Branched Flow in the Interstellar Medium

Jakob Tristan Faber

an Honors Thesis presented to the Oberlin College
Department of Physics and Astronomy

Advisors: Professors Rob Owen and Dan Stinebring

May 12, 2021

To my parents, Kim and Sebastiaan, and my sister, Maya.

Abstract

In this thesis, I introduce a universal wave behavior known as *branched flow* into the field of radio astronomy, along with a wealth of theoretical work from disparate realms of classical and quantum dynamics related to wave propagation in random media. Through analytic theory and numerical ray-tracing simulations, we confirm that the interstellar medium is very likely conducive to branched flow. Until now, branched flow has been a foreign concept to high-energy astrophysics and radio astronomy, yet the phenomenon lends a unique and enlightening opportunity to leverage ray optics and chaos theory in gaining a better understanding of radio wave scintillation in the interstellar medium.

Acknowledgments

I would like to thank my research advisor and mentor, Professor Dan Stinebring, for his unwavering support and invaluable guidance. The opportunities he has so generously given me and the time we have spent together over the past four years will undoubtedly stay with me as some of the best experiences of my life. I would also like to thank my other research advisor and mentor, Professor Rob Owen, for his enlightening teaching, thoughtful advice, and the fascinating conversations we have shared about physics and life. I also have my academic advisor, Professor Jason Stalnaker, to thank for his constant dedication to helping me maneuver life as a Physics major at Oberlin and succeed in my academics. Lastly, I would like to thank the entire Physics Department for making my experience at Oberlin truly wonderful, and for the support each and every professor has given me over the past four years.

Contents

Chapter 1

Wave Propagation in Random Media

1.1 Introduction

Wave propagation in random media is ubiquitous in nature. A familiar example is the twinkling of stars in the night sky. As starlight passes through the atmosphere, it is refracted by minute variations in temperature and pressure (Tatarskii, 1961; Roddier, 1981). Twinkling, or *scintillation*, is merely the repeated focusing (brightening) and defocusing (dimming) of starlight that occurs as a consequence of this refraction. For many of us, watching stars twinkle overhead instills a feeling of youthful whimsy. However, for optical astronomers whose telescope resolutions are severely limited by the atmospheric distortions of starlight, it is a nuisance. For many instances of wave propagation in random media, it seems that there are two sides from which the consequences can be perceived. On the one hand, they are often visibly beautiful and easy to appreciate. Yet, on the other, they challenge intuition, and present a deeply fascinating puzzle of physics. My hope is that the work presented in this thesis will emphasize the beauty that lies in trying to untangle the complexities of this natural phenomenon.

Regardless of the wave type, the scattering that occurs during propagation is dependent on the irregularities and refractive index fluctuations in the surrounding medium. As will be made clear in Chapter 2, however, the details of this dependence are trivial to an extent. A system for which scintillation has been extensively studied at optical frequencies by both natural and artificial sources, such as stars and lasers, is the atmosphere (Strohbehn, 1978). The upper ionized layer of the atmosphere (the ionosphere), however, can only be probed at radio frequencies—for instance, by radio galaxies and supernova remnants (Fallows et al., 2015, 2020). Radio waves in the ionosphere and optical waves in the lower atmosphere are distorted in similar ways; only radio waves are affected by plasma density irregularities rather than temperature and pressure fluctuations. Scintillation, however, is by no means limited to electromagnetic waves. Sounds waves traveling through the ocean are, for instance, a perfect analogue to electromagnetic wave propagation

(Keller & Papadakis, 1977). The instantiation of wave propagation in random media that will be studied in this thesis is the scintillation of radio waves emitted by pulsars (rapidly rotating neutron stars) and other astrophysical radio sources passing through plasma density irregularities in the interstellar medium (ISM).

Studies of wave propagation in random media usually seek to accomplish one (or more) of three things: (1) understanding the geometry of the medium, (2) understanding the nature or physics of the source, or (3) correcting for phase and amplitude distortions accumulated by the wave along its path through the medium. In this thesis, (1) and (2) are addressed, with the intention of addressing (3) in future work.

1.2 Thesis Summary

This thesis consists of seven chapters. Chapter 2 presents a diverse set of natural systems in which *branched flow* is known to occur. Chapter 3 explores the physics of branched flow, including the criteria that need to be met for it to exist, as well as the underlying optics with which it can be described and modeled. Chapter 4 explores the modeling of branched flow in continuous or extended media through Hamiltonian optics, as well as discretized "kick-drift" potentials through basic geometrical optics and phase space mapping. Chapter 5 outlines an argument that concludes branched flow ought to exist in the ISM. Chapter 6 introduces a three-dimensional kick-drift model with which branched flow can be simulated on astrophysical scales. Finally, Chapter 7 outlines the preliminary conclusions of our work, as well as directions for further research. In summary, we find that branched flow can almost certainly be expected to manifest in the propagation of radio waves (emitted by compact astrophysical objects such as pulsars) through the ISM.

Chapter 2

Branched Flow: A Universal Wave Behavior Unveiled

Branched flow has quietly pervaded wave propagation in random media for centuries, though it has only recently been recognized as a physical phenomenon. The term "branched flow" was first coined by Topinka et al. (2001) in reference to a remarkable electron diffusion pattern that was observed emanating from a Quantum Point Contact (QPC) in 2D electron gas (2DEG) (see Fig. 2.1a). The pattern showed a network of branch-like structures of heightened electron flux that remained stable over distances much greater than the correlation length of the 2DEG, despite its stochasticity. The conclusion was that the 2DEG was conducive to a form of ensemble diffusion. Since this groundbreaking result, the phenomenon has been replicated repeatedly in experiments and simulations (Jura et al. (2007); Shaw (2002); Patsyk et al. (2020); see Fig 2.1b).

One possible reason for branched flow's long overdue recognition is that it manifests in a transitional or transient regime between so-called "ballistic" (projectile-like) and "diffusive" or chaotic motion—in media that meet the criteria for its existence in the first place, of course. Transient dynamical regimes are generally complex and easy to miss. Furthermore, the complexity of branched flow makes it significantly more difficult to appreciate without the computational capabilities required to model it. The computational advancements of the 21st century have fortunately made such capabilities accessible on a standard laptop.

2.1 Universality in Dynamical Systems

Universality is a well-known concept in the domain of dynamical systems that classifies, broadly speaking, behaviors that do not depend on the details of the system in which they take place. A central question in both classical and quantum dynamics that motivates the study of universal phenomena is: how does nature mirror itself across a wide range of physical scales? Branched flow offers a uniquely astounding example of commonalities that can exist between physical systems with length scales spanning over 12 orders of

magnitude, from electron flow through semiconductor heterostructures (Topinka et al. (2001); see Fig. 2.1) to tsunamis racing across the Pacific Ocean (Berry (2007); see Fig 2.3). A looming question still remains, however: is branched flow an astrophysical phenomenon? The eddies in the interstellar medium can span up to 10^3 parsecs (1 pc = 10^{16}m). Thus, if the answer to this question is yes, then the introduction of interstellar wave propagation to the set of physical scales on which branched flow is already known to occur would contribute an additional 13 (25 in total) orders of magnitude to the problem.

2.2 Branched Flow In Nature

2.2.1 Electron Flow from QPCs

The experimental result by Topinka et al. (2001) (see Fig. 2.1a), later replicated by Jura et al. (2007) (see Fig. 2.1b), marked the official discovery of quantum branched flow. The imaging of electron transport in the nanostructures of the 2DEGs using scanning probe microscopy, narrowly constricted by a QPC, shows focusing due to ripples in the underlying potential landscape. Additionally, the filamentary branches show interference fringes that are separated by half the Fermi wavelength, which mark the preservation of phase coherence in the electron transport. The concept of phase coherence is of great interest to other studies of wave propagation in random media as well, including acoustic oceanography (see Section 2.2.2).

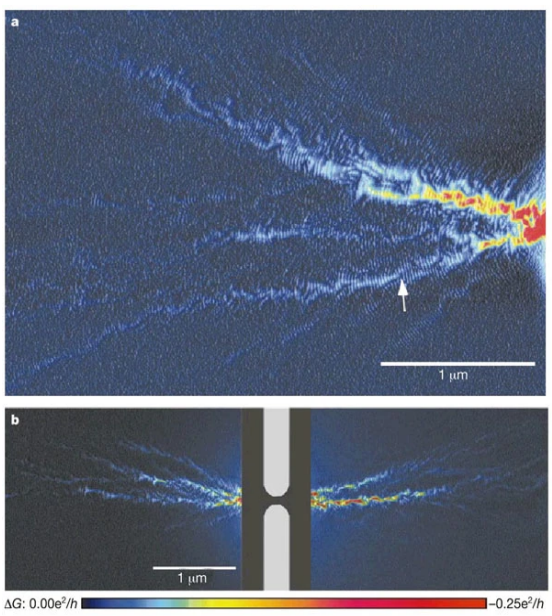

(a) The original experiment by Topinka et al. (2001).

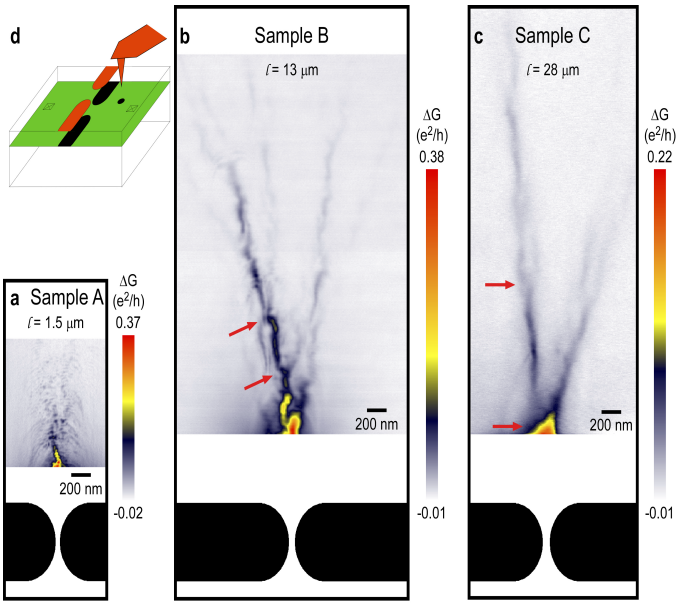

(b) A follow-up experiment by Jura et al. (2007).

Figure (2.1) Scanning probe microscopy image of electron flow from a quantum point contact (QPC) through two dimensional electron gas (2DEG).

2.2.2 Deep Ocean Acoustics

Acoustic oceanography has made extensive theoretical and experimental contributions to the understanding of long-range wave propagation through random media and its relationship to chaos theory. When sound is projected at 2 kilometers below the ocean surface, for instance, it experiences a loss of coherence due to multiple scattering, which is caused by the interaction of sound waves with internal ocean waves and mesoscale eddies. Geometrical acoustics simulations have shown that sound waves in these environments remain stable over many kilometers and eventually diverge chaotically beyond a so-called "predictability horizon" (Wolfson & Tappert (2000); Wolfson & Tomsovic (2001); see Figures 2.2a and 2.2b).

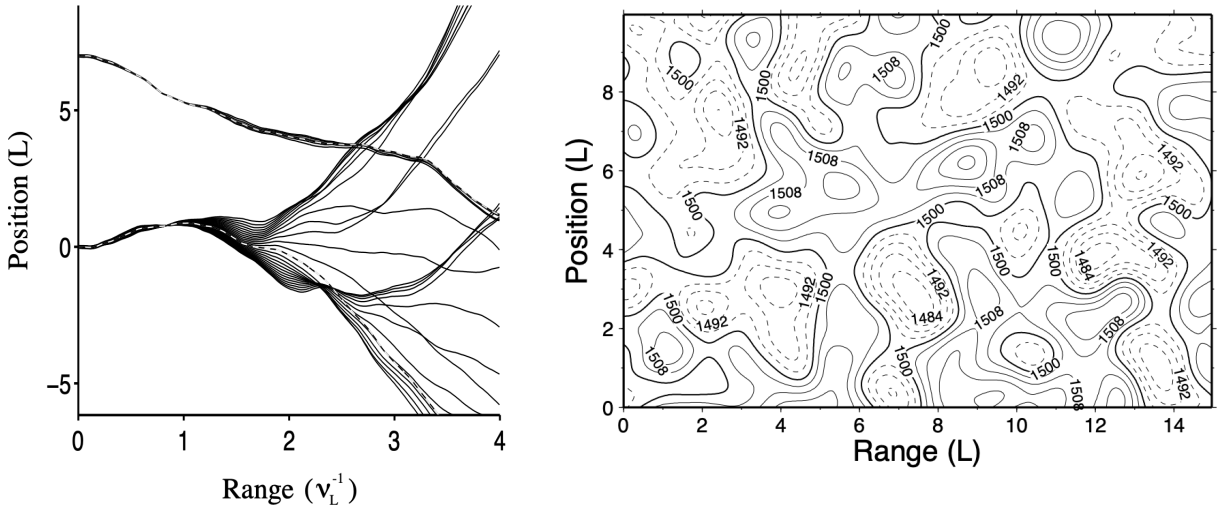

(a) A geometrical acoustics ray-tracing simulation in the ocean mesoscale.

(b) Two-dimensional contour plot of the ocean mesoscale as a Gaussian random potential.

Figure (2.2) A simulation of sound wave propagation in the ocean by Wolfson & Tomsovic (2001).

2.2.3 Rogue Waves at Sea

Rogue waves or "freak waves" have posed a great danger to ships sailing the high-seas for centuries, though they have only recently been causally linked to extreme stochastic focusing (Heller, 2005). Extreme ("rogue") events of this nature occur in optical systems as well. Waves that attain the status of "extreme" are generally defined as those which have a height greater than twice the "significant height"—i.e., the average of the upper third of wave heights in a given population (Dudley et al., 2019). As we will see in later sections (see Chapter 4), the *intersection* of branches has been shown to produce the extreme amplitudes exhibited by rogue waves (Metzger et al., 2014).

2.2.4 Tsunamis in Shallow Water

The focusing (or lensing) of tsunami pressure waves by small-scale topological undulations in the ocean floor has been studied theoretically for many years (Berry (2007); see Fig. 2.3a). Over the last few decades, efforts have been made by organizations like the NOAA Center for Tsunami Research[1] to collect precise data on the locations of the seismic events responsible for these natural disasters (see Fig. 2.3b). By pinpointing the locations of these events, oceanographers are able to use bathymetry (ocean floor mapping) to evaluate the trajectories of randomly focused pressure waves that strike coastlines (Degueldre et al., 2016). Other systems of relevance to branched flow can hardly compete with the urgency of this one, given the devastating consequences that tsunamis can have and the lives that can be saved by predicting their contact points on land prior to their arrival.

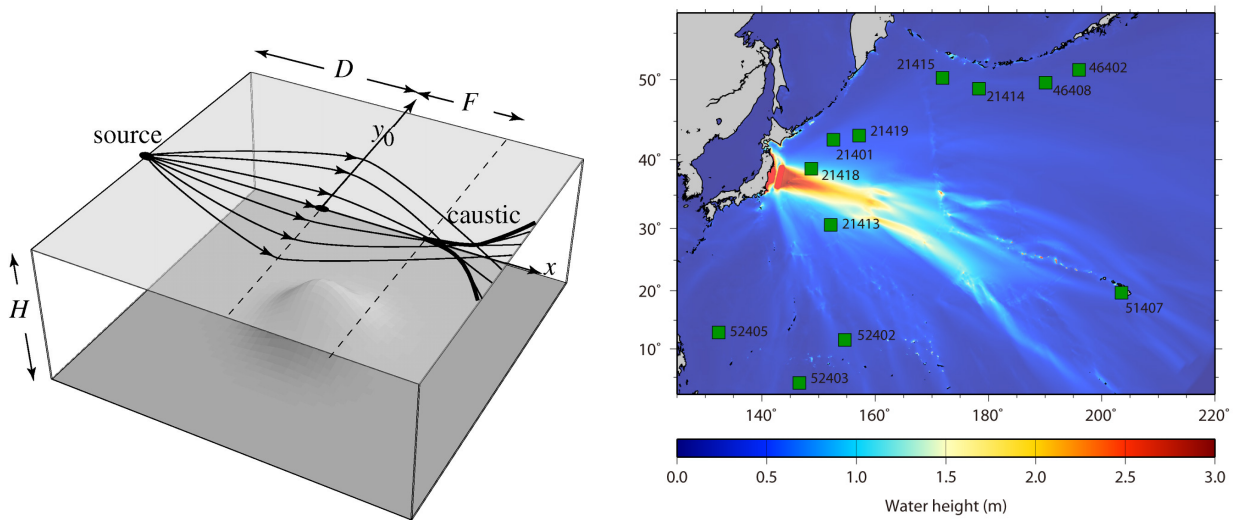

(a) A bulge (lens) in the ocean floor that leads to a focusing of tsunami pressure waves (Berry, 2007).

(b) The wave-height map of a tsunami near Tohoku, Japan in 2011 (NOAA Center for Tsunami Research).

Figure (2.3) The process by which tsunamis are focused and amplified by ocean floor topology (a) and the branched flow that results (b).

2.2.5 A Laser In a Soap Bubble

The first experimental confirmation of branched flow in optics was announced by Patsyk et al. (2020) in June 2020. The experiment involved the injection of a 532-nm laser beam, generated by a single-mode fiber, into a soap film. The pattern that emerged was a beautifully familiar network of scattered filamentary structures (see Fig. 2.4). It is worth noting that this breakthrough was a purely experimental one since branched flow had been studied theoretically and simulated in optics prior to this result (Brandstötter et al., 2019).

[1] Visit https://nctr.pmel.noaa.gov/honshu20110311/ for additional figures and animations.

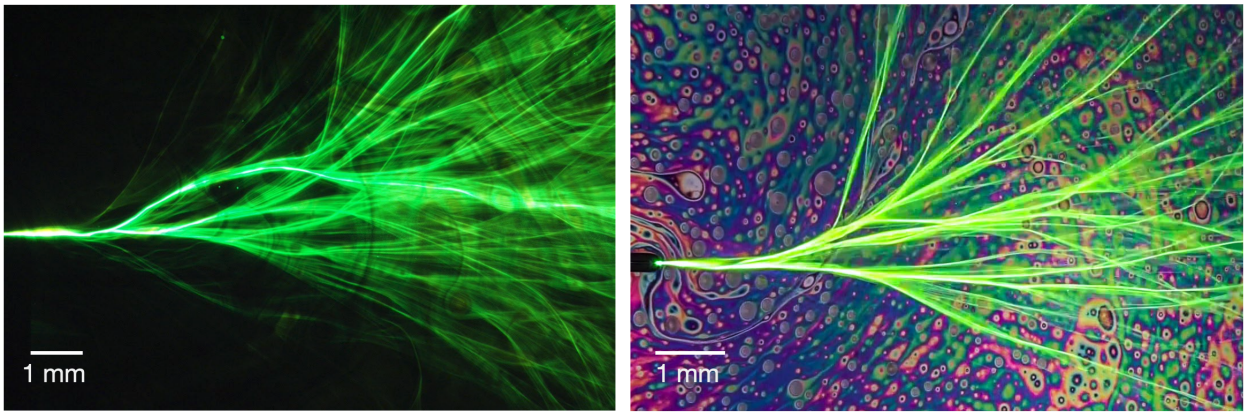

Figure (2.4) The injection of a laser beam into a soap film, marking the first experimental discovery of branched flow in optics (Patsyk et al., 2020). The right-most figure shows the soap film thickness variations (illuminated by white light) that weakly refract the laser beam and produce stable branches.

Chapter 3

The Physics of Branched Flow

3.1 Waves to Rays: Background, Definitions, and Assumptions

3.1.1 The Geometric "Ray" Limit

Branched flow is fundamentally a theory of wave propagation that can be described in a three-dimensional Cartesian coordinate system $\mathbf{x} = (x, y, z)$ by the linear wave equation:

$$\frac{1}{c^2}\frac{\partial^2\psi(\mathbf{x};t)}{\partial t^2} = \nabla^2\psi(\mathbf{x};t), \tag{3.1}$$

where $\partial^2\psi(\mathbf{x};t)$ represents an arbitrary complex scalar wave function, and z marks the range downstream from the wave source. As we know from countless experiments and simulations, the basic ingredients needed to produce branched flow are two-fold: (1) a wave source and (2) a weak, correlated, random medium. With this in mind, we can make a number of approximations that will greatly simplify our problem. The first is the *paraxial* (or *small-angle*, or *quasi-two-dimensional*) approximation, which extends from the presupposition that any medium in which branched flow occurs will be weakly refracting. Consequently, all scattering that occurs upon interaction with that medium will be almost entirely longitudinal (along the propagation axis) and in the forward direction (towards the observer). Branched flow is thus classified as a *small-angle forward scattering* problem. For problems of this kind, the wavefield can be simplified by making the parabolic approximation to the one-way Helmholtz equation (Tatarskii, 1961), which takes the form:

$$\frac{i}{k}\frac{\partial}{\partial z}A(\mathbf{x},z) = \left[-\frac{1}{2k^2}\nabla_\perp^2 - \delta n(\mathbf{x},z)\right]A(\mathbf{x},z), \tag{3.2}$$

for wavefield amplitude $A(\mathbf{x},z)$, refractive index $n(\mathbf{x},z) = \sqrt{\varepsilon}$ in a homogeneous isotropic medium[1], wavenumber k, and location downstream z. Our final assumption, which is warranted in all physical contexts treated here (including the upcoming astrophysical case), is that the propagating waves oscillate rapidly with respect to the oscillations within the medium, or equivalently that the wavelength λ is much smaller than

[1] This assumption can be relaxed.

the correlation length l_c (average irregularity size) of the medium. With this approximation, entered by taking the so-called "ray limit" $\lambda \to 0$ ($k \to \infty$), we find ourselves in a purely geometrical formulation of wave propagation known as *geometrical optics* in which waves are treated as rays (Kline & Kay, 1965; Born & Wolf, 1999; Stavroudis, 2006). One consequence of formulating wave propagation this way is that both diffraction and interference effects are ignored. Still, for our purposes, this approximation is valid.

3.1.2 Geometrical Optics

As stated previously, refractive wave propagation can modeled by tracing the paths of rays as they propagate through a medium $n(\mathbf{x}, z)$. Ray propagation is governed by the following ray equations:

$$\begin{aligned} \frac{dx}{dz} &= \theta, \\ \frac{d\theta}{dz} &= \nabla \delta n(x(z); z), \end{aligned} \tag{3.3}$$

where, in the paraxial approximation, the cross-range or transverse positions of rays can be parameterized by the longitudinal coordinate z. In continuous or *extended* media, a single ray's angle and position between the source ($z = 0$) and observer at z are calculated by integrating over n along the path:

$$\theta(z) = \int_0^z dz' \nabla \delta n \left(\mathbf{x}\left(z'\right); z' \right) \tag{3.4}$$

$$\mathbf{x}(z) = \int_0^z dz' \left(z - z'\right) \nabla \delta n \left(\mathbf{x}\left(z'\right); z' \right). \tag{3.5}$$

When a ray is refracted away from the straight-line path connecting the source to the observer, it accumulates a geometrical phase delay $\phi(\mathbf{x})$, defined as:

$$\phi(\mathbf{x}) = \int_0^z k \delta n \left(\mathbf{x}\left(z'\right); z' \right) dz'. \tag{3.6}$$

This is easily generalized to a wavefront of initially parallel rays that propagate from their source to an observer, each encountering entirely different refractive indices in the intervening medium that lead to different phase delays (path lengths) and distort the wavefront over time.

3.1.3 Hamiltonian Optics

Hamiltonian optics is a formulation of geometrical optics that offers a particularly intuitive picture of waves as ray paths by way of a mechanical analog to the motion (flow) of particles. There are many useful features of this formalism, including the ability to represent caustic geometries in phase space, analyze the stabilities of respective ray paths, and characterize wave propagation using chaos theory. In the Hamiltonian formulation, the accumulated geometrical phase in Eq. 3.6 can be related directly to the integrated potential $V(\mathbf{x}, z)$ traversed by a charged particle:

$$U(\mathbf{x}) = \int_0^z V(\mathbf{x}; z')dz', \tag{3.7}$$

$$\phi(\mathbf{x}) = \frac{k}{mv_z e^2} U(\mathbf{x}). \tag{3.8}$$

For the moment, we will work in one transverse dimension x rather than two ($\mathbf{x}$), so we can get a better handle on the formulation. We begin with a Hamiltonian (assuming a "particle" or ray of unit mass $m = 1$):

$$\mathcal{H} = \frac{p_x^2}{2} + V(x; z), \tag{3.9}$$

and derive Hamilton's equations:

$$\begin{aligned} \frac{dx}{dz} &= \frac{\partial \mathcal{H}(x, p_x; z)}{\partial p_x}, \\ \frac{dp_x}{dz} &= -\frac{\partial \mathcal{H}(x, p_x; z)}{\partial x}. \end{aligned} \tag{3.10}$$

These can be straightforwardly simplified as functions of the potential $V(x; z)$:

$$\begin{aligned} \frac{dx}{dz} &= p_x, \\ \frac{dp_x}{dz} &= -\frac{\partial V(x; z)}{\partial x}, \end{aligned} \tag{3.11}$$

where x is the transverse location of the ray and p_x is the ray momentum such that each trajectory is defined by $(x(z), p_x(z))$. Geometrically speaking, the ray momentum can be defined as $p_x = \tan(\theta)$, or simply θ in the small-angle approximation.

3.1.4 Lagrangian Manifolds

Another useful tool with which we can describe ray momentum is the classical action $\mathcal{S}(x)$, which is effectively equivalent to the ray path length and defined as the integral along the path:

$$\mathcal{S}(x) = \int_0^z \mathcal{L}(x; z)dz, \tag{3.12}$$

where the $\mathcal{L}$ is the classical Lagrangian. In this representation, we can introduce the Hamilton-Jacobi Equation (Metzger et al., 2010a), a nonlinear partial differential equation (PDE) that is fundamental to all Hamiltonian flows:

$$\frac{\partial}{\partial z}\mathcal{S}(x; z) + \mathcal{H}\left(x, \frac{\partial \mathcal{S}}{\partial x}; z\right) = 0, \tag{3.13}$$

within which $\mathcal{H}$ takes the form:

$$\mathcal{H}\left(x, \frac{\partial \mathcal{S}}{\partial x}; z\right) = \frac{1}{2}\left(\frac{\partial \mathcal{S}}{\partial x}\right)^2 + V(x; z). \tag{3.14}$$

Finally, we see that $S(x)$ is related to the ray momentum by:

$$p(x) = \frac{\partial \mathcal{S}(x)}{\partial x}, \tag{3.15}$$

and that the Lagrangian $\mathcal{L}$ effectively represents a graph of $p(x)$ for a surface of points $(x, p(x))$ in phase space (Metzger et al., 2010b; Littlejohn, 1992). The gradual deformation of such a manifold can be seen in Fig 3.2.

3.2 Caustics and Branches

3.2.1 What Constitutes a Branch?

When waves propagate through random media, consecutive weakly refracting events can accumulate over many correlation lengths and lead to highly intense focusing regions downstream from the source. These focusing regions are known as caustics. Classically speaking, they are *singularities* where the flow density diverges. In geometrical optics, caustics are represented by the converging and bundling of rays. In two dimensions, there are often only two types of relevant caustic geometries: cusps and folds. Cusps are, broadly speaking, focusing regions at which large numbers of rays converge (usually imperfectly). Additionally, cusps generate or *organize* pairs of fold caustics (Nye, 1999) that propagate onward from the cusp. The formation of fold and cusp caustics can be seen in Fig. 3.1.

The most elementary definition of a *branch* is a region of heightened flow density that is bounded by two caustics on either side (Metzger et al., 2010b; White & Fornberg, 1998; Kaplan, 2002). The caustics that bound the branch are fold caustics, which are represented by turning points of the Lagrangian manifold, the evolution of which is shown in Fig. 3.2. The Lagrangian manifold shows how caustics can effectively act as wave-guides in phase space.

It is worth noting that, while branches also exist in three dimensions (shown in Fig. 3.3 and described in Chapter 6), the classes of stable caustics that produce them are broader and significantly more complex.

Caustics in Geometrical Optics

One complication that arises in the ray limit is the inability to properly model the gain of the wavefield at caustics (since they exhibit classically infinite amplitudes). For this reason, it is still not well understood how exactly ray amplitude statistics compare to wavefield intensity statistics (Wolfson & Tomsovic (2001); Colosi (2015)).

There are two good reasons for continuing with the ray description of wave propagation, however, particularly in the context of chaos theory. The first is that, in the physical contexts that bring about branched

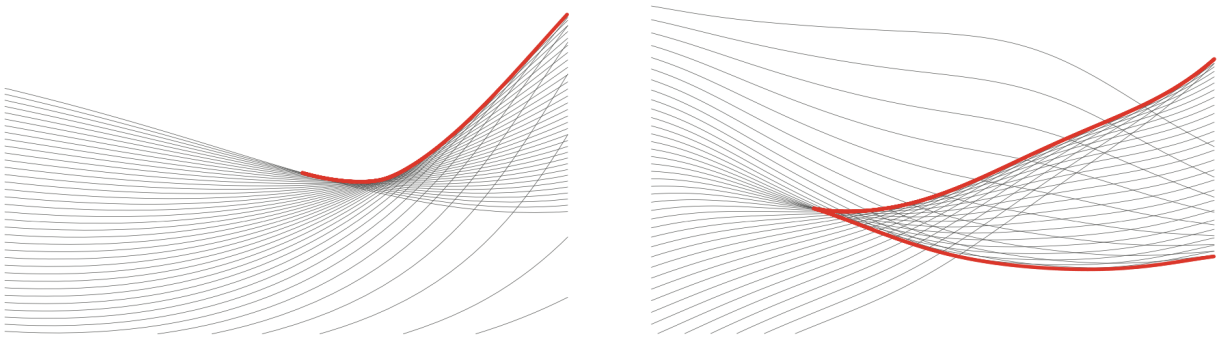

Figure (3.1) The caustic geometries relevant to wave propagation in two dimensions are the fold and cusp—which produces two folds (borrowed from J.J. Metzger's PhD Thesis (2010)).

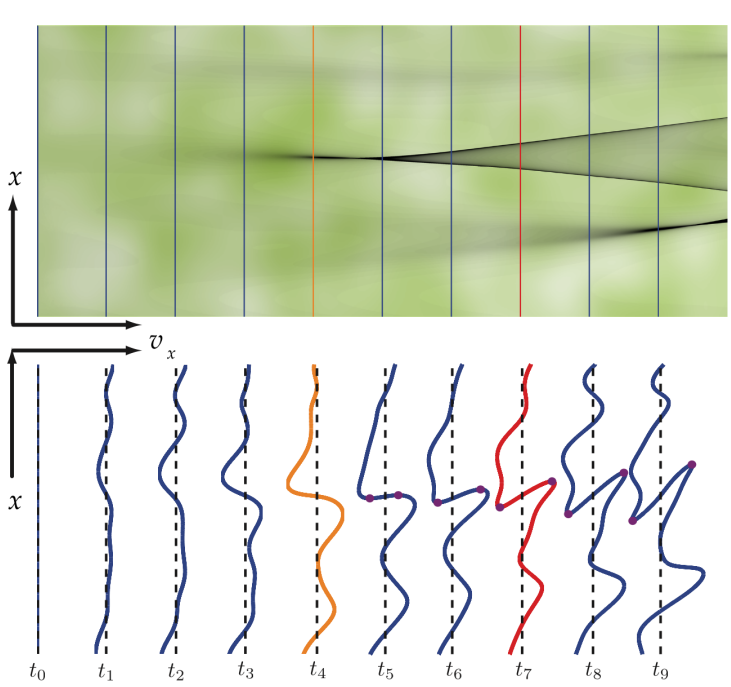

(a) The evolution of a Lagrangian manifold along two fold caustics emanating from a cusp.

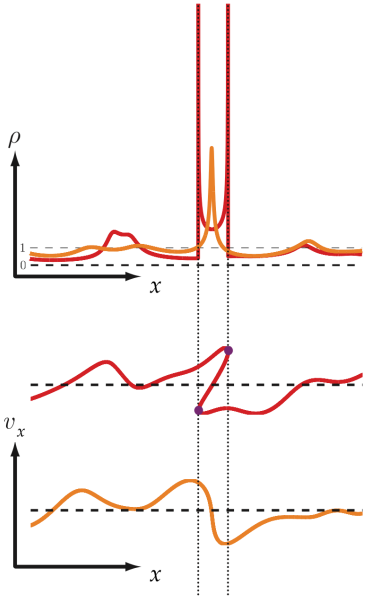

(b) The caustics (points of classically infinite ray density) occur at the turning points of the Lagrangian manifold.

Figure (3.2) The configuration space representation against the phase space representation of fold and cusp caustic formation (borrowed from J.J. Metzger's PhD Thesis (2010)).

flow, waves are consistently found to exhibit a ray-like character, suggesting that the approximation is generally valid. The second is that ray-tracing can provide crucial insights into the physics underlying chaotic wave propagation, which is much more difficult to extract through wave-optics (or wavefield) approaches.

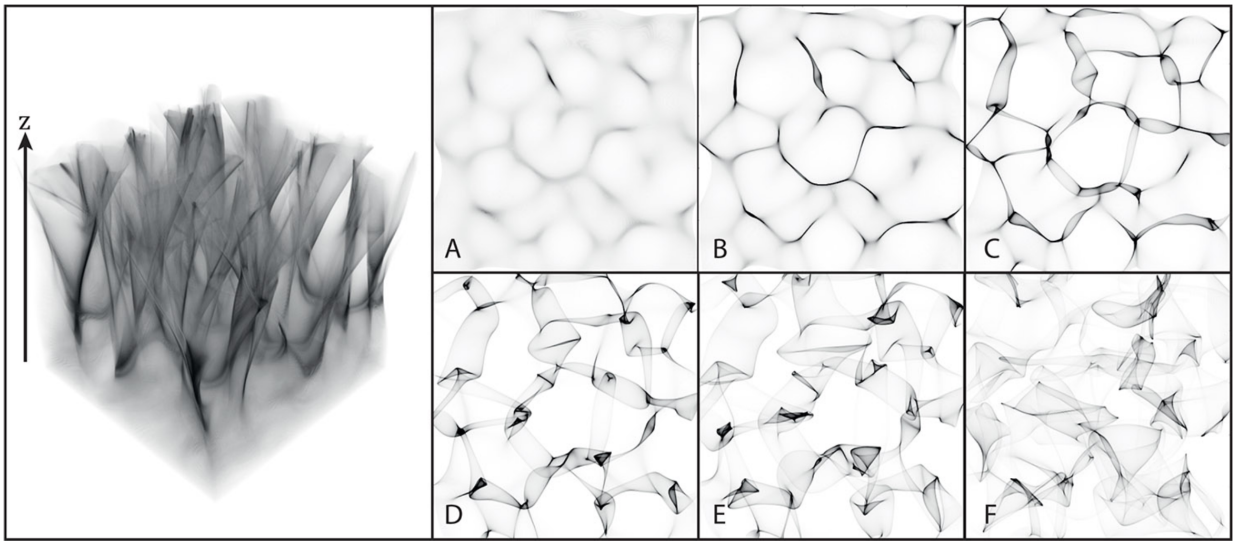

Figure (3.3) Branch flow in three dimensions (1 longitudinal and 2 transverse) (Heller et al., 2019a).

3.2.2 A Fundamental Spatio-Temporal Length Scale

The distance from a wave source to point at which the first cusp caustics are formed is a fundamental statistical property of wave propagation in weak, correlated random media (Metzger et al., 2010b, 2014), and is defined by the correlation length l_c and fluctuation strength (standard deviation) σ of the medium:

$$z_{peak} = \frac{l_c}{\sigma^{2/3}}. \tag{3.16}$$

While this length scale can be derived in more mathematically rigorous ways (Barkhofen (2013); Metzger et al. (2010b)) it does not require a tedious derivation.

Let's imagine the random walk of a ray away from a source with propagation direction z and cross-range x. If we assume a random walk with a binomial distribution, we can claim that, at each step, the ray will be deflected by a scattering angle $+\theta_s$ or $-\theta_s$—each with equal weighting. By the paraxial approximation, we assume that $\theta_s \ll 1$. For random walks, we know that the accumulated angular deflection after N steps will be $\theta_N \sim \sqrt{N}\theta_s$. Now, assuming that each step taken has length l, we can define the total distance traversed in the cross-range dimension at location z downstream as $x_z \sim (Nl)(\sqrt{N}\theta_s) \sim N^{3/2} l\theta_s$. We know that focusing occurs at a location z_0 when a ray covers a lateral distance of one correlation length, such that $x_z = l_c$. Finally, by setting $l = l_c$, we arrive at the following relation:

$$z_0 \sim \frac{l_c^{2/3} l^{1/3}}{\theta_s^{2/3}} = \frac{l_c}{\theta_s^{2/3}}. \tag{3.17}$$

Thus, we have successfully arrived at the 2/3 value in Eq. 3.16. The final step in tying together Equations 3.16 and 3.17 is recognizing that $\theta_s \propto \sigma$. Derivations quite similar to this one have been outlined by Voronovich (1994) and Heller (2018). Furthermore, as will be shown in Chapter 4, the onset of branching

at z_0 coincides with the first peak in the so-called "scintillation index", which measures the normalized intensity variance in the flow.

3.3 Branched Flow in Phase Space: Stability and Chaos

3.3.1 Stability and the Stability Matrix

Stability in chaotic wave propagation has been extensively explored within acoustic oceanography. An especially interesting phenomenon that occurs in the ray limit is that, while rays on their own tend towards chaos rather quickly, bundled rays (emanating from a point source with wide angular aperture, for instance) generally remain stable. This has been termed "manifold stability" (Cerruti & Tomsovic, 2002).

The stability matrix is a useful tool for studying dynamical systems, which describes the way a small neighborhood $\{\delta x, \delta p\}$, in phase space evolves for all ranges. We derive it by locally linearizing the ray dynamics (Heller, 1991). At a point z downstream, we have:

$$\begin{pmatrix} \delta p_z \\ \delta x_z \end{pmatrix} = M \begin{pmatrix} \delta p_0 \\ \delta x_0 \end{pmatrix}, \tag{3.18}$$

which leads to the stability matrix in terms of partial derivatives:

$$M = \begin{pmatrix} m_{11} & m_{12} \\ m_{21} & m_{22} \end{pmatrix} = \begin{pmatrix} \left.\frac{\partial p_z}{\partial p_0}\right|_{x_0} & \left.\frac{\partial p_z}{\partial x_0}\right|_{p_0} \\ \left.\frac{\partial x_z}{\partial p_0}\right|_{x_0} & \left.\frac{\partial x_z}{\partial x_0}\right|_{p_0} \end{pmatrix}. \tag{3.19}$$

As pointed out by Wolfson & Tomsovic (2001), the matrix element m_{21} comes up in the prefactor of Green's function of Eq. 3.2, and thus informs us of the varying wave amplitudes. The second-order partial derivative of V governs the evolution of M:

$$\frac{d}{dz}M = KM, \tag{3.20}$$

where:

$$K = \begin{pmatrix} -\frac{\partial^2 \mathcal{H}}{\partial x \partial p} & -\frac{\partial^2 \mathcal{H}}{\partial x^2} \\ \frac{\partial^2 \mathcal{H}}{\partial p^2} & \frac{\partial^2 \mathcal{H}}{\partial x \partial p} \end{pmatrix} = \begin{pmatrix} 0 & -\frac{\partial^2 V}{\partial x^2} \\ 1 & 0 \end{pmatrix}, \tag{3.21}$$

and the initial condition of M is the identity matrix I. The stability matrix can be used to distinguish between stable and unstable (chaotic) motion and is, therefore, closely related to the Lyapunov exponent.

3.3.2 Chaos and Lyapunov Exponents

In chaos theory, the *Lyapunov exponent* ν_L is an important metric for determining the onset of chaos, or preservation of stable motion, and is related to the trace of the stability matrix. The trace of the stability matrix $\mathrm{Tr}(M)$ is range dependent and defined as the sum of the entries along the diagonal of M:

$$\mathrm{Tr}(M) = \sum_{i=1}^{2} m_{ii}. \tag{3.22}$$

If we diagonalize M by a linear similarity transform $M = L^{-1}\Lambda L$, where Λ is a matrix containing the eigenvalues of M along its diagonal:

$$\Lambda = LML^{-1} \Longrightarrow \begin{pmatrix} \lambda & 0 \\ 0 & \lambda^{-1} \end{pmatrix}, \tag{3.23}$$

$\mathrm{Tr}(M)$ remains invariant, so:

$$\mathrm{Tr}(M) = \mathrm{Tr}\left(L^{-1}\Lambda L\right) = \frac{\mathrm{Tr}(\Lambda)\,\mathrm{Tr}(L)}{\mathrm{Tr}(L)} = \mathrm{Tr}(\Lambda) = \lambda + \lambda^{-1}. \tag{3.24}$$

The Lyapunov exponent is a limit of the so-called *stability exponent* ν. The stability exponent characterizes the exponential rate at which infinitely close rays diverge from each other. Furthermore, ν can be defined using the $\mathrm{Tr}(M)$ as:

$$\nu = \frac{\ln|\,\mathrm{Tr}(M)|}{z}. \tag{3.25}$$

The stability exponent describes individual rays at points z downstream. The Lyapunov exponent is then defined as the limit of ν at infinite range:

$$\nu_L = \lim_{z\to\infty} \frac{\ln|\,\mathrm{Tr}(M)|}{z}, \tag{3.26}$$

at which point, $\mathrm{Tr}(M)$ is controlled by one eigenvalue $\lambda = \max\{\lambda_i\}$. Alternatively, we can define ν_L in terms of the momentum $p(z)$ as:

$$\nu_L = \lim_{z\to\infty} \frac{1}{z} \log_2 \frac{p(z)}{p(0)}. \tag{3.27}$$

It is important to note that the $|\,\mathrm{Tr}(M)|$ is a useful value for distinguishing stable from chaotic propagation. If $|\,\mathrm{Tr}(M)|$ is less than 2, the ray is considered stable; otherwise, it is unstable (or tending towards chaos). This can be visualized by sharp dips in the plots of $\mathrm{Ln}(|\,\mathrm{Tr}(M)|)$ shown in Figures 4.2 through 4.5. Furthermore, the Lyapunov exponent can be distinguished from the stability exponent by the fact that it is independent both of individual rays and individual realizations of potential $V(x;z)$ (Wolfson & Tomsovic, 2001).

3.3.3 The Kick-Drift Model

The kick-drift model is a discretized representation of wave propagation in phase space, alternatively referred to as the "thin lens" model, which will be explored in more depth in Section 6.1.2. Propagation is divided into discrete steps for a randomly changing potential (i.e. a δ-correlated interaction at every n^{th} step). This model is used most frequently in physical systems with longitudinal length scales that greatly exceed the transverse lengths scales. The mapping in phase space is a step-by-step (point-to-point) area-preserving map of the phase plane onto itself, defined by the kick step:

$$p_{n+1} = p_n - \left.\frac{dV_n(x)}{dx}\right|_{x=x_n}, \tag{3.28}$$

which is then immediately followed by a drift step:

$$q_{n+1} = q_n + p_{n+1}. \tag{3.29}$$

This model has been visually represented as a stretching and tangling of equidistant points over the course of many drifts in Photoshop by Tobias Kramer[2] (see Fig. 3.4). This is equivalent to the tangling of a Lagrangian manifold, as shown in the previous figures (see Fig. 3.2b), as well as upcoming simulations in Chapter 4 (see Figures 4.8 and 4.9). As will be discussed in Chapter 6, the kick-drift model is a phase space analog to the so-called "thin screen" model often used in astronomy.

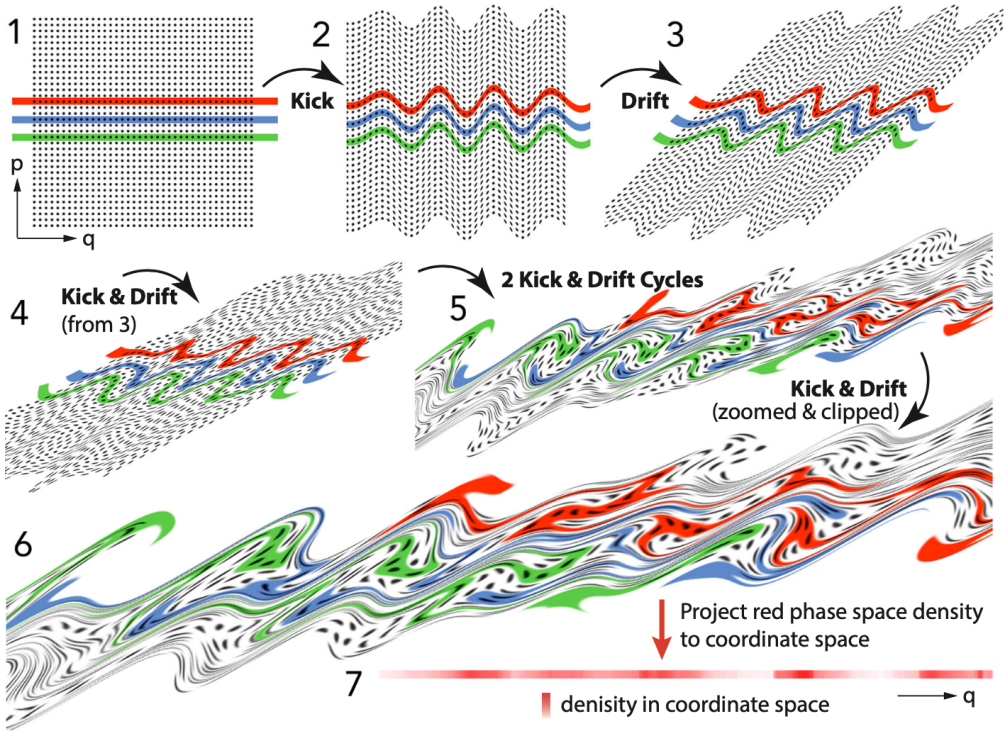

Figure (3.4) The Kick-Drift model visualized in Photoshop (see Heller et al. (2019b)).

[2] https://github.com/tobiaskramer/branch.git

Chapter 4

Branched Flow in Two Dimensional Random Media

Now that we have the necessary mathematical framework, we can begin to model branched flow in both extended and kick-drift (discretized) random media.

4.1 Two Dimensional Extended Potential

4.1.1 Gaussian Random Potential

To model a Gaussian random potential, we employed a method similar to that outlined by Bernhardt (2009), who generated a potential by superimposing a large number of Gaussian random plane waves. First, we defined a single plane wave:

$$V\left(x',y'\right) = Re\left[\frac{e^{i\left(\vec{k}\cdot\vec{x}'+\phi_j\right)}}{\sqrt{2\pi L}}\right] = \frac{1}{\sqrt{2\pi L}}\cos\left(kx'\cos\theta_j + ky'\sin\theta_j + \phi_j\right), \tag{4.1}$$

with wavenumber $k = 2\pi/L$, the wavelength (correlation length) L oriented at angle θ w.r.t. the x' axis, and a phase ϕ_j at the origin. The locations (x, y) throughout the potential are in dimensionless units of $L/2\pi$. Then, by superimposing N plane waves (where N is $\mathcal{O}(10^5)$ to satisfy the Central Limit Theorem (Bernhardt, 2009)), each defined by Eq. 4.1, with random angle θ_j and phase ϕ_j uniformly distributed on $[0, 2\pi]$, and uniform correlation length L, we obtained the potential:

$$V\left(x',y'\right) = \frac{\sigma}{\sqrt{2\pi L}}\sum_{j=1}^{N}\cos\left(kx'\cos\theta_j + ky'\sin\theta_j + \phi_j\right). \tag{4.2}$$

For dimensionless variables x and y, such that $x = kx', y = ky'$, and $k = 2\pi/L$, we can rewrite this expression as:

$$V(x,y) = \frac{\sigma}{\sqrt{2\pi L}} \sum_{j=1}^{N} \cos\left(x\cos\theta_j + y\sin\theta_j + \phi_j\right). \tag{4.3}$$

A realization of $V(x,y)$ for $L = 1$ is shown in Fig. 4.1.

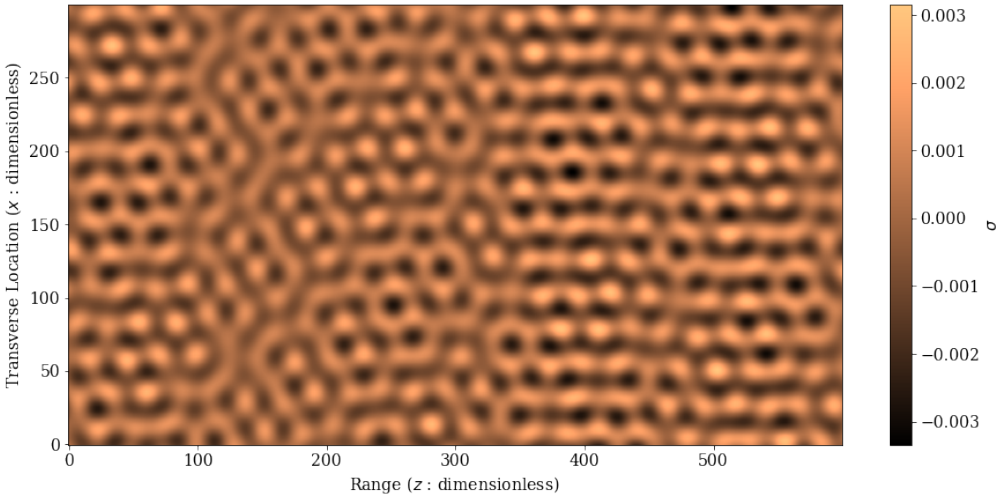

Figure (4.1) Gaussian random potential with zero mean and a standard deviation of $\sigma = 10^{-3}$, generated by superimposing 10^5 Gaussian random plane waves.

4.1.2 Scintillation Index

In the upcoming discussion, we will be referencing an important parameter known as the *scintillation index* (also referred to as the *modulation index* in contexts related to radio astronomy), which is defined as the normalized variance of the intensity:

$$\sigma_I^2(x) = \frac{\langle I^2 \rangle}{\langle I \rangle^2} - 1 = \frac{\text{variance of the intensity}}{(\text{ mean intensity })^2}. \tag{4.4}$$

4.1.3 Ray Propagation for a Point Source and a Plane Wave

To model branched flow in continuous potentials, we use Scipy's `odeint` to numerically solve Hamilton's equations as well as the elements in the stability matrix in Eq. 3.19. The results for a point source with zero-momentum and small spatial separation initial conditions are shown in Figures 4.2, 4.3, and 4.4. Each realization exhibits a potential of different scattering strength σ. The results for a plane wave are shown in Fig. 4.5, which also reveal the fundamental length scale $z_0 = l_c\sigma^{-2/3}$ that we derived in Section 3.2.2, which marks the onset of branching in weak, correlated random media.

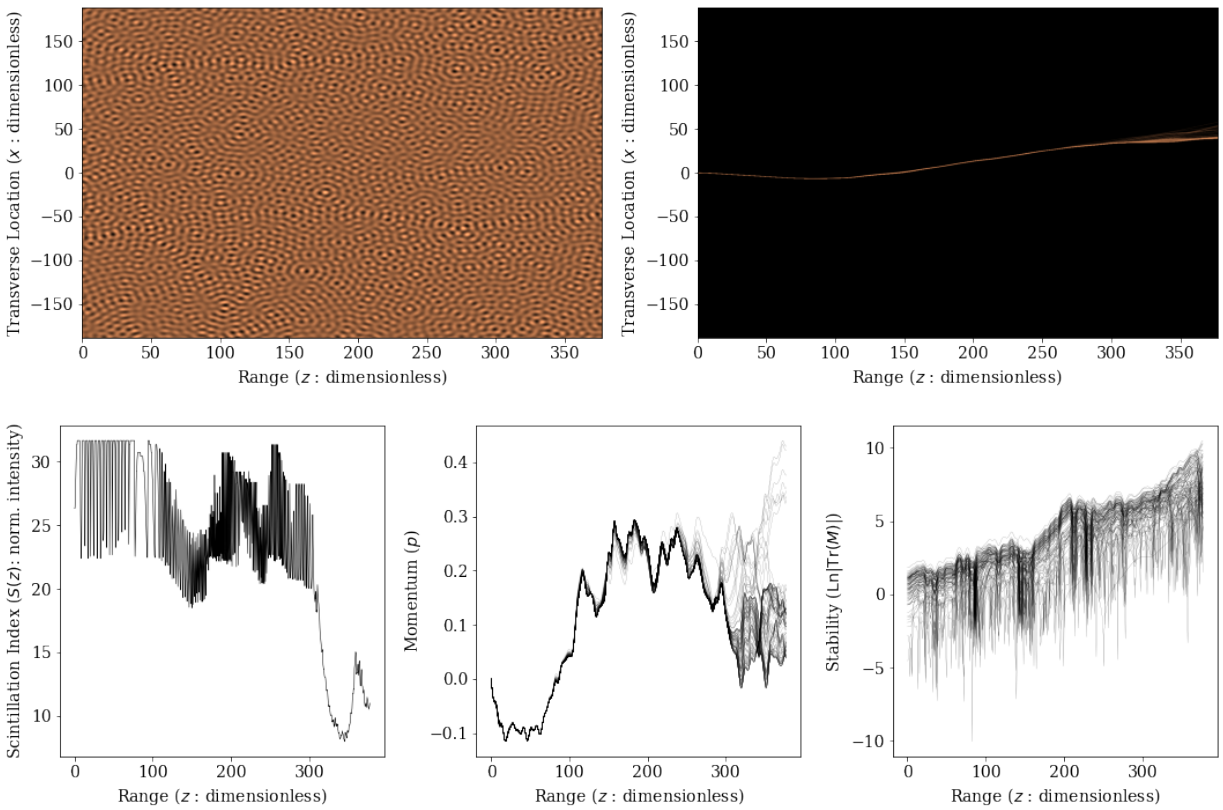

Figure (4.2) A refracted beam in a potential of $\sigma = 0.01$ exhibiting a scintillation index with peaks corresponding to focusing regions.

4.2 Two Dimensional Kick-Drift (Thin Screen) Potential

Alternatively, we can use straightforward geometric optics to propagate rays by collapsing the extended random medium into a series of delta-correlated scattering phase screens. To do so, we go a step further with the formulation in Section 3.1.2. We imagine that the intervening medium between the wave source and observer can be divided into N slabs, each of which is collapsed to an infinitely thin phase screen, such that the accumulated phase (see Eq. 3.6) of the i^{th} screen, which is added to the incident wave amplitude, is given by:

$$\phi_i(\mathbf{x}) = \int_{z_i'}^{z_{i+1}'} k\delta n(\mathbf{x}, z)dz. \tag{4.5}$$

The advantage of having multiple phase screens is that information on the ray paths, which are approximated in a piecewise linear fashion, is given by the multiple scattering from multiple screens. Importantly, however, when approximating a medium using phase screens, it is assumed that the approximation is improved by increasing the number of screens.

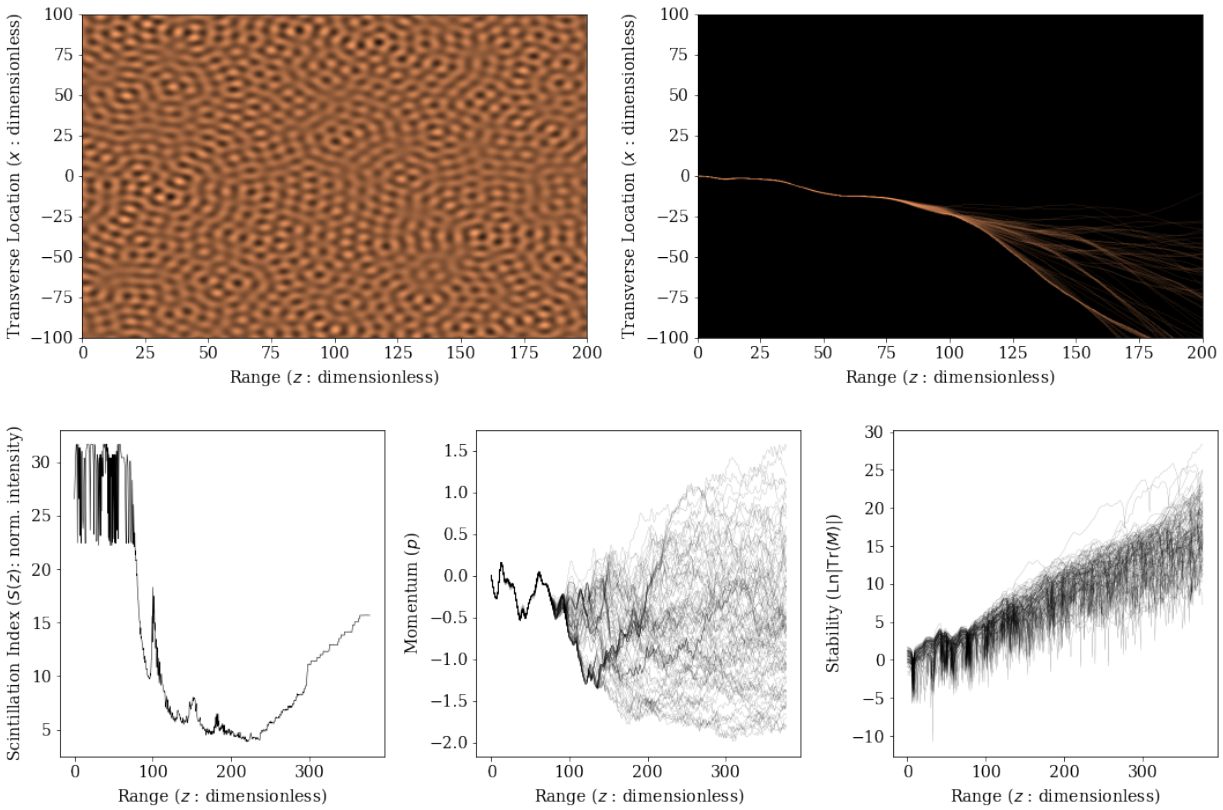

Figure (4.3) A refracted beam in a potential of $\sigma = 0.05$ exhibiting a scintillation index with peaks corresponding to focusing regions.

4.2.1 The Thin Phase Screen Model

As we know, wave propagation in random media is entirely governed by variations in the refractive index. These variations are what define the phase screen such that the statistics of the extended medium are also contained within its discretized elements. As we have noted, since the wave is expected to encounter a large number of uncorrelated density fluctuations, the phase delay is taken to be a Gaussian random variable by the Central Limit Theorem.

The correlation lengths of an extended medium divided into individual screens are preserved in the transverse dimensions and integrated out along the wave path for the longitudinal dimensions—thus, making the phase screens statistically independent.

The most common and useful measure of the phase statistics is the phase structure function $D_\phi(\mathbf{r})$, which measures how a quantity differs at separation $\mathbf{r}$:

$$D_\phi(\mathbf{r}) \equiv \left\langle [\phi(\mathbf{x}) - \phi(\mathbf{x}+\mathbf{r})]^2 \right\rangle = 2\phi_{rms}^2 \left(1 - \rho_\phi(\mathbf{r})\right). \tag{4.6}$$

This expression contains within it the phase autocorrelation function $\rho_\phi\,(r_x, r_y)$, defined as:

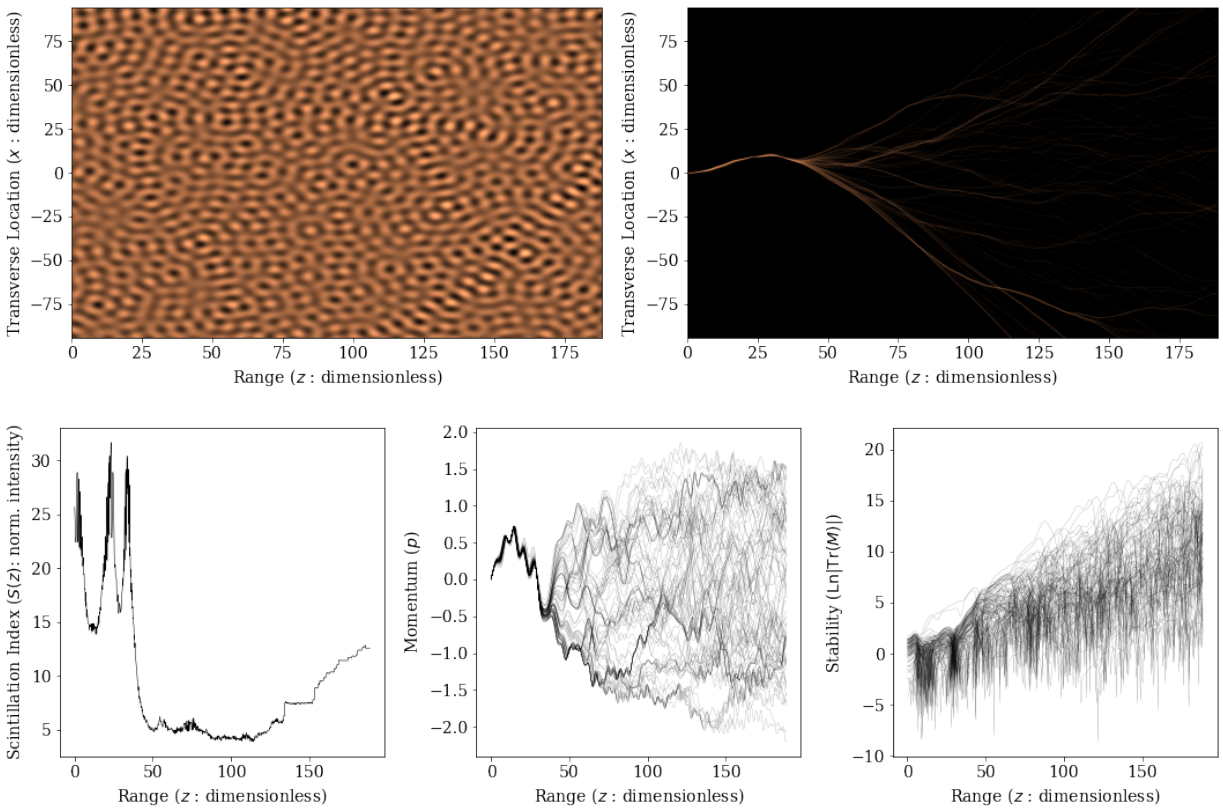

Figure (4.4) A refracted beam in a potential of $\sigma = 0.1$ exhibiting a scintillation index with peaks corresponding to focusing regions.

$$\phi^2_{rms}\rho_\phi\left(r_x, r_y\right)\delta_{ij} \equiv \left\langle \phi_i(x,y)\phi_j\left(x+r_x, y+r_y\right)\right\rangle, \tag{4.7}$$

where the angular brackets denote an ensemble averaging. It is also important to note that the direct relationship between phase delay and refractive index fluctuations $\delta n(x)$ (which have zero-mean) informs us that $\langle\phi\rangle = 0$.

The structure function is theoretically equivalent to the autocorrelation function but does not require a stationary mean (Ishimaru, 1978). Furthermore, it can be defined in terms of the wavenumber spectrum of the phase fluctuations, $\mid \phi(q)^2 \mid$:

$$D_\phi(\mathbf{r}) = \int d\mathbf{q}(1-\cos(\mathbf{q}\cdot\mathbf{r}))|\phi(\mathbf{q})|^2. \tag{4.8}$$

When dealing with a single irregularity scale, the phase power spectrum is Gaussian:

$$|\phi(q)|^2 \approx \exp\left(\frac{-q^2a^2}{2}\right), \tag{4.9}$$

where a is the size of a typical irregularity or "turbule". The associated phase structure function is defined as $D_\phi(r) = 2\phi_0^2(1-\rho(r))$ where $\rho(r) = \exp\left(-r^2/2a^2\right)$. A realization of a one-dimensional Gaussian phase

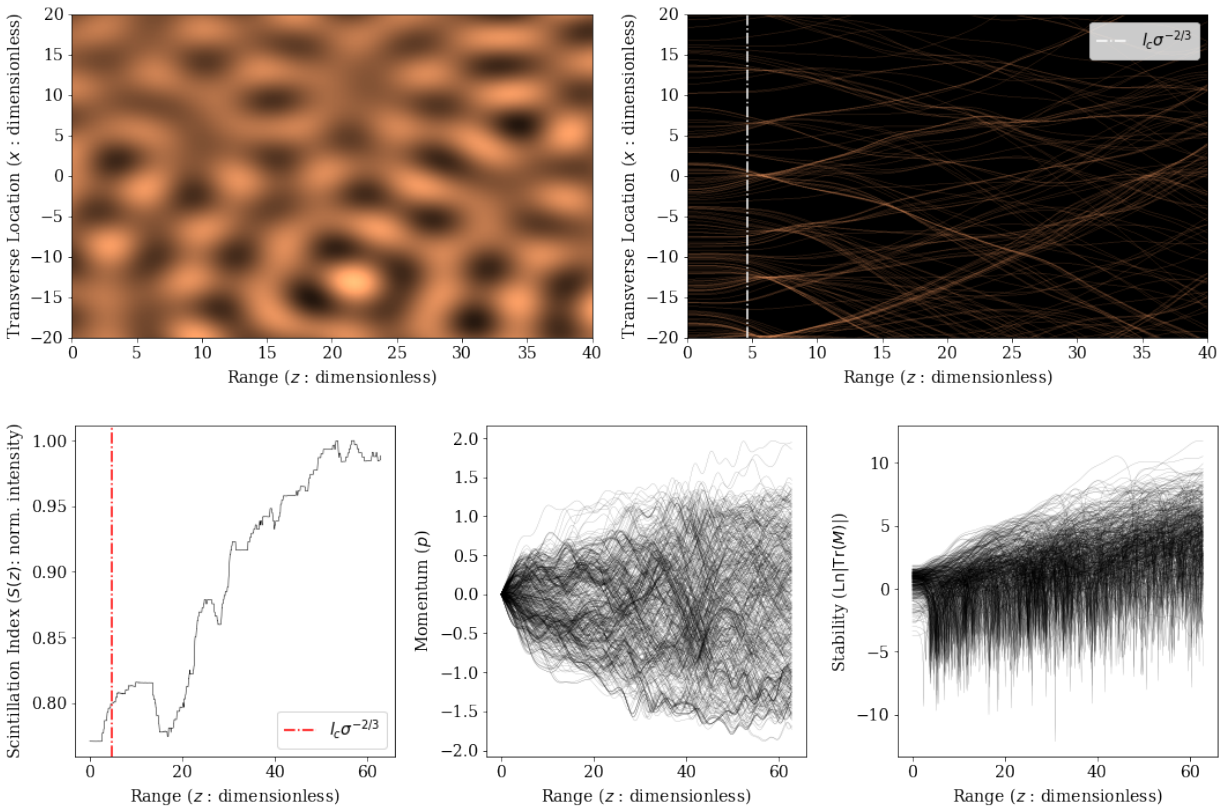

Figure (4.5) A refracted plane wave in a potential of $\sigma = 0.1$ exhibiting a scintillation index with peaks corresponding to focusing regions. The red line shows that the onset of branching indeed corresponds to the fundamental length scale derived in Section 3.2.2.

screen and its corresponding autocorrelation function are shown in Fig. 4.6.

As we will see in Chapter 5, random media like the Interstellar Medium, or even the Earth's atmosphere, are not always accurately modeled by a single irregularity scale a due to their turbulent nature. For media that exhibit a multi-scale range of irregularities, the phase power spectrum is taken to be a power-law with some spectral index (exponent) β:

$$|\phi(q)|^2 \approx q^{-\beta} \quad q_0 < q < q_1, \tag{4.10}$$

where q_0 represents a lowest wavenumber corresponding to a largest irregularity size, or *outer scale*, $L_0 = 2\pi/q_0$ and q_1 is the highest wavenumber corresponding the smallest irregularity size, or *inner scale*, $l_0 = 2\pi/q_1$. Power-law wavenumber spectra generally span between 3 and 5 orders of magnitude (Cordes et al., 1986). The ISM and other turbulent media in nature are often characterized by Kolmogorov turbulence, which is described by a power-law of $\beta = -5/3$ in one dimension. Realizations of a one-dimensional Kolmogorov phase screen and its corresponding wavenumber spectrum are shown in Fig. 4.7.

Other spectra that contain multiple spectral indices (so-called "hybrid" spectra) can also be considered

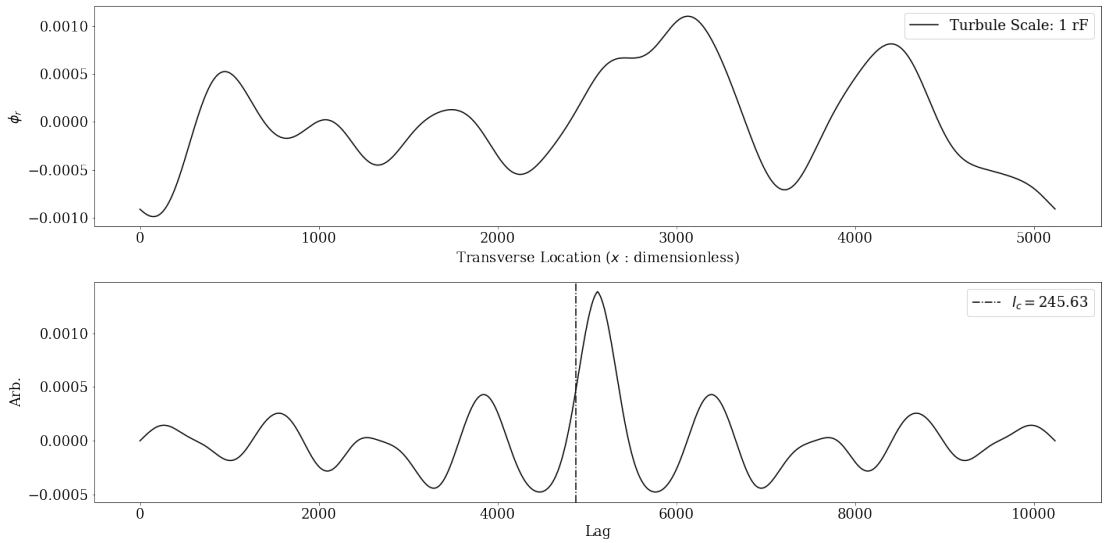

Figure (4.6) Realization of a one-dimensional Gaussian phase screen and auto-correlation function, which reveals the correlation length ("turbule size").

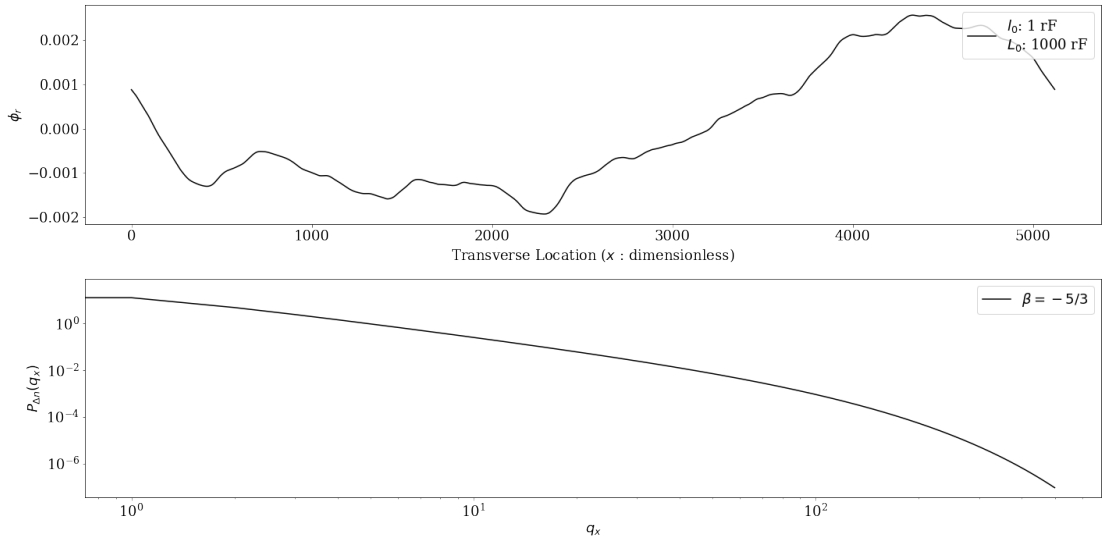

Figure (4.7) Realization of a one-dimensional Kolmogorov phase screen with spectral index $\beta = -5/3$, as shown by its wavenumber spectrum.

and have often been proposed as more accurate descriptions of turbulent media like the ISM. For the purposes of this work, however, we will focus solely on Gaussian and Kolmogorov media.

4.2.2 Ray Propagation for a Plane Wave

For the kick-drift model, a plane wave yields the most interesting phase space structure, so that is what we will focus on. When dealing with simple, one-dimensional screens, the propagation of rays is entirely

analogous to the kick-drift model shown in Section 3.3.3. The ray propagation is defined solely by the transverse phase gradient $\nabla_\perp \phi_i(\mathrm{x})$ of the screen, such that the scattering angle can be written as:

$$\Delta\theta_{\mathrm{i}}(\mathrm{x}) = \frac{1}{k}\nabla_\perp \phi_i(\mathrm{x}). \tag{4.11}$$

Here, both $\phi_i(\mathbf{x})$ and $\Delta\theta_{\mathbf{i}}(\mathbf{x})$ are assumed to be nonlinear functions of x. The rays are then propagated by the kick of the i^{th} screen:

$$\theta_{\mathrm{i}}(\mathrm{x}) = \theta_{\mathrm{i}-1}(\mathrm{x}) + \Delta\theta_{\mathrm{i}}(\mathrm{x}), \tag{4.12}$$

and subsequent drift to the next screen:

$$x_{i+1} = x_i + \Delta z_i \theta_i\left(x_i\right), \tag{4.13}$$

which scatters the ray again. The above kick-drift steps are iterated until the observer plane is reached. This can be thought of as a discretization of the analytic solution to the ray paths presented in Section 3.1.2 (see Equations 3.4 and 3.5).

By using the kick-drift model, we traced rays through 20 consecutive one-dimensional Gaussian and Kolmogorov phase screens, the results of which are shown in Figures 4.8 and 4.9, respectively. In addition to calculating the ray paths, we examined the progressive morphological tangling of their Lagrangian manifolds in phase space at each step to observe the onset of focusing (manifold folding). Due to the smaller scale irregularities in the Kolmogorov phase screens, the consequent phase space tangling appeared significantly more irregular than for the smoother Gaussian screens. In each figure, the phase screens are plotted over the ray paths as well to give an indication of their respective features and downstream affects.

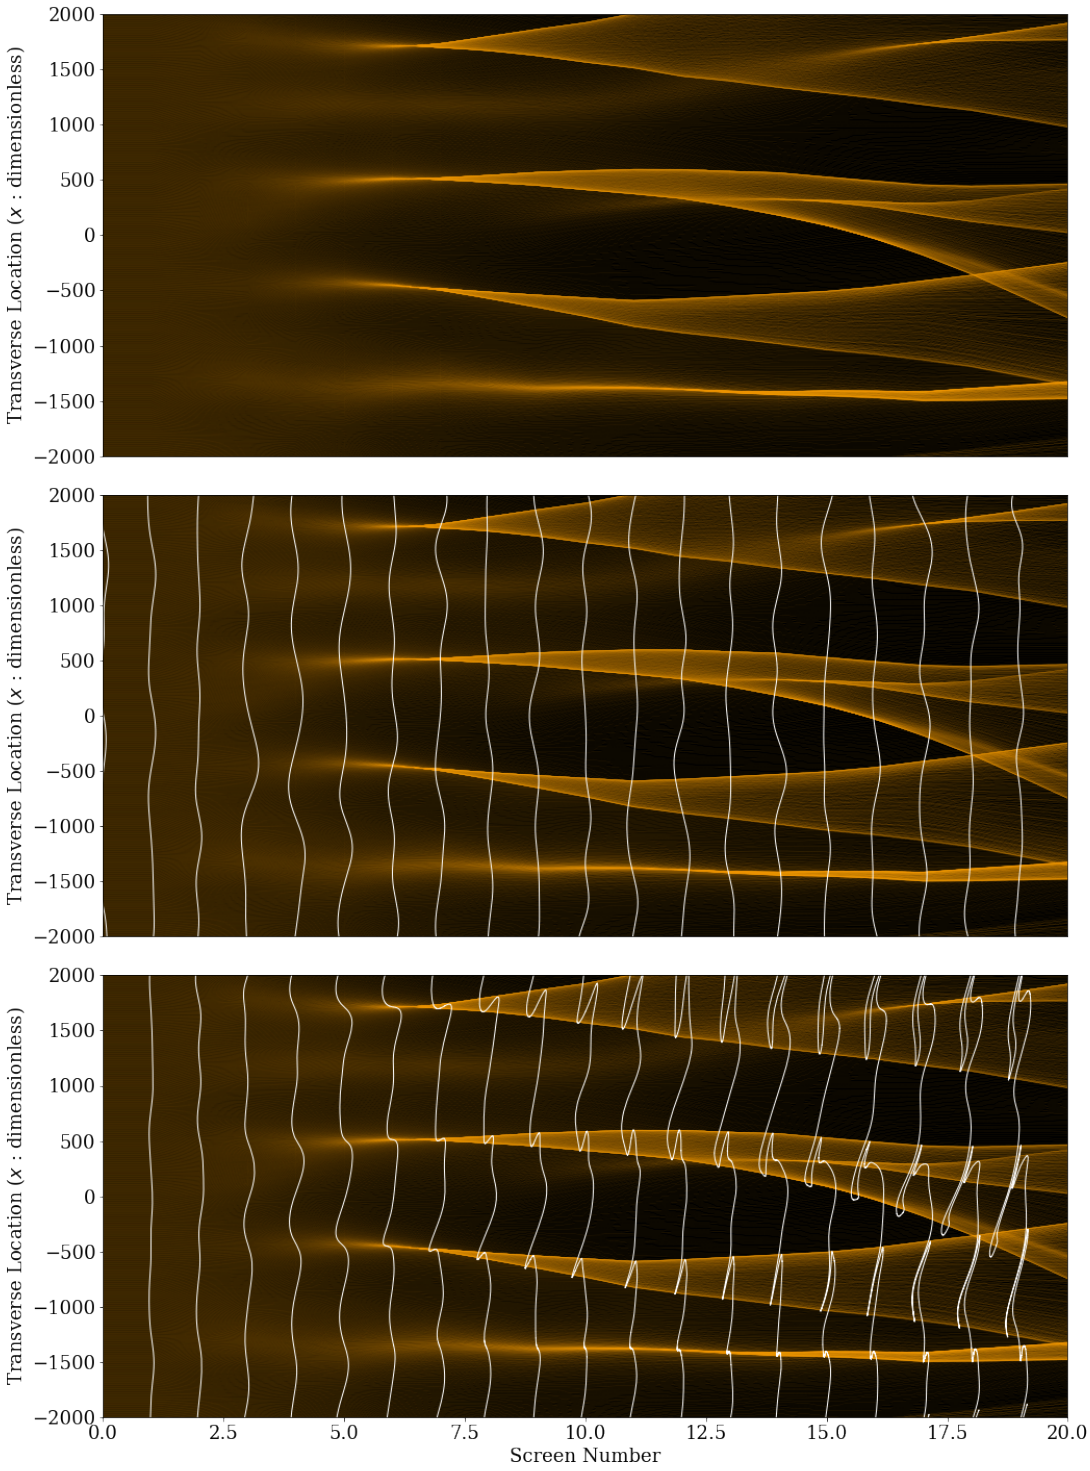

Figure (4.8) Ray tracing through 20 consecutive Gaussian random phase screens. The phase screens are shown in the middle figure and the progressive tangling of the Lagrangian manifolds in phase space can be seen in the lowest figure.

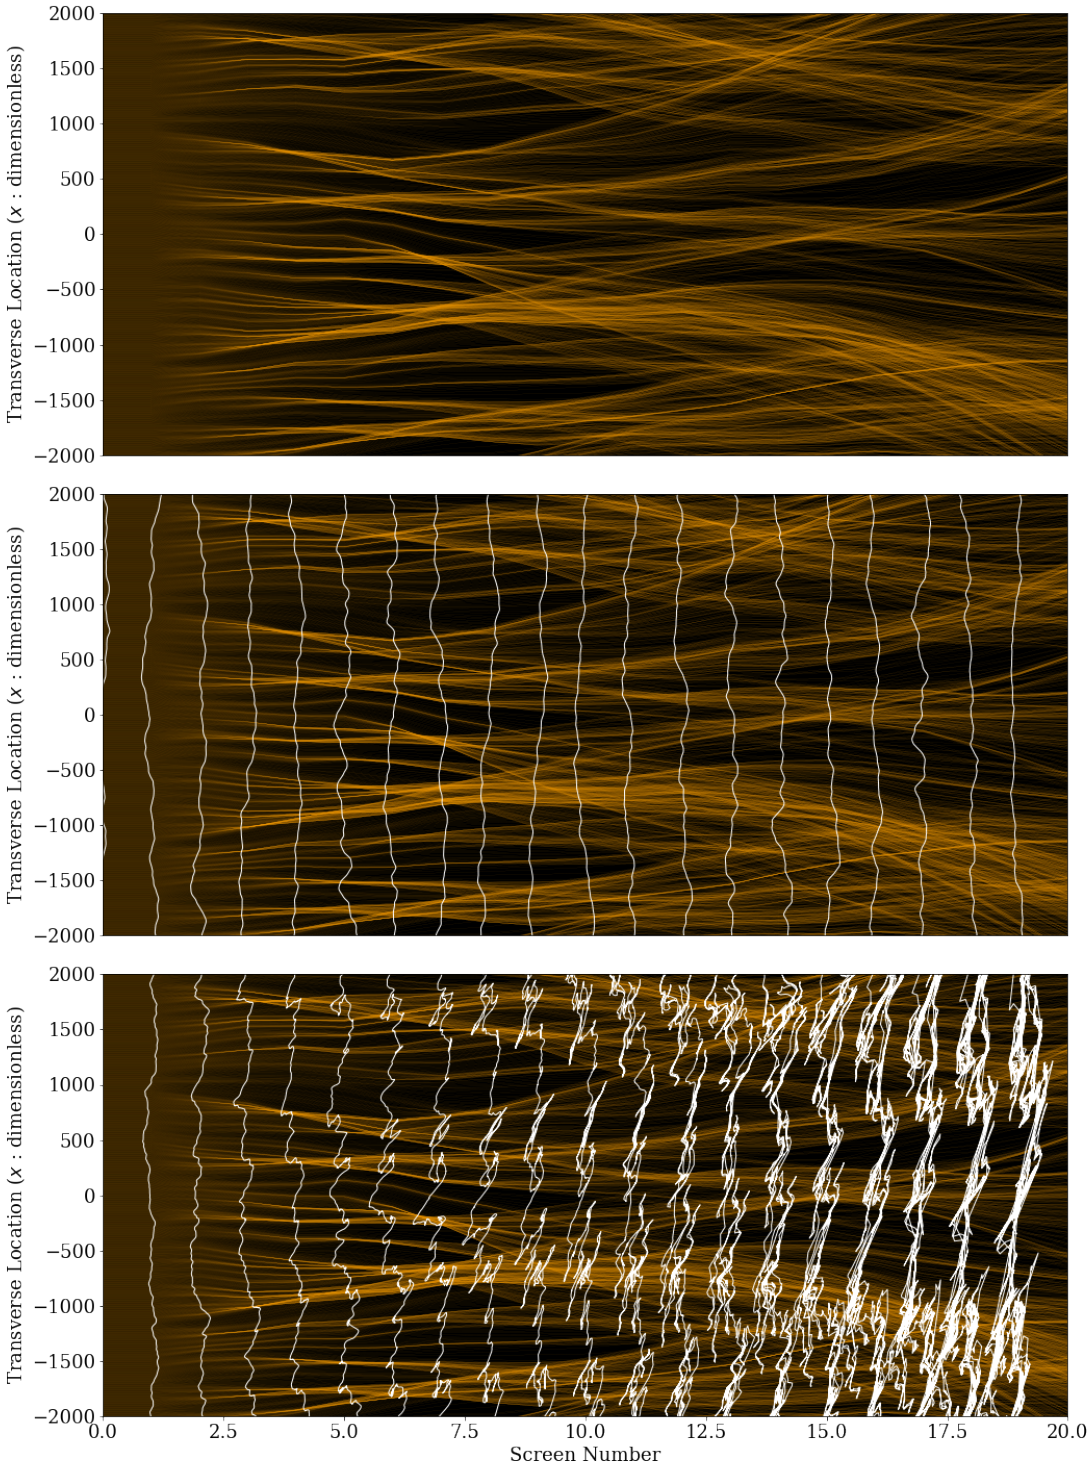

Figure (4.9) Ray tracing through 20 consecutive Kolmogorov random phase screens. The phase screens are shown in the middle figure and the progressive tangling of Lagrangian manifolds in phase space can be seen in the lowest figure. The phase space morphologies are far more complex for Kolmogorov media due to multi-scale irregularities.

Chapter 5

The Astrophysical Picture: The Case for Interstellar Branched Flow

5.1 The Interstellar Medium

Radio waves from compact sources are influenced by a number of turbulent ionized media on their journey to telescopes on Earth. These media include the interstellar medium (ISM), the interplanetary medium, and the Earth's ionosphere. The interstellar medium (ISM) is a term that refers to all ordinary matter that exists between stars in galaxies (such as the Milky Way). This includes plasma, dust, gas, and cosmic rays. In this discussion, we will focus on the ISM—specifically, the regions (phases) containing free electrons and magnetic fields that comprise the ionized plasma in the ISM, known as the ionized interstellar medium (IISM) (henceforth, implied by the use of ISM).

There are three effects of the interstellar medium that are important to note. The first is dispersion, which is the frequency-dependent group-velocity perturbations that radiation experiences in an ionized medium. The second is Faraday rotation, which is a polarization-dependent fluctuation in the refractive index of a magnetized ionized medium. The third and final effect is interstellar scattering, which can be either diffractive or refractive and occurs as a consequence of fluctuations in electron density of the ionized plasma. A simple schematic of the large scale (refractive) and small-scale (diffractive) scattering of a plane wave emitted by a coherent source in the ISM is shown in Fig. 5.1. As a wavefront is scattered, the phase and amplitude modulations that occur can lead to a variety of phenomena including interstellar scintillation, angular broadening, and pulse smearing (for pulsed sources such as pulsars). The observational consequences of these effects offer unique insights into the geometry and composition of the ISM. For a more in-depth discussion of interstellar scattering, the reader is referred to Rickett (1977).

The electron density fluctuations in the ISM have been described by both Gaussian and power-law (e.g.,

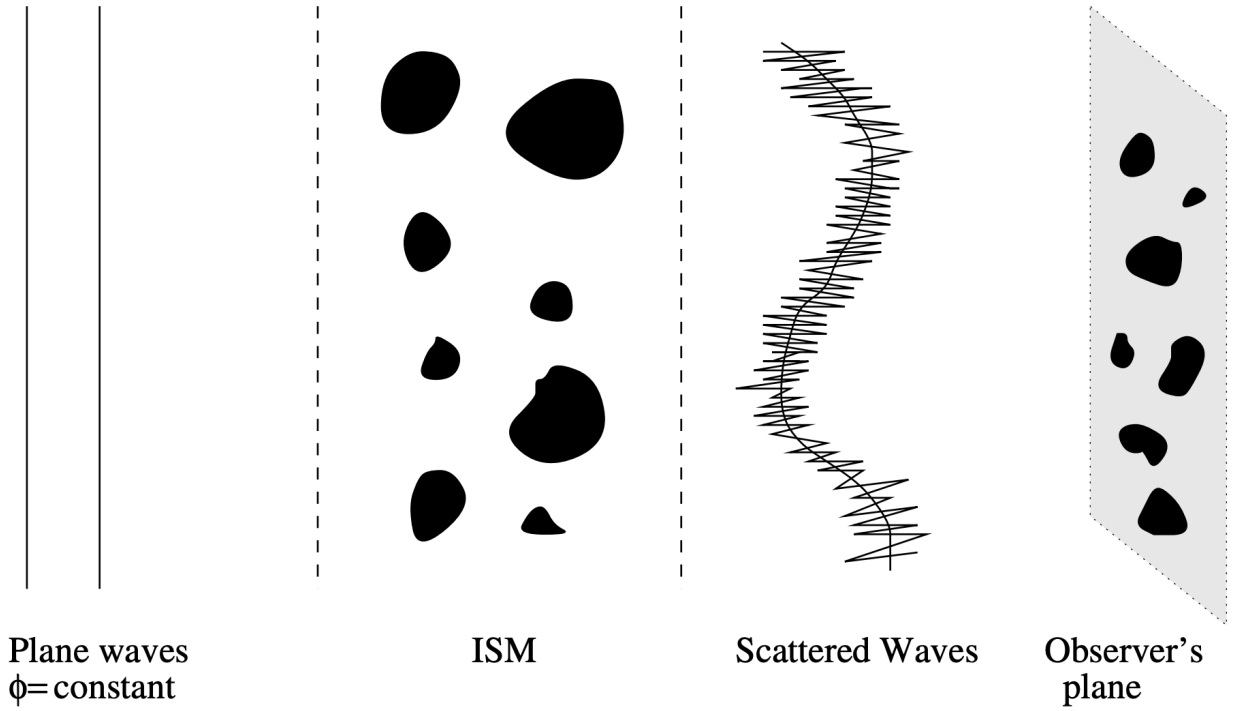

Figure (5.1) Schematic of diffractive (small-scale) and refractive (large-scale) scattering of a plane wave propagating through the ISM towards an observer (Bondi, 2005).

Kolmogorov) irregularity spectra. The most straightforward example of a random medium is described by a Gaussian irregularity (power) spectrum. As we know from previous studies, this has been repeatedly confirmed to be a suitable medium for branched flow if the statistics allow. The ISM has at times been modelled by a Gaussian irregularity spectrum, although most consider it to be better characterized by a power-law irregularity spectrum—specifically, that of Kolmogorov turbulence (Coles et al., 1987).

5.2 Pulsars in the Milky Way

Pulsars, in addition to being sources with pulsed emissions, can be idealized as point sources. As such, they have been identified by Heller et al. (2019b) as ideal candidates for branched flow. Here, I will argue that pulsars in the ISM do indeed meet the criteria necessary for branched flow to occur.

Effectively, what we are trying to do is apply certain principles of pulsar astrophysics to the heuristic outlined in Section 3.2.2. We begin with an expression for the scattering strength of a medium, defined by Patsyk et al. (2020) as the ratio between the standard deviation of the potential and the energy:

$$\nu_0 = \frac{1}{2}\sqrt{\frac{\langle n^4 \rangle}{\bar{n}^4} - 1}, \tag{5.1}$$

where n is the refractive index of the medium (in this case, the ISM) with a non-zero mean, $\bar{n}$ is the mean

refractive index, and the brackets represent ensemble averaging.

If we assume that the refractive index profile of the ISM is Gaussian, such that $n = 1 + \sigma_n N(0,1)$, then we know that, for a normal distribution: $\langle n^4 \rangle = \mu^4 + 6\mu^2\sigma^2 + 3\sigma^4$. If we claim that $\mu = \bar{n} = 1$, we arrive at an expression for the ISM scattering strength in terms of the standard deviation σ_n of its fluctuation (irregularity) scale:

$$\nu_0 = \frac{1}{2}\sqrt{6\sigma_n^2 + 3\sigma_n^4} \approx \sqrt{3/2}\sigma_n, \tag{5.2}$$

where the final approximation is arrived at by assuming $\sigma_n \ll 1$. For now, as we assumed in the heuristic in Section 3.2.2, we will claim $\nu_0 = \sigma_n$.

For a plasma in the ISM, we know that the refractive index is approximately (Lyne & Graham-Smith, 2012):

$$n \approx 1 + \frac{n_e r_e \lambda^2}{2\pi}, \tag{5.3}$$

where n_e is the electron density, r_e is the classical electron radius, and λ is the observing wavelength. We can now redefine our scattering strength as:

$$\nu_0 \approx \sigma_n = \frac{\delta n_e r_e \lambda^2}{2\pi}. \tag{5.4}$$

5.3 Focusing Power of the Interstellar Medium

Now we are in a position to determine whether the cumulative scattering power in the ISM is sufficient to achieve $D_0 < D$ along a particular line of sight, which we know to be a prerequisite for branched flow.

Consider, as we have before, the total angle that is reached after propagating through the entire extended, homogeneous ISM. The r.m.s. scattering angle will be:

$$\Delta\theta \approx \sqrt{N}\,\theta_0 = \left(\frac{D}{a}\right)^{1/2}\theta_0. \tag{5.5}$$

But, since $D = Na$ by assumption (i.e., number of irregularities encountered), $D_0 < D$ translates to:

$$\Delta\theta > \theta_0^{2/3}. \tag{5.6}$$

Note that $N = \theta_0^{-2/3}$ in this model.

For the ISM, we know from Lyne & Graham-Smith (2012), that an irregularity or "turbule" of size a is related to the r.m.s. phase delay by:

$$l_c = \frac{\Delta\phi_{rms}}{(r_e \epsilon \bar{n}_e)^{2/3}\,\lambda^{1/3}(2\pi)^{1/3}}. \tag{5.7}$$

We would like to put this in terms of a scattering angle θ_s, so we first rewrite D_0 in terms of l_cL

$$D_0 = \frac{l_c}{\left(\epsilon \bar{n}_e r_e \lambda^2 / 2\pi\right)^{2/3}}, \tag{5.8}$$

and see that θ_s can be written as:

$$\Delta\theta_s = \frac{1}{2\pi}\left(\frac{D}{a}\right)^{1/2} r_e \Delta n_e \lambda^2. \tag{5.9}$$

Rearranging, we arrive at an expression for D:

$$D = a\left[\frac{2\pi\Delta\theta_s}{r_e \epsilon \bar{n}_e \lambda^2}\right]^2. \tag{5.10}$$

As we already established, for focusing to occur, the inequality $D_0 < D$ must hold. Using this, we can remove a from the problem:

$$\left(\epsilon \bar{n}_e r_e \lambda^2 / 2\pi\right)^{2/3} < \left[\frac{2\pi\Delta\theta_s}{r_e \epsilon \bar{n}_e \lambda^2}\right]^2, \tag{5.11}$$

And, finally, we arrive at an expression for θ_s in terms of the refractive index:

$$\Delta\theta_s > \frac{\epsilon \bar{n}_e r_e \lambda^2}{2\pi}\left[\frac{2\pi}{\epsilon \bar{n}_e r_e \lambda^2}\right]^{1/3}, \tag{5.12}$$

which simplifies to:

$$\Delta\theta_s > \left(\frac{\epsilon \bar{n}_e r_e \lambda^2}{2\pi}\right)^{2/3}. \tag{5.13}$$

The final step is to introduce a concept known as *scattering measure* (SM), which is a measure of the scattering strength of the ISM frequently used in radio astronomy. SM is defined differently for Galactic and extragalactic sources (Cordes et al., 1991). Assuming a three-dimensional Kolmogorov spectral index of $\beta = 11/3$, SM is a frequency dependent measure that scales with distance and angular size θ as follows:

$$\mathrm{SM} = 1\mathrm{kpcm}^{-20/3}\nu_{\mathrm{GHz}}^{11/3}\begin{cases}\left(\dfrac{\theta_{\mathrm{FWHM}}}{128\mathrm{mas}}\right)^{5/3}, & \text{extragalactic source;}\\ \left(\dfrac{\theta_{\mathrm{FWHM}}}{71\mathrm{mas}}\right)^{5/3}, & \text{galactic source.}\end{cases} \tag{5.14}$$

Now, with SM defined, we can leverage our expression for θ_s by substituting it in for the angular size of the source (i.e., setting $\theta_s = \theta_{\mathrm{FWHM}}$):

$$\left(\frac{\Delta\theta_s}{71\,\mathrm{mas}}\right) = \frac{\mathrm{SM}^{3/5}}{\nu_{\mathrm{GHz}}^{11/5}}. \tag{5.15}$$

Plugging this back into Eq. 5.13, we find:

$$\frac{\mathrm{SM}^{3/5}}{\nu_{\mathrm{GHz}}^{11/5}} > \left(\frac{\epsilon \bar{n}_e r_e \lambda^2}{2\pi}\right)^{2/3}. \tag{5.16}$$

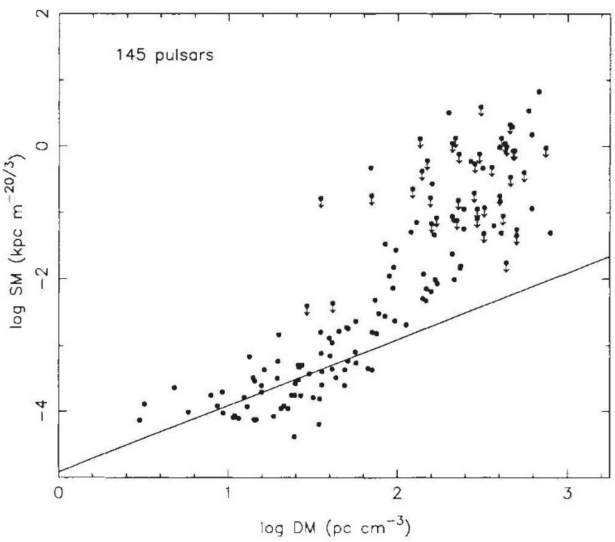

Figure (5.2) Scattering measures for 145 pulsars plotted against their corresponding dispersion measures (electron column densities) (Cordes et al., 1991).

By substituting in the relevant values for the ISM, we find that, for a radio waves at $\nu = 0.3$ GHz (or 300 MHz), in order for focusing to occur, the SM of the ISM must satisfy:

$$\mathrm{SM} > 0.158 \left(\nu_{\mathrm{GHz}}^{13/9} \left(\epsilon n_{e,\mathrm{cm}^3} \right)^{10/9} \right). \tag{5.17}$$

More specifically, if we consider our medium to be weakly varying such that ϵ is $\mathcal{O}(10^{-3})$ and our electron number density is a typical value of $\bar{n}_e = 0.03\mathrm{cm}^{-3}$, then for focusing occur, the scattering measure must satisfy:

$$\mathrm{SM} > \left(1.49 \times 10^{-6}\right) \nu_{\mathrm{GHz}}^{13/9}. \tag{5.18}$$

Fortunately, SM values have been documented for a host of pulsars in the ISM by Cordes et al. (1991) (see Fig. 5.2). Since most pulsars exhibit SM values greater than 10^{-4} kpc m$^{-20/3}$, we are well within the downstream range needed to observe branching in radio waves due to the ISM! While there are still many considerations left to be discussed concerning the distribution of matter in the ISM, this result is immensely encouraging.

Chapter 6

Branched Flow in Three Dimensional Random Media: Radio Waves in the Interstellar Medium

6.1 Three Dimensional Kick-Drift Model

In order to model wave propagation in the ISM, we extend our kick-drift model to three dimensions and generate a series of consecutive two-dimensional scattering screens with spectral indices of $\beta = -8/3$, corresponding to a Kolmogorov wavenumber spectrum in two dimensions. A realization of a single screen along with its wavenumber spectrum is shown in Fig. 6.1. The kick-drift model is standard in astronomical contexts due to the significant difference in length scale between the transverse and longitudinal dimensions of wave propagation.

6.1.1 A Thin Screen (Lens) Geometry

To model wave propagation in a three-dimensional ISM, we simply generalize our thin phase screen model from Section 4.2.1 and propagate a plane of initially equidistant and parallel rays between consecutive two-dimensional phase screens. A schematic showing the two-dimensional kick-drift geometry can be seen in Fig. 6.2. The argument pertaining to the statistical independence of consecutive screens and the preservation of medium statistics in the transverse dimensions still holds. The treatment here is an extension of one proposed by Grillo & Cordes (2018), and is largely inspired by the work of Clegg et al. (1998). We begin by defining planes in which the source and observer will "live" (their specific locations denoted by $\vec{x}_s$ and $\vec{x}_{obs}$ respectively), and represent the medium between them using an arbitrary number of lens planes (or

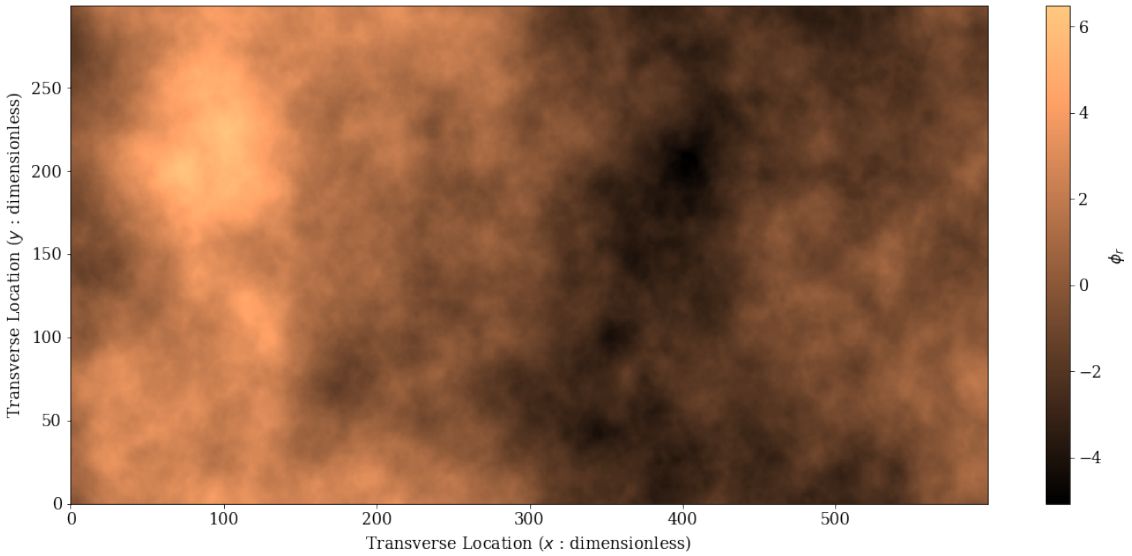

(a) Kolmogorov turbulent scattering screen.

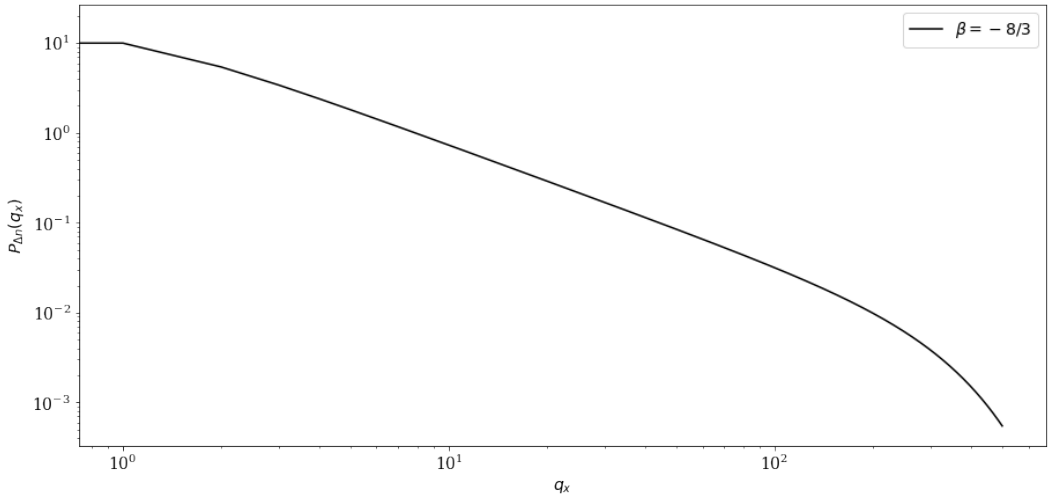

(b) Kolmogorov wavenumber spectrum.

Figure (6.1) Two-dimensional Kolmogorov scattering screen with spectral index $\beta = -8/3$ and its corresponding wavenumber spectrum.

phase screens). The simplest geometry (shown in Fig. 6.2) consists of only one phase screen containing the integrated phase fluctuations for the entire medium. For our purposes, we are interested in a large number of consecutive kick-drift steps which allow for branched flow to emerge.

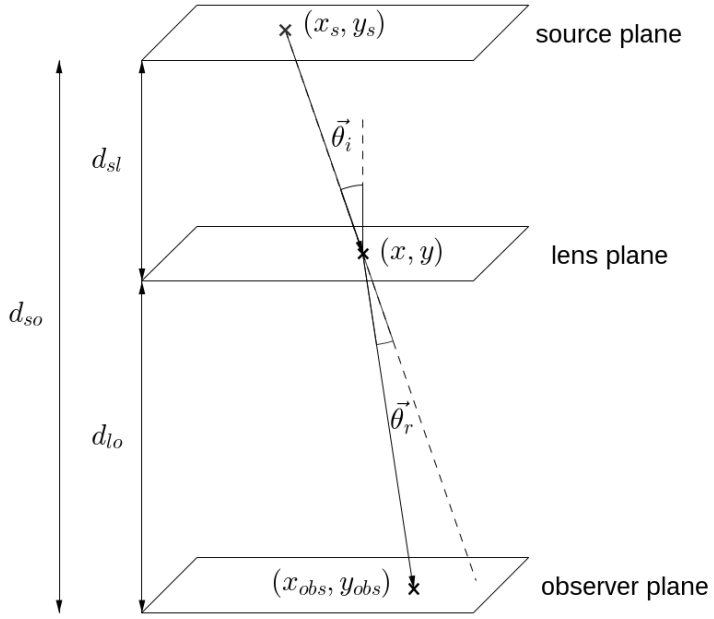

Figure (6.2) A thin screen (lens) geometry (borrowed from Grillo & Cordes (2018)).

6.1.2 Refractive Interstellar Scattering and the Lens Equation

In the case of pulsars, refractive interstellar scattering and scintillation play important roles. Refractive scintillation is the only type of scattering that can be modeled by ray optics, so we will ignore diffractive effects for now. To model these refractive effects, we will use the lens equation, which allows us to map ray coordinates through phase screens that lie between source and observer.

To more straightforwardly present the thin screen formalism in the case of interstellar optics, we will, for the moment, assume the existence of a single phase screen. As argued in Section 3.1.2, we will operate in the paraxial approximation, where the now two-dimensional angles of incidence (as shown in Fig. 6.2) and scattering can be written as:

$$\begin{aligned} \vec{\theta}_i &= \frac{\vec{x}_s - \vec{x}}{d_{sl}}, \\ \vec{\theta}_s &= \frac{\vec{x}_{obs} - \vec{x}}{d_{lo}} - \vec{\theta}_i. \end{aligned} \tag{6.1}$$

By combining these expressions, we arrive at the lens equation as a function of the relevant distances shown in Fig. 6.2, as well asx source and observer coordinates, and θ_s:

$$\vec{x}_s \left(\frac{d_{lo}}{d_{so}}\right) + \vec{x}_{obs} \left(\frac{d_{sl}}{d_{so}}\right) = \vec{x} + \vec{\theta}_s \left(\frac{d_{sl} d_{lo}}{d_{so}}\right). \tag{6.2}$$

To simplify the lens equation, we can operate in a new set of coordinates $\vec{x}'$ that combine the source and observer locations, which are, in turn, scaled by the locations of the phase screen:

$$\vec{x}' \equiv \vec{x}_s \left(\frac{d_{lo}}{d_{so}}\right) + \vec{x}_{obs} \left(\frac{d_{sl}}{d_{so}}\right). \tag{6.3}$$

Using this new set of coordinates, our lens equations can be rewritten as:

$$\vec{x}' = \vec{x} + \vec{\theta}_s \left(\frac{d_{sl} d_{lo}}{d_{so}}\right). \tag{6.4}$$

If we assume that the gradient of our phase screen is small and that the medium we are modeling is effectively homogeneous, we can come up with a general expression for the scattering angle θ_s. Before doing so, however, we need to address some preliminary optics—much of which we have already covered in some capacity. The first is the geometric phase delay ϕ_{scr}, defined as:

$$\phi_{\rm scr} = \omega\tau = kc\tau, \tag{6.5}$$

which is the phase delay (or phase advancement, if the refractive index is greater than unity) that is introduced on a ray by the screen through which it is passing. Here, τ is the difference in propagation time between a scattered and unscattered ray, k is the wavenumber, and ω is the angular frequency of the radiation ($\omega = 2\pi c/\lambda$).

Furthermore, in the case of the plasma responsible for interstellar scattering (which we will assume is unmagnetized), the frequency-dependent refractive index for a radio wave is defined as:

$$\begin{aligned} n_r &= \sqrt{1 - \left(\frac{\omega_e}{\omega}\right)^2} \\ &= \sqrt{1 - \frac{\lambda^2 r_e n_e}{\pi}} \\ &\approx 1 - \frac{\lambda^2 r_e n_e}{2\pi}, \end{aligned} \tag{6.6}$$

where $\omega_e = \sqrt{4\pi n_e e^2/m_e}$ is the electron plasma frequency, and e, m_e, r_e, n_e are the charge, mass, radius and number density of the electrons. The final expression for n_r is arrived at by assuming that, for a radio spectrum, $\omega_e \ll \omega$.

Drawing back on Section 4.2.2, we know that the scattering angle of a ray can be written as (Born & Wolf, 1999):

$$\vec{\theta}_s = \frac{1}{k} \nabla_\perp \phi_{scr}, \tag{6.7}$$

which can then be recast in terms of a physically relevant quantity known as the *dispersion measure* (DM), or the column density of electrons through which the ray is passing, as:

$$\phi_{\rm scr}(\vec{x}) = -\lambda r_e {\rm DM}(\vec{x}). \tag{6.8}$$

As suggested by Grillo & Cordes (2018), we can rewrite the DM($\vec{x}$) as a constant (with the electron density in typical units of pc cm^{-3}) multiplied by a function with unit maximum $V(\vec{x})$, which represents

the potential variations of the screen scaled by a particular DM value. The scattering angle can thus be rewritten as:

$$\vec{\theta}_s = -\frac{c^2 r_e \mathrm{DM}_l}{2\pi\nu^2}\nabla V(\vec{x}). \tag{6.9}$$

Once again using the formulation proposed by Grillo & Cordes (2018), we identify the two frequency-dependent parameters required to model wave propagation in the ISM, namely the Fresnel scale (r_F):

$$r_F = \sqrt{\frac{c d_s d_{lo}}{2\pi\nu}}, \tag{6.10}$$

and the scattering strength of the screen (ϕ_0):

$$\phi_0 = \frac{-c r_e \mathrm{DM}}{\nu}, \tag{6.11}$$

and combine them into one parameter $\alpha = r_F^2\phi_0$ which allows us to rewrite our lens equation as:

$$\vec{x}' = \vec{x} + \alpha\nabla V(\vec{x}). \tag{6.12}$$

Lastly, we make our terms dimensionless and divide by the characteristic transverse length scale of the phase screen (in this case, some small fraction of an Astronomical Unit (AU)). In doing so we redefine our scattering and observer plane coordinates as $u'_x = x/a_x$ and $u'_y = y/a_y$. Our mapping of rays between scattering screens can thus be written as the mapping between planes u and u' (where, at each propagation step, the u plane is the location of the "kick" screen and the u' plane is the location of the screen to which the ray is drifting). This mapping is repeated for all screens until the ray reaches the observer plane:

$$\begin{aligned}\begin{bmatrix} u'_x \\ u'_y \end{bmatrix} &= \begin{bmatrix} u_x + \dfrac{\alpha}{a_x^2}\dfrac{\partial V(u_x,u_y)}{\partial u_x} \\ u_y + \dfrac{\alpha}{a_y^2}\dfrac{\partial V(u_x,u_y)}{\partial u_y} \end{bmatrix} \\ &= \begin{bmatrix} u_x + \alpha_x\dfrac{\partial V(u_x,u_y)}{\partial u_x} \\ u_y + \alpha_y\dfrac{\partial V(u_x,u_y)}{\partial u_y} \end{bmatrix}\end{aligned} \tag{6.13}$$

To test this mapping formalism, we propagated a plane of initially parallel rays through 4 consecutive Gaussian phase screens with large-scale topological variations over a distance of 2 kiloparsecs (6.171×10^{19} m). The results are shown in Figures 6.3 and 6.4. In these images, the wave intensity corresponds to the ray density in a particular region, and the formation of caustics is indicated by the folding or corrugation of the wavefront.

It is important to emphasize that caustic formation does not necessarily lead to branched flow, as is evidenced by Figures 6.3 and 6.4. For four consecutive strong scattering screens, the wavefronts are prevented from entering into a sufficiently stable regime of propagation.

6.1.3 Branched Flow in the ISM

Using the formalism outlined in Section 6.1.2, we were able to simulate radio wave propagation in the ISM by propagating initially equidistant and parallel rays (paraxial wavefronts, or plane waves) over a distance of 1 kiloparsec through many scattering screens with Gaussian and Kolmogorov statistics. Our results suggest the existence of branching in both Gaussian and Kolmogorov media with length scales and electron column densities that reflect those of the ISM.

Preliminary results of this simulation for a Gaussian medium (modeled by 100 consecutive weakly scattering Gaussian random phase screens) are shown in Figures 6.5 and 6.6. The nodal structures exhibited in the two transverse dimensions are familiar and, as shown in Fig. 6.6, remain stable over the full propagation distance to the observer. Fig. 6.5 gives a more informative view of the propagation by showing the trajectories in two transverse dimensions (x and y) as well as one transverse and one longitudinal dimension (x or y and z). Viewing the ray paths in the longitudinal dimension offers a particularly clear view of the onset of branching (around screen 20) and the stability that is maintained until the rays reach the observer plane. Additionally, normalized ray intensities (peaking at caustic locations) are shown for a single horizontal slice along the observer plane. As a final confirmation of true branching, coherence is shown in momentum space as well, which we expect.

The results for a Kolmogorov medium (modeled by 100 consecutive weakly scattering Kolmogorov random phase screens) are shown in Figures 6.7 and 6.8. While the nodal structure that is observed due to the smoother Gaussian screens is not as obvious, there is still clear branching in both the full transverse (x, y) and transverse-longitudinal $(x/y, z)$ representations. The scattering strengths of the Kolmogorov screens are intended to be roughly equivalent to the Gaussian screens and are defined such that the outer scale (an analogue to the correlation length for power-law (multi-scale) media) is comparable to the correlation length of the Gaussian medium. Hence, the onset of branching occurs at a roughly similar point downstream near screen 20. While the coherence in momentum space in Fig. 6.7 is not quite as telling as it is in Fig. 6.5, the other configuration space representations are sufficient for identifying distinct branching in ray propagation.

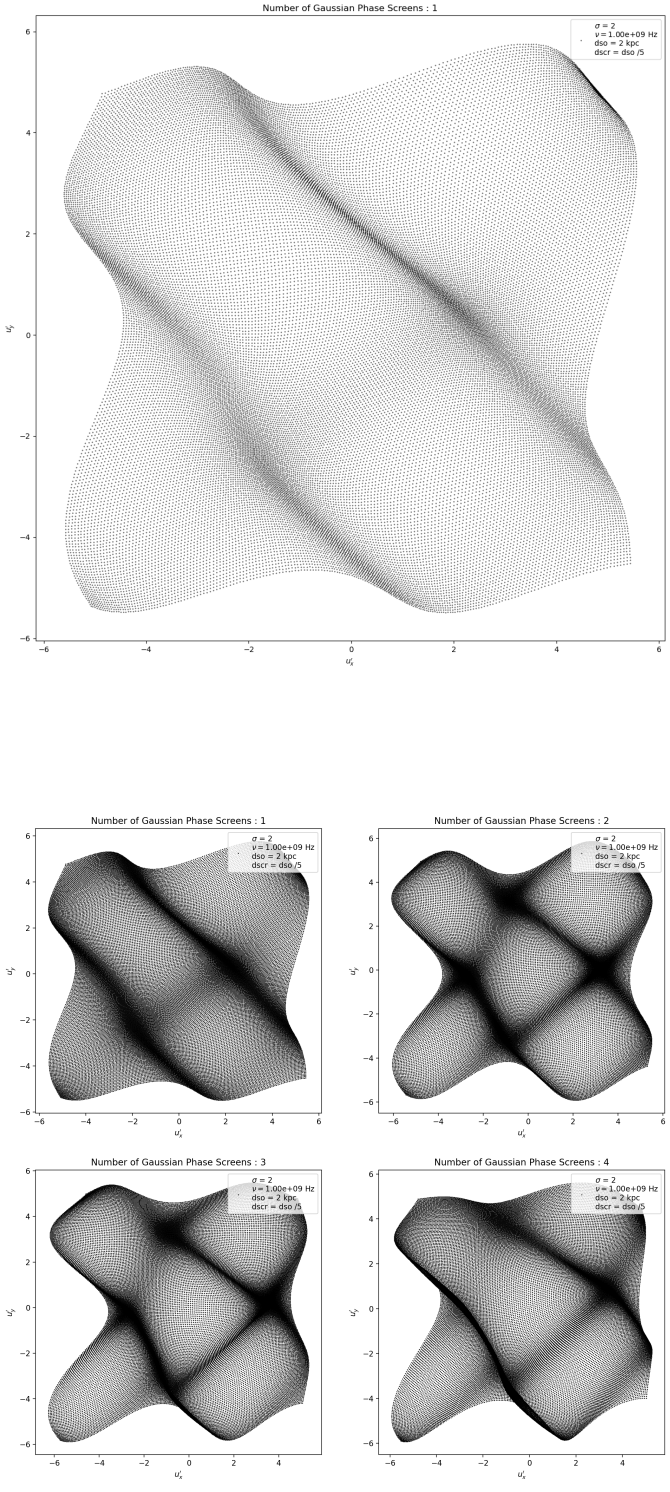

Figure (6.3) Scattering of a plane wave due to four consecutive Gaussian random phase screens with scattering strengths of $\sigma = 2$, spaced equally over a distance of 2 kiloparsecs. The wavefront shows no caustic formation, but does exhibit intensity variations (increased and descreased ray density).

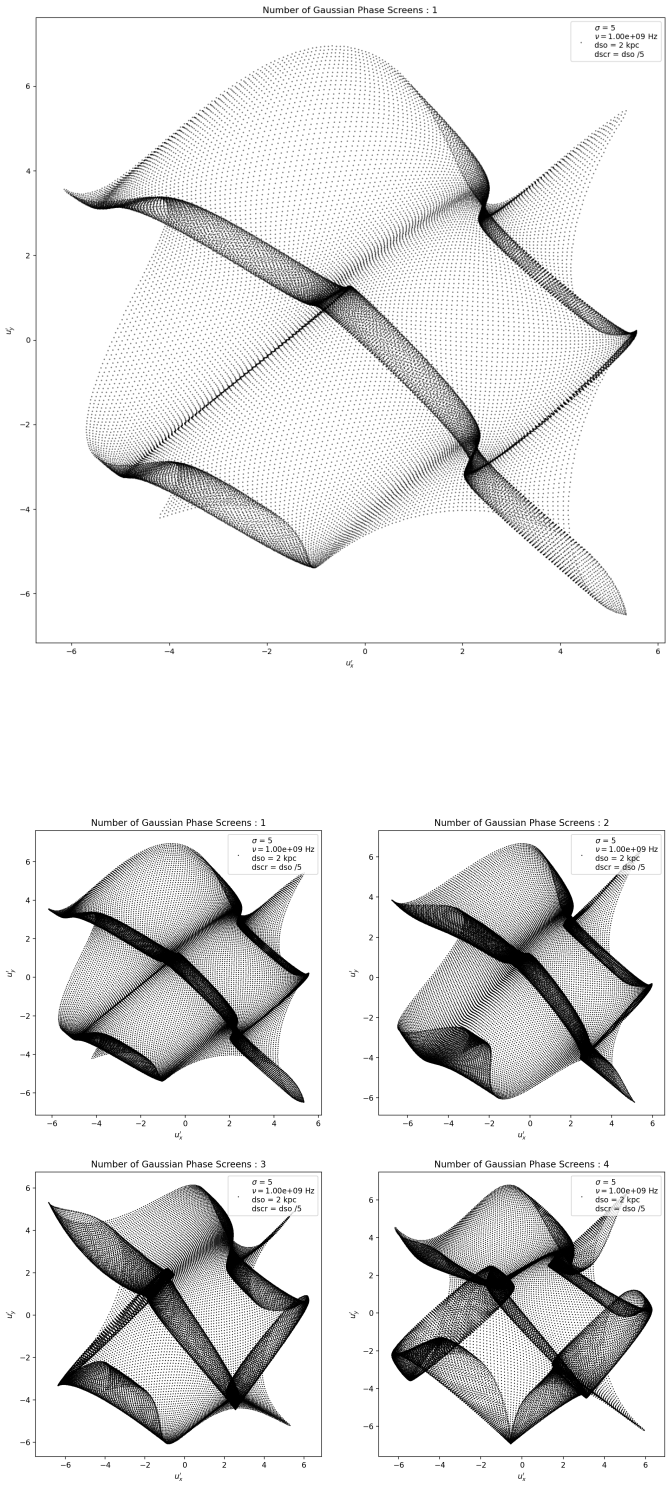

Figure (6.4) Scattering of a plane wave due to four consecutive Gaussian random phase screens with scattering strengths of $\sigma = 5$, spaced equally over a distance of 2 kiloparsecs. The wavefront is folded over itself and is therefore showing caustic formation, but no branching.

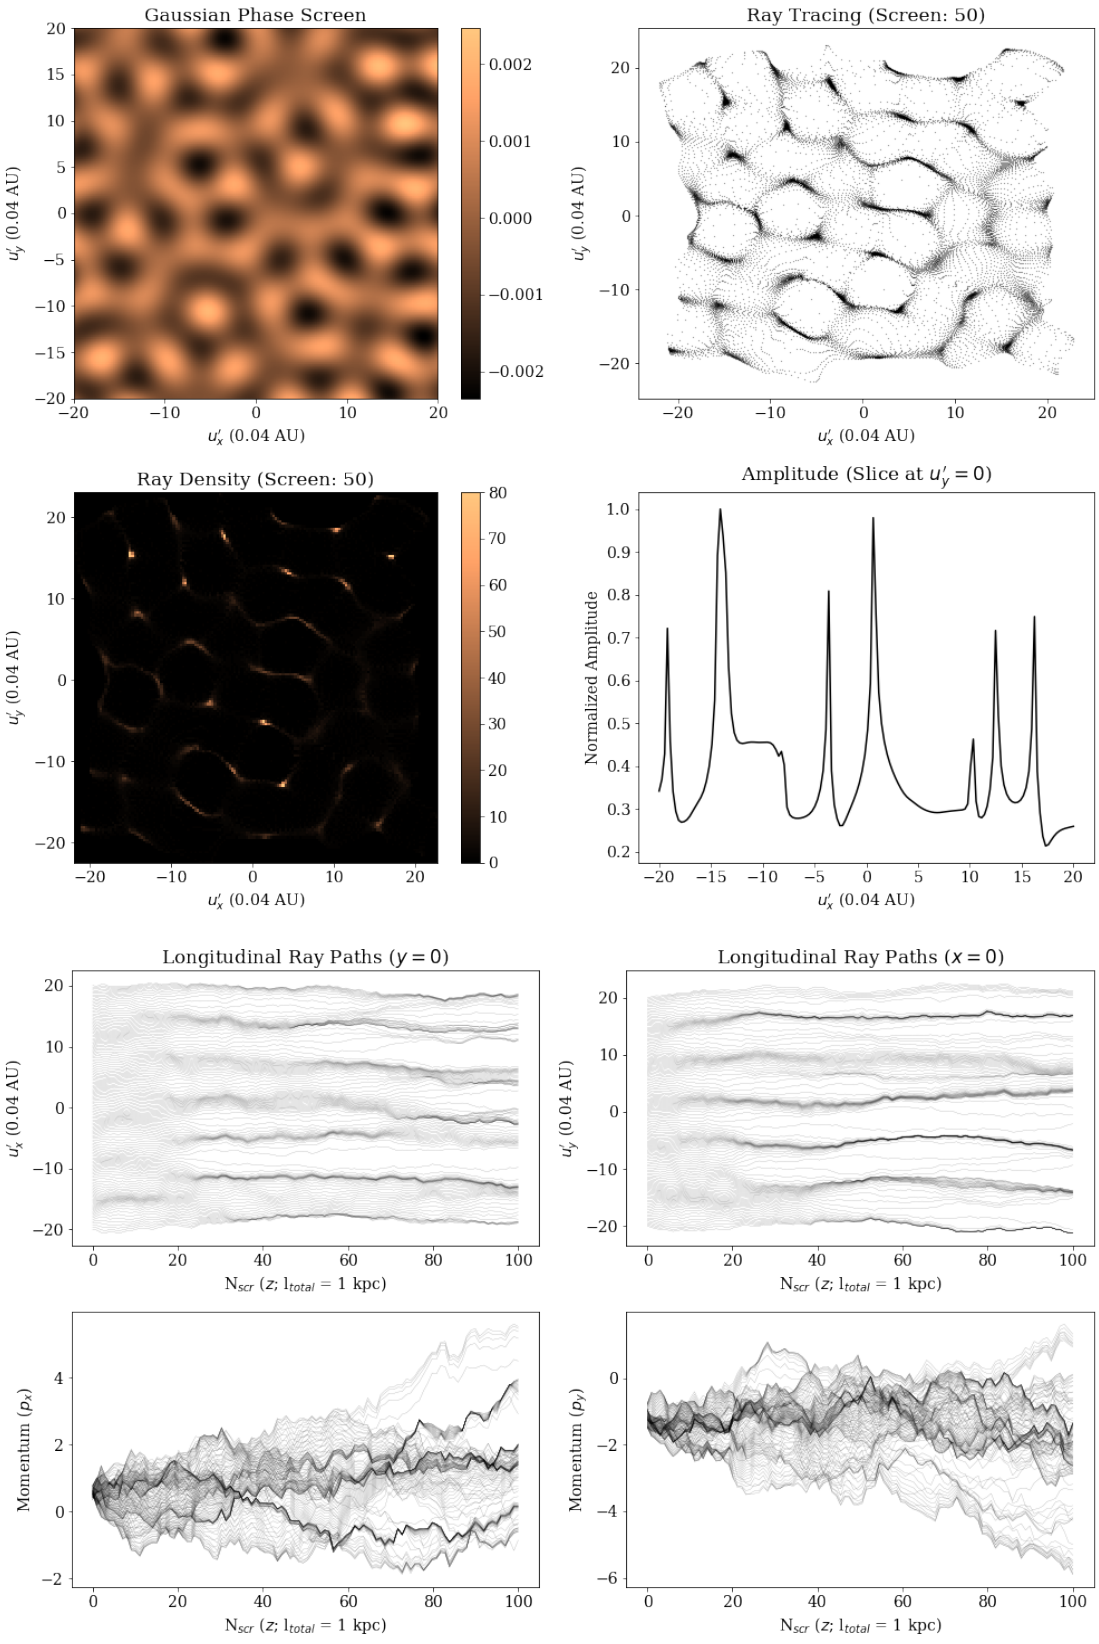

Figure (6.5) Interstellar scattering in a discretized three-dimensional Gaussian ISM, represented by 100 consecutive random phase screens with characteristic length scales of $a = 0.04$ AU and strengths of $\sigma = 10^{-3}$.

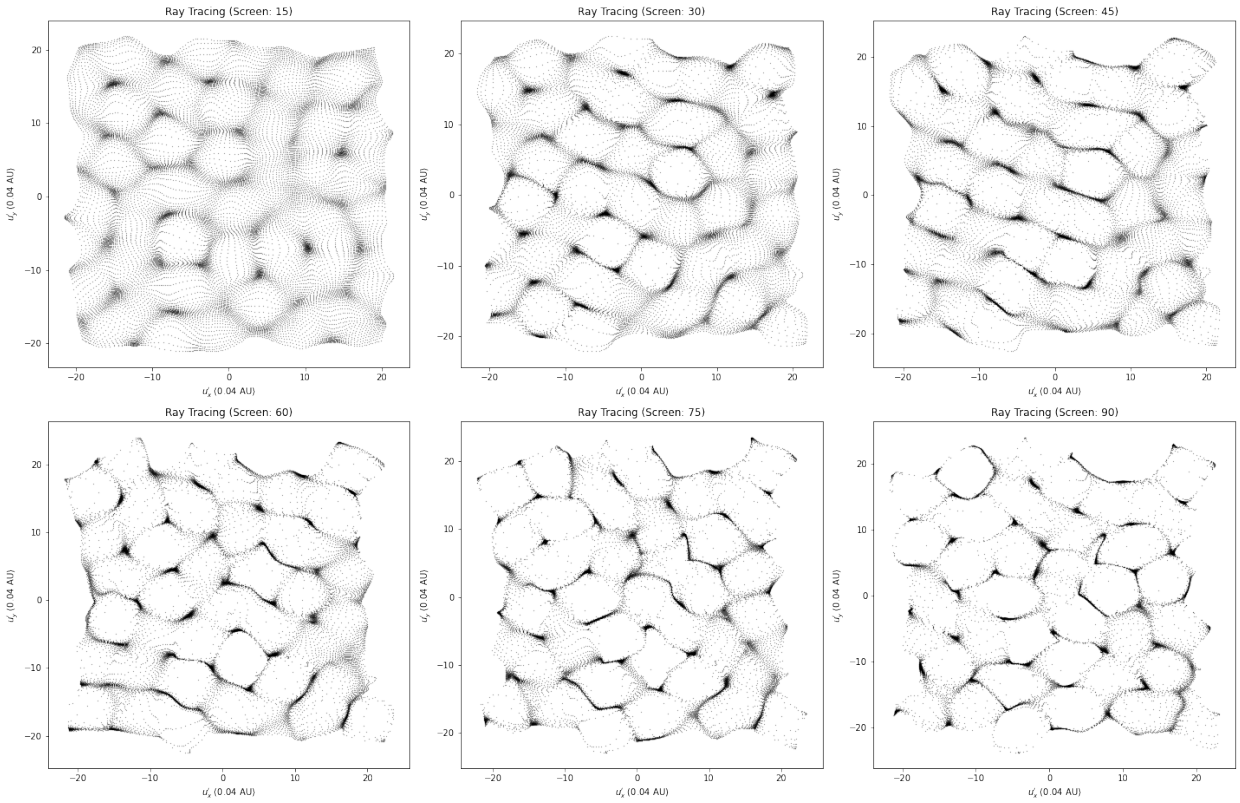

Figure (6.6) Evolving wavefront through 100 Gaussian random phase screens (6 of which are shown at various intermediate steps) with characteristic length scales of $a = 0.04$ AU and strengths of $\sigma = 10^{-3}$.

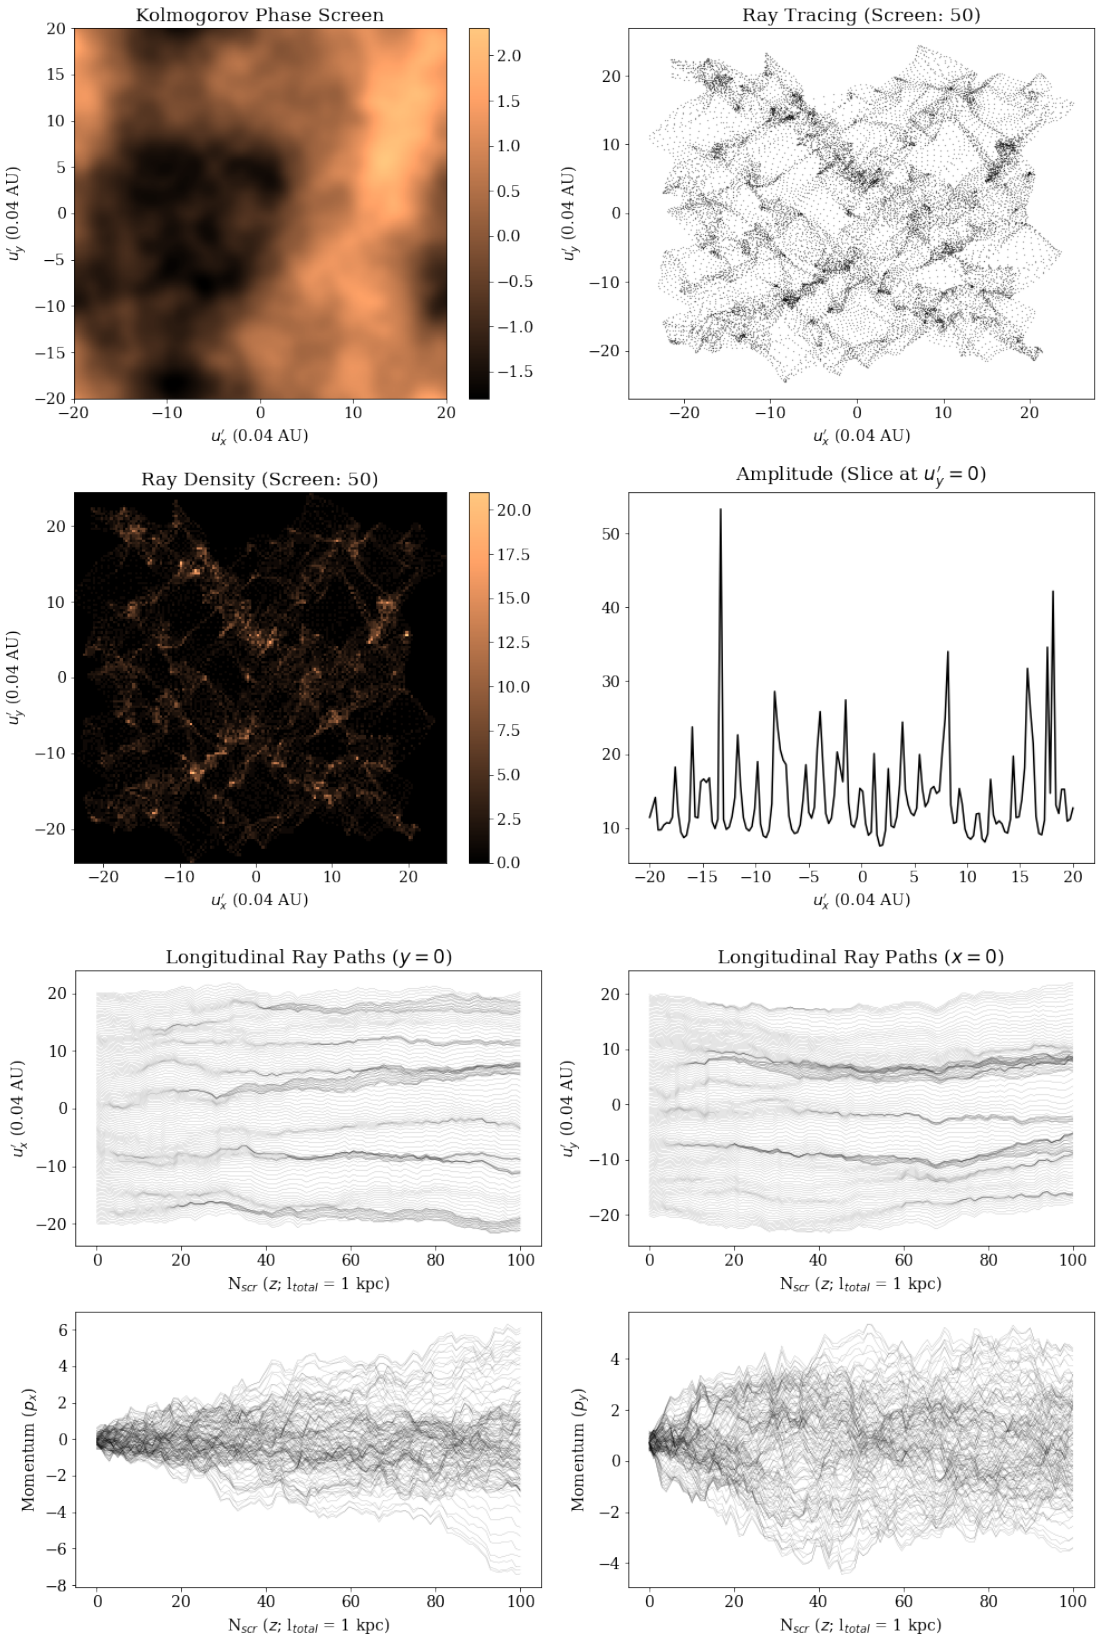

Figure (6.7) Interstellar scattering in a discretized three-dimensional Kolmogorov ISM, represented by 100 consecutive random phase screens with characteristic length scales of $a = 0.04$ AU and spectral indices of $\beta = -8/3$.

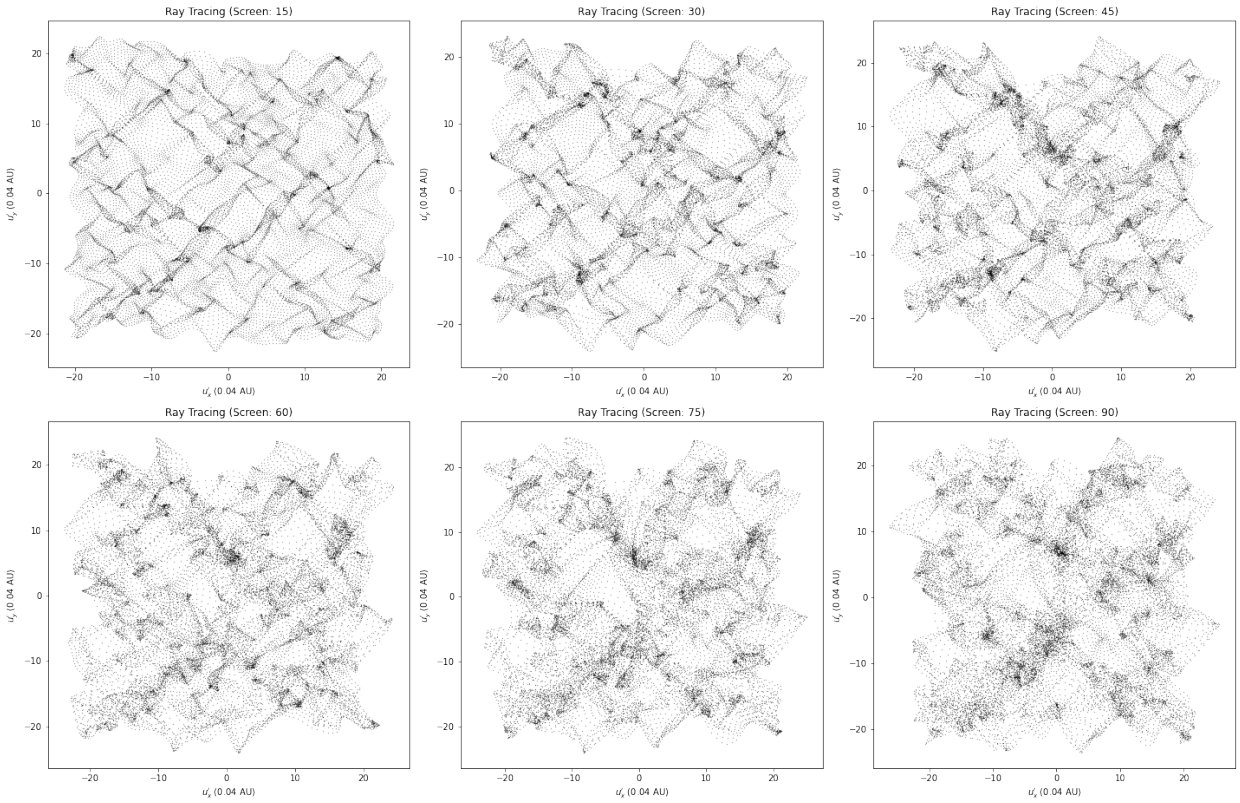

Figure (6.8) Evolving wavefront through 100 Kolmogorov random phase screens (6 of which are shown at various intermediate steps) with characteristic length scales of $a = 0.04$ AU and spectral indices of $\beta = -8/3$.

Chapter 7

Preliminary Conclusions and Future Work

We have shown, through analytic theory and numerical simulations, that the ISM does indeed appear to meet the criteria for branched flow in radio wave propagation from compact astrophysical sources such as pulsars. So far, we have limited our simulation of three-dimensional wave propagation to plane waves rather than spherical waves (emitted by point sources). Spherical waves are, however, more appropriate for modeling the radio wave scintillation from pulsars and will therefore serve as the next step for expanding the simulation presented in Chapter 6. At this stage, however, our simulation is sufficient for modeling the dynamics of greatest importance to our investigation.

The work presented here is preliminary. In pursuit of providing conclusive evidence for the existence of branched flow in the ISM, we must compare our theoretical results and simulations with the standard observables that are studied in radio astronomy—namely, dynamic and secondary spectra (i.e., spectrograms and their Fourier transforms). The simulation presented in Chapter 6 has neared the point where these observables can be generated by incorporating phase information into the geometrical optics formalism. By incorporating phase information, we can move closer to a more comprehensive model of (physical) wave optics that will allow us to model diffractive and interference effects, in addition to chaotic wave propagation, between coherent radio waves that undergo multipath propagation (scintillation) in the ISM.

References

Barkhofen, S. 2013, Microwave Measurements on n-Disk Systems and Investigation of Branching in correlated Potentials and turbulent Flows, ed. U. P. D. Kuhl (Philipps-Universität Marburg)

Bernhardt, E. 2009, Unpublished

Berry, M. V. 2007, Proceedings of the Royal Society A, 463, 3055

Bondi, M. 2005, in The Role of VLBI in Astrophysics, Astrometry and Geodesy, ed. F. Mantovani & A. Kus (Dordrecht: Springer Netherlands), 233–242

Born, M., & Wolf, E. 1999 (Cambridge University Press), doi:10.1017/cbo9781139644181

Brandstötter, A., Girschik, A., Ambichl, P., & Rotter, S. 2019, Proceedings of the National Academy of Sciences, 116, 13260

Cerruti, N. R., & Tomsovic, S. 2002, Physical Review Letters, 88, doi:10.1103/physrevlett.88.054103

Clegg, A. W., Fey, A. L., & Lazio, T. J. W. 1998, The Astrophysical Journal, 496, 253–266

Coles, W. A., Frehlich, R. G., Rickett, B. J., & Codona, J. L. 1987, , 315, 666

Colosi, J. A. 2015, Sound Propagation through the Stochastic Ocean (Cambridge University Press), doi:10.1017/cbo9781139680417

Cordes, J. M., Pidwerbetsky, A., & Lovelace, R. V. E. 1986, The Astrophysical Journal, 310, 737

Cordes, J. M., Weisberg, J. M., Frail, D. A., Spangler, S. R., & Ryan, M. 1991, Nature, 354, 121

Degueldre, H., Metzger, J. J., Geisel, T., & Fleischmann, R. 2016, Nature Physics, 12, 259

Dudley, J. M., Genty, G., Mussot, A., Chabchoub, A., & Dias, F. 2019, Nature Reviews Physics, 1, 675

Fallows, R. A., Coles, W. A., McKay, D., et al. 2015, arXiv:1511.00937

Fallows, R. A., Forte, B., Astin, I., et al. 2020, A LOFAR Observation of Ionospheric Scintillation from Two Simultaneous Travelling Ionospheric Disturbances, arXiv:2003.04013

Grillo, G., & Cordes, J. 2018, arXiv:1810.09058

Heller, E. 2005, Europhysics News, 36, 159

Heller, E. J. 1991, Session LII-Chaos and quantum physics (Les Houches, France, 1989)

—. 2018, The Semiclassical Way to Dynamics and Spectroscopy (Princeton University Press)

Heller, E. J., Fleischmann, R., & Kramer, T. 2019a, Branched Flow, arXiv:1910.07086

—. 2019b, arXiv:1910.07086 [physics], arXiv: 1910.07086

Ishimaru, A. 1978, Wave propagation and scattering in random media. Volume 1 - Single scattering and transport theory, Vol. 1, doi:10.1016/B978-0-12-374701-3.X5001-7

Jura, M. P., Topinka, M. A., Urban, L., et al. 2007, Nature Physics, 3, 841–845

Kaplan, L. 2002, Physical Review Letters, 89, doi:10.1103/physrevlett.89.184103

Keller, J. B., & Papadakis, J. S., eds. 1977, Wave Propagation and Underwater Acoustics (Springer Berlin Heidelberg), doi:10.1007/3-540-08527-0

Kline, M., & Kay, I. W. 1965, Electromagnetic Theory and Geometrical Optics

Littlejohn, R. G. 1992, Journal of Statistical Physics, 68, 7

Lyne, A., & Graham-Smith, F. 2012, Pulsar Astronomy

Metzger, J. J., Fleischmann, R., & Geisel, T. 2010a, Phys. Rev. Lett., 105, 020601

—. 2010b, Physical Review Letters, 105, 020601

—. 2014, Physical Review Letters, 112, 203903, 00000

Nye, J. F. 1999, Natural focusing and fine structure of light: caustics and wave dislocations

Patsyk, A., Sivan, U., Segev, M., & Bandres, M. A. 2020, Nature, 583, 60

Rickett, B. J. 1977, Annual Review of Astronomy and Astrophysics, 15, 479

Roddier, F. 1981, Progess in Optics, 19, 281

Shaw, S. E. J. 2002, PhD thesis, Harvard University

Stavroudis, O. N. 2006, The Mathematics of Geometrical and Physical Optics: The k-function and its Ramifications (Weinheim: Wiley-VCH)

Strohbehn, J. W., ed. 1978, Laser Beam Propagation in the Atmosphere (Springer Berlin Heidelberg), doi:10.1007/3-540-08812-1

Tatarskii, V. I. 1961, Wave Propagation in Turbulent Medium

Topinka, M. A., LeRoy, B. J., Westervelt, R. M., et al. 2001, Nature, 410, 183–186

Voronovich, A. G. 1994, Wave Scattering from Rough Surfaces (Springer Berlin Heidelberg), doi:10.1007/978-3-642-97544-8

White, B. S., & Fornberg, B. 1998, Journal of Fluid Mechanics, 355, 113

Wolfson, M. A., & Tappert, F. D. 2000, The Journal of the Acoustical Society of America, 107, 154

Wolfson, M. A., & Tomsovic, S. 2001, The Journal of the Acoustical Society of America, 109, 2693–2703