated from the single-image regions by two fold caustics. At the fold caustic the normalized intensity profile diverges.

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

The geometric optics limit is attained in two stages: at short wavelengths, each real stationary point corresponds to a distinct image. As the wavelength is increased, each image itself forms an interference pattern, as illustrated in the Young's double slit experiment examined in Appendix B. The limit of short wavelengths, in which phase coherence is maintained, is often called the eikonal approximation. However, when phase coherence is lost - for example, when the source size becomes larger than the spacing of its fringes, interference effects disappear altogether. This assumption of loss of coherence is implicit in the geometric optics limit. However, objects smaller than the Fresnel scale are still seen to scintillate, as a result of coherent interference effects on unresolved scales. This is reflected in the expression

stars twinkle, planets don't. Interstellar scintillation typically occurs for sources smaller than about a micro arcsecond, corresponding to the Fresnel angle $\theta_F = \sqrt{\lambda/d}$ (with λ the wavelength and d the distance from the lens) on the sky. This condition is true for most FRBs and pulsars. Interplanetary scintillation due to the solar wind is commonly seen for many compact extragalactic radio sources at low frequencies [45]. In this case, the characteristic Fresnel angle for wavelengths of a few meters and distances of an astronomical unit is a fraction of an arcsecond. Ionospheric scintillation is strongest at the lowest frequencies, and is commonly seen at solar maximum or at equatorial locations near sunrise or sunset [46], and causes loss of lock in GPS. The Fresnel angular scale for a screen at a distance of 200km at wavelengths of a meter is 8 arc minutes, causing all celestial sources except the sun and the moon to scintillate.

B. Catastrophe theory

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

In the catastrophe theory of higher-dimensional functions, degenerate critical points are included because they are stable. René Thom (1972) proved [49] that the stable critical points with co-dimension⁴ K less than or equal to 4 are classified by the seven “elementary catastrophes.” These seven singularities suffice to classify the full range of caustics emerging in three-dimensional lenses. Thom named the seven catastrophes: *the fold, cusp, swallowtail, butterfly, and the elliptic, hyperbolic and parabolic umbilic.* The caustics were in the subse-

⁴ The co-dimension of a caustic is roughly the dimensionality of the singularity. The stable critical points of a n -dimensional function are completely classified by the caustics with co-dimension smaller or equal to n .

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

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

quent years connected and labeled by the Coxeter reflection groups (Arnol'd [50, 51]). The theory was subsequently applied to optical interference patterns by Berry and collaborators, and beautiful experiments were performed [6]. For a more recent theoretical investigation of catastrophe theory and caustics with applications to large-scale structure formation see [52]. Here we briefly review catastrophe theory and its application to oscillatory integrals.

Table I lists the seven “elementary catastrophes” and their unfoldings $\phi(\mathbf{x}; \boldsymbol{\mu})$. The unfolding $\phi(\mathbf{x}; \boldsymbol{\mu})$ evaluated at $\boldsymbol{\mu} = \mathbf{0}$ is the normal form of the catastrophe, representing the archetypical form of the critical point near $\mathbf{x} = \mathbf{0}$. We observe that the fold and the cusp respectively correspond to a cubic and quartic critical point of x . The parameter $\boldsymbol{\mu}$ represents the ways in which the caustic can decompose into lower-order caustics. In the case of the fold, we see that a linear perturbation decomposes the fold into a minimum and a maximum for $\mu < 0$ and no critical point at all for $\mu > 0$. The seven catastrophes belong to two families, classified by their co-rank⁵. The A -family is of co-rank $N = 1$, while the D -family is of co-rank $N = 2$. Critical points with higher co-rank have a co-dimension higher than 4, and for this reason are not included here. The co-rank N and the co-dimension K characterize the critical point. It generally takes N variables to describe the critical point, and it takes K parameters to describe its unfolding. In more prosaic terms, N is the dimension of the space of \mathbf{x} 's and K is the dimension of the space of $\boldsymbol{\mu}$'s.

⁵ The co-rank is the number vanishing eigenvalues of the Hessian matrix.

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

Table II: The intensity and fringe separation scaling relations for the catastrophes shown shown in Table I. At large ν the maximum intensity (14) is given by $I_0\nu^{2\beta}$ (see the discussion following Eq. (21)) and the fringe scaling exponents are defined in (22).

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

$$\Psi(\boldsymbol{\mu}; \nu) = \left(\frac{\nu}{\pi}\right)^{N/2} \int e^{i\phi(\boldsymbol{x}; \boldsymbol{\mu})\nu} d\boldsymbol{x}, \quad (20)$$

forms a caustic at the critical point $\boldsymbol{\mu} = 0$. For a detailed analysis including illustrations of the intensities obtained in each case, see chapter 36 of [53]. As ν is increased to large values, the normalized intensity $I(\boldsymbol{\mu}; \nu) = |\Psi(\boldsymbol{\mu}; \nu)|^2$ diverges and the scale of the associated diffraction fringes shrinks to zero according to scaling laws which are specific for each catastrophe. At large ν , the maximum of the intensity is attained near $\boldsymbol{\mu} = 0$ as illustrated, for example, by the fold singularity shown in Fig. 13. The maximum intensity obeys the following scaling law at large ν :

$$I(\mathbf{0}, \nu) = I_0\nu^{2\beta}. \quad (21)$$

The constant β , termed the *singularity index* by Arnold (Arnold [51] and Varchenko [54]), is universal, being invariant under diffeomorphisms and depending only on the topological class of the catastrophe. It is given, for each case, in the fourth column of Table II. The scaling with ν is easily seen by examining the corresponding normal form. Setting the unfolding parameter to zero, *i.e.*, $\boldsymbol{\mu} = \mathbf{0}$, in the phases listed in Table I, one can render the phase of the integrand independent of ν by rescaling the integration variables \boldsymbol{x} . For example, for A_2 we set $x = \nu^{-\frac{1}{3}}y$. Taking into account the ν -dependence arising from the Jacobian in the integration measure as well from the prefactor in (20), one infers that the amplitude at the caustic scales as $\nu^{\frac{1}{6}}$ for A_2 and hence that $\beta = \frac{1}{6}$ for this case. For the two dimensional

lenses, one has to rescale both x_1 and x_2 in order to remove ν from the exponent but the argument is otherwise the same.

For each of the normal forms of the phase listed in Table I, one may also analytically compute the constant I_0 , and we provide its numerical value in Table II. When considering a class of lens models for modeling purposes (such as the localised models we consider later), it may be helpful to notice an additional scaling property. At large ν , the amplitude is determined by the form of the phase near the critical point. Indeed, this is how universality arises. For any lens model which includes a given catastrophe, the leading terms in the Taylor expansion of the phase about the associated critical point will, after coordinate redefinitions, take the form of one of the “normal forms” listed in Table I. Coordinate rescalings are one of the simplest such transformations, which have a simple effect on any model and on its Taylor expansion about any critical point. One can easily derive the scaling behavior of the intensity under such transformations of the lens model. For an A_n catastrophe, for example, we may consider a set of lens models whose phase $\phi(x, 0) \sim \frac{a}{n+1}x^{n+1}$ near $x = 0$, with a a constant. By rescaling x we can remove the a dependence from the phase and hence infer that the intensity scales as $a^{-\frac{2}{n+1}}$. Physically, decreasing a means increasing the size of the lens, so it makes sense that the corresponding intensity grows. For a type D_n catastrophes we may likewise have a phase in the form $ax_1^{n+1} \pm bx_1x_2^2$. Again, the a and b dependence can be removed from the phase by rescaling x_1 and x_2 . Hence one can infer that the intensity scales as $a^{-\frac{1}{n-1}}b^{-1}$. (As an aside, note that for D_4^+ the normal form used here differs, from that given in Table I. The two forms may be shown to be equivalent under a linear coordinate transformation $x'_i = A_{ij}x_j$, $i, j = 1, 2$, and hence are in the same equivalence class. In the rest of this paper we always use the normal forms listed in Table I.)

The pattern of fringes may likewise be shown to scale as

$$\Psi(\boldsymbol{\mu}; \nu) = \nu^\beta \Psi((\nu^{\sigma_1}\mu_1, \dots, \nu^{\sigma_K}\mu_K), \nu), \quad (22)$$

where the fringe exponents σ_i , defined by Berry [4], are also listed in Table II. The sum of the fringe exponents, $\gamma = \sum_{i=1}^K \sigma_i$, represents the scaling exponent for the K -dimensional hypervolume of the diffraction pattern known as the *fringe index*. All of these exponents are invariant under diffeomorphisms, making them topological features.

While these catastrophes provide an exhaustive list, the precise forms in Table I are




Figure 3: The unfolding diagram of the seven elementary catastrophes.

unlikely to arise, since apart from A_1^\pm , they are completely de-localized, with the strength of the lens diverging away from the critical point. Later in this paper we shall consider more realistic, localized lenses which generate catastrophes within them of the listed form. In the vicinity of such a catastrophe, one can expect the behaviors indicated in Table II to hold. Note, however, that to compute the maximum intensity at such a catastrophe, one must first redefine the coordinate \mathbf{x} to put the exponent locally into the normal form of the catastrophe listed in Table I (this is guaranteed to be possible by the theorems mentioned above), and take into account the associated Jacobian factor when evaluating the integral.

The seven elementary catastrophes listed above form an intricate hierarchy which unfold under perturbations according to the unfolding diagram (see Fig. 3). As we saw before, the fold caustic (A_2) splits under a small perturbation into a maximum and a minimum each corresponding to an A_1 . Analogously, butterfly caustic (A_5) unfolds into a swallowtail caustic (A_4), which in its turn unfolds into the cusp (A_3) and the fold caustic (A_2). The parabolic umbilic caustic (D_5) has a more intricate structure as it can unfold into both the elliptic (D_4^-), hyperbolic umbilic (D_4^+) and the swallowtail caustic (A_4). The elliptic (D_4^-) and hyperbolic umbilic caustic (D_4^+) always unfold into cusp caustics (A_3).

IV. EXAMPLE: A ONE-DIMENSIONAL LENS

Picard-Lefschetz theory, as described for example in [27], enables us to deform the real integration domain \mathbb{R}^N of the Fresnel-Kirchhoff integral (13) into a set of Lefschetz thimbles, *i.e.*, steepest descent contours, $\mathcal{J}_i \subset \mathbb{C}^N$ each corresponding to a relevant saddle point \bar{x}_i ,

$$\Psi(\boldsymbol{\mu}; \nu) = \left(\frac{\nu}{\pi}\right)^{N/2} \sum_i \int_{\mathcal{J}_i} e^{i\phi(\mathbf{x}; \boldsymbol{\mu})\nu} d\mathbf{x}. \quad (23)$$

The exponent $i\phi(\mathbf{x}; \boldsymbol{\mu})\nu$ evaluated along a steepest descent contour \mathcal{J}_i has a constant imaginary part while its real part $h = \text{Re}[i\phi(\mathbf{x}; \boldsymbol{\mu})\nu]$ is monotonically decreasing. As a consequence, the conditionally convergent oscillatory integral is transformed into a sum of convex

integrals. This is remarkable, as the originally conditionally convergent integral is generally sensitive to regularization and also, if the integral is performed iteratively, to the order in which the partial integrals are taken (see Appendix A for an instructive example). The integrals over Lefschetz thimbles have none of these ambiguities. It is for this reason that we will interpret the integral over the sum of Lefschetz thimbles as the *definition* of the integral taken over the real integration domain. Once we have identified the correct set of thimbles, we can use conventional numerical methods to evaluate the integral on each thimble.

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

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




Figure 4: A comparison of the Gaussian lens (red) and the rational approximation (blue).

A. Geometric optics approximation

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

$$\varphi(x) = \frac{\alpha}{1 + x^2}, \quad (24)$$

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

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

Picard-Lefschetz theory, however, applies to analytic lenses in general. The lens (24) is a rational approximation to the Gaussian lens

$$\varphi(x) = \alpha e^{-x^2}, \quad (25)$$

which has an essential singularity at infinity on the Riemann sphere and an infinite number of saddle points in the complex plane. See Fig. 4 for a comparison between the two lenses. It

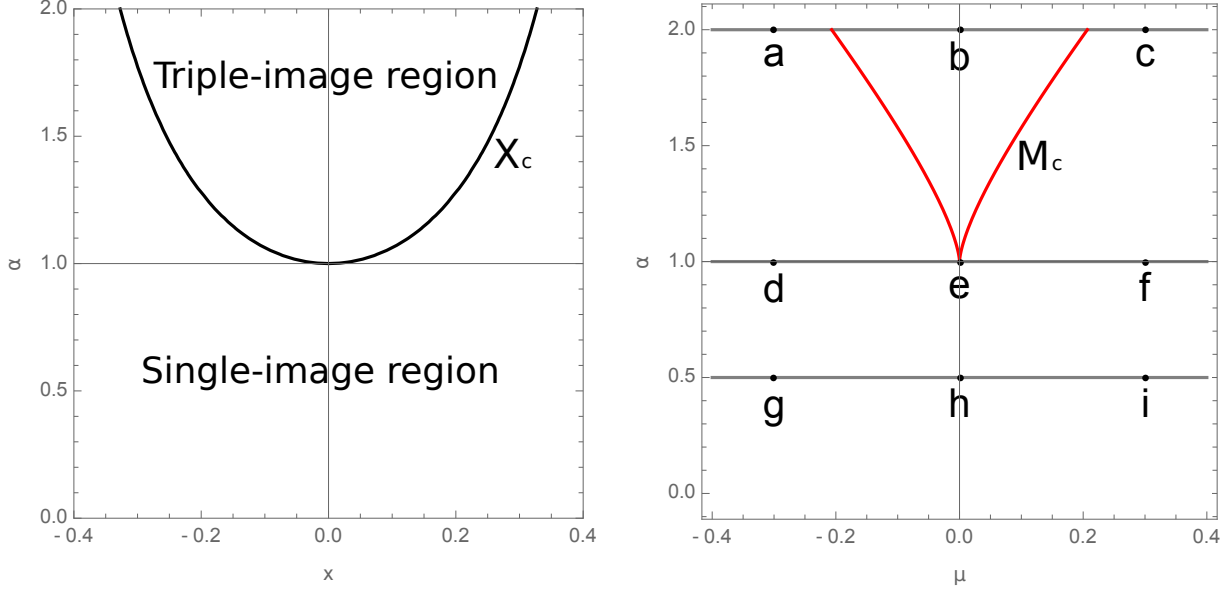
(a) Caustic in the x - α plane.(b) Caustic in the μ - α plane.

Figure 5: The caustics of the one-dimensional rational lens in the x - α and the μ - α plane. The points in the *left panel* correspond to the Picard-Lefschetz diagrams in Fig. 7. The lines correspond to the panels in Fig. 11.

is a wonderful fact that many real-valued functions with intricate structure in the complex plane, can be well-approximated with a Padé approximation, whose analytic continuation possesses only a finite number of poles.

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

$$\xi(x) = x - \frac{\alpha x}{(1+x^2)^2}, \quad (26)$$

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

$$\frac{\partial^2 \phi(x)}{\partial x^2} = 2 \frac{\partial \xi(x)}{\partial x} = 2 + 2\alpha \frac{3x^2 - 1}{(1+x^2)^3} = 0. \quad (27)$$

See Fig. 5 for the caustic surface in the x - α and the μ - α planes. For $\alpha < 1$ no such real root exists. The lensed image consists of a single-image region. For $\alpha = 1$ there is a single real-valued root at $x_c = 0$ with multiplicity two and the corresponding point $\mu_c = \xi(x_c) = 0$ in the parameter space M . In the μ - α plane, this point is non-differentiable on the caustic set. This is an example of a cusp caustic. For $\alpha > 1$ there are two symmetric real roots.




Figure 6: The normalized intensity in geometric optics, $I(\mu; \infty)$ is plotted as a function of μ .

These are examples of fold caustics. For further reference, for $\alpha = 2$, the two caustics are located at

$$X_c = \{-0.327334\dots, 0.327334\dots\}, \quad (28)$$

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

$$M_c = \xi(X_c) = \{+0.206751\dots, -0.206751\dots\}. \quad (29)$$

The relative normalized intensity of the lens in the geometric optics limit (see Section III A) is plotted in figure 6. For $\alpha = 1/2$, the lens does not form a caustic. The normalized intensity map is finite. For $\alpha = 1$, we see a cusp caustic at $\mu = 0$. For $\alpha = 2$, we observe two fold caustics at M_c enclosing a triple-image region. The black curves in the triple-image region are the three contributions corresponding to the three images. The red curve is the sum over the multi-image regions.

B. Finding the thimbles

We now turn to evaluating the full expression (23). First, we need to determine which Lefschetz thimbles contribute. We shall describe two distinct methods, the second of which is more efficient for numerical purposes.

1. *Method 1: following steepest ascent contours*

The exponent

$$i\phi(x; \mu)\nu = i \left[(x - \mu)^2 + \frac{\alpha}{1 + x^2} \right] \nu \quad (30)$$

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

$$\frac{\partial \phi(x; \mu)}{\partial x} = 2(x - \mu) - \frac{2\alpha x}{(1 + x^2)^2} = 0. \quad (31)$$

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

$$i\phi(\mathbf{x}; \boldsymbol{\mu})\nu = h(\mathbf{u} + i\mathbf{v}; \mu) + iH(\mathbf{u} + i\mathbf{v}; \mu), \quad (32)$$

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

$$\frac{\partial \gamma_\lambda(\mathbf{z})}{\partial \lambda} = -\nabla_{\mathbf{u} + i\mathbf{v}} h[\gamma_\lambda(\mathbf{z})] \quad (33)$$

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

$$\nabla_{\mathbf{u} + i\mathbf{v}} h = \nabla_{\mathbf{u}} h + i\nabla_{\mathbf{v}} h. \quad (34)$$

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

Depending on μ and α either one or three of the saddle points are real-valued. The lens thus has both single- and triple-image regions. See Fig. 7 for the five saddle points \bar{x}_i and the corresponding steepest ascent and descent contours. By intersecting the steepest ascent contours with the real line, we obtain the Lefschetz thimble (plotted in blue). The thimbles run from $x = -\infty$ to $x = +\infty$, while passing through the poles at $x = \pm i$.

From the caustic structure in Fig. 5 we can distinguish three regimes:

- In the regime $\alpha < 1$, the lens forms a single image. The corresponding Picard-Lefschetz analysis yields a single real-valued saddle point. For large $|\mu|$ there is, in addition, a relevant complex saddle point. When $|\mu|$ decreases to 0, the complex saddle point becomes irrelevant due to a Stokes transition. This phenomenon is discussed in detail in the next section. For $\mu = 0$, only the real saddle point is relevant. Note that the thimble can for all μ be deformed to the original integration domain \mathbb{R} . See the lower panels of Fig. 5.
- For $\alpha = 1$, the lens contains a cusp caustic at $\mu_c = 0$. For $\mu \neq \mu_c$, the Picard-Lefschetz analysis is similar to the $\alpha < 1$ regime. The thimble passes through one real-valued and one complex-valued saddle point. At the caustic $\mu = \mu_c$, three non-degenerate saddle points merge forming a degenerate saddle point. This is the signature of the cusp caustic, whose normal form is the quartic function x^4 . See the middle panels of Fig. 5.
- In the regime $\alpha > 1$, the Picard-Lefschetz analysis splits into three intervals. In the single-image region, *i.e.*, μ in $(-\infty, -\mu_c)$ or (μ_c, ∞) , the Picard-Lefschetz analysis consists of two relevant saddle points; one real and one complex. At the caustic, the complex saddle point approaches the real line and merges with its complex conjugate saddle point. This is the signature of the fold caustic. In the triple-image region, *i.e.*, $\mu \in (-\mu_c, \mu_c)$, the analysis consists of three real-valued relevant saddle points. See the upper panels of Fig. 5.

2. Method 2: flowing the integration domain

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

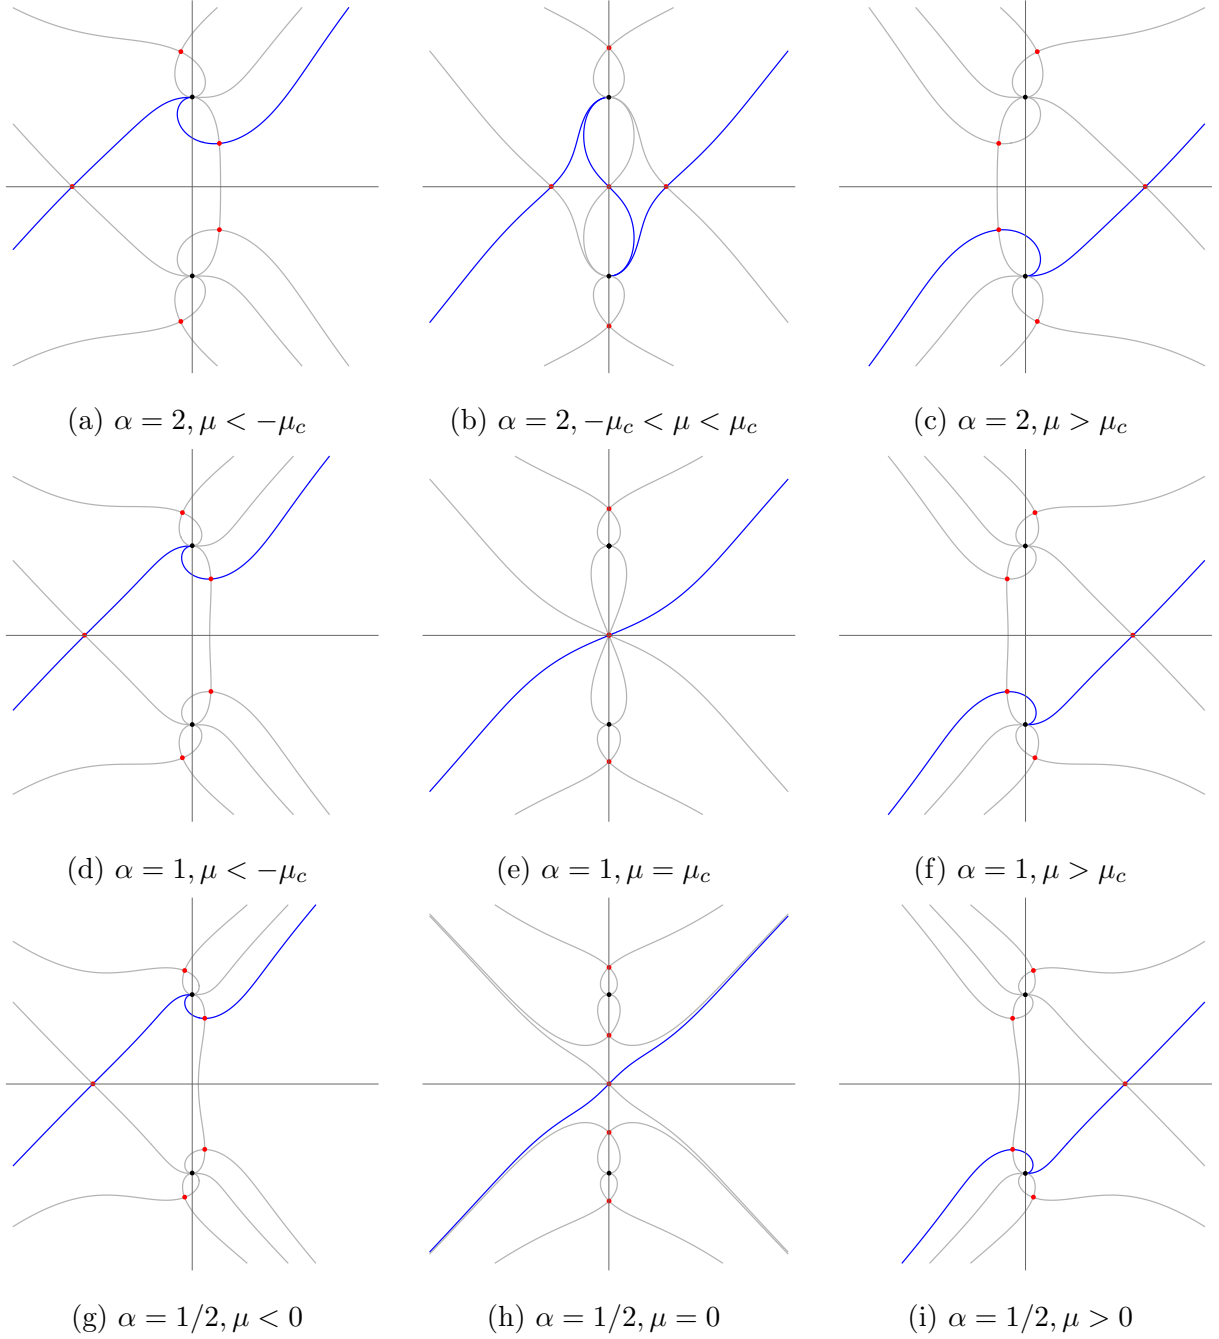


Figure 7: The Picard-Lefschetz thimbles for $\alpha = 1/2, 1, 2$ as a function of μ . The red and black points are the saddle points and poles. The curves are paths of steepest descent and ascent. The blue ones are relevant, while the grey ones are irrelevant.

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

$$X_\lambda = \gamma_\lambda(X) \subset \mathbb{C}. \quad (35)$$

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

$$\gamma_\lambda(\mathcal{J}_i) = \mathcal{J}_i \quad (36)$$

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

$$\lim_{\lambda \rightarrow \infty} X_\lambda = \mathcal{J}. \quad (37)$$

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

1. If for some μ , a few non-degenerate saddle points merge to form a higher order saddle point, the number of relevant critical points will change. At these points, the integral $\Psi(\mu; \nu)$ forms a caustic. This phenomenon can be described by catastrophe theory (see sections III B and V).
2. When the imaginary part H evaluated in two saddle points coincides for some parameter μ_s , the two corresponding steepest-descent contours can coincide. At such a parameter μ_s , the Lefschetz thimbles flip changing the number of relevant saddle points (see Fig. 8 for an illustration). This is known as a Stokes transition. The parameters μ_s for which this happens form so-called Stokes lines.

We study both phenomena in detail in the Section V.

We numerically evaluate the flow X_λ by approximating X by a set of line-segments and flowing the endpoints. Since the real part of the analytic continuation of an analytic function does not have local extrema (this follows from the Cauchy-Riemann equation), all points $z \in \mathbb{C}$ flow to poles as $\lambda \rightarrow \infty$. The limit $\lim_{\lambda \rightarrow \infty} X_\lambda$ should not be interpreted as a pointwise limit. We, for this reason, trace the length of the line-segments and add points




Figure 8: The Stokes phenomenon at which the steepest descent contours coincide and the relevant saddle points change. In the *left* panel, we have one real relevant saddle point. In the *central* panel, we see a Stokes phenomenon. The Lefschetz thimble passes to both a real and a complex saddle point. In the *right* pane, we have one real and one complex relevant saddle point.

The Stokes phenomenon occurs when the steepest descent contour of a saddle point connects with another saddle point.

when neighboring points move too far apart. We moreover remove line-segments in the neighborhoods of the poles of the h -function. The contour X_λ has converged to the thimble when the imaginary part H is approximately constant along the line-segments.

This idea is implemented by the algorithm:

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

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

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

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

See Fig. 9 for the flow of the original integration domain corresponding the rational lens for $\alpha = 2$ and $\mu = 0$. For $\lambda = 0$ the contour X_λ coincides with the real line. As λ is increased to 1, the original integration domain smoothly flows to the Lefschetz thimble \mathcal{J} consisting




Figure 9: The downward flow of the integration domain. The contour X_λ for $\lambda = 0, 0.2, 0.4, 0.6, 0.8, 1$ is plotted by the black, grey and blue curves. The five saddle points are plotted in red and the poles are plotted in black.

of three steepest descent contours \mathcal{J}_i corresponding to three relevant saddle points \bar{x}_i . By evaluating the flow for varying α and μ , we obtain the Picard-Lefschetz analysis of the lens.

For multi-dimensional oscillatory integrals, the flow algorithm can be generalized by flowing the cells of a tessellation of the original integration domain. In this paper, we start our calculations with the tessellation of a rectilinear lattice. For an two-dimensional illustration see Fig. 10.

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

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

do

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

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

 remove the corresponding cells

end if

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

 subdivide the cell into smaller cells

end if

end while

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

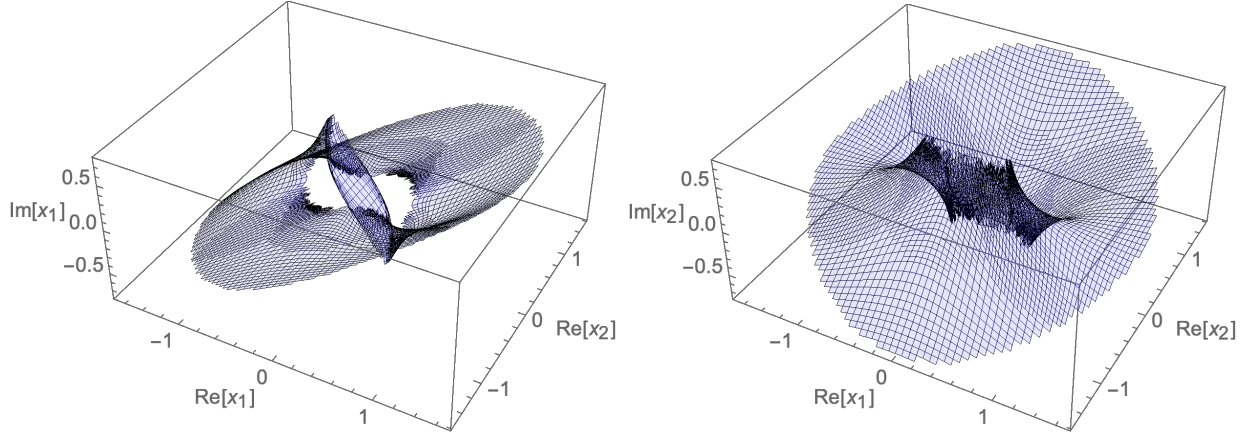


Figure 10: Two projections of the numerically obtained two-dimensional thimble \mathcal{J} in \mathbb{C}^2 for a two-dimensional oscillatory integral.

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

C. Integrating along the thimbles

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

$$\Psi(\mu; \nu) \approx \sum_i \frac{e^{i\phi(p_{i,1}; \mu)\nu} + e^{i\phi(p_{i,2}; \mu)\nu}}{2} (p_{i,2} - p_{i,1}) \quad (38)$$

summed over the line segments. For multi-dimensional oscillatory integrals, we evaluate the integral on a linear approximation of the integrand on the tessellation of the thimble. Naively, one might expect to have to compute the Lefschetz contour \mathcal{J} for every μ for which one wishes to perform the integral. However, since the thimble is a smooth function of μ , it suffices to compute the thimble for a range of μ . When integrating, we instead evaluate the integral on the thimble corresponding to the closest μ for which we have evaluated the

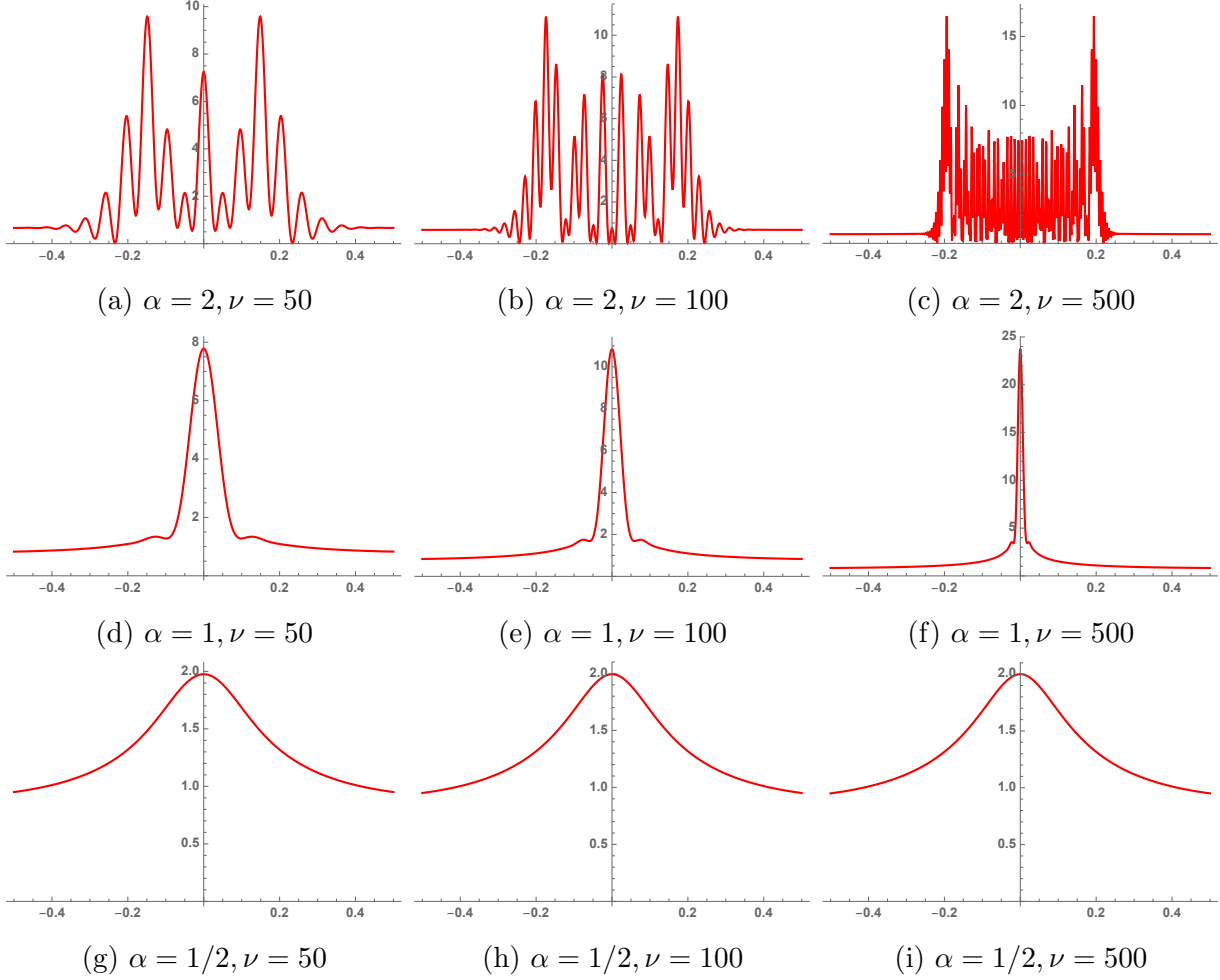


Figure 11: The normalized intensity $I(\mu; \nu) = |\Psi(\mu; \nu)|^2$ for $\alpha = 1/2, 1, 2$ as a function of μ for $\nu = 50, 100, 500$.

thimble. Finally, it should be noted that for increasing ν , the support of the integral is increasingly concentrated around the relevant saddle points. As a consequence we can, for large ν , restrict the integral to the line segments close to the saddle points. It follows from this that the numerical evaluation of the integral along the thimble becomes more and more efficient as the frequency is increased. This is in sharp contrast with conventional integration techniques which need to trace many oscillations of the integrand along the real line.

See Fig. 11 for the normalized intensity profiles of the lens evaluated along the thimble for frequencies $\nu = 50, 100, 500$. We observe the following properties of the normalized intensity profiles:

- In the regime $\alpha < 1$, the lens leads to a single-image region. The normalized intensity profile does not oscillate and is moreover independent of the frequency ν . See the

lower panels of Fig. 11.

- For $\alpha = 1$, the lens forms a cusp caustic. The caustic corresponds to the peak at $\mu_c = 0$. For increasing frequency, ν , the peak is enhanced and becomes increasingly narrow. In the eikonal limit $\nu \rightarrow \infty$, the normalized normalized intensity diverges as $\nu^{1/2}$ at the caustic μ_c (see the scaling relations in Table II). See the middle panels of Fig. 11.
- In the regime $\alpha > 1$, the lens forms a triple-image region which is bounded by two-fold caustics. We see that the triple-image region $(-\mu_c, \mu_c)$, with $\mu_c = 0.206751\dots$ for $\alpha = 2$, consists of an interference pattern bounded by two peaks at $\mu = \pm\mu_c$. The interference pattern in the triple-image region is the result of the three real saddle points. The oscillations in the single-image region result from the interplay between the relevant real and the complex saddle point. For increasing ν , the fringes of the interference pattern shrink and spikes corresponding to the fold get sharper and are increasingly enhanced. For the relevant scalings see Table II. See the upper panels of Fig. 11.

Note that the normalized intensity in the cusp exceeds the normalized intensity in the fold caustic. This related to the co-dimension of the caustic as described in Section III B. Moreover remark that the cusp caustic only exists at a single α for the one-dimensional lens, while the fold caustic appears for a range of α . Table II shows the frequency dependence of the pattern. Furthermore, note that the normalized intensity profiles at frequency $\nu = 500$, for $\alpha = 1/2, 1$, and 2 , are close to the normalized intensity maps predicted by geometric optics (see Fig. 6).

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

V. THE ELEMENTARY CATASTROPHES

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

A. The fold A_2

The fold singularity is the simplest degenerate critical point and can be viewed as the superposition of two non-degenerate saddle points. The Picard-Lefschetz analysis of the unfolding of the fold singularity is illustrated in Fig. 12. For negative μ , there are two relevant real saddle points (see Fig. 12a). As μ approaches the caustic at $\mu_c = 0$, the two saddle points merge and form the fold singularity (see Fig. 12b). Note that the fold saddle point emanates three steepest ascent and three descent curves. The thimble is non-differentiable at the degenerate saddle point. When μ is increased passed the caustic μ_c , the two saddle points move off the real axis and into the complex plane (see Fig. 12c). In this regime only one of them remains relevant.




Figure 12: The saddle points and the Lefschetz thimbles in the complex plane \mathbb{C} of the unfolding of the fold singularity A_2 .




Figure 13: The integral $\Psi(\mu, \nu)$ for the fold singularity as a function of μ for $\nu = 50, 100, 500$.

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

$$\Psi(\mu; \nu) = \sqrt{\frac{\nu}{\pi}} \int_{-\infty}^{\infty} e^{i\left(\frac{x^3}{3} + \mu x\right)\nu} dx = 2\sqrt{\pi}\nu^{1/6} \text{Ai}[\nu^{2/3}\mu]. \quad (39)$$

Note the appearance of the singularity and fringe indices $1/6$ and $2/3$ as listed in Table II. It is straightforward to derive the scaling of the amplitude and the fringes, with the change of coordinates $z = \nu^{1/3}x$. The other scaling relations are derived analogously. The Airy function is a good illustration of the interference phenomenon present in multi-image regions (seen in Fig. 13). The range $\mu < 0$, for which the two relevant saddle points reside on the real line corresponds to a double-image region, where two saddle points lead to an interference pattern. The range $\mu > 0$ corresponds to a zero-image region in which the amplitude asymptotes to zero as $\nu \rightarrow \infty$.

1. Asymptotics

Using the Picard-Lefschetz diagrams (Fig. 12) we can derive asymptotics for the integral $\Psi(\mu; \nu)$. For $\mu > 0$ the Picard-Lefschetz analysis consists of a single relevant saddle point located at $i\sqrt{\mu}$. The exponent can be approximated around the saddle point by

$$\phi(x; \mu) = i\frac{2}{3}\mu^{3/2} + i\sqrt{\mu}(x - i\sqrt{\mu})^2 + \mathcal{O}\left((x - i\sqrt{\mu})^3\right). \quad (40)$$

The saddle point approximation for this point gives an exponential falloff

$$\Psi(\mu; \nu) \approx \frac{e^{-\frac{2}{3}\mu^{3/2}\nu}}{\sqrt{2}\mu^{1/4}}. \quad (41)$$

This matches the the behaviour in Fig. 13.

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

$$\Psi(\mu; \nu) \approx \frac{e^{-\frac{2i}{3}\mu^{3/2}\nu} + ie^{\frac{2i}{3}\mu^{3/2}\nu}}{\mu^{1/4}}, \quad (42)$$

seen in Fig. 13. Observe that wave function becomes increasingly oscillatory and falls off as a power law $\Psi(\mu; \nu) \propto \frac{1}{|\mu|^{1/4}}$ in the geometric limit $\nu \rightarrow \infty$.

B. The cusp A_3

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

$$\Psi(\boldsymbol{\mu}; \nu) = \sqrt{\frac{\nu}{\pi}} \int_{-\infty}^{\infty} e^{i\left(\frac{x^4}{4} + \mu_2 \frac{x^2}{2} + \mu_1 x\right)\nu} dx. \quad (43)$$

See Fig. 14 for an illustration of unfolding of the cusp caustic and the Picard-Lefschetz analysis in the (μ_1, μ_2) -plane.

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

$$x^3 + \mu_2 x + \mu_1 = 0. \quad (44)$$

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

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

$$\mu_2 = -\frac{3}{2^{2/3}}|\mu_1|^{2/3}. \quad (45)$$

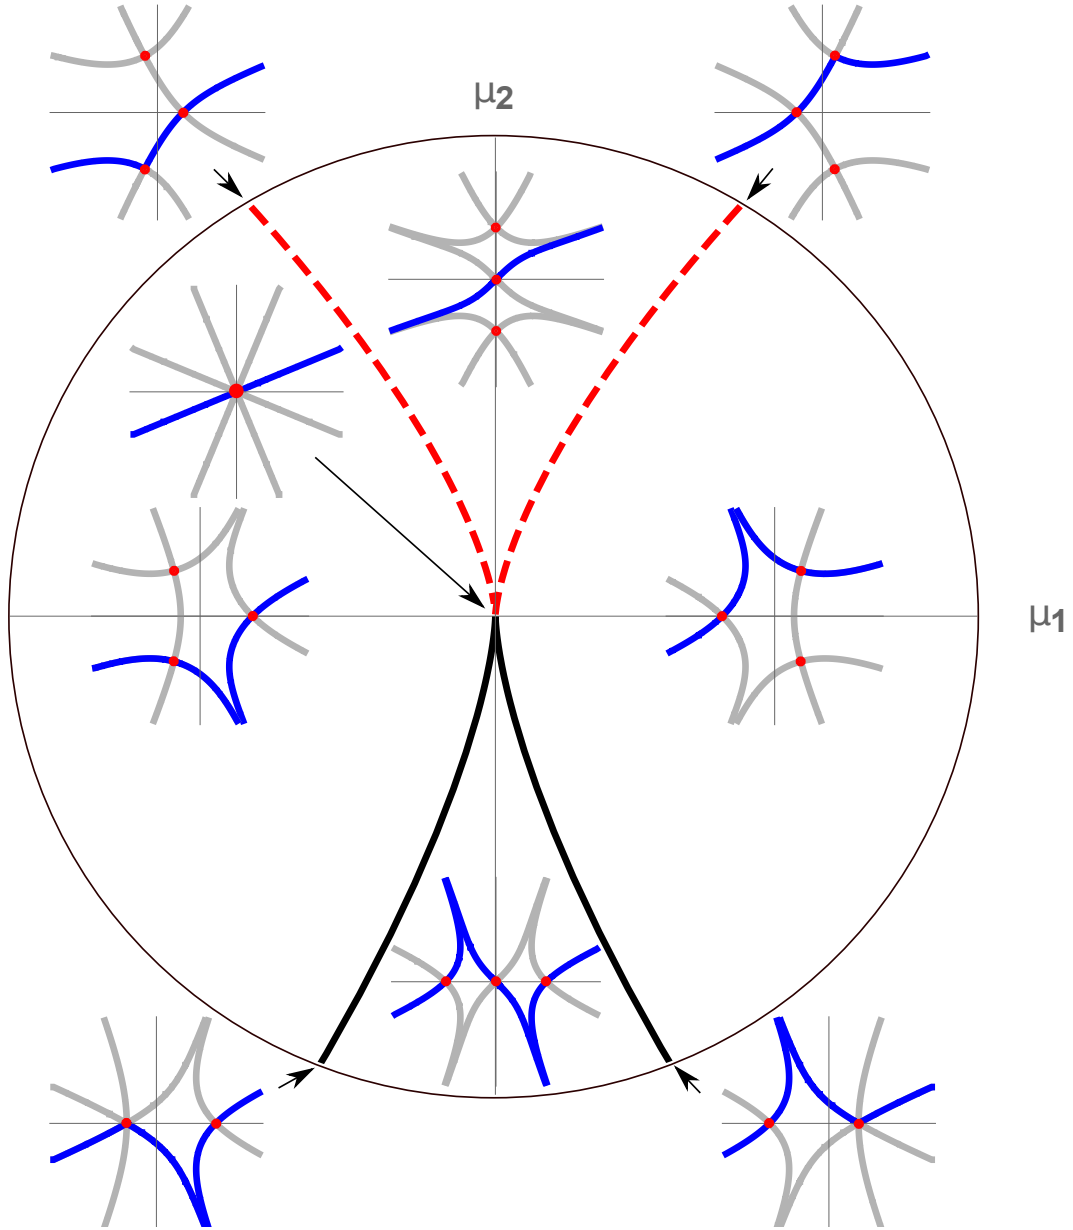


Figure 14: The cusp singularity in the unfolding μ -plane. The black curve is the fold line separating the single-image region (upper region) from the triple-image region (lower region). The dashed red line is the Stokes line. We observe that the Stokes transition corresponds to one complex saddle point becoming (ir)relevant, and that the caustics correspond to the superposition of real non-degenerate saddle points. The upper region is a single-image region since there is only one real relevant saddle point. The lower region is a triple-image region since there are three real relevant saddle points.

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

In the triple-image region enclosed by the two fold-lines, the thimble passes through three real-valued saddle points. When approaching one of the fold lines, we see that two

of the real saddle points merge and move in the complex plane. Only one of the two complex saddle points remains relevant to the integral. This is analogous to the behavior observed in the analysis of the fold caustic. At the cusp saddle point at $(\mu_1, \mu_2) = (0, 0)$ all three saddle points merge at the origin. Finally, note that the single-image region consists of three subregions, for which the Picard-Lefschetz analysis either consists of one or two relevant saddle points. These subregions are separated by two Stokes lines (red dashed lines in Fig. 14). Along these lines, the Lefschetz thimbles flip while the saddle points remain separated. The Stokes lines can be found by equating the imaginary parts of the exponents evaluated at the saddle points, *i.e.*,

$$\text{Im}[i\phi(\bar{x}_i; \boldsymbol{\mu})\nu] = \text{Im}[i\phi(\bar{x}_j; \boldsymbol{\mu})\nu] \quad (46)$$

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

$$\mu_2 = 3\sqrt[3]{\frac{3\sqrt{3}-5}{2}}|\mu_1|^{2/3}, \quad (47)$$

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

1. Numerics

Given the Lefschetz thimble, we can numerically evaluate the amplitude (see Fig. 15). In the eikonal limit $\nu \rightarrow \infty$ we observe the emergence of a fold-line (A_2) with a sharp exponential falloff in most of the single-image region and the power-law falloff in the triple-image region. We also see the emergence of a cusp caustic at the origin with a power-law falloff along the line $\{\mu_1 = 0\}$.




Figure 15: The normalized intensity $I(\boldsymbol{\mu}; \nu)$ for the cusp caustic for $\nu = 50, 100$ and 500 .

2. Asymptotics

The Picard-Lefschetz diagrams (Fig. 14) allow us to derive limiting behaviour for $\Psi(\boldsymbol{\mu}; \nu)$. From the functional form of $\phi(x; \boldsymbol{\mu})$ along the line $\{\mu_1 = 0\}$, *i.e.*,

$$\phi(x; \boldsymbol{\mu}) = \frac{\mu_2 x^2}{2} + \frac{x^4}{4}, \quad (48)$$

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

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

$$\Psi(x; \boldsymbol{\mu}) \approx \sqrt{\frac{\nu}{\pi}} \int_{-\infty}^{\infty} e^{i \frac{\mu_2 \nu}{2} x^2} dx = \sqrt{2} (-i \mu_2)^{-1/2} \quad (49)$$

which in the normalized intensity corresponds to the falloff

$$I(x; \boldsymbol{\mu}) = |\Psi(x; \boldsymbol{\mu})|^2 \approx \frac{2}{\mu_2}, \quad (50)$$

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

For $\mu_2 < 0$, the Picard-Lefschetz analysis consists of three real relevant saddle points

located at $\pm\sqrt{\mu_2}$ and 0. The exponent at the saddle point $\pm\sqrt{-\mu_2}$ can be approximated by

$$\phi(x; \boldsymbol{\mu}) = -\frac{\mu_2^2}{4} - \mu_2(x \pm \sqrt{-\mu_2})^2 + \mathcal{O}((x \pm \sqrt{-\mu_2})^3). \quad (51)$$

In the saddle point approximation,

$$\Psi(\boldsymbol{\mu}; \nu) \approx \sqrt{2} \frac{-(-1)^{3/4} + (1+i)e^{-\frac{i}{4}\mu_2^2\nu}}{\sqrt{-\mu_2}}. \quad (52)$$

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

$$I(\boldsymbol{\mu}; \nu) \propto \frac{2}{\mu_2}. \quad (53)$$

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

$$\Psi(\boldsymbol{\mu}; \nu) \approx \sqrt{2} \frac{(-1)^{1/4} e^{-i\mu_1^{4/3}\nu}}{-\sqrt{3}\mu_1^{1/3}}, \quad (54)$$

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

$$I(\boldsymbol{\mu}; \nu) = |\Psi(\boldsymbol{\mu}; \nu)|^2 \approx 2\mu_1^{-2/3}. \quad (55)$$

C. The swallowtail A_4

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

$$\Psi(\boldsymbol{\mu}; \nu) = \sqrt{\frac{\nu}{\pi}} \int_{-\infty}^{\infty} e^{i\left(\frac{x^5}{5} + \mu_3 \frac{x^3}{3} + \mu_2 \frac{x^2}{2} + \mu_1 x\right)\nu} dx. \quad (56)$$

See figures 17 and 18 for an illustration of unfolding of the swallowtail caustic and the Picard-Lefschetz analysis in the (μ_1, μ_2, μ_3) -space.

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

$$x^4 + \mu_3 x^2 + \mu_2 x + \mu_1 = 0. \quad (57)$$

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

Geometric optics applied to this integral shows that the swallowtail caustic at $\boldsymbol{\mu} = (0, 0, 0)$ emanates a cusp-line and a fold-surface (see Fig. 16). The fold-surface (the yellow surface in Fig. 16) is given by

$$A_2 = \{(3u^4 + u^2v, -4u^3 - 2uv, v) | (u, v) \in \mathbb{R}^2\} \subset M \quad (58)$$

satisfying the two constraints

$$\frac{d\phi(x; \boldsymbol{\mu})}{dx} = 0, \quad \frac{d^2\phi(x; \boldsymbol{\mu})}{dx^2} = 0. \quad (59)$$

The cusp-line (the black curve in Fig. 16) lays on the fold-surface and is given by

$$A_3 = \{(-3t^4, 8t^3, -6t^2) | t \in \mathbb{R}\} \subset M \quad (60)$$

satisfying the three constraints

$$\frac{d\phi(x; \boldsymbol{\mu})}{dx} = 0, \quad \frac{d^2\phi(x; \boldsymbol{\mu})}{dx^2} = 0, \quad \frac{d^3\phi(x; \boldsymbol{\mu})}{dx^3} = 0. \quad (61)$$

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

In figures 17 and 18 we plot three slices of the fold-surface and cusp-line for $\mu_3 = -1, 0$ and $+1$. For $\mu_3 = -1$ we obtain the characteristic swallowtail shape in the fold-surface with the cusp-line intersecting at the tips, which gives the singularity its name. For $\mu_3 = 0$ we see




Figure 16: The swallowtail singularity in the unfolding space (μ_1, μ_2, μ_3) . The yellow surface is the fold surface separating the single- double- and triple-image regions. The black line is the cusp line, along which we find the cusp saddle points.

the actual swallow caustic. The slice $\mu_3 = +1$, is simpler as it does not contain intersections with the cusp-line and only consists of the fold-surface separating two regions.

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

- We start by analysing the saddle points in the $\mu_3 = -1$ plane (Fig. 17). The Picard-Lefschetz analysis for the enclosed region in the middle of the circle consists of four relevant real saddle points. This is a quadruple-image region. Note that multiple-image regions for localized lenses always consist of an odd number of images. In such lenses, the swallowtail will in practice always appear near another caustic such as a fold.

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

Again, starting from the quadruple-image region and moving in the vertical direction, we observe that two of the outer saddle points merge to form a fold saddle point and




Figure 17: The Picard-Lefschetz analysis of the unfolding of the swallowtail (A_4) singularity at $\mu_3 = -1$. The number of real saddles gives the number of images in geometric optics.

subsequently move into the complex plane. The resulting Picard-Lefschetz analysis again consists of three relevant saddle points; two real and one complex. This again is a double-image region. If we, however, move from this double-image region to the double-image region on the left of the quadruple-image region, we pass through two Stokes lines, at which the complex saddle point switches from relevant to irrelevant.

The Stokes lines are defined by

$$\text{Im}[i\phi(\bar{x}_i; \boldsymbol{\mu})\nu] = \text{Im}[i\phi(\bar{x}_j; \boldsymbol{\mu})\nu] \quad (62)$$

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

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

- For $\mu_3 = 0$, the geometry of the fold-line is simpler as the quadruple-image region has merged into the swallowtail caustic at the origin (see Fig. 18a). The Picard-Lefschetz analysis of this slice is largely similar to the one at $\mu_3 = -1$. The double-image region (including the Stokes lines) has been deformed but is otherwise the same. The zero-image region is also unchanged. However, the intersection of the two fold-lines is replaced by the swallowtail saddle point at the origin of in the (μ_1, μ_2) -plane. Since this saddle point is the superposition of four non-degenerate saddle points, the amplitude integral is enhanced.
- For $\mu_3 = +1$, the geometry of the caustics is depicted in Fig. 18b. The fold-line separates the zero-image region on the right from the double-image region on the left. Since the Picard-Lefschetz diagram in the zero-image region consists of four complex saddle points – two of them being relevant – there exist two distinct ways in which we can pass to the double-image region; either by merging the two saddle points on the left or on the right (see upper and lower diagram). The transition between these takes place at the origin, where the four saddle points are located on the imaginary

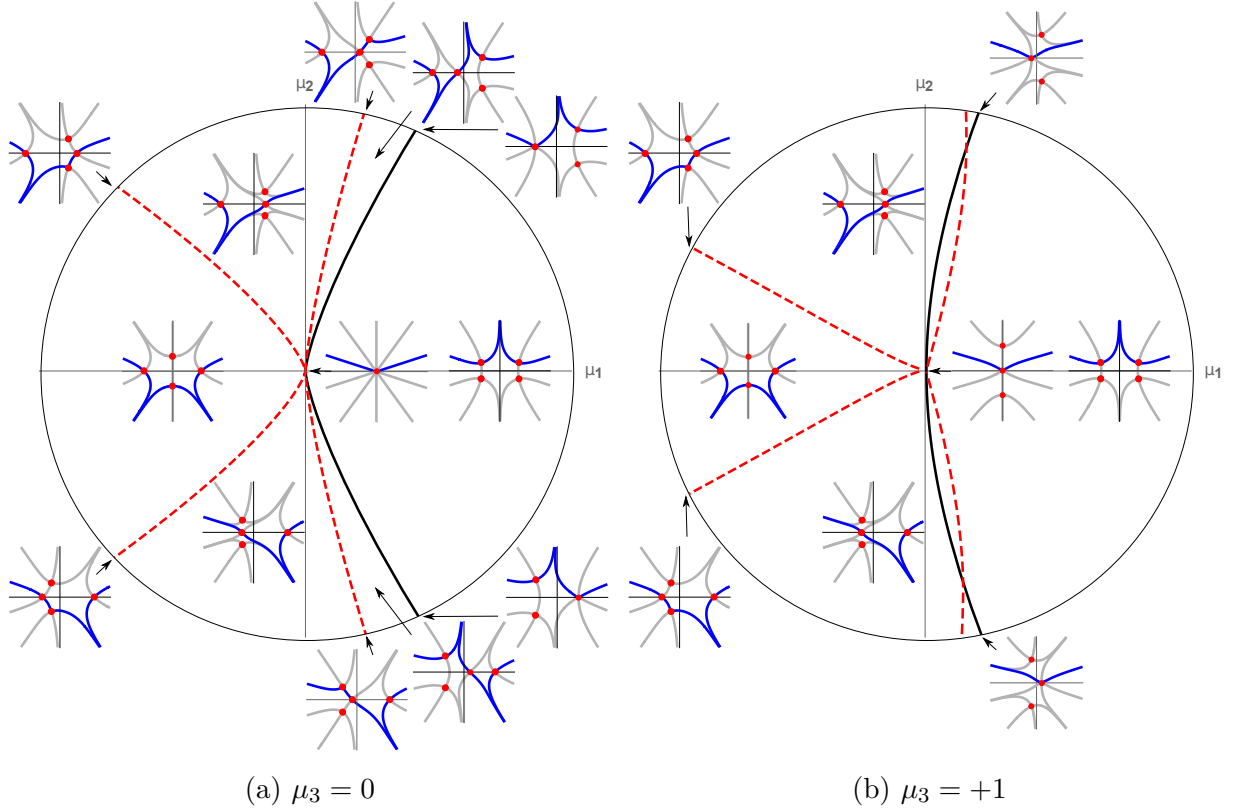


Figure 18: The Picard-Lefschetz analysis of the unfolding of the swallowtail (A_4) singularity at $\mu_3 = 0$ and $+1$. The number of real saddles gives the number of images in the geometric optics approximation.

axis. The double-image region consists of three subregions. The rightmost Stokes lines at $\mu_3 = 0$ (see Fig. 18a) have partly moved into the zero-image region.

By patching the Picard-Lefschetz analysis at $\mu_3 = -1, 0$ and $+1$ together, we obtain a complete description of the unfolding of the swallowtail singularity in the (μ_1, μ_2, μ_3) -space. Note that the Stokes lines obtained in figures 17 and 18 are intersections of Stokes-surfaces, which together with the fold-surface partition the μ -space.

1. Numerics

Given the Lefschetz thimble, we can numerically compute the normalized intensity map of the lens (see Fig. 19). The left, central and right panels depict the normalized intensity $I(\boldsymbol{\mu}; \nu)$ for $\mu_3 = -1, 0$ and $+1$. The upper, middle and lower panels depict the different frequencies $\nu = 50, 100$ and 500 .

We observe that for $\nu = 50$, interference is a dominant feature of the geometry of the caustic. The images are blurry and the geometry of the swallowtail is not resolved (Fig. 19a). We do observe the power-law falloff associated with the cusp singularities, which contrasts with the exponential falloff of the fold singularities.

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

D. The elliptic umbilic D_4^-

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

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

$$\Psi(\boldsymbol{\mu}; \nu) = \frac{\nu}{\pi} \int_{\mathbb{R}^2} e^{i(x_1^3 - 3x_1x_2^2 - \mu_3(x_1^2 + x_2^2) - \mu_2x_2 - \mu_1x_1)\nu} dx_1 dx_2. \quad (63)$$

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

$$3x_1^2 - 3x_2^2 - 2\mu_3x_1 - \mu_1 = 0 \quad (64)$$

$$-6x_1x_2 - 2\mu_3x_2 - \mu_2 = 0. \quad (65)$$

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




Figure 19: The normalized intensity, $I(\boldsymbol{\mu}; \nu)$, of the unfolding of the swallowtail caustic (A_4) sliced by the surfaces $\{\mu_3 = -1\}$, $\{\mu_3 = 0\}$, $\{\mu_3 = +1\}$ (respectively the left, central and right panels) for the frequencies $\nu = 50, 100$ and 500 (respectively the upper, the middle and lower panels).

this set of equations for μ_1 and μ_2 we obtain the Lagrangian map as a function of μ_3 ,

$$\xi_{\mu_3}(x_1, x_2) = (3x_1^2 - 3x_2^2 - 2x_1\mu_3, -2x_2(3x_1 + \mu_3), \mu_3). \quad (66)$$

In the geometric limit, we form a fold-surface and three cusp lines. The fold-surface in

base space $X = \mathbb{R}^2$ is given by

$$A_2^X(\mu_3) = \left\{ \left(\frac{\mu_3}{3} \cos \theta, \frac{\mu_3}{3} \sin \theta \right) \mid \theta \in [0, 2\pi) \right\} \quad (67)$$

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

$$|\mathcal{M}| = 0, \quad (68)$$

where the deformation tensor is given by

$$\mathcal{M} = \left[\frac{\partial^2 \phi(\mathbf{x}; \boldsymbol{\mu})}{\partial x_i \partial x_j} \right]_{i,j=1,2} \quad (69)$$

$$= \begin{pmatrix} 6x_1 - 2\mu_3 & -6x_2 \\ -6x_2 & -6x_1 - 2\mu_3 \end{pmatrix}. \quad (70)$$

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

$$A_3^X(\mu_3) = \left\{ \left(\frac{\mu_3}{3}, 0 \right), \left(\frac{\mu_3}{3} \cos \frac{2\pi}{3}, \frac{\mu_3}{3} \sin \frac{2\pi}{3} \right), \left(\frac{\mu_3}{3} \cos \frac{4\pi}{3}, \frac{\mu_3}{3} \sin \frac{4\pi}{3} \right) \right\} \quad (71)$$

in the X space.

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

The fold-surface is given by

$$A_2 = \left\{ \left(\frac{\mu_3^2}{3} (\mp 2 \cos \theta + \cos(2\theta)), -\frac{2\mu_3^2}{3} (\pm 1 + \cos(\theta)) \sin(\theta), \pm \mu_3 \right) \mid \theta \in [0, 2\pi), \mu_3 \in \mathbb{R} \right\} \quad (72)$$

where the two branches corresponding to \pm correspond to two disconnected pieces corresponding to the two eigenvalue fields of \mathcal{M} . The cusp lines are given by

$$A_3 = \left\{ (t^2, 0, t), (-t^2/2, \sqrt{3}t^2/2, t), (-t^2/2, -\sqrt{3}t^2/2, t) \mid t \in \mathbb{R} \right\}. \quad (73)$$

The fold-surface and cusp lines are illustrated in Fig. 20a. The red and the blue surfaces denote the fold surfaces corresponding to the eigenvalue fields λ_1 and λ_2 . The fold surface has a harp edge at the cusp lines (in black).

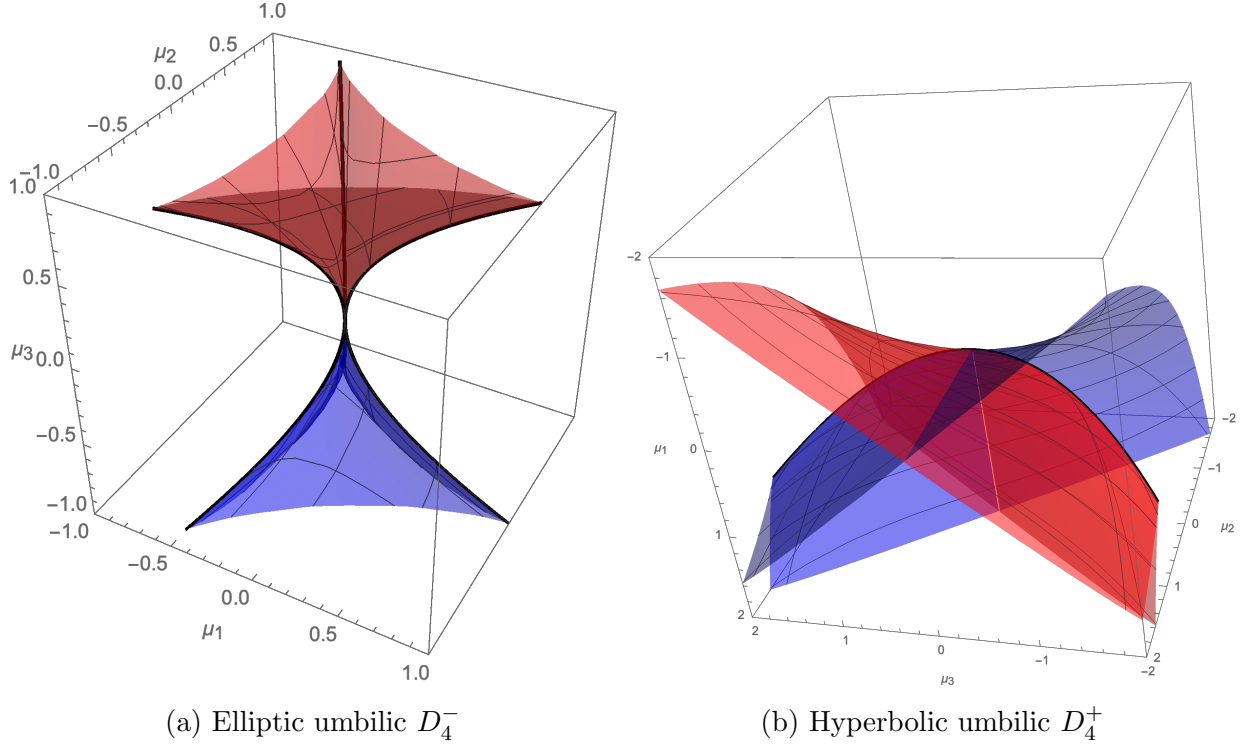


Figure 20: The elliptic (D_4^-) and hyperbolic umbilic (D_4^+) singularity in the unfolding space (μ_1, μ_2, μ_3) . The red and blue surfaces are the fold-surfaces corresponding to the eigenvalue field λ_1 and λ_2 separating the single- double- and triple-image regions. The black lines are the cusp-lines, along which we find the cusp saddle points.

Note the symmetry of the triangular singularity and point symmetry of the caustic. By performing the Picard-Lefschetz analysis for the two slices $\mu_3 = \pm 1$ and $\mu_3 = 0$ we can obtain the Picard-Lefschetz diagram of the unfolding of the singularity. See Fig. 21 for the Picard-Lefschetz analysis of the two slices. The small diagrams are the real parts of the four saddle points in the (x_1, x_2) -plane. The black circle is the caustic in the base space at the corresponding μ_3 .

- We first consider the case $\mu_3 \neq 0$. At the origin, the four saddle points are real-valued (see Fig. 21a). As a consequence, we conclude that they are all relevant. This is a quadruple-image region. One of the four saddle points is located inside the black circle. The other three are symmetrically distributed around the circle.

When crossing the fold-line, the saddle point in the circle merges with one of the outer saddle points on the circle. After passing the fold-line, the two saddle points become complex. The saddle point with the smallest real part of the exponent $i\phi(\bar{\mathbf{x}}; \boldsymbol{\mu})$ will remain relevant whereas the other saddle point becomes irrelevant. The outside of the




Figure 21: Intersection of the elliptic umbilic D_4^- by the surface $\{\mu_3 = -1\}$. The black line is the fold-line and the red line is the Stokes line. The number of real saddles gives the number of images in the geometric optics approximation.

triangle is a double-image region. Note that the real parts of the two complex saddle points always coincides with the black circle.

When approaching one of the three the cusp points, three of the four saddle points merge at a single point on the circle. Note that the four saddle points are collinear in the cusps.

Finally, note that the double-image region consists of six subregions divided by six Stokes lines. In the regions on the left, the upper right and the lower right, the Lefschetz thimble passes through two real and one complex saddle point. In the regions to the right, upper left and lower right, the Picard-Lefschetz analysis consist of only two real saddle points (the ones outside the circle. The Stokes lines are again associated with the three cusps.

- In the case, $\mu_3 = 0$, the central region is replaced by the elliptic umbilic saddle point (see Fig. 21b). The rest of the μ_1 - μ_2 -plane is divided into six distinct regions by the six Stokes lines. The upper left, lower left, and the upper right regions consist of two relevant real saddle points. These regions correspond to the upper left, lower left, and right region in Fig. 21a. The three remaining regions consist of two real and one

complex relevant saddle points.

These slices form a complete description of the Lefschetz thimble of the unfolding of the elliptic umbilic in the μ -space.

1. Numerics

Given the Lefschetz thimble, we can numerically evaluate the normalized intensity map of the lens (Fig. 22). The upper and lower panels depict the normalized intensity $I(\boldsymbol{\mu}; \nu)$ for $\mu_3 = \pm 1$ and 0. The left, central and right panels depict the frequencies $\nu = 50, 100$ and 500.

The normalized intensity map corresponding to the unfolding of the elliptic umbilic (D_4^-) has a triangular symmetry. As the frequency increases, the normalized intensity profile steepest and increases in amplitude. In the plane $\mu_3 = \pm 1$ we observe a fold-line in a triangular configuration with three cusp caustics at the corners. For the frequency, $\nu = 50$ the fold-line is relatively blurry. We again observe outward stripes emanating from the cusp caustics. These again follow a power-law falloff independent of the frequency. As the frequency is raised to $\nu = 100$ and $\nu = 500$ we observe that the fold lines become sharper and the fringes in the quadruple image region shrink. The normalized intensity at frequency $\nu = 500$ is very close to the normalized intensity map predicted by geometric optics.

E. The hyperbolic umbilic D_4^+

The hyperbolic umbilic D_4^+ completes the set of caustics appearing in two-dimensional lenses. It is again a singularity with co-rank 2 and co-dimension $K = 3$. The unfolding is described in terms of the three unfolding parameters (μ_1, μ_2, μ_3) . We consider the integral

$$\Psi(\boldsymbol{\mu}; \nu) = \frac{\nu}{\pi} \int_{\mathbb{R}^2} e^{i(x_1^3 + x_2^3 - \mu_3 x_1 x_2 - \mu_2 x_2 - \mu_1 x_1)\nu} dx_1 dx_2. \quad (74)$$

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

$$3x_1^2 - \mu_3 x_2 - \mu_1 = 0, \quad (75)$$

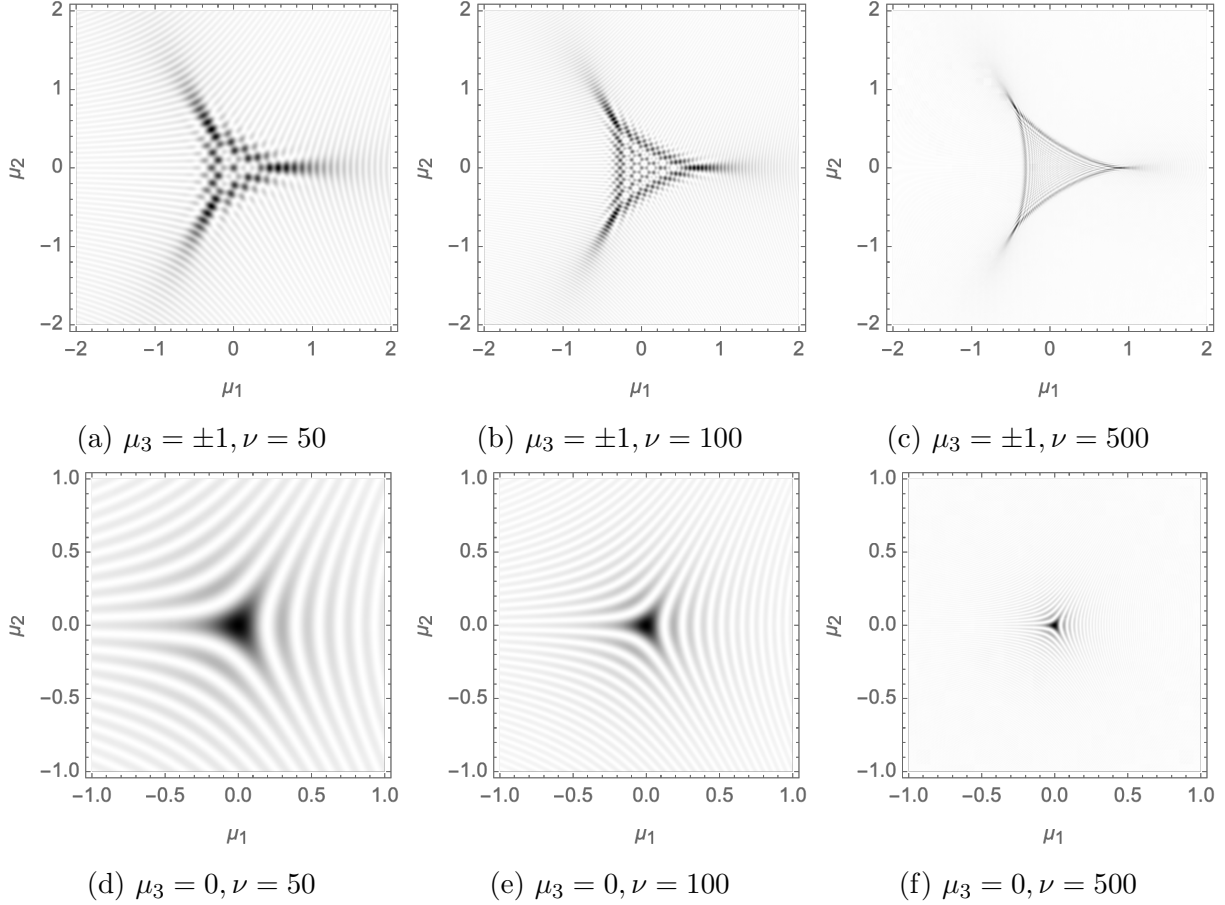


Figure 22: The normalized intensity, $I(\boldsymbol{\mu}; \nu)$, of the unfolding of the elliptic umbilic caustic (D_4^-) sliced by the surfaces $\{\mu_3 = \pm 1\}$, $\{\mu_3 = 0\}$ (respectively the upper and lower panels) for the frequencies $\nu = 50, 100$ and 500 (respectively the left, the centre and right panels).

$$3x_2^2 - \mu_3 x_1 - \mu_2 = 0. \quad (76)$$

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

$$\xi_{\mu_3}(x_1, x_2) = (3x_1^2 - x_2\mu_3, 3x_2^2 - x_1\mu_3). \quad (77)$$

In the geometric limit, we form a fold-surface and a cusp lines. The fold-surface in X

space is given by

$$A_2^X(\mu_3) = \left\{ \left(\pm \frac{\mu_3^2}{36t}, t \right) \mid t \in \mathbb{R} \right\} \quad (78)$$

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

$$|\mathcal{M}| = 0 \quad (79)$$

where the deformation tensor is given by

$$\mathcal{M} = \left[\frac{\partial^2 \phi(\mathbf{x}; \boldsymbol{\mu})}{\partial x_i \partial x_j} \right]_{i,j=1,2} \quad (80)$$

$$= \begin{pmatrix} 6x_1 & -\mu_3 \\ -\mu_3 & 6x_2 \end{pmatrix}. \quad (81)$$

The three cusp-lines are linear lines laying on the fold-surface,

$$A_3^X(\mu_3) = \{(-\mu_3/6, -\mu_3/6)\} \quad (82)$$

in the $X = \mathbb{R}^2$ space.

In the parameter space M , the elliptic umbilic point is located at the origin. The fold-surface is given by

$$A_2 = \{(3u^4 \pm 6uv^3, \pm 6u^3v + 3v^4, \mp 6uv) \mid u, t \in \mathbb{R}\} \quad (83)$$

$$A_2 = \{(3u^4 \mp 6uv^3, \mp 6u^3v + 3v^4, \mp 6uv) \mid u, t \in \mathbb{R}\} \quad (84)$$

where the two solutions correspond to two disconnected pieces corresponding to the two eigenvalue fields of \mathcal{M} . The cusp line in the parameter space is given by

$$A_3 = \{(t^2/4, t^2/4, t) \mid t \in \mathbb{R}\}. \quad (85)$$

The fold-surface and cusp-line are illustrated in Fig. 20b. The red and the blue surfaces denote the fold surfaces corresponding to the eigenvalue fields λ_1 and λ_2 of the deformation tensor \mathcal{M} . The fold surface has a sharp edge at the cusp lines (in black).

Note the symmetry of the triangular singularity and point symmetry of the caustic. By performing the Picard-Lefschetz analysis for the two slices $\mu_3 = \pm 1$ and $\mu_3 = 0$ we can obtain an understanding of the relevant saddle points.




Figure 23: Intersection of the hyperbolic umbilic D_4^+ with the surface $\{\mu_3 = -1\}$. The black line is the fold-line and the red line is the Stokes line. The number of real saddles gives the number of images in the geometric optics approximation.

See Fig. 23 for the Picard-Lefschetz analysis of the two slices. The small diagrams are the real parts of the four saddle points in the (x_1, x_2) plane. The black circle is caustic in X space at $\mu_3 = \pm 1$ and 0.

- Consider the slice $\mu_3 = \pm 1$ (see Fig. 23a). In the upper right corner, the four saddle points are real-valued. In the corresponding Picard-Lefschetz analysis, they are all relevant. This is a quadruple-image region. When we pass the left or lower fold-line, two of the four saddle points merge at the hyperbola in X space, to form a fold singularity. Afterward, both saddle points become complex. The one with the smallest real part of the exponent $i\phi(\bar{x}; \boldsymbol{\mu})\nu$ remains relevant whereas the other saddle point becomes irrelevant. Just like in the elliptic umbilic, the real part of the complex saddle points remains on the hyperboloid. This is a double-image region. Depending on whether we cross the fold line to the left or below the quadruple region, two different

saddle points merge.

If we move from the quadruple-image region to the cusp, we obtain a singularity due to the merger of three saddle points. After passing through the cusp, only the two real saddle points will be relevant. The two complex saddle points are irrelevant.

From the double-image region, we can pass the second fold-line. At this fold-line, the two remaining real saddle points merge to form a fold saddle point after which they move in the complex plane. Note that the real parts of these two saddle points remain on the second branch of the hyperbolic. Since the Picard-Lefschetz analysis does not contain any real-valued saddle points after passing the second fold-line, this is a zero-image region. The zero-image region is again subdivided into three subregions. In the upper left and lower right regions, the Picard-Lefschetz analysis consists of two relevant complex saddle points. In the lower left region, the Picard Lefschetz analysis consists of one relevant complex saddle point.

- In the case, $\mu_3 = 0$ the analysis is similar to the one obtained for $\mu_3 = \pm 1$, since the regions are trivially deformed (see Fig. 23b). In the upper right region, again four saddle points are real. All of them are relevant. This is still a quadruple-image region. The fold line along the positive μ_1 and μ_2 axis is double fold lines, as the two fold lines at $\mu_3 = \pm 1$ have merged. The left and lower right regions are zero-image regions. In the upper left and lower right regions, the Picard-Lefschetz analysis consists of two relevant complex saddle points. In the lower left region, the Picard-Lefschetz analysis again consists of one relevant complex saddle point. This concludes the Picard-Lefschetz analysis.

1. Numerics

Given the Lefschetz thimble, we can numerically compute the normalized intensity map (see Fig. 24). The upper and lower panels depict the normalized intensity $I(\boldsymbol{\mu}; \nu)$ for $\mu_3 = \pm 1$ and 0. The left, central and right panels depict the different frequencies $\nu = 50, 100$ and 500.

For both unfoldings at $\mu_3 = \pm 1$ and $\mu_3 = 0$, the normalized intensity map closely follows

the caustics structure represented in Fig. 23. In the zero-image region, the normalized intensity vanishes. In the double-image regions, for $\mu_3 = \pm 1$, the normalized intensity oscillates forming lines of equal normalized intensity as should be expected from the presence of the left fold line. In the quadruple-image regions, the normalized intensity oscillates in two directions, for $\mu_3 = \pm 1$ forming the structure we observed for the cusp caustic, and for $\mu_3 = 0$ forming an interference pattern with rectangular symmetry.

In the eikonal limit $\nu \rightarrow \infty$, the normalized intensity becomes sharper and the caustics become more pronounced. It should, in particular, be noted that the normalized intensity at the hyperbolic umbilic (in the origin in the plots corresponding to $\mu_3 = 0$), the normalized intensity rises rapidly with ν . This is in correspondence with the scaling relations we found above.

VI. TWO DIMENSIONAL LOCALIZED LENSES

The seven elementary singularities form a dictionary of the local behavior of the lens integral

$$\Psi(\boldsymbol{\mu}; \nu) = \left(\frac{\nu}{\pi}\right)^{N/2} \int_{\mathbb{R}^N} e^{i\phi(\boldsymbol{x}; \boldsymbol{\mu})\nu} d\boldsymbol{x}, \quad (86)$$

$$\phi(\boldsymbol{x}; \boldsymbol{\mu}) = (\boldsymbol{x} - \boldsymbol{\mu})^2 + \varphi(\boldsymbol{x}), \quad (87)$$

near caustics. Their corresponding normalized intensity map completely describes the local properties of lensed images. However, the global structure of the caustic is in general different. Since the normal forms of the elementary singularities are polynomials, the corresponding phase φ has support throughout the base space $X = \mathbb{R}^N$. The catastrophes with an even co-dimension K , lead to an image with an even number of images. In contrast, localized lenses lead to n -image regions with n an odd integer. We now turn to the study of interference patterns appearing in localized lenses near caustics. We evaluate three two-dimensional lenses, which simulate the behavior of a localized lens and include the five elementary catastrophes appearing in two-dimensional lenses. In the process, we also demonstrate the accuracy of the integration scheme along the Lefschetz thimble.

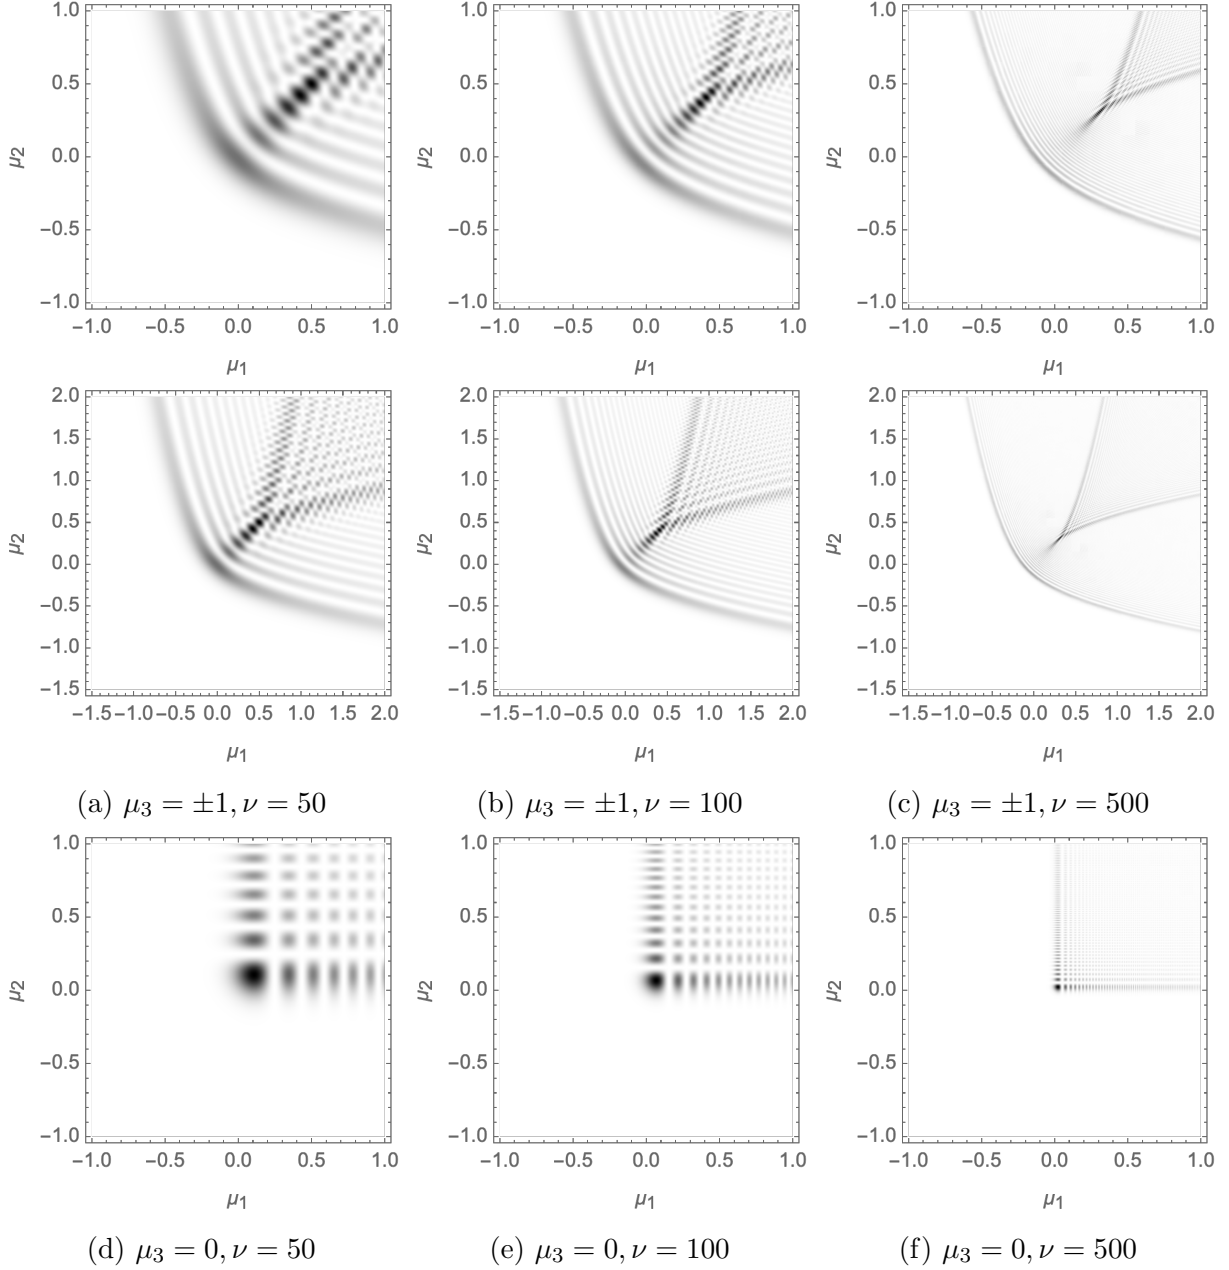


Figure 24: The normalized intensity, $I(\boldsymbol{\mu}; \nu)$, of the unfolding of the hyperbolic umbilic caustic (D_4^+) sliced by the surfaces $\{\mu_3 = \pm 1\}$, $\{\mu_3 = 0\}$ (respectively the upper and lower panels) for the frequencies $\nu = 50, 100$ and 500 (respectively the left, the centre and right panels).

A. A generic peak

In general, lensing effects are strongest near the extrema of the variation of the phase φ . It is for this reason natural to study the effect of an asymmetric peak in the phase variation

φ , with

$$\varphi(\mathbf{x}) = \frac{\alpha}{1 + x_1^2 + 2x_2^2}, \quad (88)$$

the two-dimensional generalization of the one-dimensional lens studied in Section IV. For astrophysical plasma lenses, the parameter α scales according to the dispersion relation $\alpha \propto \nu^{-2}$.

The Lagrangian map is given by

$$\xi(\mathbf{x}) = \mathbf{x} + \frac{1}{2}\nabla\varphi(\mathbf{x}) \quad (89)$$

$$= \mathbf{x} - \frac{\alpha}{(1 + x_1^2 + 2x_2^2)^2}(x_1, 2x_2). \quad (90)$$

The map forms a caustic where the deformation tensor

$$\mathcal{M}_{ij} = \frac{\partial^2\phi(\mathbf{x};\boldsymbol{\mu})}{\partial x_i\partial x_j}, \quad (91)$$

with the eigenvalue and eigenvector fields $\lambda_i(x), v_i(x)$, is singular, *i.e.*,

$$|\mathcal{M}(\mathbf{x})| = \lambda_1(\mathbf{x})\lambda_2(\mathbf{x}) = 0. \quad (92)$$

For convenience, we order the eigenvalue and eigenvector fields by $\lambda_1(\mathbf{x}) \leq \lambda_2(\mathbf{x})$.

The first caustic forms at the origin $(\mu_1, \mu_2) = (0, 0)$ for the parameter $\alpha = \frac{1}{2}$ (see Fig. 25). This is a cusp singularity. Note that by construction this caustic corresponds to the eigenvalue field λ_1 . For $\frac{1}{2} < \alpha < \frac{64}{49}$ the A_3 point forms an outgoing fold-line (A_2) with two cusps (A_3) on the left and the right. At $\alpha = 1$, a new A_3 point is created, this time corresponding to the second eigenvalue field λ_2 . For $1 < \alpha < \frac{64}{49}$ the A_3 point forms a fold-line (A_2) with two cusps (A_3) at the top and the bottom. At $\alpha = \frac{64}{49}$ the two fold lines merge in a hyperbolic umbilic (D_4^+) at $(\mu_1, \mu_2) = (0, \pm 1/\sqrt{14})$. For $\alpha > \frac{64}{49}$ the two fold lines continue to move outwards, where the fold-line corresponding to λ_1 has four cusps while the fold-line corresponding to λ_2 does not contain a cusp. Outside the fold-line of the caustics, the image consists of a single-image region. Inside the blue fold line, we find a triple-image and a five-image region enclosed by the red fold line.

The analytic continuation of the exponent, $\phi(\mathbf{x};\boldsymbol{\mu})$, possesses a pole on the two-




Figure 25: The caustics corresponding to λ_1 (blue) and λ_2 (red) as a function of α . We observe the formation of a triple- and a five-image region.

dimensional surface $x_1^2 + 2x_2^2 + 1 = 0$. Note that poles are never isolated in multi-dimensional complex analysis [57]. The exponent has nine saddle points \bar{x}_i . By evaluating the gradient of the h -function and flowing the original integration domain, we obtain a numerical representation of the thimble $\mathcal{J} \subset \mathbb{C}^2$.

Given the two-dimensional thimble \mathcal{J} , we numerically evaluate the integral $\Psi(\boldsymbol{\mu}; \nu)$. In Fig. 26, we plotted the normalized intensity of the sensed signal for $\alpha = 7/10, 64/49$ and 4 as a function of the frequency. Observe that when the wavelength is comparable to the size of the caustic structure, the normalized intensity is blurred. The caustics emerge when the wavelength becomes shorter. At the frequency $\nu = 500$, we accurately recover the image corresponding to geometric optics. Remark the stripes emanating from the cusp singularities. This is the frequency independent power-law falloff we observed in the elementary singularities.

In Fig. 27, we plot the cross-section of the normalized intensity map along the diagonal $\mu_1 = \mu_2$ for the lens with $\alpha = 4$ for $\nu = 50, 100$ and 500. Observe the four spikes while passing through the fold catastrophe. Note that the spikes increase in magnitude as ν is raised. In the astronomical context, these spikes correspond to amplification in the light-curve of the lensed source.

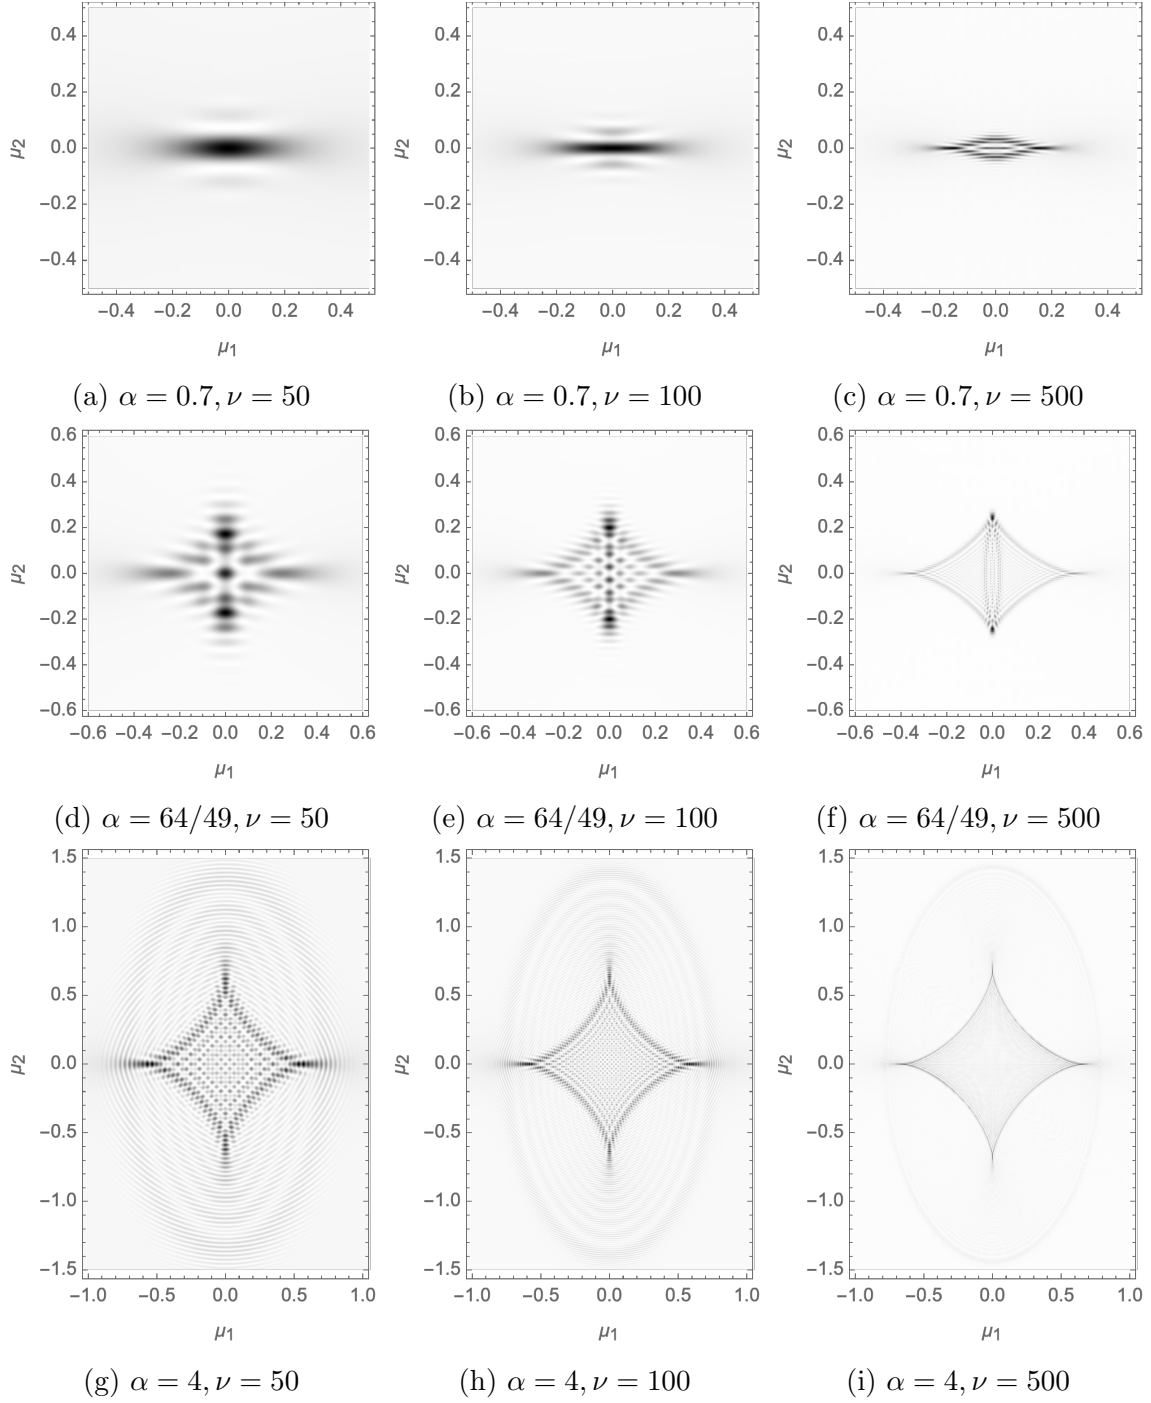


Figure 26: Intensity $I(\boldsymbol{\mu}; \nu)$ of the local lens at $\alpha = 0.7, 64/49, 4$ for $\nu = 50, 100$ and 500 .




Figure 27: The normalized intensity evaluate along the diagonal in the (μ_1, μ_2) -plane for $\alpha = 4$ for $\nu = 100$ and 500 . The black curve is the envelope predicted by geometric optics (see equation (18)).

B. A degenerate peak

A more intricate structure arises for the lens corresponding to the degenerate peak in the phase,

$$\varphi(\mathbf{x}) = \frac{\alpha}{1 + x_1^4 + x_2^2}, \quad (93)$$

with the Lagrangian map

$$\xi(\mathbf{x}) = \mathbf{x} + \frac{1}{2} \nabla \varphi(\mathbf{x}) \quad (94)$$

$$= \mathbf{x} - \frac{\alpha}{(1 + x_1^4 + x_2^2)^2} (2x_1^3, x_2). \quad (95)$$

The caustics structure of the Lagrangian map for varying α is plotted in Fig. 28. For $\alpha = 1$ we find two disconnected components, which are joined at $\alpha = 1.5$ and form an intricate pattern at $\alpha = 2$ and $\alpha = 2.5$. At $\alpha = 2$ we again find a hyperbolic umbilic caustic (D_4^+) at the two points where the cusps corresponding to the first and second eigenvalue fields λ_1, λ_2 coincide. We thus see that not only the structure at the peak but also the falloff of the variation in the phase φ is important in the study of caustics in lensed images. The caustic structure is generally sensitive to the Hessian of the phase φ , *i.e.*, the second order derivatives.

After flowing the original integration contour to the Lefschetz thimble \mathcal{J} , we numerically evaluate the amplitude $\Psi(\boldsymbol{\mu}; \nu)$ and the corresponding normalized intensity $I(\boldsymbol{\mu}; \nu)$. The resulting normalized intensity maps are plotted in figures 29 and 30. For the frequency $\nu =$




Figure 28: The caustics of the Lagrangian map for varying α . The caustics corresponding to the first and second eigenvalue fields λ_1, λ_2 in red and blue.

50, the image is again rather blurry. We can see the general shape, but cannot distinguish the detailed line structure. For the frequency $\nu = 100$, the lines are better resolved. However, the length scale of the caustics is comparable to the length scales of the interference patterns in the multi-image regions. For $\nu = 500$, we see the complete geometric structure of the caustics. The oscillations in the multi-image regions are now very fine. For this frequency, we are very close to the geometric optics approximation.

C. The swallowtail caustic

In the previous two examples of lenses corresponding to the simple peaks, we found both fold (A_2) and cusp caustics (A_3) corresponding to a single eigenvalue field, and the interaction between two eigenvalue fields via the hyperbolic umbilic (D_4^+). The two remaining caustics, *i.e.*, the swallowtail (A_4) and the elliptic umbilic (D_4^-), appear in slightly more involved lenses. For the swallowtail caustic, consider the lens

$$\varphi(\mathbf{x}) = \frac{\alpha x_1}{1 + x_1^4 + x_2^2}. \quad (96)$$

Again, in the astrophysical context, α follows the dispersion relation $\alpha \propto \nu^{-2}$.

The corresponding integrand $i\phi(\mathbf{x}; \boldsymbol{\mu})\nu$, consists of 23 saddle point in the complex plane. By deforming the integration domain to the thimble, we evaluate the two-dimensional lens integral numerically. See figures 31 and 32 for the caustics obtained from geometric optics and the corresponding normalized intensity maps for the frequency $\nu = 50, 100, 500$.

- For $\alpha = 2$, the lens forms a caustic corresponding to a single eigenvalue field (see the

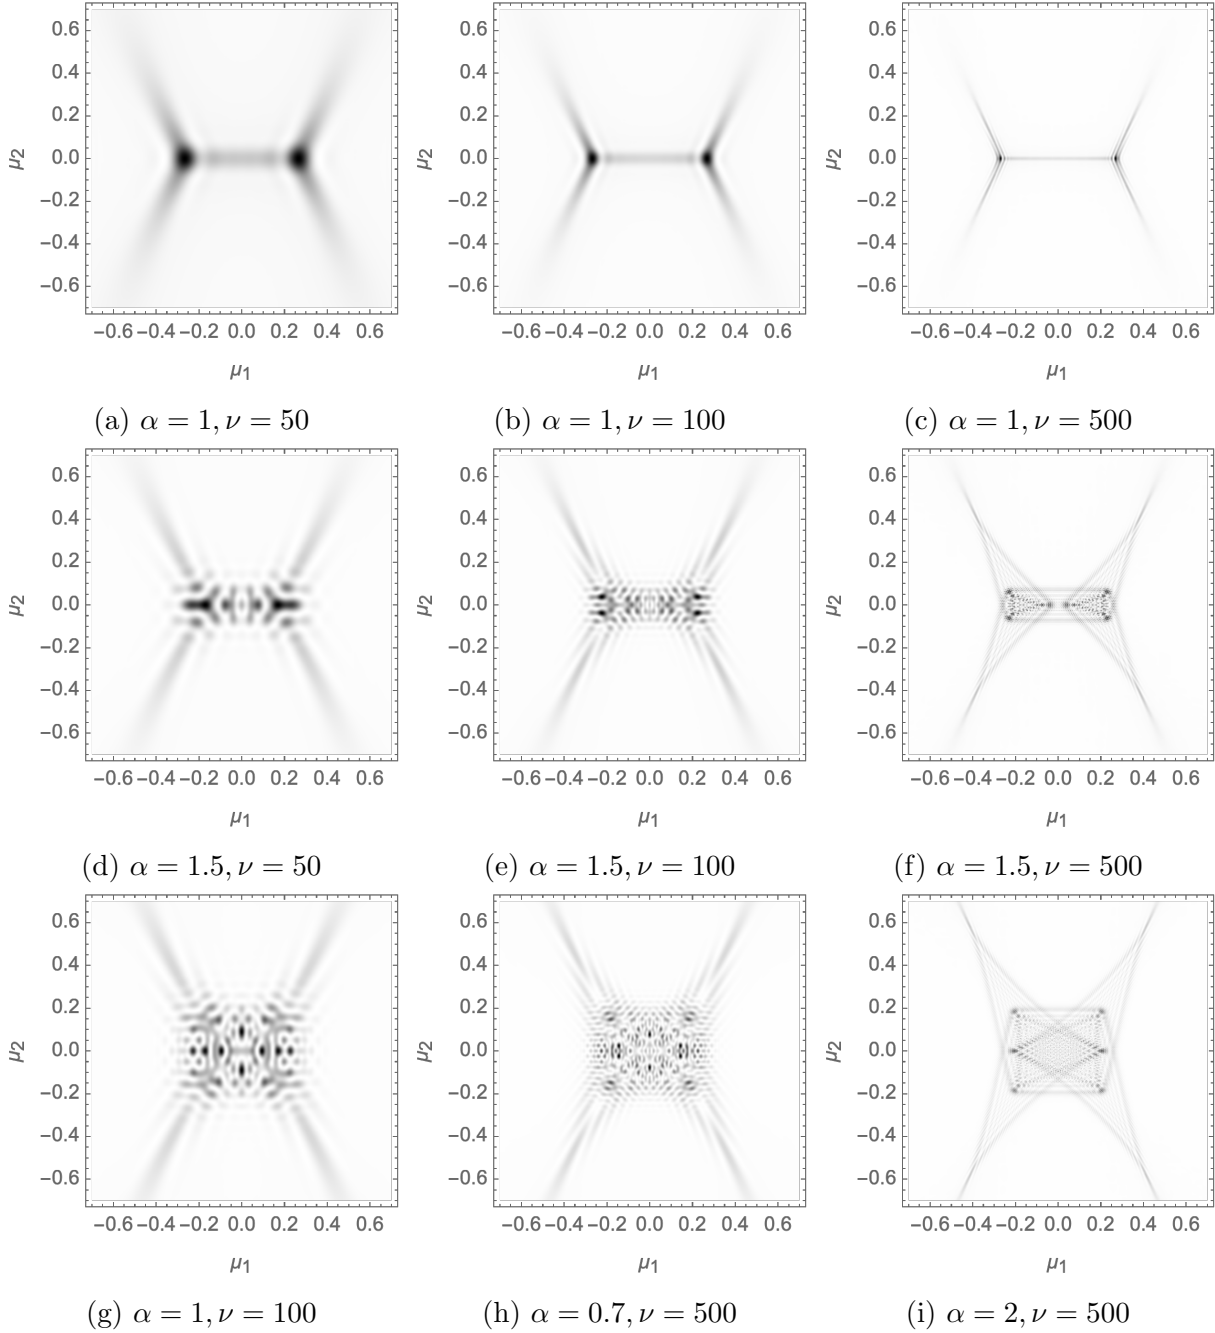


Figure 29: The normalized intensity map, $I(\boldsymbol{\mu}; \nu)$, for different frequencies.

upper panels of figure 31.). The profile is a pancake with two cusps at the tips. In the corresponding normalized intensity field, we see an interference pattern in the triple-image region, two stripes emanating from the cusps and more strikingly two diagonal stripes going to the left in the single image region. These stripes are a precursor of the swallowtail caustic emerging at later α .




Figure 30: The normalized intensity map, $I(\boldsymbol{\mu}; \nu)$, for different frequencies.

- As α is raised to 3, a second caustic emerges in the triple-image region (see the lower panels of figure 31). This caustic corresponds to the second eigenvalue field of the deformation tensor. At $\alpha = 3$ one of the two cusps of the second fold line merges with the outer fold line and transfers the cusp singularity via an elliptic umbilic caustic (D_4^-). For larger α , the blue line will thus have three cusps whereas the red line has only one.

However, more importantly, the lens forms a swallowtail caustic (A_4) in the blue line at $\alpha = 2$. This phenomenon cannot be observed in the blue fold-line but is apparent in the normalized intensity map. The two stripes already visible for $\alpha = 2$ are amplified. At the location where the swallowtail stripe coincides with the fold-line, we see an amplification of the normalized intensity in the swallowtail point.

In the normalized intensity map, we see that the geometry becomes sharper and sharper as we increase the frequency and approach the geometric optics limit. Note that the normalized intensity of the hyperbolic umbilic (D_4^+) outshines the other caustics at frequency $\nu = 500$.

- Finally, for $\alpha = 4$, we see that the swallowtail caustic has unfolded into its characteristic shape in the blue fold-line (see Fig. 32). We see the same structure emerge in the normalized intensity map. However, in addition, we how to obtain a large number of stripes emanating from the cusp caustics.

We also see that the lens at $\alpha = 4$, consists of a second hyperbolic umbilic (D_4^+) appearing at the origin, where the blue and the red fold-lines meet. As the frequency




Figure 31: The normalized intensity map, $I(\boldsymbol{\mu}; \nu)$, for $\alpha = 2, 3$ and frequencies $\nu = 50, 100, 500$.

is raised, we again see that the normalized intensity spikes for this caustic.

D. The elliptic umbilic caustic

We conclude this section by studying the elliptic umbilic (D_4^-) caustic in a localized lens. The elliptic umbilic forms when the deformation tensor is singular due to two eigenvalues vanishing simultaneously. The geometry of the caustic however differs from the hyperbolic umbilic (D_4^+), in that it includes the merger of three cusp caustics. We here study the




Figure 32: The normalized intensity map, $I(\boldsymbol{\mu}; \nu)$, for $\alpha = 4$ and frequencies $\nu = 50, 100, 500$.




Figure 33: The caustics of the Lagrangian map for varying α .

localized lens

$$\varphi(\mathbf{x}) = \frac{\alpha(x_1^3 - 3x_1x_2^2)}{1 + x_1^2 + x_2^2}. \quad (97)$$

From geometric optics, we observe the caustic structure of the lens (see Fig. 33).

- For small, $\alpha < 1.4$, the lens consists of three Zel'dovich pancakes with a triangular symmetry. Three of the cusp caustic point to the origin of the parameter space.
- At $\alpha = 1.4$, we observe that the three Zel'dovich pancakes are joined by three fold-lines forming a triangular structure.
- As $\alpha > 1.4$, the triangle decouples from the three Zel'dovich pancakes. The three




Figure 34: The normalized intensity map, $I(\boldsymbol{\mu}; \nu)$, for $\alpha = 1, 1.4, 5$ and frequencies $\nu = 50, 100$.

resulting fold lines move away from the origin and the triangle shrinks to a point. The point at which the triangle is contracted to a point is the elliptic umbilic caustic. The region enclosed by the large triangle is a 5-image region. The region enclosed by the small triangle is a 7-image region.

Note that since the elliptic umbilic caustic only forms after three cusp caustics have formed a triangular fold line, the caustic will be rare in simple lenses. It is nonetheless a stable configuration, as a small deformation of the lens preserves the structure.

Using the Picard-Lefschetz analysis, we evaluate the normalized intensity map for the configurations $\alpha = 1, 1.4$, and 5 for the frequencies $\nu = 50, 100$ (see Fig. 34).

- For $\alpha = 1$, we observe that even though the triangular structure is not yet present in the geometric optics analysis, it is present in the normalized intensity map at finite frequency (see the left panels of Fig. 34). That is to say, the normalized intensity is enhanced at the triangle, however as $\nu \rightarrow \infty$ the normalized intensity at the triangle

will remain finite.

- At $\alpha = 1.4$, the triangle has formed in the geometric optics analysis (see the central panels of Fig. 34). In the normalized intensity maps, the triangle is enhanced. The normalized intensity will now diverge in the geometric optics limit.
- As α is further increased to $\alpha = 5$, the triangle shrinks to a point and interference effects between the different fold lines start to appear (see the right panels of Fig. 34). At $\alpha = 5$, we do no longer observe the fold lines but rather observe a triangular blob at the origin of the parameter space. This closely resembles the normalized intensity map of the elementary elliptic umbilic catastrophe. It is however a bit more intricate as a close inspection demonstrates that caustic structure oscillates at a high frequency due to the interference of the elliptic umbilic with the surrounding multi-image region.

Note that there are a few small numerical artefacts present in the normalized intensity map for the lens at $\alpha = 5$. The lens outside of the triangle, is a 5-image region in which some of the real saddle points are located far away from the origin in the lens plane. The inside of the triangular region is a 7-image region. The Lefschetz thimble has a complicated shape and the tessellation of the thimble can occasionally miss a few points.

VII. SIGNATURES OF CAUSTICS IN FAST RADIO BURSTS

A Fast Radio Burst (FRB) is a millisecond transient radio pulse, caused by some yet to be identified high-energy astrophysical process. The first burst was found by Duncan Lorimer and his student David Narkevic in 2007 while scanning through archival pulsar survey data [58]. The burst in question had been detected in 2001 by the Parkes Observatory in Australia. In subsequent years, several other bursts were observed, among which the first repeating source (named FRB 121102) [8] was detected in 2012 by the Arecibo Observatory in Puerto Rico. In the last few months, several new detections have been announced by the Canadian Hydrogen Intensity Mapping Experiment (CHIME) collaboration, including the second repeating FRB source (named FRB 180814) [59]. FRBs are now known to be relatively common, with approximately 10,000 bright fast radio bursts occurring per day over the entire sky. Telescopes capable of detecting a significant fraction of these bursts should become possible in coming decades, an exciting prospect indeed.

The source of fast radio bursts is yet to be identified. Many different models have been proposed but none is yet compelling. They range from rapidly spinning neutron stars or black holes and regions of very high electromagnetic fields, to more exotic sources [60, 61]. It seems likely that the bursts are extragalactic in origin, as the first observed repeater, FRB 121102, has been identified with a galaxy at a distance of approximately 3 billion light years [62–64]. As mentioned in the introduction, it is likely that the phenomenology of fast radio bursts is strongly affected by astrophysical plasma lensing. They have a characteristic time-frequency profile, their frequency typically falling during the pulse, or series of pulses. This profile is probably due to the fact that lower frequencies are more strongly lensed and thus follow longer geometrical paths, and also because they propagate more slowly.

The methods and results we have reported here should be helpful in modeling the effects of plasma lensing on observed FRBs. The lensing may take place in a variety of places – near the source, near the observer or in between. If the line of sight encounters a caustic due to a plasma lens, the FRB may be amplified, enhancing the chances of detection. For reasons we have explained, caustics are likely to be localized in frequency, leading to the observed spectral shape. The “marching down” features could also be due to asymmetric structures in the lens, leading to angled caustics. This requires a preferred time asymmetry, which could in turn provide hints about the structure of the lens itself. In the lensing example of B1957+20 [13], the lens is due to a companion wind. In this specimen, the time-frequency caustics march both up and down. This symmetry could be broken if the wind contained shock waves, which could preferentially move retrograde in the rotating frame. Quantitative lens modeling can be tested on the pulsar binary system, and then applied to FRB data. This could be the scope of a future paper.

Since the observed radio waves have a relatively long wavelength, the corresponding diffraction catastrophes are likely to fill a significant volume in the parameter space of the normalized intensity maps. Therefore it is important to study the complete interference pattern. It follows from Table II that the elliptic (D_4^-), the hyperbolic (D_4^+) umbilic and to a lesser extent the swallowtail (A_4) caustic lead to the largest spikes in the normalized intensity map. Of these three caustics, the swallowtail (A_4) and the hyperbolic (D_4^+) umbilic caustics are most likely to be realized in simple lenses, of which the hyperbolic caustic gives the greatest amplification. However, these caustics will not generically occur in time-frequency data, as they are formed at point in three-dimensional functions. The line of sight, is,

however, reasonably likely to pass close to them, as they fill a finite volume of the parameter space. In principle, we do expect to see the cusp (A_3) points and the fold (A_2) lines caustics, in the data. However, note that these caustics lead to a lesser amplification of the source.

As we observed in the previous sections, caustics due to multi-dimensional lenses never occur as isolated events. The caustics of co-dimension four, *i.e.*, the umbilics D_4^\pm , and the swallowtail A_4 caustics, are always accompanied by cusp (A_3) points and fold (A_2) lines. It thus follows that when a fast radio burst is indeed amplified by a lens, that the corresponding peak in time-frequency space will be of characteristic shape. More concretely, after identifying the time and the frequency with the two of the unfolding parameters μ , we expect the peak to resample the normalized intensity map of the corresponding elementary catastrophe computed in Section V. That is to say, the peak corresponding to elliptic (D_4^-) umbilic caustic should exhibit a triangular symmetry and the peak corresponding to the swallow (A_4) caustic will exhibit the characteristic swallowtail geometry in the fold-line and two cusps caustics.

Further investigation is required to estimate the number density of the different caustics for generic two-dimensional lenses and the most likely normalized intensity profiles along the line of sight.

VIII. CONCLUSIONS

Conditionally convergent oscillatory integrals play a central role in modern physics. However, these integrals are often difficult to define as their definition, in the multi-dimensional case, can depend on the order of integration or the regularization scheme. They are, moreover, generically impossible to evaluate analytically and too expensive to evaluate with conventional numerical methods. In this paper we have brought Picard-Lefschetz theory to bear. We have shown how in a multi-dimensional oscillatory integral, the integrand generically defines a set of relevant Lefschetz thimbles in the complexified integration domain, along which the integral is absolutely convergent. These thimbles can be thought of as an ‘integrand-dependent Wick rotation’. The integral evaluated along the set of relevant thimbles in fact provides an unambiguous definition of the original integral itself. We moreover have presented a new, efficient numerical scheme both to find the thimbles and to efficiently evaluate the integral along them in polynomial time. The virtue of this new method that

the efficiency actually *increases* as the integrand becomes more oscillatory.

In particular, we have studied the Lefschetz thimbles for caustic catastrophes and the Stokes phenomenon occurring in two-dimensional lenses. Given the thimbles, we numerically evaluate the normalized intensity maps over all frequencies study the resulting interference patterns. We have shown that the normalized intensity maps smoothly converge to the caustics predicted by geometric optics, without introducing numerical artifacts.

Our method renders feasible the calculation of interference patterns in a wide variety of interesting astrophysical contexts, in particular to model the effect of plasma lenses on radio sources. So far, such modeling has been restricted to the simplest examples of fold and cusp singularities, produced by one dimensional lenses. More realistic, two-dimensional models, including the swallowtail, elliptic umbilic and hyperbolic umbilic caustics are now accessible. We have computed the normalized intensity maps for a few representative examples, and commented briefly on likely observational signatures. A statistical analysis of the normalized intensity profiles for the diffraction catastrophes generated by a realistic plasma lens ensembles will be the subject of further investigations.

Finally, we analyzed a simple model of Young's double slit experiment, representing an initial exploration of the use of these methods for describing interference in quantum mechanics.

Acknowledgments

We thank Roger Blandford, Claudio Bunster, Neal Dalal, Angelika Fertig, Sterl Finney, Steven Gratton, James Hartle, Estelle Inack, Nick Kaiser, Nynke Niezink, Laura Sberna and Doug Scalapino for interesting comments and encouragement. Research at Perimeter Institute is supported by the Government of Canada through Industry Canada and by the Province of Ontario through the Ministry of Research and Innovation. Ue-Li Pen holds Associate positions at the Dunlap Institute for Astronomy and Astrophysics and at Perimeter Institute. Ue-Li Pen and Neil Turok are Associate Fellows in the Canadian Institute for Advanced Research (CIFAR) Gravity and the Extreme Universe program. The authors also

gratefully acknowledge support from the Centre for the Universe at Perimeter Institute.

-
- [1] R. Feynman and A. Hibbs, *Quantum mechanics and path integrals*. International series in pure and applied physics. McGraw-Hill, 1965.
- [2] J. Klauder, *A Modern Approach to Functional Integration*. Applied and Numerical Harmonic Analysis. Birkhäuser Boston, 2010. https://books.google.ca/books?id=QYA_AAAAQBAJ.
- [3] M. V. Berry and J. F. Nye, “Fine structure in caustic junctions,” *Nature* **267** (May, 1977) 34–36.
- [4] M. V. Berry, “Focusing and twinkling: critical exponents from catastrophes in non-Gaussian random short waves,” *Journal of Physics A Mathematical General* **10** (Dec, 1977) 2061–2081.
- [5] M. V. Berry, J. F. Nye, and F. J. Wright, “The Elliptic Umbilic Diffraction Catastrophe,” *Philosophical Transactions of the Royal Society of London Series A* **291** (Apr, 1979) 453–484.
- [6] M. V. Berry and C. Upstill, “IV Catastrophe Optics: Morphologies of Caustics and Their Diffraction Patterns,” *Progress in Optics* **18** (Jan, 1980) 257–346.
- [7] M. V. Berry, “Focused tsunami waves,” *Proceedings of the Royal Society of London Series A* **463** (Nov, 2007) 3055–3071.
- [8] L. G. Spitler, J. M. Cordes, J. W. T. Hessels, D. R. Lorimer, M. A. McLaughlin, S. Chatterjee, F. Crawford, J. S. Deneva, V. M. Kaspi, R. S. Wharton, B. Allen, S. Bogdanov, A. Brazier, F. Camilo, P. C. C. Freire, F. A. Jenet, C. Karako-Argaman, B. Knispel, P. Lazarus, K. J. Lee, J. van Leeuwen, R. Lynch, S. M. Ransom, P. Scholz, X. Siemens, I. H. Stairs, K. Stovall, J. K. Swiggum, A. Venkataraman, W. W. Zhu, C. Aulbert, and H. Fehrmann, “Fast Radio Burst Discovered in the Arecibo Pulsar ALFA Survey,” *Astrophys. J.* **790** (Aug, 2014) 101, [arXiv:1404.2934](https://arxiv.org/abs/1404.2934) [astro-ph.HE].
- [9] CHIME/FRB Collaboration, M. Amiri *et al.*, “Observations of fast radio bursts at frequencies down to 400 megahertz,” *Nature* **566** (2019) no. 7743, 230–234, [arXiv:1901.04524](https://arxiv.org/abs/1901.04524) [astro-ph.HE].
- [10] A. Josephy *et al.*, “CHIME/FRB Detection of the Original Repeating Fast Radio Burst Source FRB 121102,” [arXiv:1906.11305](https://arxiv.org/abs/1906.11305) [astro-ph.HE].
- [11] D. B. Melrose and P. G. Watson, “Scintillation of Radio Sources: The Role of

- Caustics,” *Astrophys. J.* **647** (Aug, 2006) 1131–1141.
- [12] J. M. Cordes, I. Wasserman, J. W. T. Hessels, T. J. W. Lazio, S. Chatterjee, and R. S. Wharton, “Lensing of Fast Radio Bursts by Plasma Structures in Host Galaxies,” *Astrophys. J.* **842** (Jun, 2017) 35, [arXiv:1703.06580 \[astro-ph.HE\]](#).
- [13] R. Main, I. S. Yang, V. Chan, D. Li, F. X. Lin, N. Mahajan, U.-L. Pen, K. Vanderlinde, and M. H. van Kerkwijk, “Pulsar emission amplified and resolved by plasma lensing in an eclipsing binary,” *Nature* **557** (May, 2018) 522–525, [arXiv:1805.09348 \[astro-ph.HE\]](#).
- [14] U.-L. Pen and L. King, “Refractive convergent plasma lenses explain extreme scattering events and pulsar scintillation,” *Mon. Not. R. Astron. Soc.* **421** (Mar, 2012) L132–L136, [arXiv:1111.6806 \[astro-ph.GA\]](#).
- [15] B. J. Rickett, “Interstellar scattering and scintillation of radio waves.,” *Annual review of astronomy and astrophysics* **15** (Jan, 1977) 479–504.
- [16] A. Ishimaru, “Optical scattering and diffusion in turbulence and scatterers (A),” *Journal of the Optical Society of America (1917-1983)* **68** (Oct, 1978) 1368.
- [17] U.-L. Pen and Y. Levin, “Pulsar scintillations from corrugated reconnection sheets in the interstellar medium,” *Mon. Not. R. Astron. Soc.* **442** (Aug, 2014) 3338–3346, [arXiv:1302.1897 \[astro-ph.GA\]](#).
- [18] A. V. Bilous, S. M. Ransom, and P. Demorest, “Unusually Bright Single Pulses from the Binary Pulsar B1744-24A: A Case of Strong Lensing?,” *Astrophys. J.* **877** (Jun, 2019) 125, [arXiv:1811.05766 \[astro-ph.HE\]](#).
- [19] R. L. Fiedler, B. Dennison, K. J. Johnston, and A. Hewish, “Extreme scattering events caused by compact structures in the interstellar medium,” *Nature (London)* **326** (Apr, 1987) 675–678.
- [20] U.-L. Pen and L. King, “Refractive convergent plasma lenses explain extreme scattering events and pulsar scintillation,” *Mon. Not. R. Astron. Soc.* **421** (Mar, 2012) L132–L136, [arXiv:1111.6806 \[astro-ph.GA\]](#).
- [21] L. Dai and T. Venumadhav, “On the waveforms of gravitationally lensed gravitational waves,” *arXiv e-prints* (Feb, 2017) [arXiv:1702.04724](#), [arXiv:1702.04724 \[gr-qc\]](#).
- [22] X. Er and A. Rogers, “Two families of astrophysical diverging lens models,” *Mon. Not. R. Astron. Soc.* **475** (Mar, 2018) 867–878, [arXiv:1712.06900 \[astro-ph.GA\]](#).
- [23] W. A. Coles, J. P. Filice, R. G. Frehlich, and M. Yadlowsky, “Simulation of wave

- propagation in three-dimensional random media,” *Applied Optics* **34** (Apr, 1995) 2089.
- [24] W. A. Coles, B. J. Rickett, J. J. Gao, G. Hobbs, and J. P. W. Verbiest, “Scattering of Pulsar Radio Emission by the Interstellar Plasma,” *The Astrophysical Journal* **717** (Jul, 2010) 1206–1221, [arXiv:1005.4914 \[astro-ph.GA\]](#).
- [25] G. Grillo and J. Cordes, “Wave asymptotics and their application to astrophysical plasma lensing,” *arXiv e-prints* (Oct, 2018) arXiv:1810.09058, [arXiv:1810.09058 \[astro-ph.CO\]](#).
- [26] E. Witten, “A New Look At The Path Integral Of Quantum Mechanics,” *arXiv e-prints* (Sep, 2010) arXiv:1009.6032, [arXiv:1009.6032 \[hep-th\]](#).
- [27] J. Feldbrugge, J.-L. Lehners, and N. Turok, “Lorentzian Quantum Cosmology,” *Phys. Rev. D* **95** (2017) no. 10, 103508, [arXiv:1703.02076 \[hep-th\]](#).
- [28] J. Feldbrugge, J.-L. Lehners, and N. Turok, “No Smooth Beginning for Spacetime,” *Physical Review Letters* **119** (Oct, 2017) 171301, [arXiv:1705.00192 \[hep-th\]](#).
- [29] J. Feldbrugge, J.-L. Lehners, and N. Turok, “No rescue for the no boundary proposal: Pointers to the future of quantum cosmology,” *Phys. Rev. D* **97** (2018) no. 2, 023509, [arXiv:1708.05104 \[hep-th\]](#).
- [30] J. Feldbrugge, J.-L. Lehners, and N. Turok, “Inconsistencies of the New No-Boundary Proposal,” *Universe* **4** (2018) no. 10, 100, [arXiv:1805.01609 \[hep-th\]](#).
- [31] A. Di Tucci, J. Feldbrugge, J.-L. Lehners, and N. Turok, “Quantum Incompleteness of Inflation,” *arXiv e-prints* (June, 2019) arXiv:1906.09007, [arXiv:1906.09007 \[hep-th\]](#).
- [32] J. Feldbrugge, A. Fertig, L. Sberna, , and N. Turok, “In preparation: The Spacetime Amplitude Approach to Relativistic Quantum Mechanics: the Schwinger Effect,”.
- [33] L. Dai and W. Lu, “Probing Motion of Fast Radio Burst Sources by Timing Strongly Lensed Repeaters,” *Astrophys. J.* **847** (Sep, 2017) 19, [arXiv:1706.06103 \[astro-ph.HE\]](#).
- [34] M. Born and E. Wolf, *Principles of Optics*. Cambridge university press, 1999.
- [35] J. Jackson, *Classical Electrodynamics*. John Wiley, New York, 1975.
- [36] J. Feldbrugge and N. Turok, “In preparation,”.
- [37] G. Basar, G. V. Dunne, and M. Ünsal, “Resurgence theory, ghost-instantons, and analytic continuation of path integrals,” *Journal of High Energy Physics* **2013** (Oct, 2013) 41, [arXiv:1308.1108 \[hep-th\]](#).
- [38] A. Behtash, G. V. Dunne, T. Schaefer, T. Sulejmanpasic, and M. Unsal, “Toward Picard-Lefschetz Theory of Path Integrals, Complex Saddles and Resurgence,” *arXiv e-prints*

- (Oct, 2015) arXiv:1510.03435, [arXiv:1510.03435 \[hep-th\]](#).
- [39] G. V. Dunne and M. Ünsal, “What is QFT? Resurgent trans-series, Lefschetz thimbles, and new exact saddles,” *arXiv e-prints* (Nov, 2015) arXiv:1511.05977, [arXiv:1511.05977 \[hep-lat\]](#).
- [40] A. Behtash, G. V. Dunne, T. Schäfer, T. Sulejmanpasic, and M. Ünsal, “Complexified Path Integrals, Exact Saddles, and Supersymmetry,” *Phys. Rev. Lett.* **116** (Jan, 2016) 011601, [arXiv:1510.00978 \[hep-th\]](#).
- [41] A. Behtash, G. V. Dunne, T. Schäfer, T. Sulejmanpasic, and M. Ünsal, “Critical points at infinity, non-Gaussian saddles, and bions,” *Journal of High Energy Physics* **2018** (Jun, 2018) 68, [arXiv:1803.11533 \[hep-th\]](#).
- [42] M. Serone, G. Spada, and G. Villadoro, “The power of perturbation theory,” *Journal of High Energy Physics* **2017** (May, 2017) 56, [arXiv:1702.04148 \[hep-th\]](#).
- [43] Y. Tanizaki and T. Koike, “Real-time Feynman path integral with Picard-Lefschetz theory and its applications to quantum tunneling,” *Annals of Physics* **351** (Dec, 2014) 250–274, [arXiv:1406.2386 \[math-ph\]](#).
- [44] G. Brooker, *Modern Classical Optics*. Oxford Master Series in Physics. OUP Oxford, 2003. <https://books.google.ca/books?id=Z1bY1eW4znIC>.
- [45] W. A. Coles, “Interplanetary Scintillation,” *Space Science Reviews* **21** (Feb, 1978) 411–425.
- [46] M. v. d. Kamp, P. S. Cannon, and M. Terkildsen, “Effect of the ionosphere on defocusing of space-based radars,” *Radio Science* **44** (Feb, 2009) 1–15.
- [47] V. I. Arnol’d, “Normal forms for functions near degenerate critical points, the Weyl groups of A_k, D_k, E_k and Lagrangian singularities,” *Functional Anal. Appl* **6** (1972) 1972.
- [48] V. I. Arnol’d, “Wave front evolution and equivalent Morse lemma,” *Communications in Pure Applied Mathematics* **29** (Nov., 1976) 557–582.
- [49] R. Thom, *Structural Stability And Morphogenesis*. CRC Press, 2018. <https://books.google.ca/books?id=YVNPdWAAQBAJ>.
- [50] V. I. Arnol’d, “Remarks on the Stationary Phase Method and Coxeter Numbers,” *Russian Mathematical Surveys* **28** (Oct, 1973) 19–48.
- [51] V. I. Arnol’d, “Critical Points of Smooth Functions and Their Normal Forms,” *Russian Mathematical Surveys* **30** (Oct, 1975) 1–75.
- [52] J. Feldbrugge, R. van de Weygaert, J. Hidding, and J. Feldbrugge, “Caustic Skeleton & i

- Cosmic Web,” *Journal of Cosmology and Astro-Particle Physics* **2018** (May, 2018) 027, [arXiv:1703.09598](https://arxiv.org/abs/1703.09598) [astro-ph.CO].
- [53] I. Thompson, “NIST Handbook of Mathematical Functions, edited by Frank W.J. Olver, Daniel W. Lozier, Ronald F. Boisvert, Charles W. Clark,” *Contemporary Physics* **52** (Sep, 2011) 497–498.
- [54] A. N. Varchenko, “Zeta-Function of Monodromy and Newton’s Diagram.,” *Inventiones Mathematicae* **37** (Jan, 1976) 253.
- [55] M. Morse, “Relations between the critical points of a real function on n independent variables,” *Transaction of the American Mathematical Society* **27** (1925) 345–396.
- [56] E. Stein, J. Milnor, M. Spivak, R. Wells, R. Wells, and J. Mather, *Morse Theory*. Annals of mathematics studies. Princeton University Press, 1963.
<https://books.google.ca/books?id=uj44vQEACAAJ>.
- [57] R. M. Range, “Complex analysis: A brief tour into higher dimensions,” *The American Mathematical Monthly* **110** (2003) 89–108.
- [58] D. R. Lorimer, M. Bailes, M. A. McLaughlin, D. J. Narkevic, and F. Crawford, “A Bright Millisecond Radio Burst of Extragalactic Origin,” *Science* **318** (Nov, 2007) 777, [arXiv:0709.4301](https://arxiv.org/abs/0709.4301) [astro-ph].
- [59] CHIME/FRB Collaboration, “A second source of repeating fast radio bursts,” *Nature* **566** (Jan, 2019) 235–238, [arXiv:1901.04525](https://arxiv.org/abs/1901.04525) [astro-ph.HE].
- [60] U.-L. Pen, “The nature of fast radio bursts,” *Nature Astronomy* **2** (Oct, 2018) 842–844, [arXiv:1811.00605](https://arxiv.org/abs/1811.00605) [astro-ph.HE].
- [61] E. Platts, A. Weltman, A. Walters, S. P. Tendulkar, J. E. B. Gordin, and S. Kandhai, “A Living Theory Catalogue for Fast Radio Bursts,” *arXiv e-prints* (Oct, 2018) [arXiv:1810.05836](https://arxiv.org/abs/1810.05836), [arXiv:1810.05836](https://arxiv.org/abs/1810.05836) [astro-ph.HE].
- [62] S. Chatterjee, C. J. Law, R. S. Wharton, S. Burke-Spolaor, J. W. T. Hessels, G. C. Bower, J. M. Cordes, S. P. Tendulkar, C. G. Bassa, P. Demorest, B. J. Butler, A. Seymour, P. Scholz, M. W. Abruzzo, S. Bogdanov, V. M. Kaspi, A. Keimpema, T. J. W. Lazio, B. Marcote, M. A. McLaughlin, Z. Paragi, S. M. Ransom, M. Rupen, L. G. Spitler, and H. J. van Langevelde, “A direct localization of a fast radio burst and its host,” *Nature* **541** (Jan, 2017) 58–61, [arXiv:1701.01098](https://arxiv.org/abs/1701.01098) [astro-ph.HE].
- [63] S. Chatterjee, R. Wharton, C. J. Law, J. Hessels, S. Burke-Spolaor, G. C. Bower, M. W.

- Abruzzo, C. Bassa, B. J. Butler, J. M. Cordes, D. Paul, V. M. Kaspi, M. McLaughlin, S. M. Ransom, P. Scholz, A. Seymour, L. Spitler, S. P. Tendulkar, PALFA Survey Team, VLA+AO FRB121102 Simultaneous Campaign Team, and EVN FRB121102 Campaign Team, “Localizing the Fast Radio Burst 121102,” in *American Astronomical Society Meeting Abstracts #229*, vol. 229 of *American Astronomical Society Meeting Abstracts*, p. 330.01. Jan, 2017.
- [64] D. Michilli, A. Seymour, J. W. T. Hessels, L. G. Spitler, V. Gajjar, A. M. Archibald, G. C. Bower, S. Chatterjee, J. M. Cordes, K. Gourdji, G. H. Heald, V. M. Kaspi, C. J. Law, C. Sobey, E. A. K. Adams, C. G. Bassa, S. Bogdanov, C. Brinkman, P. Demorest, F. Fernandez, G. Hellbourg, T. J. W. Lazio, R. S. Lynch, N. Maddox, B. Marcote, M. A. McLaughlin, Z. Paragi, S. M. Ransom, P. Scholz, A. P. V. Siemion, S. P. Tendulkar, P. van Rooy, R. S. Wharton, and D. Whitlow, “An extreme magneto-ionic environment associated with the fast radio burst source FRB 121102,” *Nature* **553** (Jan, 2018) 182–185, [arXiv:1801.03965](https://arxiv.org/abs/1801.03965) [astro-ph.HE].
- [65] J. Doob, *Measure Theory*. Graduate Texts in Mathematics. Springer New York, 2012. <https://books.google.ca/books?id=HOPhBwAAQBAJ>.

Appendix A: Defining oscillatory integrals

Oscillatory integrals, which do not converge absolutely, are sometimes claimed to be ill-defined since the key theorems of measure theory, *e.g.* the dominated convergence theorem and Fubini’s theorem, do not apply [65]. We here study conditionally convergent oscillatory integrals for the one- and multi-dimensional case and propose a definition using Picard-Lefschetz theory in terms of absolutely convergent ones.

1. One-dimensional integral

The Fresnel integral

$$F(\infty) = \int_{-\infty}^{\infty} e^{ix^2} dx = (1+i)\sqrt{\frac{\pi}{2}} \quad (\text{A1})$$




Figure 35: The Euler or Cornu spiral. The black line is the real and imaginary part of $F(R)$ as a function of R . The black point is the real and imaginary part of the limit

$$\lim_{R \rightarrow \infty} F(R) = (1 + i)\sqrt{\frac{\pi}{2}}.$$

exists, even though the integral is only conditionally convergent. The integral is usually defined as a limit of the partial integral

$$F(R) = \int_{-R}^R e^{ix^2} dx, \quad (\text{A2})$$

i.e., $\lim_{R \rightarrow \infty} F(R) = (1 + i)\sqrt{\frac{\pi}{2}}$ following the Euler or Cornu spiral (see Fig. 35). This definition is as important to the integral as the integrand, as different regularization schemes – which do not approach the real line by adding points incrementally – lead to different answers.

The definition of the conditionally convergent integral in terms of the limit $R \rightarrow \infty$ is equivalent to the assumption of analyticity, since Cauchy’s integral theorem only applies to integrals over \mathbb{R} defined this way. We can alternatively define the integral by deforming the integration contour \mathbb{R} in the complex plane \mathbb{C} to the Lefschetz thimble

$$\mathcal{J} = \{(1 + i)u | u \in \mathbb{R}\} = (1 + i)\mathbb{R}, \quad (\text{A3})$$

for which the integrand is convex and the integral is absolutely convergent, *i.e.*,

$$\int_{\mathbb{R}} e^{ix^2} dx = \int_{\mathcal{J}} e^{ix^2} dx \quad (\text{A4})$$

$$= (1 + i) \int_{\mathbb{R}} e^{-2u^2} du \quad (\text{A5})$$

$$= (1 + i) \sqrt{\frac{\pi}{2}}. \quad (\text{A6})$$

Note that this definition does not depend on a limit. The regularization is completely determined by the assumption of analyticity.

2. Multi-dimensional integrals

Multi-dimensional conditionally convergent oscillatory integrals such as

$$\int_{\mathbb{R}^N} e^{if(x_1, \dots, x_N)} dx_1 \dots dx_N \quad (\text{A7})$$

for $N \in \mathbb{N}$ and appropriate functions f , play an important role in optics but cannot be uniquely defined using an extension of the regularization scheme described above for the one-dimensional case. To show this, let's consider the two-dimensional analogue of the Fresnel integral

$$\int_{\mathbb{R}^2} e^{i(x^2+y^2)} dx dy. \quad (\text{A8})$$

Since this integral factorizes, it is reasonable to require the integral to converge to

$$F(\infty)^2 = \left((1 + i) \sqrt{\frac{\pi}{2}} \right)^2 = i\pi. \quad (\text{A9})$$

However, for general $f(x_1, \dots, x_n)$ we are not able to write the integral as a product of one-dimensional integrals. This thus should not be considered as a desirable definition of the integral.

To see the dependence on the regularization scheme, consider the integral in polar coordinates. We write

$$I(R) = \int_{\mathbb{D}_R} e^{i(x^2+y^2)} dx dy \quad (\text{A10})$$

$$= 2\pi \int_0^R r e^{ir^2} dr \quad (\text{A11})$$

$$= i\pi \left(1 - e^{iR^2}\right), \quad (\text{A12})$$

with \mathbb{D}_R the disk of radius R centred at the origin. We thus find that the limit $\lim_{R \rightarrow \infty} I(R)$ does not exist! The function $I(R)$ instead circles the ‘correct answer’ $i\pi$ with increasing angular velocity.

It is instead appropriate to define the integral in terms of the Lefschetz thimble

$$\mathcal{J} = \{(1+i)(u, v) | (u, v) \in \mathbb{R}^2\} \quad (\text{A13})$$

$$= (1+i)\mathbb{R}^2. \quad (\text{A14})$$

Along the thimble, the integral is absolutely convergent

$$\int_{\mathbb{R}^2} e^{i(x^2+y^2)} dx dy = \int_{\mathcal{J}} e^{i(x^2+y^2)} dx dy \quad (\text{A15})$$

$$= (1+i)^2 \int_{\mathbb{R}^2} e^{-2(u^2+v^2)} du dv \quad (\text{A16})$$

$$= 2i \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} e^{-2(u^2+v^2)} du dv \quad (\text{A17})$$

$$= i\pi. \quad (\text{A18})$$

On the thimble we can safely convert the integral over the real plane \mathbb{R}^2 into the iterative integral using Fubini’s theorem, since the integral in u and v over \mathbb{R}^2 is absolutely convergent. This definition straightforwardly generalizes to general multi-dimensional conditionally convergent integrals.

Appendix B: Young’s double-slit experiment

In this appendix, we generalize our treatment of interference in order to tackle Young’s famous double slit experiment. In spite of the extreme simplicity of this example, and its centrality to introductory discussions of quantum physics, detailed interference patterns are surprisingly hard to compute. By generalizing our treatment of the Fresnel-Kirchhoff integral we shall be able to efficiently study the pattern created by a pair of smooth, finite size slits in detail. In particular, we shall see how quantum interference effects disappear in the classical limit, as \hbar is taken to zero.




Figure 36: The transition amplitude of the wall $T(x)$ with two slits at $s_1 = -1$, $s_2 = 1$ with width $\epsilon = 0.1$.

The generalization required is to make the interference “phase” complex in order to damp out the amplitude away from two narrow slits. Modeling this complex phase with a simple rational function, our numerical techniques allow us to efficiently find the relevant Lefschetz thimbles and compute the detailed interference pattern at all values of the parameters.

Consider a distant point source emitting particles towards a screen, with a thin barrier separating the screen from the source. The barrier is opaque to the particles except in the neighbourhood of two slits. In dimensionless coordinates (which we shall define below), the transmission amplitude takes the form

$$T(x) \propto \exp \left[\frac{\epsilon}{\epsilon^2 + (x - s_1)^2} + \frac{\epsilon}{\epsilon^2 + (x - s_2)^2} - \frac{1}{\epsilon} \right], \quad (\text{B1})$$

consisting of two peaks each of strength unity, centered respectively at $x = s_1$ and $x = s_2$. Here, $\epsilon > 0$ is a small number representing both the width of the slits (see Fig. 36) and the opacity of the barrier: away from the slits, the latter is given by $T \sim \exp(-1/\epsilon)$.

Assuming the incident amplitude for the particles to be coherent and constant across the slits, we may then compute the path integral amplitude just as in Section II. Here, however, we deal with a particle of fixed mass m , energy E and momentum $p = \sqrt{2mE}$. The last formula in (8), in the same small displacement-approximations made in Eq. (12) above, yields a Pythagorean contribution to the phase, $p(x - \mu)^2/(2d\hbar)$ where d is the distance from the slits to the screen. Setting $x \rightarrow xa$ where a is the characteristic dimension of the slits and x is dimensionless, we take the quantity $\hbar 2d/(pa^2)$ to be our new, dimensionless \hbar . In terms of these dimensionless quantities, the amplitude for the particle to arrive at

position μ on the screen is therefore given by the oscillatory integral

$$\Psi(\mu) = \mathcal{N} \int e^{\frac{i}{\hbar}(x-\mu)^2} T(x) dx \quad (\text{B2})$$

with the normalization constant \mathcal{N} , ensuring unitarity $\int |\Psi(\mu)|^2 d\mu = 1$. The probability for the particle to arrive at μ on the screen is given by the absolute square of the wavefunction

$$I(\mu) = |\Psi(\mu)|^2. \quad (\text{B3})$$

Note that the dimensionless version of Planck's constant \hbar appears in this nonrelativistic problem, whereas it cancelled out of our earlier formulae for a massless particle, as a result of the latter's scale covariance.

We evaluate the wavefunction (equation (B2)), by analytically continuing the exponent

$$\phi(x; \mu) = \frac{i}{\hbar}(x - \mu)^2 + \frac{\epsilon}{\epsilon^2 + (x - s_1)^2} + \frac{\epsilon}{\epsilon^2 + (x - s_2)^2} - \frac{1}{\epsilon}, \quad (\text{B4})$$

in the complex plane and evaluating the Lefschetz thimble. The exponent, ϕ , has four poles and nine saddle points. The poles at $x = s_i \pm i\epsilon$ correspond to the slit centered at s_i . The saddle points are roots of a ninth order polynomial. We can associate four saddle points to each slit. The remaining saddle point is shared and moves between the the poles corresponding to the two slits as a function of the position on the screen μ .

Fig. 37 shows the corresponding Picard-Lefschetz diagrams for various positions μ for $\hbar = 1$. In the description we will for simplicity assume the left slit to be at s_1 and the right slit to be at s_2 , *i.e.*, $s_1 < s_2$:

- For positions on the screen far to the left of the slits, $\mu \ll s_1$, the thimble consists of five steepest descent contours. The thimble runs from the lower left to the upper right via a complex saddle point. The thimble subsequently loops around the upper left and the upper right poles. For positions $\mu \leq \frac{s_1 + s_2}{2}$, the wavefunction is dominated by the left slit. It is for this reason not surprising to see that the thimble corresponding to the right slit is representative in this regime.
- As μ approaches s_1 , we observe a Stokes transition after which only four saddle points are relevant. The thimble moves from the lower-left via a saddle point to the upper

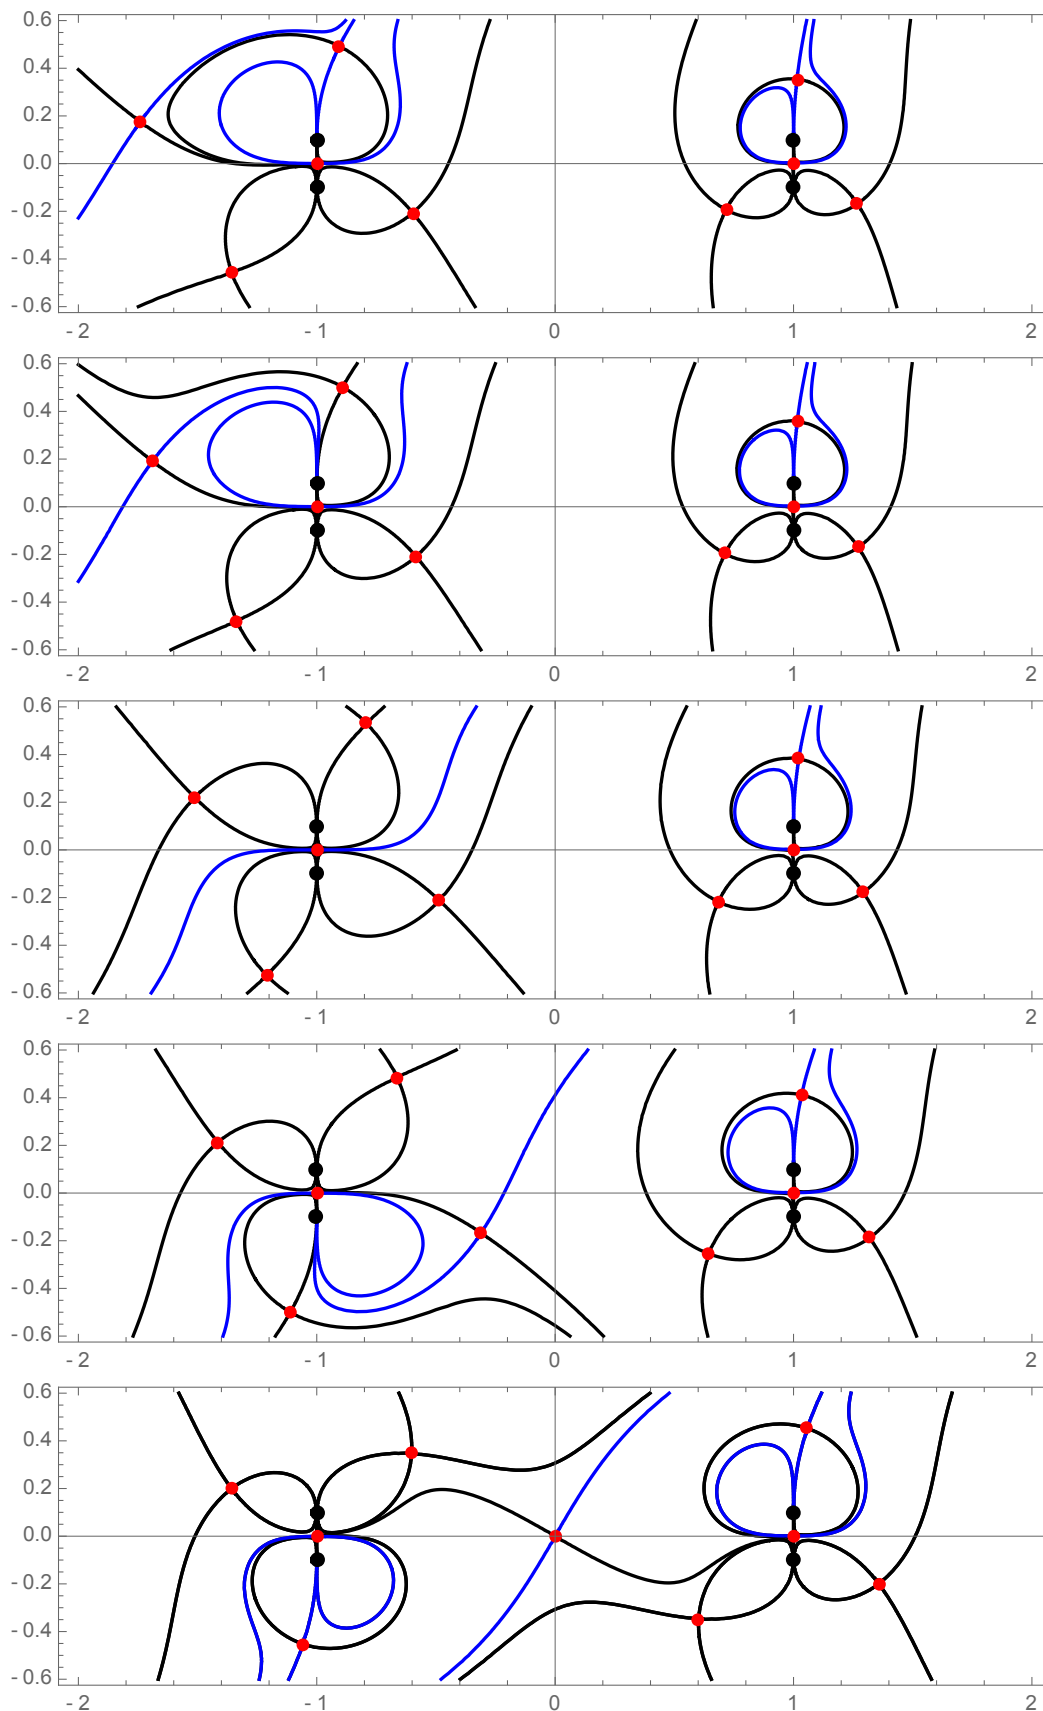


Figure 37: The Picard-Lefschetz diagram for the Young experiment for $\epsilon = 0.1$ from $\mu = -1.6, -1.5, -1, -0.5$, and $\mu = 0$, with the steepest ascent, descent contours (black) and thimbles (blue) corresponding to the saddle points (red) for $\hbar = 1$.

left pole, after which it passes through the saddle point between the two left poles. The right part of the thimble is largely unchanged.

- For μ near s_1 , we observe yet another Stokes transition after which only three saddle points remain relevant. The thimble runs from the lower left via the saddle point between the two left poles to the upper right.
- When μ approaches the mid-point $\frac{\mu_1+\mu_2}{2}$, we observe that a complex saddle point becomes relevant after a Stokes phenomenon. The thimble now consists of four steepest descent contours.
- For μ near the mid-point $\mu = \frac{\mu_1+\mu_2}{2}$, we observe that after yet another Stokes transition, we obtain a thimble consisting of five steepest descent contours. Note that the middle saddle point has moved to the origin $x = 0$. When the position μ is increased further, this saddle point will move to the poles corresponding to the right slit. The corresponding Picard-Lefschetz diagrams are mirror images of the ones discussed above.

In the semi-classical limit $\hbar \rightarrow 0$, the geometry of the Lefschetz thimble is to an increasing extent determined by the Pythagorean term in equation (B4). As a consequence, after a few Stokes transitions, the eight saddle points which can be associated to the poles corresponding to the two slits become tighter bound to the poles representing the geometry of the right part of the thimble in figure 37. The remaining saddle point still moves between the poles corresponding to the two slits. However, note that the integral is increasingly dominated by the two saddle points between the four poles. These two saddle points approach the real line at $x = s_1$ and $x = s_2$ in this limit.

Given the thimble, we can efficiently evaluate the oscillatory integral for various \hbar (see figure 38). For relatively large \hbar , the intensity on the screen is dominated by interference effects. For both $\hbar = 1$ and $\hbar = 1/2$ we do not observe the classical intensity peaks corresponding to the two slits. In the semi-classical limit, $\hbar \rightarrow 0$, the interference pattern is slowly replaced by the classical peaks. Note that this transition from the quantum to the classical regime cannot be studied in the traditional thin slit approximation.

Observe that, while the behavior of strong lenses is dominated by caustics where the saddle points become degenerate, the qualitative behavior of the double-slit experiment is

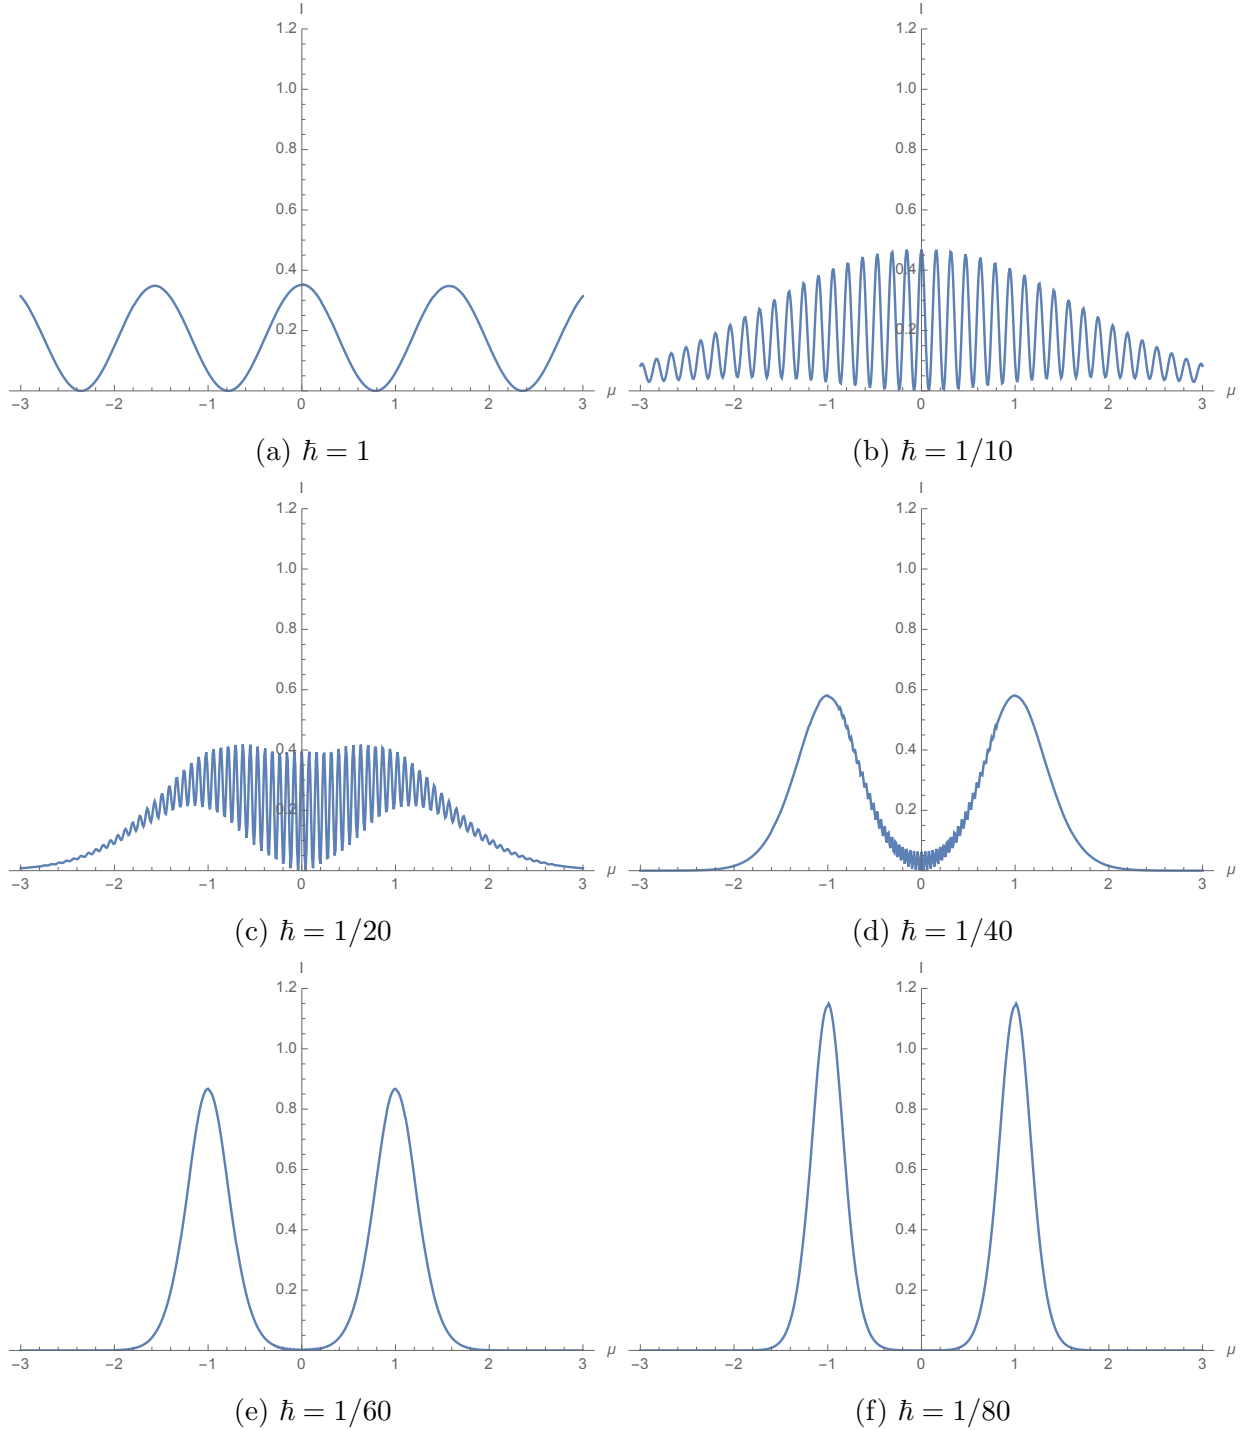


Figure 38: The intensity I as a function of position μ for various \hbar

completely determined by the Stokes transitions. The saddle points are everywhere non-degenerate and the h -function is a Morse function. In both instances, the saddle point approximation fails and the integral should be evaluated along the complete Lefschetz thimble. We expect this to be a generic feature in quantum mechanical interference phenomena.