



Provided by the author(s) and University of Galway in accordance with publisher policies. Please cite the published version when available.

Title	Gradient-index lenses: Symplectic ray tracing and optical testing
Author(s)	McKeon, Ben
Publication Date	2024-02-22
Publisher	NUI Galway
Item record	http://hdl.handle.net/10379/18062

Downloaded 2024-05-02T06:41:42Z

Some rights reserved. For more information, please see the item record link above.





OLLSCOIL NA GAILLIMHE

UNIVERSITY OF GALWAY

Gradient-index lenses
Symplectic ray tracing and optical testing

Ben McKeon
School of Natural Sciences
University of Galway

Supervisor
Dr. Alexander Goncharov

A document submitted in partial fulfilment of
the requirements for the degree of
Master of Science
February 14, 2024

Abstract

The chief objective of this thesis is to provide an introduction to symplectic numerical methods and how they may be applied to optical problems, particularly for tracing rays within gradient-index (GRIN) optics. Specifically, we investigate how symplectic methods compare in terms of accuracy with well-established numerical integration techniques such as Euler's method and the fourth-order Runge-Kutta method (RK4). As a near-term application, symplectic methods are used to render a test image which requires nonlinear ray tracing. The accuracy of implicit numerical methods is also considered, in addition to the derivation of algebraic iteration schemes for lenses with separable index profiles thereby removing the need for root solvers when using implicit methods. Finally, the pyramid wavefront sensor, a component commonly employed in adaptive optics systems, is considered as a means of measuring aberrations present within GRIN elements and is proposed as a tool to undertake the characterisation and optical testing of same.

Keywords: *Symplectic numerical methods, GRIN optics, optical testing.*

Contents

Abstract	i
Table of contents	ii
List of figures	iv
1 Introduction	1
2 Mathematical preliminaries	4
2.1 Hamiltonian optics: the state of the art	4
2.2 A review of the Hamiltonian formalism	5
2.3 Phase space and symplectic integrators	7
2.4 Ray tracing in homogeneous media	14
3 Numerical experiments with spherical GRIN lenses	17
3.1 An alternative derivation of the optical Hamiltonian	17
3.1.1 A brief remark on alternatives to ray tracing	19
3.2 Results of numerical experiments	20
3.2.1 A note of caution on higher-order methods	29
3.3 Applying symplectic methods to image rendering	32
3.4 Summary	36
4 An implicit method for separable GRIN profiles	37
4.1 Introduction	37
4.2 The implicit midpoint method	38
4.2.1 Construction by composition	38
4.2.2 Deriving the necessary iteration scheme for each lens	42
4.3 Numerical ray tracing	45
5 GRIN aberrometry with a pyramid wavefront sensor	54

5.1	A lack of testing techniques	54
5.2	The pyramid wavefront sensor	56
5.3	The test procedure	59
6	Conclusions and recommendations	63
	Appendices	64
A.1	Numerical recipes for explicit methods	64
A.2	Proof that the implicit midpoint is indeed symplectic	67
A.3	Proof of the nonexistence of explicit symplectic Runge-Kutta methods	69
	Bibliography	70

List of Figures

1.1	Using a gradient-index element to achieve a perfect focus.	3
2.1	The phase space of a one-dimensional simple harmonic oscillator.	9
2.2	The symplectic structure.	10
2.3	Differences between symplectic and nonsymplectic methods.	12
3.1	A schematic of a spherical GRIN lens of radius $R_0 = k\Delta t$	20
3.2	Numerical traces through a Lüneburg lens of radius R_0	21
3.3	Numerical traces through a Maxwell fish-eye lens of radius R_0	22
3.4	Numerical traces through an Eaton lens of radius R_0	25
3.5	Numerical traces through an optical black hole of radius R_0	26
3.6	Error in ray tracing each lens with numerous step sizes	27
3.7	Numerical traces through a Mikaelian lens of length l	30
3.8	Error in ray tracing the Mikaelian lens with numerous step sizes	31
3.9	1080p test images of a Schwarzschild black hole.	33
3.10	Normalised pixel-wise difference maps for the velocity Verlet and Ruth’s methods.	34
3.11	A plot of trace time versus image height for images with 16:9 and 4:3 aspect ratios.	35
4.1	Exact traces available for the lenses given in Table 4.1.	41
4.2	A constant-index meniscus lens with $n = 1.40$ exhibiting coma.	47
4.3	Implicit midpoint tracing a Lüneburg lens of radius R_0	48
4.4	Implicit midpoint tracing an Ilinsky lens whose front surface is of radius r_1	49
4.5	Implicit midpoint tracing a Tarkhanov lens whose back surface is of radius r_2	50
4.6	Transverse coma calculated by each method	51
4.7	Differences in the optical path travelled by forward and reverse traces	52
5.1	Spot diagrams for Ilinsky and Tarkhanov lenses	55
5.2	Schematic of a pyramid wavefront sensor.	57
5.3	Detector plane images of transmission and phase masks.	58
5.4	An Ilinsky lens being tested with the pyramid wavefront sensor.	59

5.5	Wavefront profiles for Ilinsky and Tarkhanov lenses.	60
5.6	A Zernike wavefront decomposition for each lens.	61

Declaration

I, the Candidate, certify that the Thesis is all my own work and that I have not obtained a degree in this University or elsewhere on the basis of any of this work.

Signed: *Ben McKeon.*

Date: February 14, 2024

To Sasha,
whose patience and encouragement made this possible.

To my parents and grandparents,
for lighting the spark and fostering a desire for knowledge.

“We’re fabulous today.”

Chapter 1

Introduction

When students first begin their study of optics, the refractive index is often presented as a material property with a fixed value. Birefringent media, with various constant index values along a given axis or set of axes are then typically introduced, though the process of determining ray trajectories through isotropic or birefringent media is essentially identical. However, the idea of a position-dependent refractive index is often considered a specialist topic and is seldom taught in great detail, since the determination of analytical trajectories through media with non-constant refractive indices becomes significantly more challenging and is even impossible in a great number of cases.

That said, it is still worth examining spatially-varying gradient-index (GRIN) elements as they offer an additional degree of freedom in the design of optical components. For instance, Figure 1 shows how it is possible to use a GRIN lens with spherical surfaces to direct an axially-collimated bundle of rays to an ideal focus. Thus, we no longer need to manufacture complex aspheric surfaces in order to correct for spherical aberration. Various techniques exist for the fabrication of GRIN elements. Photopolymerisation, for example, involves the exposure of monomers to UV light with varying intensity, causing the monomers to partially polymerise, creating a gradient-index profile [64]. Another technique involves the Fickian diffusion of exchange molecules across a flat boundary, reporting that lenses manufactured via this method experience an order of magnitude reduction in spherical aberration when compared with constant-index lenses of the same geometry [28].

Gradient index optics are not just synthetic curiosities; rather, many natural examples also exist. Earth's atmosphere has a spatially- and temporally-varying refractive index [66], causing a notable reduction in the quality of images obtained by Earth-based astronomical observatories [83]. Also, in the field of ophthalmology, developing an accurate GRIN model of the human eye

which accounts for ageing is an active area of research [84]. Finally, the most exotic example of a natural GRIN lens, is perhaps the Solar Gravitational Lens, where detector spacecraft are placed along the gravitational focal line of the Sun in order to obtain high-resolution images of distant exoplanets. However, many technological obstacles prevent it from being realised [96].

This thesis primarily investigates symplectic numerical methods, a class of numerical integrators designed for solving Hamiltonian systems, comparing their performance with a range of common numerical methods used in the solution of ordinary differential equations (ODEs). We also show how implicit numerical methods may be employed without the need for a root-solving algorithm, provided that the refractive index of a lens is of a particular form. The pyramid wavefront sensor is then proposed as an aberrometer for GRIN optics.

It is perhaps worth stating at the outset that the simulations presented here are purely phenomenological; rays are simply treated as the trajectories of corpuscles, calculated based on the eikonal approximation of Maxwell's equations [11]. Hence, any effects owing to the wave nature of light (i.e. diffraction), are neglected completely. So too are the singularities of light wave amplitudes which occur at caustics surfaces which are tangential to each of the rays [88]. Furthermore, wavefronts are considered to be surfaces orientated perpendicular to each ray at a given instance during its propagation within an optical system, that is, all points of the same phase. While the phase function governing a particular wavefront is typically periodic, the pyramid wavefront sensor struggles to reconstruct wavefronts with phase discontinuities [9]. Hence, wavefront branches are disregarded here for simplicity.

This thesis is laid out as follows: Chapter 2 provides an overview of the mathematics required for the Hamiltonian formulation of geometric optics and symplectic methods before applying them to a selection of spherical GRIN lenses in Chapter 3, where image rendering is also presented as an alternative use for symplectic numerical techniques. Chapter 4 then demonstrates how lenses with certain separable index profiles enable the use of implicit symplectic methods without the need for computationally expensive nonlinear root-solvers. Simulations of aberrometry with the pyramid wavefront sensor are given in Chapter 5. Concluding remarks and recommendations then follow in Chapter 6. Three appendices accompany the text. The first provides further details of each numerical method used in the thesis. The second contains a proof that the implicit midpoint method is symplectic. The third appendix then proves the non-existence of explicit symplectic Runge-Kutta methods.

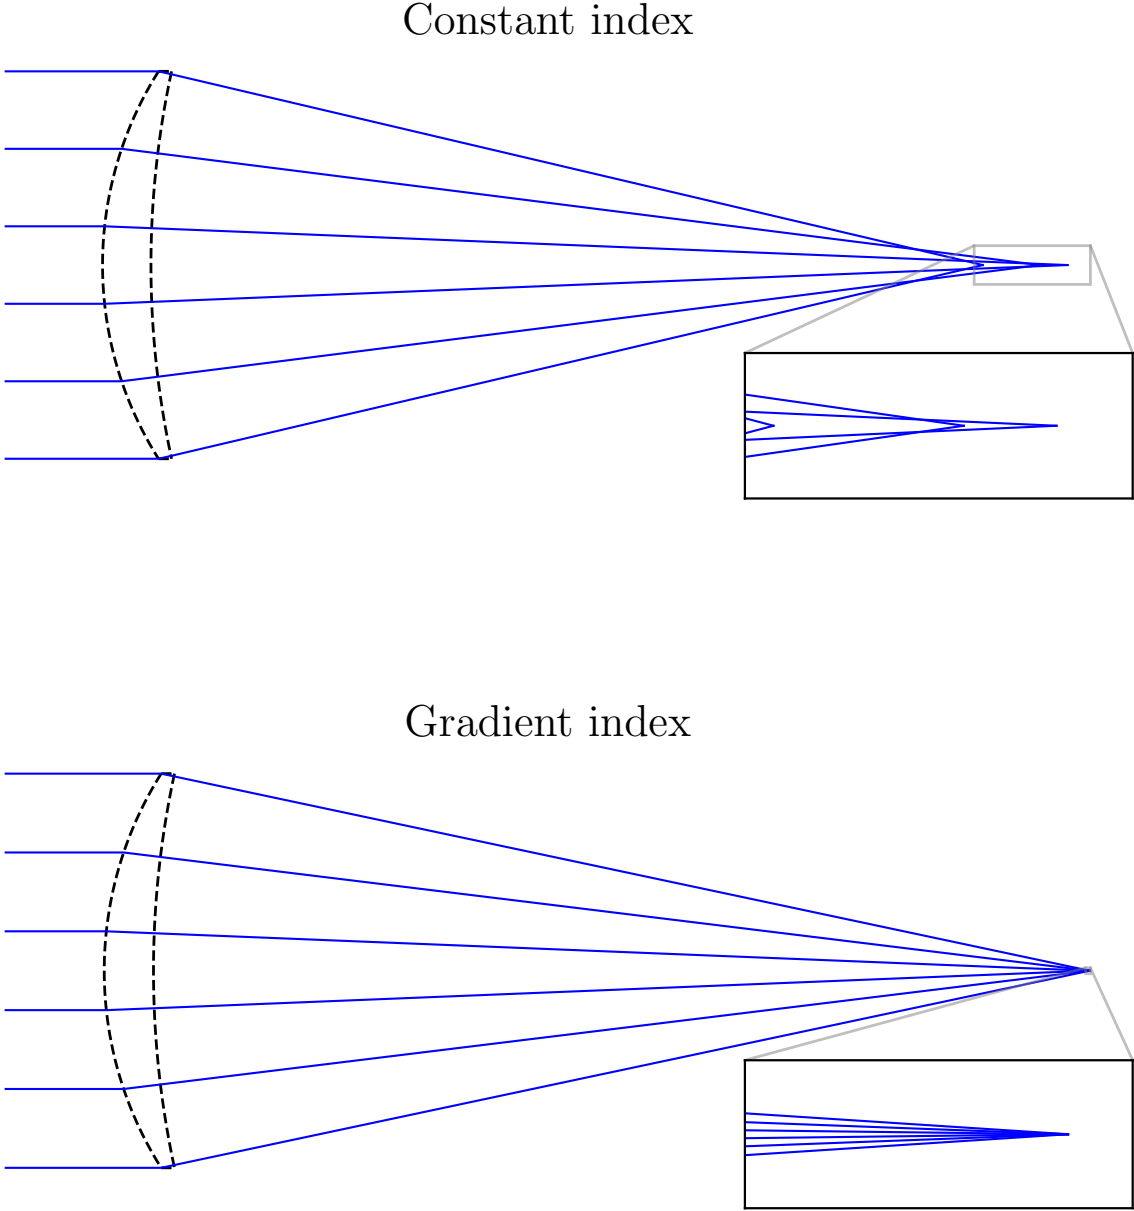


Figure 1.1: Using a gradient-index element to achieve a perfect focus. This particular element is known as the Ilinsky lens and is examined in further detail in Chapter 4

Chapter 2

Mathematical preliminaries

2.1 Hamiltonian optics: the state of the art

Hamiltonian dynamics in its most general sense has its roots in optics, first being presented in William Rowan Hamilton's *Theory of systems of rays* [39]. However, Hamilton's initial work on optics was quickly extended to mechanics [40], laying the foundation for quantum mechanics and forming a large portion of mechanics courses at university level. While several volumes have been published specifically on Hamiltonian optics [14, 70, 90], Hamilton's analysis is seldom exploited to its full potential, typically being neglected in favour of an equivalent Lagrangian framework, which has already proved useful for ray tracing in GRIN media [51, 65, 73].

Nonetheless, by using the more symmetrical Hamiltonian framework, a system's phase space may be examined, thereby offering a unique perspective. The utility of phase space in optics is well documented [92, 94, 100] and has several advantages. Firstly, the inverse problem determining the refractive index of a medium from the optical path of a given ray is simply solved by a Legendre transform between coordinate and momentum spaces [63]. Secondly, aberrations present in a freeform optical system can be easily quantified by polynomial curve-fitting of phase space data, allowing a more natural alternative framework to Seidel aberration theory [4, 5]. Moreover, since the determination of ray trajectories requires the solution of a nonlinear partial differential equation, which is only analytic in specific cases [58], numerical methods prove to be indispensable for nonlinear ray tracing [35]. Furthermore, by adopting the Hamiltonian perspective, we may construct numerical techniques for ray tracing in GRIN media that are both accurate and computationally inexpensive. Thus, by turning once again to phase space, we can construct numerical methods capable of preserving certain phase space properties which only become apparent when considered from a Hamiltonian point of view. Chief among them is the symplectic structure, representing the spread of light within an optical system [18]. Numerical

methods which conserve the symplectic structure are referred to as symplectic integrators and have been well studied [37, 78, 103]. Popular applications include celestial mechanics [45, 77] and molecular dynamics [34, 87]. By comparison, they have seen limited use in optics [68, 80], with symplectic ray tracing receiving little consideration over the past two decades, despite numerous advances being made in the relevant underlying theory [37, 38, 80].

Consequently, the objective of this chapter is to provide the necessary mathematical preliminaries for Hamiltonian optics without the need for extensive prior knowledge of topics such as differential geometry or group theory. Notation is deliberately kept similar to existing literature so as to allow for assimilation of relevant results from pre-existing manuscripts, ultimately highlighting the merits of applying symplectic numerical methods to optical problems.

2.2 A review of the Hamiltonian formalism

To begin, we define a dynamical system in k -dimensional configuration space by its generalised coordinates $\mathbf{q} = [q_1(t), q_2(t), \dots, q_k(t)]$, parameterised by some scalar t . For mechanical systems, t typically represents time, though other parameterisations (e.g. arclength) may also be used as any parameterisation is not generally unique [93]. The system's corresponding generalised velocities are $\dot{\mathbf{q}} = [\dot{q}_1(t), \dot{q}_2(t), \dots, \dot{q}_k(t)]$, where throughout the remainder of this thesis, the notation \dot{q}_i represents the total derivative with respect to t , that is

$$\dot{q}_i = \frac{dq_i}{dt}. \quad (2.1)$$

The action \mathcal{S} is then defined as the integral of the Lagrangian $L = L(\mathbf{q}(t), \dot{\mathbf{q}}(t), t)$ along a given trajectory between a pair of fixed endpoints parameterised by t_1 and t_2

$$\mathcal{S} = \int_{t_1}^{t_2} L dt. \quad (2.2)$$

For a Lagrangian that is not explicitly t -dependent, [i.e. $L = L(\mathbf{q}(t), \dot{\mathbf{q}}(t))$], multiple trajectories exist between t_1 and t_2 , such that the difference between any two trajectories is minimised to the first order and higher-order terms are assumed to be negligible [48]. In other words, we have defined the stationary action principle, expressed mathematically as follows:

$$\delta\mathcal{S} = \delta \int_{t_1}^{t_2} L dt = 0. \quad (2.3)$$

The system's generalised momenta (or conjugate momenta) are then defined as

$$p_i \equiv \frac{\partial L}{\partial \dot{q}_i}, \quad (2.4)$$

where the energy function of the system is given by

$$E_L = \sum_{i=1}^k \dot{q}_i \frac{\partial L}{\partial \dot{q}_i} - L. \quad (2.5)$$

The Hamiltonian may then be defined by an inverse Legendre transform of the energy function:

$$H = \sum_{i=1}^k \dot{q}_i p_i - L. \quad (2.6)$$

Then, by making an infinitesimal perturbation of the Hamiltonian, we notice

$$\begin{aligned} \delta H &= \sum_{i=1}^k (\dot{q}_i \delta p_i + p_i \delta \dot{q}_i) - \delta L \\ &= \sum_{i=1}^k \left[\dot{q}_i \delta p_i + p_i \delta \dot{q}_i - \left(\frac{\partial L}{\partial q_i} \delta q_i + \frac{\partial L}{\partial \dot{q}_i} \delta \dot{q}_i \right) \right]. \end{aligned} \quad (2.7)$$

However, recalling Eq. (2.4), we find

$$\delta H = \sum_{i=1}^k \left(\dot{q}_i \delta p_i - \frac{\partial L}{\partial q_i} \delta q_i \right) \quad (2.8)$$

$$= \sum_{i=1}^k \left(\frac{\partial H}{\partial p_i} \delta p_i + \frac{\partial H}{\partial q_i} \delta q_i \right), \quad (2.9)$$

then, comparing Eqs. (2.8) and (2.9)

$$\frac{\partial H}{\partial p_i} = \dot{q}_i, \quad (2.10)$$

$$\frac{\partial H}{\partial q_i} = -\frac{\partial L}{\partial q_i}. \quad (2.11)$$

From the Euler-Lagrange equation, we obtain the following useful result

$$\frac{\partial L}{\partial q_i} = \frac{d}{dt} \left(\frac{\partial L}{\partial \dot{q}_i} \right) = \dot{p}_i, \quad (2.12)$$

which gives us Hamilton's equations

$$\frac{\partial H}{\partial p_i} = \dot{q}_i, \quad (2.13)$$

$$\frac{\partial H}{\partial q_i} = -\dot{p}_i. \quad (2.14)$$

The primary advantage of Hamilton's formalism is that it reduces k second-order Euler-Lagrange equations into $2k$ first-order equations. Furthermore, certain symmetries and invariants might not be obvious in configuration space, but become apparent in phase space [41], which is the natural setting for Hamiltonian dynamics. One such invariant is the symplectic structure [3]. Although the Hamiltonian formulation may not necessarily reduce the difficulty in finding exact solutions when compared with its Lagrangian counterpart, we can exploit the symplectic structure to find highly accurate numerical solutions.

2.3 Phase space and symplectic integrators

In the search for a full description of a dynamical system, the configuration space picture of Lagrangian dynamics is necessarily incomplete. The system considered in the previous section exists in a k -dimensional configuration space, describing its position for a given value of the parameter t . Configuration space does, however, fail to take account of the generalised velocities. Hence, we require an extended space which describes both the position and velocity (or momentum) of the system at a given value of t and also how the system evolves with changes in t ; the extended space we require is phase space.

Our system, with its k -dimensional configuration space has an associated phase space of dimension $2k$, consisting of the configuration space and the system's momentum space, which may be thought of as an analogous configuration space for the system's generalised momenta [41]. As a result, only the complete phase space of a one-dimensional system can be depicted easily on a single graph; representing the full phase space of higher-dimensional systems is cumbersome. To assist with visualisation, however, phase space may be decomposed into two-dimensional phase planes [86], with one for each $p_i q_i$ pair. Regardless of the dimension of the system's phase space, a given point \mathbf{z}_0 in phase space is described by

$$\mathbf{z}_0 = \mathbf{z}(t_0) = \begin{pmatrix} \mathbf{p}_0 \\ \mathbf{q}_0 \end{pmatrix} = \begin{pmatrix} \mathbf{p}(t_0) \\ \mathbf{q}(t_0) \end{pmatrix}, \quad (2.15)$$

and so Hamilton's equations may be rewritten in matrix form:

$$\dot{\mathbf{z}}_0 = \mathbf{J}^{-1} \frac{\partial H}{\partial \mathbf{z}_0}, \quad (2.16)$$

where

$$\mathbf{J} \equiv \begin{pmatrix} \mathbf{0} & \mathbf{I}_k \\ -\mathbf{I}_k & \mathbf{0} \end{pmatrix} \quad (2.17)$$

is a $2k \times 2k$ skew-symmetric block matrix. $\mathbf{0}$ is a $k \times k$ null matrix and \mathbf{I}_k is a $k \times k$ identity

matrix. For $t_1 = t_0 + \Delta t$, we define another phase space point

$$\mathbf{z}_1 = \mathbf{z}(t_1) = \begin{pmatrix} \mathbf{p}_1 \\ \mathbf{q}_1 \end{pmatrix} = \begin{pmatrix} \mathbf{p}(t_1) \\ \mathbf{q}(t_1) \end{pmatrix} \quad (2.18)$$

that also satisfies Eq. (2.16) i.e.

$$\dot{\mathbf{z}}_1 = \mathbf{J}^{-1} \frac{\partial H}{\partial \mathbf{z}_1}. \quad (2.19)$$

After some algebra, Eq. (2.19) yields an alternative matrix equation

$$\begin{pmatrix} \frac{\partial \mathbf{z}_1}{\partial \mathbf{z}_0} \end{pmatrix}^T \mathbf{J} \begin{pmatrix} \frac{\partial \mathbf{z}_1}{\partial \mathbf{z}_0} \end{pmatrix} = \mathbf{J}, \quad (2.20)$$

in general, for $t_n = t_0 + n\Delta t$,

$$\begin{pmatrix} \frac{\partial \mathbf{z}_2}{\partial \mathbf{z}_1} \end{pmatrix}^T \mathbf{J} \begin{pmatrix} \frac{\partial \mathbf{z}_2}{\partial \mathbf{z}_1} \end{pmatrix} = \begin{pmatrix} \frac{\partial \mathbf{z}_3}{\partial \mathbf{z}_2} \end{pmatrix}^T \mathbf{J} \begin{pmatrix} \frac{\partial \mathbf{z}_3}{\partial \mathbf{z}_2} \end{pmatrix} = \dots = \begin{pmatrix} \frac{\partial \mathbf{z}_{n+1}}{\partial \mathbf{z}_n} \end{pmatrix}^T \mathbf{J} \begin{pmatrix} \frac{\partial \mathbf{z}_{n+1}}{\partial \mathbf{z}_n} \end{pmatrix} = \mathbf{J}, \quad (2.21)$$

where the superscript T denotes the matrix transpose. Any matrix that satisfies Eq. (2.20) is deemed to be symplectic as it conserves the symplectic structure in phase space. $\begin{pmatrix} \frac{\partial \mathbf{z}_{n+1}}{\partial \mathbf{z}_n} \end{pmatrix}$ is the Jacobian of the system's flow [23]. This flow (so called because of its similarity to a fluid flow) traces the phase space trajectory of the system, mapping the initial condition to another state

$$\begin{pmatrix} \mathbf{p}_0 \\ \mathbf{q}_0 \end{pmatrix} \mapsto \begin{pmatrix} \mathbf{p}_n \\ \mathbf{q}_n \end{pmatrix}. \quad (2.22)$$

As a simple example, Figure 2.1 shows the phase space of a one-dimensional simple harmonic oscillator, whose Hamiltonian is given by

$$H(p, q) = \frac{p^2}{2} + \frac{q^2}{2}. \quad (2.23)$$

From Eq. (2.23), we see the trajectory in each $p_i q_i$ plane is a circle, which is further confirmed in Figure 2.1 and so, it follows that the transformations required to map the initial conditions to another point on the phase space trajectory are the parametric equations of a circle

$$\underbrace{\begin{pmatrix} p_n \\ q_n \end{pmatrix}}_{\mathbf{z}_n} = \underbrace{\begin{pmatrix} \cos t & -\sin t \\ \sin t & \cos t \end{pmatrix}}_{\text{flow}} \underbrace{\begin{pmatrix} p_0 \\ q_0 \end{pmatrix}}_{\mathbf{z}_0}. \quad (2.24)$$

Ensuring that the Jacobian of the flow is symplectic via substitution into Eq. (2.20) is a straightforward task. Recalling Eq. (2.23) once again, we observe that a symplectic transformation corresponds to some initial state being rotated anticlockwise through an angle of $\frac{\pi}{2}$ radians. Thus, we now have a means of depicting the symplectic structure. Without loss of generality, for a

pair of vectors in phase space (\mathbf{v}, \mathbf{w}) and using the notation presented in [25], the symplectic structure $[\mathbf{v}, \mathbf{w}]$ is defined to be

$$[\mathbf{v}, \mathbf{w}] \equiv (\mathbf{J} \times \mathbf{v}) \cdot \mathbf{w}, \quad (2.25)$$

where \mathbf{J} is the same skew-symmetric matrix from Eq. (2.16). For the one-dimensional case, as in Figure 2.1, $[\mathbf{v}, \mathbf{w}]$ is simply the area of the projection of the parallelogram defined by \mathbf{v} and \mathbf{w} . In higher dimensions, $[\mathbf{v}, \mathbf{w}]$ becomes a sum of k terms, where each term is the area of the parallelogram defined in the k th phase plane. We can then more easily visualise a volume in $2k$ dimensions as a series of two-dimensional “slices” through each phase plane. Irrespective of the dimensions of the system, the symplectic structure is unique to Hamiltonian systems and is thus conserved by Hamilton’s equations (i.e. $\frac{d}{dt}[\mathbf{v}, \mathbf{w}] = 0$). The formal statement describing the conservation of the symplectic structure is known as Liouville’s theorem, which is demonstrated graphically in Figure 2.2, where the area of the parallelogram remains unchanged for each point on the phase space trajectory. In optics, conserving the symplectic structure corresponds to conserving étendue [18], meaning the extent to which the rays spread out cannot decrease as the light propagates.




Figure 2.1: The phase space of a one-dimensional simple harmonic oscillator. Beginning at $(q_0, p_0) = (1, 0)$, the oscillator is then released and allowed to oscillate indefinitely. The trajectory traced by the system is indicated by the arrows. To fully describe a two-dimensional oscillator, we require two such plots. Likewise, a three-dimensional one would require three plots and so on. The maximum momentum and displacement are normalised with $q_{max} = p_{max} = 1$.




Figure 2.2: The symplectic structure $[\mathbf{v}, \mathbf{w}]$ (i.e. the area of each parallelogram) is conserved and unchanging at each point on the phase space trajectory. Arrows represent the vectors \mathbf{v} and \mathbf{w} .

In more general terms, however, the application of symplectic numerical methods to Hamiltonian systems is possible since (2.23) is of the form $H = P(p_x, p_y, p_z) + N(x, y, z)$, making it separable. This allows the terms dependent on position and momentum to be treated independently. Symplectic numerical methods for separable Hamiltonians are constructed via splitting, giving two separate iterative schemes to calculate the position and momentum related to a given trajectory. Constructing the relevant iterative schemes first requires us to solve Hamilton's equations in matrix form. First, we combine the position and momentum vectors for a given ray into a single matrix $\mathbf{z} = (\mathbf{p}, \mathbf{q})^T$ with \mathbf{p} and \mathbf{q} representing a ray's position and optical momentum vectors for a given ray, as before. and define the Liouville operator D_H [25] to act on \mathbf{z} , where

$$D_H \mathbf{z} := \sum_{i=1}^3 \left(\frac{\partial \mathbf{z}}{\partial q_i} \frac{\partial H}{\partial p_i} + \frac{\partial \mathbf{z}}{\partial p_i} \frac{\partial H}{\partial q_i} \right). \quad (2.26)$$

The system (2.16) may now be written more concisely as

$$\frac{d\mathbf{z}}{dt} = D_H \mathbf{z}. \quad (2.27)$$

Treating D_H as a vector operator, the solution to Eq. (2.27) is a matrix exponential

$$\mathbf{z}(t) = [\exp(tD_H)] \mathbf{z}(0). \quad (2.28)$$

Recalling that (2.23) is separable, if we further define

$$\begin{aligned} D_{P\mathbf{z}} &:= \sum_{i=1}^3 \left(\frac{\partial \mathbf{z}}{\partial q_i} \frac{\partial H}{\partial p_i} \right), \\ D_{N\mathbf{z}} &:= \sum_{i=1}^3 \left(\frac{\partial \mathbf{z}}{\partial p_i} \frac{\partial H}{\partial q_i} \right), \end{aligned} \tag{2.29}$$

We then observe $D_H \cdot = D_P \cdot + D_N \cdot$, where “ \cdot ” represents the vector operand. By choosing sets of coefficients $\{c_i\}$ for momentum and $\{d_i\}$ for position, we subdivide the difference between 0 and t into m equal parts, where m is known as the order of the numerical method. For our numerical method to be consistent (i.e. both useful and physically meaningful), we require $\sum_{i=1}^m c_i = \sum_{i=1}^m d_i = 1$. Assuming an infinitesimally small t we now write

$$\exp [t(D_P + D_N)] = \prod_{i=1}^m \exp(tc_i D_P) \exp(td_i D_N) + \mathcal{O}(t^{m+1}), \tag{2.30}$$

However, this does not exclude the possibility of negative coefficients in the case where the sums of $\{c_i\}$ and $\{d_i\}$ would otherwise be greater than unity. Ruth’s method [32], for instance, provides a concrete example of a fourth-order method with negative coefficients and is considered in the appendix. Next, performing first-order Taylor expansions of $\exp(tc_i D_P)$ and $\exp(td_i D_N)$

$$\exp(tc_i D_P) = I + tD_P + \mathcal{O}(t^2), \tag{2.31}$$

$$\exp(td_i D_N) = I + tD_N + \mathcal{O}(t^2). \tag{2.32}$$

In the simplest case of a first-order method (i.e. taking $c_i = d_i = 1$), if we know the state governing a trajectory \mathbf{z}_n at $t = n$, we may calculate \mathbf{z}_{n+1} at $t_{n+1} = t_n + \Delta t$ (where Δt is some small, finite difference between subsequent iterations). Multiplying (2.31) by \mathbf{z}_n before substituting the result into (2.32), we obtain the symplectic Euler method [38]:

$$\begin{aligned} p_{i_{n+1}} &= p_{i_n} - \frac{\partial}{\partial q_i} [N(q_{i_n})] \Delta t, \\ q_{i_{n+1}} &= q_{i_n} + \frac{\partial}{\partial p_i} [P(p_{i_{n+1}})] \Delta t, \end{aligned} \tag{2.33}$$

Higher-order symplectic methods can then be constructed by choosing appropriate values for c_i and d_i . However, finding the optimal values for these coefficients often involves a large system of equations which may not have an exact solution [32, 57, 103]. Nonetheless, Higher-order methods can then be constructed by composition of the symplectic Euler method using fractional steps, since the composition of two symplectic methods is itself a symplectic method [49]. The composition technique for constructing symplectic methods is presented in greater detail in Chapter 4. Moreover, by considering higher-order integrators, we can construct numerical

techniques that are both symplectic and symmetric [37]. In contrast, the explicit Euler method

$$\begin{aligned} p_{i_{n+1}} &= p_{i_n} - \frac{\partial}{\partial q_i} [N(q_{i_n})] \Delta t, \\ q_{i_{n+1}} &= q_{i_n} + \frac{\partial}{\partial p_i} [P(p_{i_n})] \Delta t, \end{aligned} \quad (2.34)$$

is not generally a symplectic method as it does not conserve the symplectic structure. Differences between symplectic and nonsymplectic numerical methods are more easily observed by comparing their behaviour during each iteration. Figure 2.3 illustrates how position and momentum values calculated during each iteration are handled differently by each method. A nonsymplectic method will update the momentum and position using data obtained exclusively from the previous iteration, while a symplectic method initially calculates the momentum, then, with this updated momentum, it obtains the position during the same iteration, resulting in a lower overall error for a greater number of iterations [30, 37, 38]. A useful consequence of this result is that symplectic numerical methods are typically more accurate than their nonsymplectic counterparts, particularly when using a relatively large step between successive iterations [37, 38].




Figure 2.3: Differences between symplectic and nonsymplectic methods, with \mathbf{p} and \mathbf{q} representing momentum and position, respectively. By making use of the momentum obtained from the same iteration to update the position, symplectic methods are often more accurate.

However, a special case exists where the explicit Euler method behaves as a symplectic method [30]. Specifically, for a system whose Hamiltonian is separable and has constant generalised momenta, the explicit Euler method is symplectic. Hence, we observe

$$\begin{aligned}\frac{\partial P}{\partial p_i} &= p_i, \\ -\frac{\partial N}{\partial q_i} &= 0.\end{aligned}\tag{2.35}$$

Now, applying the explicit Euler method

$$\begin{aligned}p_{i_{n+1}} &= p_{i_n} \quad \forall t, \\ q_{i_{n+1}} &= q_{i_n} + \Delta t p_{i_n},\end{aligned}\tag{2.36}$$

then, calculating the Jacobian of the flow

$$\begin{aligned}\frac{\partial p_{i_{n+1}}}{\partial p_{i_n}} &= 1, & \frac{\partial p_{i_{n+1}}}{\partial q_i} &= 0, \\ \frac{\partial q_{i_{n+1}}}{\partial p_{i_n}} &= \Delta t, & \frac{\partial q_{i_{n+1}}}{\partial q_{i_n}} &= 1.\end{aligned}$$

$$\frac{\partial \mathbf{z}_{n+1}}{\partial \mathbf{z}_n} = \begin{pmatrix} 1 & 0 \\ \Delta t & 1 \end{pmatrix}.$$

Finally, testing the Jacobian for symplecticity

$$\begin{aligned}\left(\frac{\partial \mathbf{z}_{n+1}}{\partial \mathbf{z}_n}\right)^T \mathbf{J} \left(\frac{\partial \mathbf{z}_{n+1}}{\partial \mathbf{z}_n}\right) &= \begin{pmatrix} 1 & 0 \\ \Delta t & 1 \end{pmatrix}^T \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ \Delta t & 1 \end{pmatrix} \\ &= \begin{pmatrix} 1 & \Delta t \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ \Delta t & 1 \end{pmatrix} \\ &= \begin{pmatrix} -\Delta t & 1 \\ -1 & 0 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ \Delta t & 1 \end{pmatrix} \\ &= \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \\ &= \mathbf{J}.\end{aligned}$$

Hence, for a separable Hamiltonian whose generalised momenta are constant, the explicit Euler satisfies Eq. (2.20) and is therefore symplectic.

2.4 Ray tracing in homogeneous media

While the previous section illustrates an interesting result within the Hamiltonian framework, it is natural to consider whether this finding has any relevance to optics. In fact, it provides a motivating example for symplectic ray tracing within homogeneous media, describing the propagation of rays through isotropic or birefringent media. Of course, exact ray paths in such systems are easily solvable. Even so, it is encouraging to see that symplectic ray tracing produces equivalent results.

From Fermat's principle of least time, the path travelled by a light ray between two points $\mathbf{A} = [q_{1A}(t), q_{2A}(t), q_{3A}(t)]$ and $\mathbf{B} = [q_{1B}(t), q_{2B}(t), q_{3B}(t)]$ is the one that takes the least amount of time [18]. This is mathematically equivalent to Eq. (2.3), thereby allowing us to define an optical Lagrangian [46] in three-dimensional space:

$$\delta S = \delta \int_{\mathbf{A}}^{\mathbf{B}} n \frac{ds}{dt} dt = \delta \int_{\mathbf{A}}^{\mathbf{B}} L dt = 0, \quad (2.37)$$

where S is the optical path (not to be confused with the action \mathcal{S}), $ds = \sqrt{dq_1^2 + dq_2^2 + dq_3^2}$ is an infinitesimal change in the geometrical path s and n is the refractive index, assumed (for now) to be constant along each coordinate axis (i.e. the medium is homogeneous). Substituting the expression for ds into Eq. (2.37), yields the optical Lagrangian:

$$L = n\sqrt{\dot{q}_1^2 + \dot{q}_2^2 + \dot{q}_3^2}, \quad (2.38)$$

Consequently, the system's optical momenta (so called due to its similarity to the concept of conjugate momentum in classical mechanics) are defined as

$$\mathbf{p} \equiv \left(\frac{\partial L}{\partial \dot{q}_1}, \frac{\partial L}{\partial \dot{q}_2}, \frac{\partial L}{\partial \dot{q}_3} \right) \quad (2.39)$$

$$= \left(\frac{n\dot{q}_1}{\sqrt{\dot{q}_1^2 + \dot{q}_2^2 + \dot{q}_3^2}}, \frac{n\dot{q}_2}{\sqrt{\dot{q}_1^2 + \dot{q}_2^2 + \dot{q}_3^2}}, \frac{n\dot{q}_3}{\sqrt{\dot{q}_1^2 + \dot{q}_2^2 + \dot{q}_3^2}} \right) \quad (2.40)$$

$$= \left(n \frac{dq_1}{ds}, n \frac{dq_2}{ds}, n \frac{dq_3}{ds} \right) \quad (2.41)$$

$$= (p_1, p_2, p_3). \quad (2.42)$$

Following algebraic manipulation of the optical Lagrangian

$$\begin{aligned} L &= \dot{q}_1 \frac{n\dot{q}_1}{\sqrt{\dot{q}_1^2 + \dot{q}_2^2 + \dot{q}_3^2}} + \dot{q}_2 \frac{n\dot{q}_2}{\sqrt{\dot{q}_1^2 + \dot{q}_2^2 + \dot{q}_3^2}} + \dot{q}_3 \frac{n\dot{q}_3}{\sqrt{\dot{q}_1^2 + \dot{q}_2^2 + \dot{q}_3^2}} \\ &= \dot{q}_1 p_1 + \dot{q}_2 p_2 + \dot{q}_3 p_3. \end{aligned} \quad (2.43)$$

Recalling the definition of the Hamiltonian from Eq. (2.6)

$$\begin{aligned} H &= \sum_{i=1}^k \dot{q}_i p_i - L \\ &= 0. \end{aligned} \quad (2.44)$$

We see that the H vanishes identically. This gives us some direction in choosing the correct form for the optical Hamiltonian. In this instance, we make use of the following ansatz:

$$H = \frac{1}{2} [\|\mathbf{p}\|^2 - n^2(\mathbf{q})] = 0. \quad (2.45)$$

Further details on the rationale behind this choice for the optical Hamiltonian are given in Chapter 3. Now, we realise that the use of a harmonic oscillator as an example in Section 2.3 was no coincidence; Eq. (2.45) and Eq. (2.23) have a similar form, save for a few crucial differences. First, we notice that the sign of the second term in Eq. (2.45) is negative, meaning its phase space trajectory is in the opposite direction to Figure 2.1. Second, the function chosen for the refractive index will determine whether our “oscillator” is “damped” or “driven”, which could cause the phase space trajectory to spiral inwards or outwards rather than remain periodic. First, however, we take the gradient of the optical path

$$\begin{aligned} \nabla S &= n \frac{d\mathbf{q}}{ds} \\ &= \mathbf{p}. \end{aligned} \quad (2.46)$$

Moreover, considering the eikonal equation [18]

$$\left(\frac{\partial S}{\partial q_1}\right)^2 + \left(\frac{\partial S}{\partial q_2}\right)^2 + \left(\frac{\partial S}{\partial q_3}\right)^2 = n^2, \quad (2.47)$$

we see $p_1^2 + p_2^2 + p_3^2 - n^2 = 0$, indicating our choice for the optical Hamiltonian in Eq. (2.45) is suitable. From Eqs. (2.46) and (2.47), the differential equation of light rays [58] is then derived.

$$\frac{d}{ds} \left(n(\mathbf{q}) \frac{d\mathbf{q}}{ds} \right) = \nabla n(\mathbf{q}), \quad (2.48)$$

from Eqs. (2.45) and (2.48), we can demonstrate

$$\frac{d\mathbf{p}}{ds/n} = -\frac{\partial H}{\partial \mathbf{q}}. \quad (2.49)$$

By defining $dt \equiv \frac{ds}{n}$, Eq. (2.49) is rewritten as

$$\frac{d\mathbf{p}}{dt} = -\frac{\partial H}{\partial \mathbf{q}}, \quad (2.50)$$

which is identical to Eq. (2.14). However, in order to avoid confusion, some caution is warranted. Here, dt does not represent a small, finite time step, but rather a small finite change in the geometric path of the ray, divided by the refractive index of the medium. Nonetheless, a useful consequence of parameterising Eq. (2.49) in terms of dt is that the optical momenta reduce to the direction cosines $\mathbf{p} = (\cos \alpha_{q_1}, \cos \alpha_{q_2}, \cos \alpha_{q_3})$, where $\alpha_{q_1}, \alpha_{q_2}$ and α_{q_3} are the angles made with axes q_1, q_2 and q_3 respectively. In a manner similar to Eq. (2.49), one can also demonstrate

$$\frac{d\mathbf{q}}{dt} = \frac{\partial H}{\partial \mathbf{p}}. \quad (2.51)$$

Returning to Eq. (2.48), since we assumed the refractive index is constant $\dot{p}_1 = \dot{p}_2 = \dot{p}_3 = 0$, indicating optical momentum is conserved in accordance with Snell's law. Hence, p_1, p_2 and p_3 are each constant. Likewise, from Eq. (2.13)

$$\dot{q}_1 = \frac{1}{2} \frac{\partial}{\partial p_1} (p_1^2 + p_2^2 + p_3^2 - n^2) = p_1, \quad (2.52)$$

$$\dot{q}_2 = \frac{1}{2} \frac{\partial}{\partial p_2} (p_1^2 + p_2^2 + p_3^2 - n^2) = p_2, \quad (2.53)$$

$$\dot{q}_3 = \frac{1}{2} \frac{\partial}{\partial p_3} (p_1^2 + p_2^2 + p_3^2 - n^2) = p_3. \quad (2.54)$$

Substituting from Eqs. (2.52)–(2.54) into Eq. (2.33)

$$q_1 = q_{10} + p_1 t, \quad (2.55)$$

$$q_2 = q_{20} + p_2 t, \quad (2.56)$$

$$q_3 = q_{30} + p_3 t, \quad (2.57)$$

or in vector form, the system can be written as

$$\mathbf{q} = \mathbf{q}_0 + \mathbf{p}t, \quad (2.58)$$

which is an exact description of light rays traversing a homogeneous medium. This formulation is applicable to isotropic media or birefringent media, where t may replace Δt in Eqs. (2.55)–(2.58). However, the power of symplectic ray tracing becomes particularly apparent when it is applied to gradient-index media, where the ray trajectories are not necessarily available in a closed form, making the use of numerical methods an absolute necessity. The following chapter explores the suitability of symplectic numerical methods for nonlinear ray tracing and compares their performance with popular nonsymplectic methods.

Chapter 3

Numerical experiments with spherical GRIN lenses

3.1 An alternative derivation of the optical Hamiltonian

The previous chapter demonstrated how to derive the optical Lagrangian as a consequence of Fermat's principle before performing a Legendre transform on this Lagrangian to obtain the associated optical Hamiltonian [46, 94, 100]. Here, however, we dispense with this procedure for the sake of brevity, showing that it is also possible to derive the necessary optical Hamiltonian directly from the eikonal equation [74]

$$\left(\frac{\partial S}{\partial x}\right)^2 + \left(\frac{\partial S}{\partial y}\right)^2 + \left(\frac{\partial S}{\partial z}\right)^2 = n^2(x, y, z), \quad (3.1)$$

Taking the total differential dS of the optical path, via the chain rule, we find

$$dS = \nabla S \cdot d\mathbf{q}, \quad (3.2)$$

where ∇S is the gradient of the optical path vector and $d\mathbf{q}$ is the differential of the position vector $\mathbf{q} = (x, y, z)^T$. Recalling that by definition $dS = n ds$, therefore we may write

$$n ds = \nabla S \cdot d\mathbf{q}, \quad (3.3)$$

then, multiplying both sides by n/ds

$$n^2 = \nabla S \cdot n \frac{d\mathbf{q}}{ds}, \quad (3.4)$$

and $n^2 = |\nabla S|^2 = \nabla S \cdot \nabla S$ from Eq. (3.1),

$$\nabla S = n \frac{d\mathbf{q}}{ds}, \quad (3.5)$$

which is identical to the optical momentum defined in Eq. (2.42). Hence,

$$\mathbf{p} = n \frac{d\mathbf{q}}{ds}. \quad (3.6)$$

Now, differentiating \mathbf{p} with respect to s

$$\begin{aligned} \frac{d}{ds} \left(n \frac{d\mathbf{q}}{ds} \right) &= \frac{d}{ds} (\nabla S) \\ &= \frac{d\mathbf{q}}{ds} \cdot \nabla (\nabla S), \end{aligned} \quad (3.7)$$

then, rearranging Eq. (3.5) $(\nabla S)/n = d\mathbf{q}/ds$

$$\begin{aligned} \frac{d}{ds} \left(n \frac{d\mathbf{q}}{ds} \right) &= \frac{\nabla S}{n} \cdot \nabla (\nabla S) \\ &= \frac{\nabla (\nabla S^2)}{2n}. \end{aligned} \quad (3.8)$$

Recalling $|\nabla S|^2 = n^2$ from Eq. (3.1) once again

$$\frac{d}{ds} \left(n \frac{d\mathbf{q}}{ds} \right) = \frac{\nabla n^2}{2n}. \quad (3.9)$$

Finally, multiplying both sides by n

$$n \frac{d\mathbf{p}}{ds} = \nabla \left(\frac{n^2}{2} \right). \quad (3.10)$$

Defining $dt = ds/n$, we now rewrite Eq. (3.1) as a system of six ordinary differential equations

$$\begin{aligned} \frac{dx}{dt} &= p_x, & \frac{dp_x}{dt} &= \frac{\partial}{\partial x} \left(\frac{n^2}{2} \right), \\ \frac{dy}{dt} &= p_y, & \frac{dp_y}{dt} &= \frac{\partial}{\partial y} \left(\frac{n^2}{2} \right), \\ \frac{dz}{dt} &= p_z, & \frac{dp_z}{dt} &= \frac{\partial}{\partial z} \left(\frac{n^2}{2} \right). \end{aligned} \quad (3.11)$$

By choosing to parameterise our system in terms of t , we transform Eq. (3.1) from a partial differential equation into a system of Hamilton's equations for some optical Hamiltonian H .

Thus, we may now write.

$$\begin{aligned}
 p_x &= \frac{\partial H}{\partial p_x}, & \frac{\partial}{\partial x} \left(\frac{n^2}{2} \right) &= -\frac{\partial H}{\partial x}, \\
 p_y &= \frac{\partial H}{\partial p_y}, & \frac{\partial}{\partial y} \left(\frac{n^2}{2} \right) &= -\frac{\partial H}{\partial y}, \\
 p_z &= \frac{\partial H}{\partial p_z}, & \frac{\partial}{\partial z} \left(\frac{n^2}{2} \right) &= -\frac{\partial H}{\partial z},
 \end{aligned} \tag{3.12}$$

Solving for H in (3.12) is then a straightforward exercise, providing us with the optical Hamiltonian. We see it is identical in form to Eq. (2.45), i.e. the ansatz made in the previous chapter.

$$H = \frac{1}{2} [p_x^2 + p_y^2 + p_z^2 - n^2(x, y, z)]. \tag{3.13}$$

3.1.1 A brief remark on alternatives to ray tracing

While ray tracing is among the highest fidelity methods for image formation [85], other methods for solving the eikonal equation have arisen outside of the optical domain in fields such as geophysics and control theory [53, 67]. Rather than taking a Hamiltonian ray-optical approach, these alternative techniques aim to solve for the optical path directly from the eikonal equation. Chief among them is the fast marching method [81], developed in the 1990s by James Sethian, which divides a scene into a grid of nodes and iteratively solves for the optical path directly, building a solution that is informed by the information available from the boundary conditions.

Another method related to the fast marching method is the fast sweeping method [104], first proposed in 2005 by Hong-Kai Zhao. Like the fast marching method, it directly solves for the optical path, but it uses an upwind differencing scheme to discretise the domain on which the equation is to be solved. Once the initial optical path values are assigned, Gauss-Seidel iteration is used to update them until the optical path values converge satisfactorily. Fast sweeping and fast marching approaches have been compared extensively [17]. The fast sweeping method tends to outperform the fast marching method on problems where the characteristics are mostly straight lines (i.e. for plane wavefronts). By contrast, the fast marching method performs better when the object we wish to image is an irregular shape.

Nevertheless, for optical design and testing of elements, numerical ray tracing methods are typically sufficient for our needs. Moreover, in our experiments with point sources and collimated rays, the benefit offered by the fast marching method for complicated object geometries simply isn't warranted. Additionally, while the numerical methods used here solve for the ray trajectory in a series of linear steps, this may be to our advantage in the simulation of lenses which can be practically fabricated. For instance, spherical lenses are constructed by repeated stacking

of increasingly larger spherical shells, as shown in Figure 3.1. Each shell is of a finite width (Δt , say) whose refractive indices differ from each other by $\Delta n \ll 1$. Thus, a numerical trace constructed by a sequence of linear steps, each of size Δt could, in effect, be exact for a real GRIN lens. Here, the ray tracing methods used are iterative methods, typically written in the form $\mathbf{q}_{n+1} = f(\mathbf{q}_n)$. In other words, each update depends solely on information from previous iterations, though Chapter 4 considers methods with update schemes that are implicit in \mathbf{q}_{n+1} .




Figure 3.1: A schematic of a spherical GRIN lens of radius $R_0 = k\Delta t$, constructed from k spherical shells. Each shell is of width Δt and its refractive index will differ from subsequent shells by Δn , with the innermost shell having a refractive index of n_0 .

3.2 Results of numerical experiments

In order to test the validity of symplectic methods for solving optical problems, simulated traces with a monochromatic light source of wavelength $\lambda = 589.3$ nm were carried out within a selection of spherically symmetric lenses, each of radius $R_0 = 1$ mm, all of which have been extensively studied in existing optical literature. The Lüneburg [52] and Maxwell fish-eye [55] lenses are classical GRIN elements which have found applications as directional antennae for radio and microwave devices [24, 54] and more recently at optical wavelengths [105]. While the Lüneburg lens will direct a collimated bundle of rays to a single focus on its surface, the Maxwell fish-eye is described as an absolute optical instrument, theoretically capable of perfect point-to-point imaging in three-dimensional space [97].

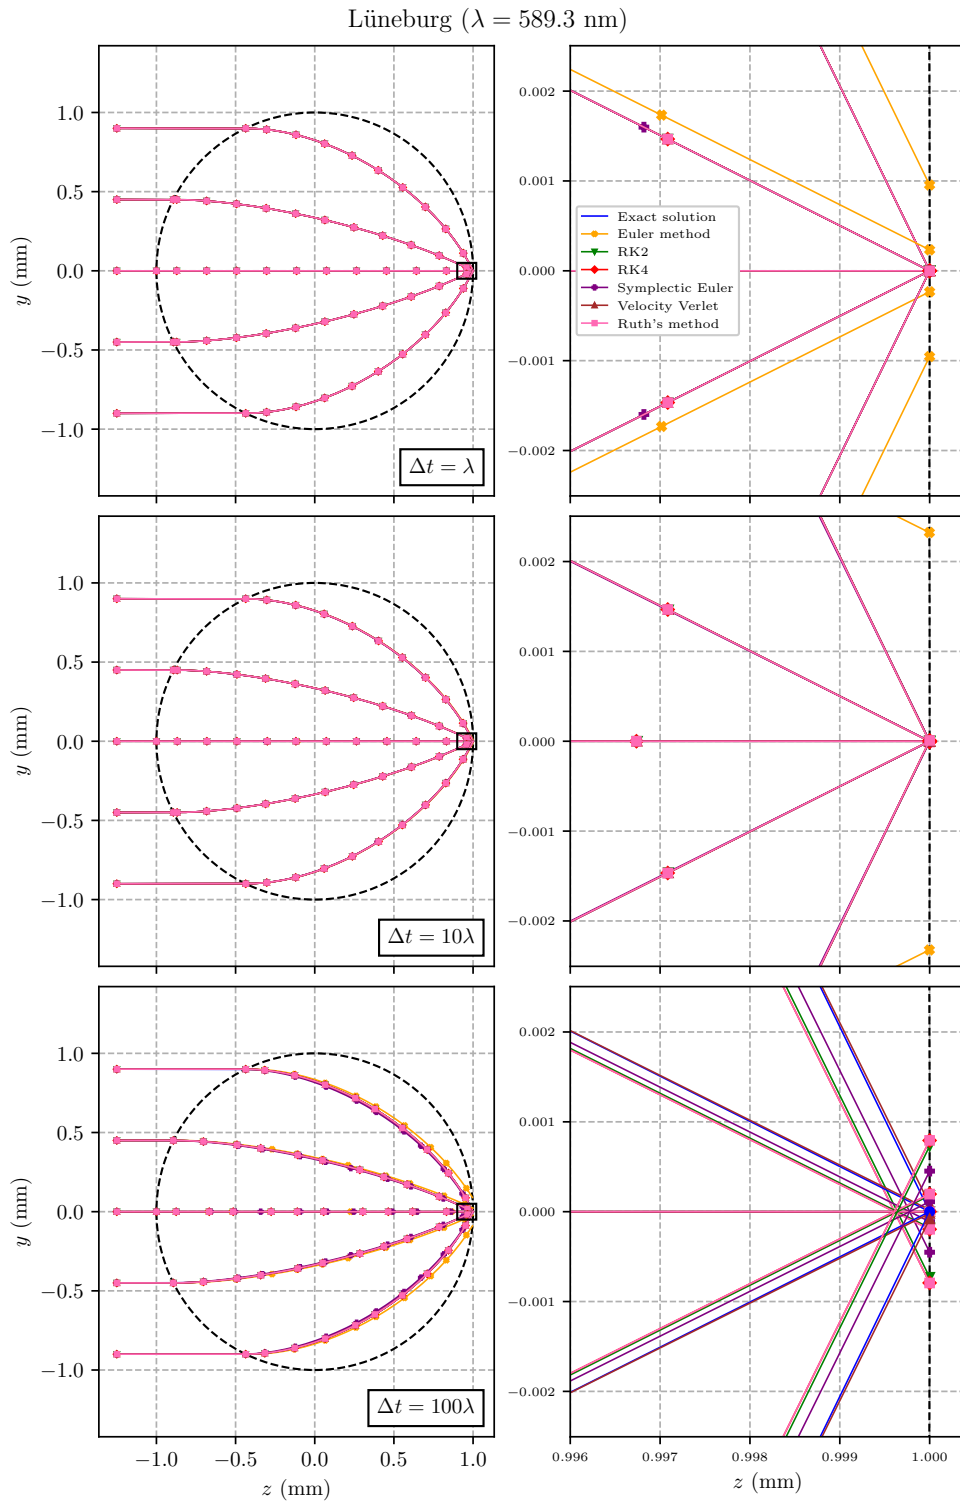


Figure 3.2: Numerical traces through a Lüneburg lens of radius R_0 . The region shown in each right-hand column is marked on each trace in the left-hand column by a black rectangle and depicts the focal point in greater detail.

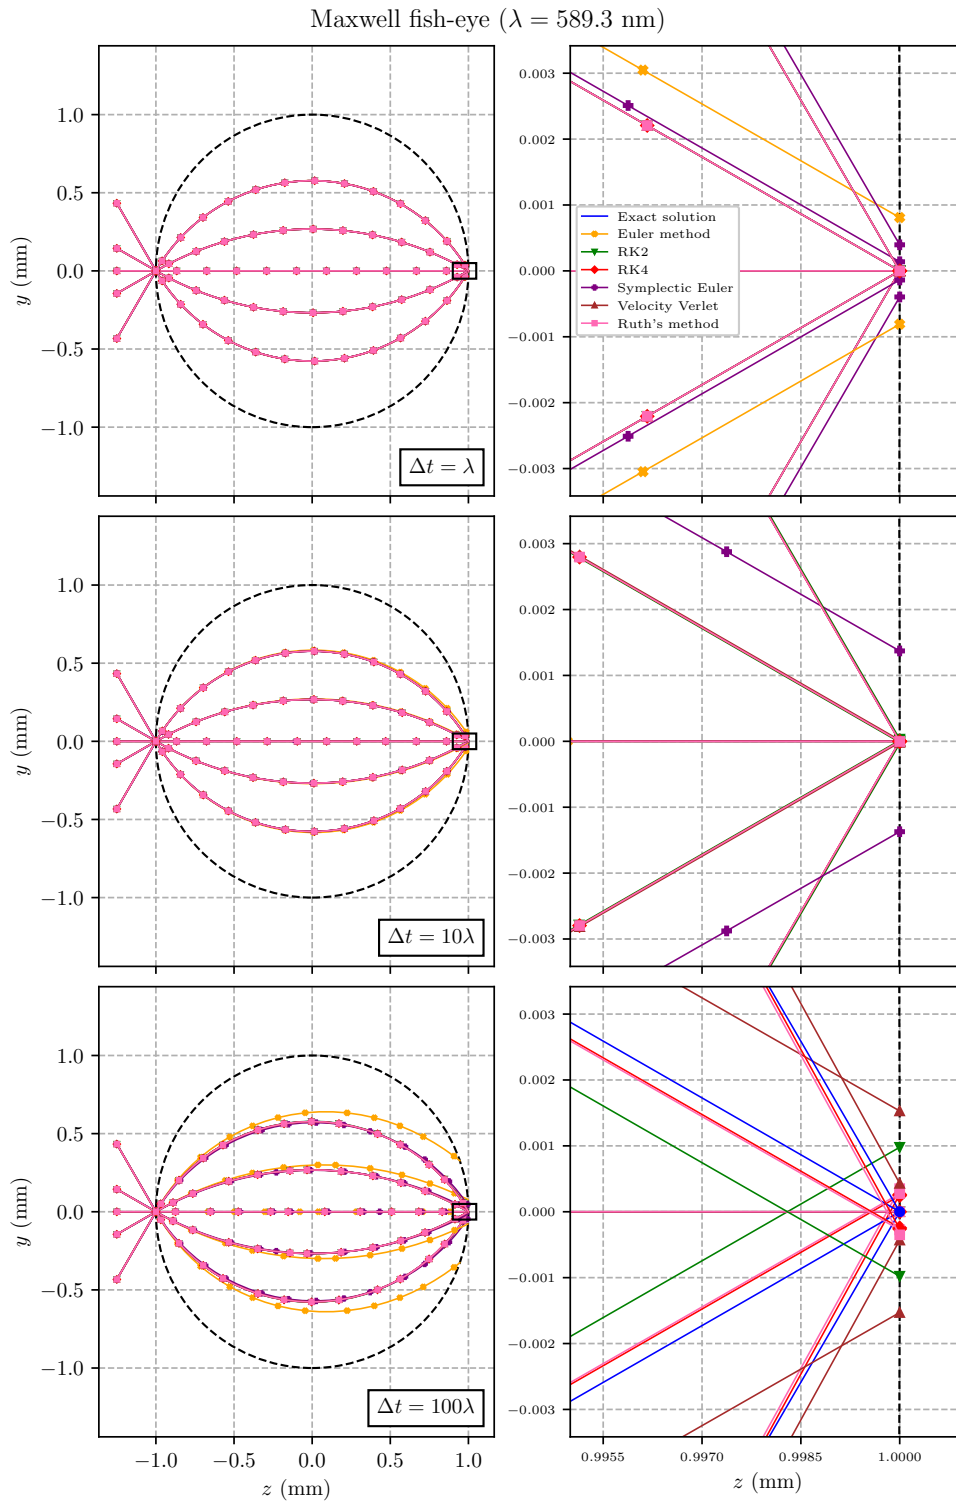


Figure 3.3: Numerical traces through a Maxwell fish-eye lens of radius R_0 . The right-hand column once again depicts the focal point in greater detail.

The Eaton lens [27] and the optical black hole (also referred to as a concentrator lens) [69] differ from the Lüneburg and Maxwell fish-eye lenses with refractive index profiles that are singular at their centres, providing a significant challenge to numerical methods which may be used to trace these lenses. Thankfully, however, we may take advantage of analogies which exist between such singular GRIN lenses and systems examined in other fields; the ray path within the Eaton lens, for instance, is effectively identical to the trajectory of a body in the Kepler problem of celestial mechanics [13]. The Eaton lens is gradually finding applications as a beam deflector [26], while the optical black hole could prove useful for laboratory astrophysics [44].

Table 3.1 provides further details about each of the four lenses used during numerical experiments where six numerical methods were chosen to trace each lens. However, for the Eaton lens, the symplectic Chin-Chen 4A method [19] developed for celestial mechanics was also tested on account of the Eaton lens' behaviour as an optical analogue for celestial mechanics. Table 3.2 provides information on the order and symplecticity of each method used, with their associated numerical recipes being presented in Appendix A.1.

Lens name	Index profile	Singularity present?
Lüneburg	$n(y, z) = \sqrt{2 - \frac{y^2+z^2}{R_0^2}}$	No
Maxwell fish-eye	$n(y, z) = 2 \left(1 + \frac{y^2+z^2}{R_0^2}\right)^{-1}$	No
Eaton	$n(y, z) = \sqrt{2\sqrt{\frac{R_0^2}{y^2+z^2}} - 1}$	Yes
Optical black hole	$n(y, z) = \sqrt{\frac{R_0^2}{y^2+z^2}}$	Yes

Table 3.1: The lenses used in these numerical experiments and their respective index profiles. R_0 represents the radius of each lens. All lenses are configured such that no refraction will take place at the surface of each lens when it is surrounded by a vacuum.

Since the Lüneburg and Maxwell fish-eye lenses directed all rays to a single focal point, it was possible to trace multiple rays and examine their behaviour simultaneously. However, as the Eaton lens and optical black hole direct each individual ray to a unique focus, only a single ray was traced in these lenses. The initial ray heights for the single rays in the Eaton lens and optical black hole were $y_0 = 0.58R_0$ and $y_0 = 0.94R_0$ respectively. These values were chosen so each numerical method would be challenged to maintain its accuracy in the vicinity of the singularity and avoid producing erroneous unphysical results.

Figures 3.2 and 3.3 depict numerical ray traces through the Lüneburg and Maxwell fish-eye lenses respectively, with the left-hand column depicting the traces in their entirety while the right-hand column presents a detailed view of the focal points. The step size values range from $\Delta t = \lambda$ to $\Delta t = 100\lambda$, where $\lambda = 589.3$ nm. The exact solution is depicted in each case by a continuous blue line. As one might expect, the disparity between the numerical methods decreases with a reduction in Δt and an increase in method order m . For the Lüneburg lens,

all methods with the exception of the Euler method are indistinguishable from each other and the exact solution for $\Delta t = \lambda$. Significant discrepancies between each of the methods are not readily visible until $\Delta t = 100\lambda$, where each of the methods can be easily identified, though significant overlap is still observed between RK2, RK4 and Ruth’s method. For Maxwell’s fish-eye, however, both the Euler method and the symplectic Euler method deviate noticeably at $\Delta t = \lambda$, with both methods disappearing from the inset when Δt is increased to 100λ . Each of the second- and fourth-order methods then follow suit, accumulating greater errors and deviating significantly from the expected focal point as Δt reaches 100λ . RK2, in particular, appears to substantially underestimate the focal distance while the Velocity Verlet method displays the opposite behaviour, overestimating the focal distance, though not to the same extent. RK4 and Ruth’s method both appear to underestimate the focal distance by approximately the same amount, indicating unsurprisingly that any differences between symplectic and nonsymplectic methods appear to diminish with an increase in method order for the same choice of step size.

Method name	Order m	Symplectic?
Euler method	1	No
Symplectic Euler method [38]	1	Yes
RK2	2	No
Velocity Verlet [98]	2	Yes
RK4 [82]	4	No
Ruth’s method [32]	4	Yes
Chin-Chen 4A [19]	4	Yes

Table 3.2: The numerical methods used to ray trace each lens. The Chin-Chen 4A method is applied to the Eaton lens only, owing to its similarity to the Kepler problem.

For both the Lüneburg and Maxwell fish-eye lenses, Figure 3.6 indicates that symplectic methods tend to exhibit lower focus errors than their nonsymplectic counterparts of the same order, with the focus error being defined as the difference between the numerical and exact focus along the lens perimeter, divided by the lens circumference, normalised in terms of the lens radius R_0 . The dashed horizontal line represents the Airy disk radius, providing a benchmark for diffraction-limited imaging. Consequently, any error values below this line may be considered exact. Most noticeably, the symplectic Euler method vastly outperforms the standard Euler method by several orders of magnitude for the Lüneburg lens and approximately one order of magnitude for the Maxwell fish-eye. The initial conditions presented in Figure 3.6 for the Lüneburg and Maxwell fish-eye lenses are associated with the rays in Figures 3.2 and 3.3 whose optical momentum experiences the greatest rate of change. Differences between the fourth-order methods are much less drastic, with both RK4 and Ruth’s method producing errors on the same order of magnitude. However, for second-order methods, the results are much more surprising; the velocity Verlet method is the only method for which the numerical trace is diffraction limited for all step sizes examined. It not only outperforms RK2, but also RK4 and Ruth’s method,

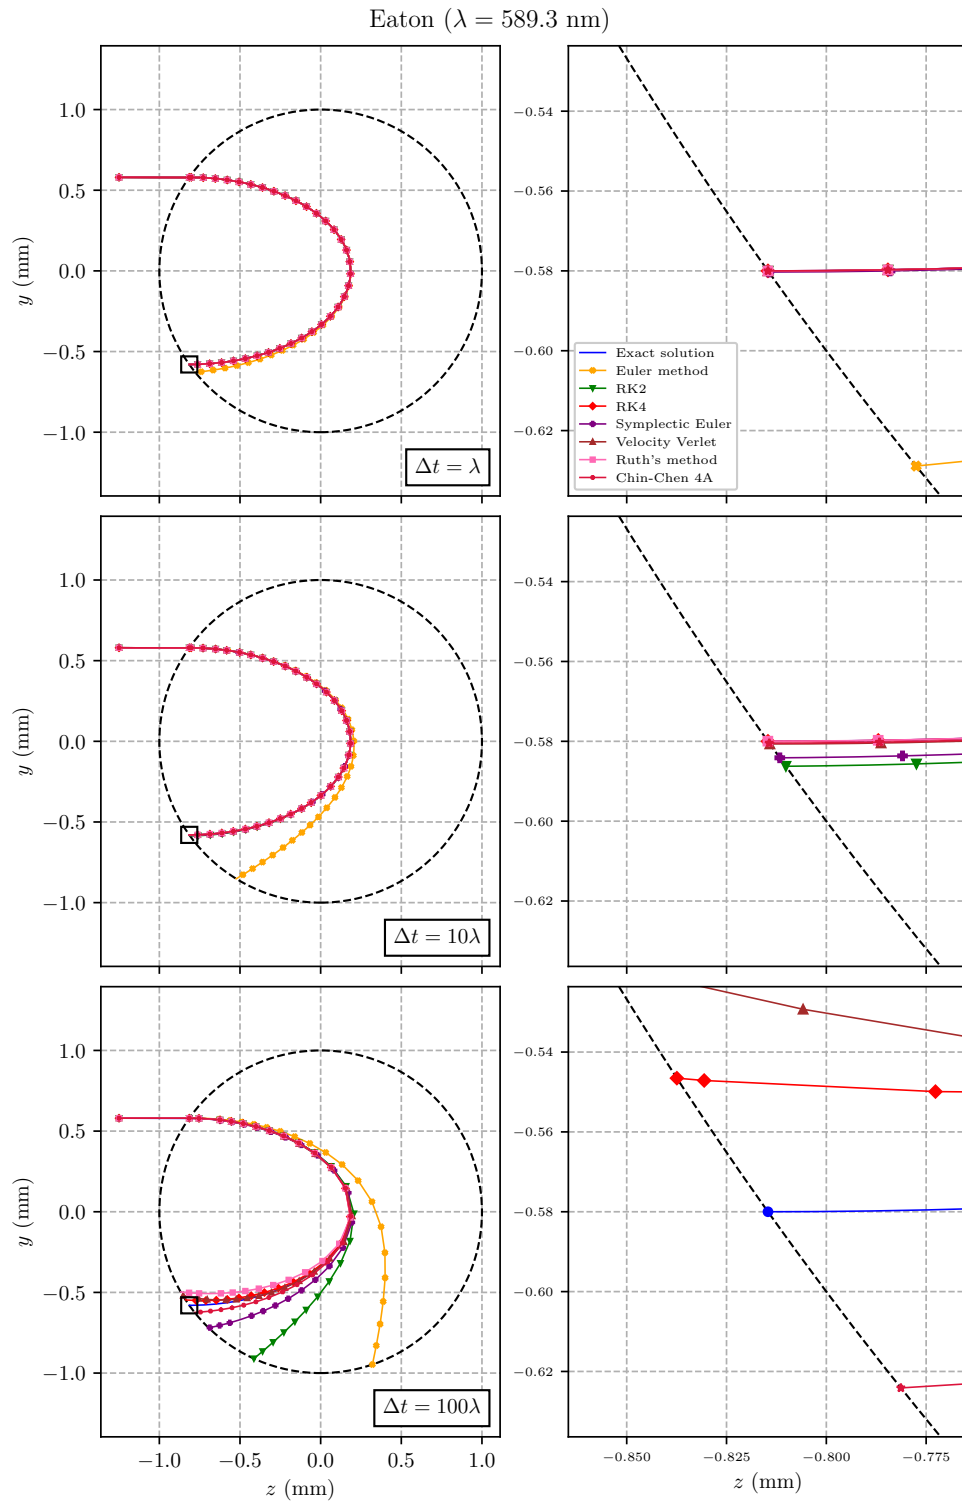


Figure 3.4: Numerical traces through an Eaton lens of radius R_0 . The left- and right-hand columns respectively show the ray trace for each Δt and a detailed view of the focus, as before.




Figure 3.5: Numerical traces through an optical black hole of radius R_0 . The shaded region on each plot represents the event horizon of a Schwarzschild black hole, providing a stop condition for the ray trace.

particularly when $\Delta t > 5\lambda$ for the Lüneburg lens and when $\Delta t = 2.5\lambda$ or 25λ for the Maxwell fish-eye, suggesting that the increased computational expense incurred by using higher-order methods is not necessarily justified by greater accuracy. This is especially true of the Lüneburg lens when $\Delta t = \lambda$, where the velocity Verlet method is exact to within machine precision.

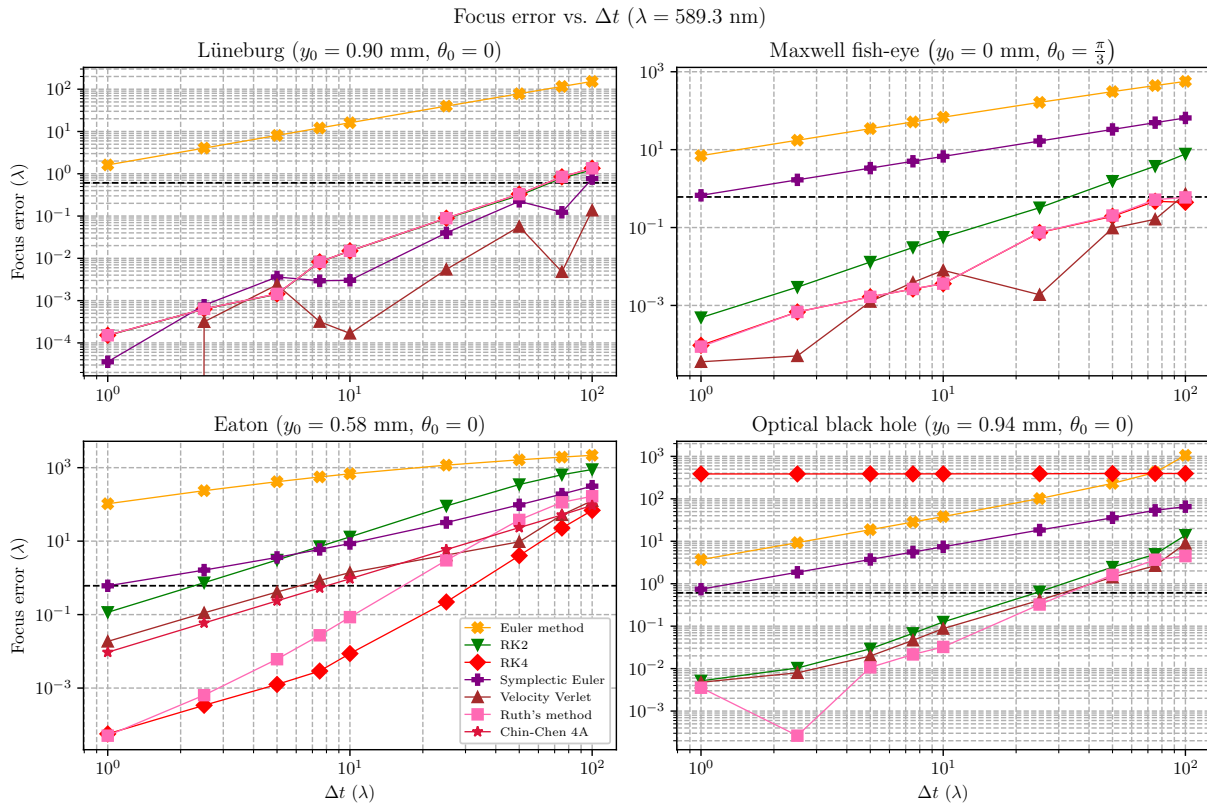


Figure 3.6: Error in ray tracing each lens with numerous step sizes. The error is defined as the distance between the exact and numerical focus along the lens perimeter, divided by the lens circumference, normalised in terms of the lens radius R_0 . The dashed line represents the diffraction limit for the wavelength $\lambda = 589.3$ nm.

A similar analysis is presented for the Eaton lens and optical black hole in Figures 3.4 and 3.5, where the symplectic methods once again generally appear to offer superior performance to nonsymplectic methods, particularly at larger step sizes. For the Eaton lens with $\Delta t = 10\lambda$, the symplectic Euler method not only outperforms the standard Euler method (which is only contained within the inset for $\Delta t = \lambda$) but also proves more accurate than RK2. For $\Delta t = 100\lambda$, both RK4 and the Chin-Chen 4A method appear to offer the most accurate solutions, suggesting that the use of methods developed outside optics could be of benefit for numerical ray tracing also. Although it does not dominate all other methods to the same extent as it did with the Lüneburg lens and Maxwell fish-eye, the velocity Verlet method, still appears to perform somewhat well, particularly when compared with RK2, with the velocity Verlet method focusing just outside the inset. In contrast, Ruth's method performs somewhat poorly here, requiring a

smaller step size in order to match the results produced by RK4 and the velocity Verlet method. Nonetheless, Ruth’s method does mimic their behaviour at least in a qualitative sense by overestimating the focusing ability of the Eaton lens. The Chin-Chen 4A method instead follows similar trajectories to all other methods, where the focusing ability of the lens is underestimated. Thus, the combination of an overestimating method with an underestimating one would perhaps form the basis for a successful predictor-corrector scheme, though this would presumably come at an elevated computational cost, especially if the methods chosen are not symplectic.

In contrast with the retroreflecting behaviour of the Eaton lens, the singular index profile of the optical black hole directs rays along an infinite inward spiral. However, the shaded grey circle in Figure 3.5 provides a practical stopping condition and represents the event horizon of a Schwarzschild black hole with radius $0.67R_0$. Despite the necessary truncation of the ray paths, some surprising results may still be observed. For instance, RK4 exhibits uncharacteristically poor performance across all step sizes. Moreover, it is surpassed by all symplectic methods and even the standard Euler method when $\Delta t = 100\lambda$. Unlike the Eaton lens, however, RK2 performs similarly to the Velocity Verlet method across all step sizes, providing a marked difference in its ray tracing capabilities for the two singular-index lenses. Finally, the symplectic Euler method outperforms the nonsymplectic Euler method by a considerable margin, particularly for $\Delta t \geq 10\lambda$, further emphasising the gains in accuracy which may be made through the use of symplectic numerical methods, particularly if a low-order method is desired. Furthermore, as was the case with the Eaton lens, The combination of methods which both over- and underestimate the focal distance (such as the velocity Verlet and Ruth’s methods, for instance) could once again form a useful predictor-corrector method, enabling a larger step size to be used.

On the whole, Symplectic methods demonstrate more consistent performance across both singular lenses. Though the higher-order methods do not necessarily dominate nonsymplectic methods to the same extent as they did for the Lüneburg and Maxwell fish-eye lenses, they tend to at least perform as well as the more established nonsymplectic algorithms at smaller step sizes. This is especially true of the optical black hole, where Ruth’s method exhibits the lowest error of all methods for $\Delta t = \lambda$, while RK4’s focus error remains almost constant for all chosen step sizes, as depicted in Figure 3.6. Differences in error between the symplectic and standard Euler methods for both the Eaton lens and optical black hole follow a similar general trend of decreasing with the step size, though the error values tend to oscillate less than those of the Lüneburg lens and Maxwell fish-eye. A notable exception, however, is Ruth’s method with $\Delta t = 2.5\lambda$ for the optical black hole, where the error increases once again for $\Delta t = \lambda$. By contrast, for the Eaton lens, RK4 still appears to be the most accurate choice for most step sizes, being the only method which offers a diffraction limited image for a step size greater than 10λ . Nonetheless, differences in accuracy between RK4 and Ruth’s method for the Eaton lens become almost negligible when $\Delta t = \lambda$. Additionally, the Chin-Chen 4A method appears to

underperform at smaller step sizes, with the velocity Verlet method making it almost redundant by offering similarly accurate solutions while demanding less computational resources.

Nevertheless, performing calculations at quadruple rather than double precision would enable us to examine any trends further and ascertain whether or not RK4 continues to hold its own versus Ruth's method for ray tracing within the Eaton lens and also check the constancy of its focus error when tracing the optical black hole. Likewise, any differences between RK2 and the velocity Verlet method within the optical black hole could also be considered in greater detail. Indeed, the impressive performance of the velocity Verlet method within the Lüneburg and Maxwell fish-eye lenses could also be further tested by quadruple-precision numerical experiments.

3.2.1 A note of caution on higher-order methods

While the velocity Verlet and Ruth's methods certainly appear to perform well for spherical GRIN profiles, they are not necessarily the most suitable methods in all circumstances. Specifically, radial profiles seem to provide a significant challenge to each of these methods. As a concrete example, we will examine the Mikaelian lens (also known commercially as the SELFOC lens) [61], a rod lens with planar end surfaces capable of focusing axially collimated rays to a point on its back surface. The necessary index profile is that of a hyperbolic secant, i.e.

$$n(y) = n_0 \operatorname{sech} \left(\frac{\pi y}{2l} \right), \quad (3.14)$$

where n_0 is the refractive index on the optical axis and l is the length of the lens. The relevant optical Hamiltonian is then written as

$$H_{\text{Mikaelian}} = \frac{1}{2} \left[p_y^2 + p_z^2 - n_0^2 \operatorname{sech}^2 \left(\frac{\pi y}{2l} \right) \right]. \quad (3.15)$$

Like each of the previous spherical lenses, an analytic trace also exists for the Mikaelian lens, with periodic ray trajectories given by

$$y(z) = \frac{2l}{\pi} \operatorname{arcsinh} \left[\sinh \left(\frac{\pi y_0}{2l} \right) \cos \left(\frac{\pi z}{2l} \right) \right], \quad (3.16)$$

where y_0 is the ray height at the front surface of the lens. Since these trajectories are periodic, the Mikaelian lens finds applications mainly in integrated optics and microwave antennae [60]. Figure 3.7 shows that even for $\Delta t = \lambda$, significant errors are present for all numerical methods, though distinguishing one method from another is rather challenging. However, by increasing the step size by an order of magnitude to 10λ , we see that even the standard Euler method performs better than the velocity Verlet method, raising serious concerns about the applicability of the velocity Verlet method to axial or radial index profiles. The method then redeems itself to some degree for $\Delta t = 100\lambda$, where it outperforms the standard Euler RK2 and RK4 methods. Ruth's




Figure 3.7: Numerical traces through a Mikaelian lens of length l .

method appears not to perform particularly well either, though its error seems not to increase substantially with the step size and performs considerably better than RK4, most noticeably for $\Delta t = 100\lambda$. The most unusual result is the excellent performance of the symplectic Euler method, which gives a significantly more accurate solution than all higher-order methods presented here.

Considering the change in focus error (defined for the Mikaelian lens as the height above the focus on the lens back surface) with step size for the marginal ray in Figure 3.8, we observe that the RK2, RK4, standard Euler and velocity Verlet methods follow repeated exponential trends, with the second-order methods growing at the fastest rate followed by RK4 and then by the standard Euler method. Still, each symplectic method is more accurate than its nonsymplectic counterpart of the same order. Curiously, the error associated with Ruth's method mains essentially flat for all step sizes, much like RK4 did when tracing the optical black hole. However, the most unusual result is given by the symplectic Euler method, further confirming our observations in Figure 3.7. However, its focus error is only below the diffraction limit for $\Delta t = 50\lambda$, suggesting that the optimum step size is not necessarily the smallest in this instance. For smaller step sizes, the error then increases rapidly, tending towards the values obtained by Ruth's method.

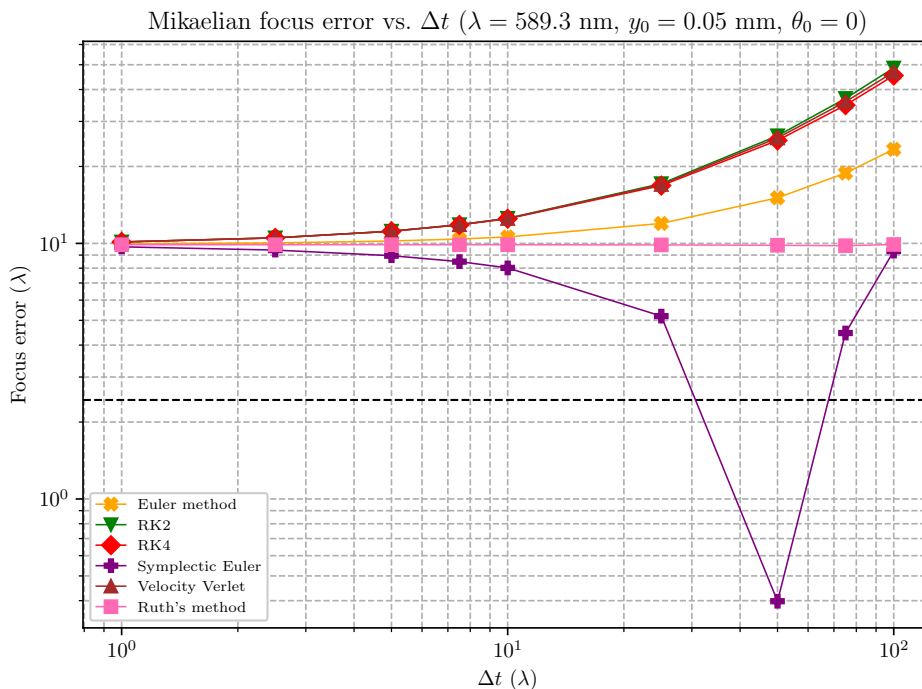


Figure 3.8: Error in ray tracing the Mikaelian lens with numerous step sizes. The error is defined as the height above the ideal focus on the lens' back surface. The dashed line once again represents the diffraction limit for the wavelength $\lambda = 589.3$ nm.

Overall, these results seem to suggest that including higher-order terms may sometimes adversely affect the accuracy of numerical traces, particularly if the error associated with those terms is

not necessarily balanced out by the lower-order terms. For instance, neglecting terms of order $(\Delta t)^2$ from the velocity Verlet method gives a solution which is exact to within machine precision for $\Delta t = \lambda$. Another valid approach could be to solve the optimisation problem of finding the most suitable step size for the symplectic Euler method when tracing the Mikaelian lens.


3.3 Applying symplectic methods to image rendering

Although it may be some time yet before the GRIN lenses considered in the previous section are incorporated into consumer optical devices, symplectic ray tracing finds a more immediate application in nonlinear ray tracing for rendering computer images. Considering the consistent performance of both the velocity Verlet method and Ruth’s method when tracing spherical index profiles, we now examine their suitability for ray tracing a black hole test image, with RK4 given as a nonsymplectic comparison. The ray trace programme, written in Python, was originally made available on GitHub [2], requiring some modifications to remove dependencies on deprecated functions and the inclusion of the symplectic methods within the existing programme.


1080p samples of the black hole test images are shown in Figure 3.9. In a Cartesian coordinate system whose origin is at the centre of the black hole’s event horizon, the observer is located at $(x, y, z) = (0, r_S, -20r_S)$, where r_S is the black hole’s Schwarzschild radius (i.e. the radius of its event horizon). All image rendering took place on a Dell™ Inspiron 5570 computer with an Intel® Core™ i7 CPU and 8 GB of DDR4 RAM. Each ray was then traced from the observer’s location for a total of 250 iterations with a step size of $0.16r_S$. The rendering workload was balanced as equally as possible between four threads, with each thread running in parallel, assigned to one of the four CPU cores. At first glance, no significant differences can be seen between each of the three images in Figure 3.9. However, Figure 3.10 presents a normalised pixel-wise difference map between both symplectic methods and RK4. Some noticeable differences between the velocity Verlet method and RK4 may then be observed, especially in the vicinity of the black hole’s event horizon and at the edge of its accretion disk. By contrast, any discrepancies between Ruth’s method and RK4 are substantially more difficult to distinguish. The radius of the event horizon in pixels for each image was then calculated, with the results displayed in Table 3.3. We see that the resolution of RK4 and Ruth’s method is identical, each resulting in an angular resolution per pixel of 64.0041 arcseconds for our observer. The angular resolution per pixel for the velocity Verlet method is marginally lower at 64.4041 arcseconds, though Figure 3.9 has already demonstrated this reduction in resolution is practically negligible.

Symplectic methods further demonstrate their advantage in Figure 3.11, where both the velocity Verlet and Ruth’s methods provide modest reductions in the computation times required to trace the test images at various resolutions with both 16:9 and 4:3 aspect ratios. The same computer, step size, number of iterations and observer position used to render the 1080p images in Figure

RK4



Ruth's method



Velocity Verlet




Figure 3.9: 1080p test images of a Schwarzschild black hole, showing no significant difference in image quality between the three numerical methods shown. The observer is located at $(x, y, z) = (0, r_S, -20r_S)$, where the origin of this coordinate system is at the centre of the black hole's event horizon and r_S represents its Schwarzschild radius.

Method	r_S (pixels)	Angular resolution per pixel (arcsec.)
RK4	161	64.0041
Ruth's method	161	64.0041
Velocity Verlet	160	64.4041

Table 3.3: Angular resolution per pixel for each of the three numerical methods used to render 1080p images as observed from $(x, y, z) = (0, r_S, -20r_S)$. The angular resolutions of RK4 and Ruth's method are identical; the velocity Verlet method's angular resolution is slightly lower.

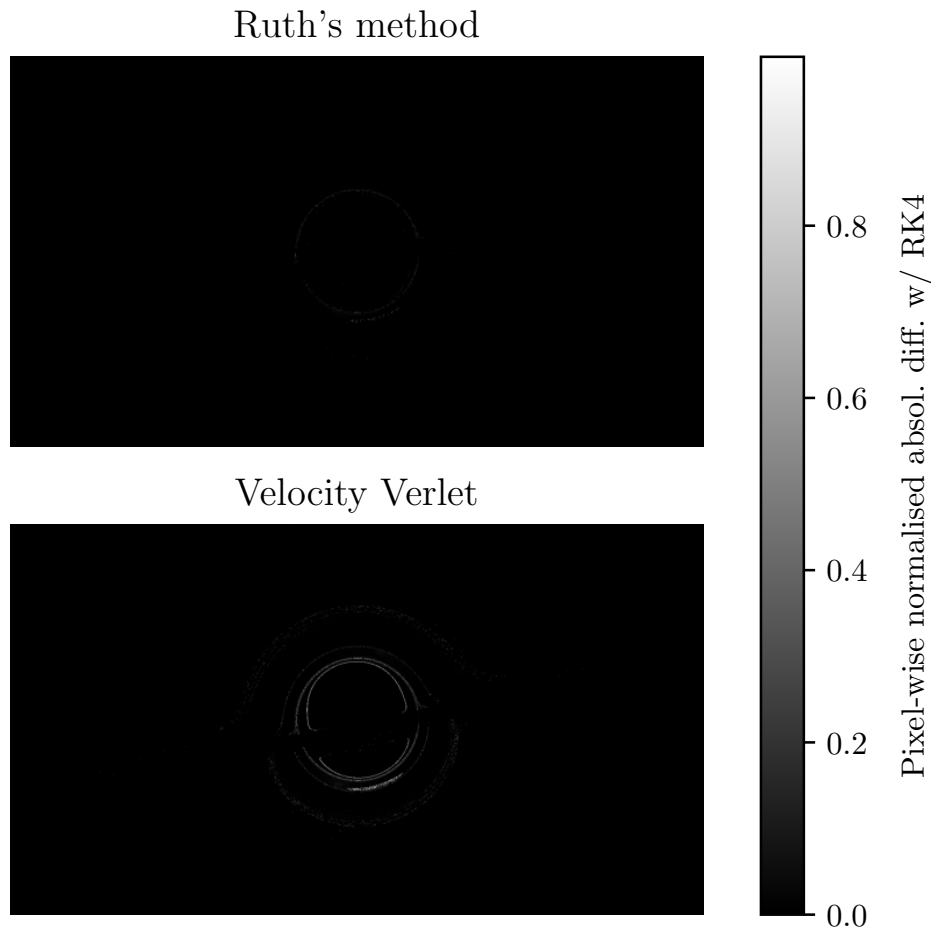


Figure 3.10: A normalised pixel-wise difference map for Ruth's method and the velocity Verlet method when compared with RK4. Any differences between RK4 and the velocity Verlet method are significantly more noticeable than those of Ruth's method.

3.9 were again used to render each image here. Owing to a greater number of pixels overall, each 4:3 ratio image took longer to render than a 16:9 image at the same approximate resolution. Independent of the aspect ratio, we notice that RK4 consistently requires more time to trace the same image when compared with Ruth's method and the velocity Verlet method, with RK4 taking over 30 minutes to trace the largest image, yet the velocity Verlet method needs only half that time to carry out the same task. Furthermore, the difference in the time required by these two methods appears to grow with a decrease in image resolution. Ruth's method also offers a modest reduction in the time required to trace the test image when compared with RK4. However, the fact that its image quality is virtually identical to that of the velocity Verlet method makes it somewhat a redundant choice for our task. Nonetheless, this still proves that symplectic methods prove to be a more optimal choice for nonlinear ray tracing, especially, as in our case, where the programme may be executed in parallel, since parallel implementations of symplectic methods should not need extensive modification of serial code.




Figure 3.11: A plot of the necessary trace time versus image height for images with 16:9 and 4:3 aspect ratios. Both Ruth's method and the velocity Verlet method require less time than RK4 to ray trace the same test image.

Although the use of the Python language for this ray tracer made its modification straightforward, additional speed increases could be achieved by using a just-in-time Python compiler such as Cython or Numba [7,47], or even rewriting the programme completely in a compiled language such as C, with multiprocessing aspects being handled by well-established parallel computing

APIs such as OpenMP or MPI [20, 21]. The ray tracer would presumably stand to benefit greatly from GPU acceleration as well. Since only a CPU was available for the numerical experiments presented here, performance benchmarking on a GPU-enabled system would provide an interesting comparison with our present results. Particularly in the context of image rendering, symplectic ray tracing methods could also be tested against fast marching and fast sweeping methods, mentioned briefly earlier in this chapter. Nevertheless, the existing Python ray tracer could still serve as a valuable educational resource for introducing students to topics including nonlinear ray tracing, numerical integration and high-performance computing.

3.4 Summary

In general, symplectic methods offer consistently accurate performance when ray tracing GRIN elements and often outperform popular nonsymplectic numerical techniques such as RK4. Although optics currently lags behind fields such as celestial mechanics and molecular dynamics in its adoption of symplectic methods, the velocity Verlet and Ruth's methods both appear to provide a useful template for the development of symplectic numerical methods specifically for ray tracing within spherical GRIN profiles. That said, the construction of a symplectic method whose performance is adequate across a range of spherical, axial and radial profiles will presumably pose a considerable challenge. Nevertheless, complex optical elements such as natural and artificial eyes, for which analytical ray traces may not exist could indeed benefit from the application of tailor-made symplectic methods. In addition to standard GRIN elements, freeform GRIN (F-GRIN) optics offer another degree of freedom in optical design, taking advantage of a lack of rotational invariance as well as a GRIN profile which may or may not be continuous [50]. A selection of bespoke symplectic numerical methods could accelerate the design and prototyping of useful F-GRIN lenses. The ability to accurately trace arbitrary F-GRIN elements would hopefully, in turn, lead to developments in fabrication and metrology, addressing in particular a current lack of useful testing methods for conventional GRIN and F-GRIN elements alike.

Moreover, each of the symplectic methods considered in this paper are explicit methods, meaning each subsequent step relies on position and optical momentum data obtained previously. Several implicit symplectic methods, often renowned for their increased numerical stability, could also be used. Unfortunately, implicit methods typically require a root-solving algorithm to calculate the relevant data during each iteration, substantially increasing the computational cost required in their implementation. However, certain index profiles exist for which implicit methods may be used without the need for a root solver, thus eliminating any extra computational expenses incurred. A selection of such index profiles are considered at length in the following chapter.

Chapter 4

An implicit method for separable GRIN profiles

4.1 Introduction

The vast majority of studies related to symplectic numerical methods focus disproportionately on explicit methods [33, 57, 62], which rely solely on data obtained during previous calculations in order to update a ray's trajectory. Implicit methods, which are known for their increased numerical stability [99], instead substitute data which is yet to be obtained into the relevant equations, requiring the desired quantities to be isolated in order to be calculated. In most instances, however, the relevant equations are nonlinear in the quantities of interest (e.g. the ray's position or direction cosine), meaning a root-finding algorithm such as Newton's method must also be employed during the iteration process, greatly increasing the computational cost associated with their use [37]. Thus, the use of implicit numerical methods is often limited to situations where the need for additional numerical stability justifies the increased memory overhead. This result is particularly relevant to symplectic methods, as symplectic Runge-Kutta methods are implicit by definition [79]. A proof of this result is available in Appendix A.3.

Nevertheless, it was previously shown that separable index profiles (i.e. profiles for which the terms dependent on each coordinate can be treated as separate functions) reduce the difficulty in solving the eikonal equation [12]. Subject to certain constraints, we find that separable GRIN profiles enable the use of implicit numerical methods without the need for a root-solver, thereby greatly reducing the computational cost involved. Here, we identify a general expression for the refractive index, allowing the inexpensive use of implicit symplectic methods while also performing numerical experiments on a selection of lenses whose index profiles are of the required form, allowing us to compare an implicit symplectic method with a selection of explicit methods,

both symplectic and nonsymplectic. First, however, we recall Eq. (2.45)

$$H = \frac{1}{2} [p_x^2 + p_y^2 + p_z^2 - n^2(x, y, z)], \quad (4.1)$$

where we notice that the form of the refractive index will greatly influence the level of difficulty in computing ray trajectories. Thus, in some instances, we may benefit from the increased numerical stability of an implicit symplectic method [37]. However, implicit methods often require a root-finding scheme to solve for each subsequent iterate, making their implementation computationally expensive. Yet, in some cases, we may solve for each iterate algebraically, taking advantage of the additional stability offered by implicit methods at no extra computational cost, provided the relevant GRIN profile can be written in the following form

$$n^2(x, y, z) = c_1 \left[\frac{c_2 + c_3(x + c_4)^{\alpha_1}}{c_5 + c_6(x + c_7)^{\alpha_2}} + \frac{c_8 + c_9(y + c_{10})^{\alpha_3}}{c_{11} + c_{12}(y + c_{13})^{\alpha_4}} + \frac{c_{14} + c_{15}(z + c_{16})^{\alpha_5}}{c_{17} + c_{18}(z + c_{19})^{\alpha_6}} + c_{20} \right], \quad (4.2)$$

where c_1, \dots, c_{20} are constants with $\alpha_1, \alpha_3, \alpha_5 \in \{0, 1, 2, 3, 4, 5\}$ and $\alpha_2, \alpha_4, \alpha_6 \in \{0, 1, 2\}$ subject to the constraints $\alpha_1 + \alpha_2 \leq 5$, $\alpha_3 + \alpha_4 \leq 5$ and $\alpha_5 + \alpha_6 \leq 5$. Thus, any iterative scheme constructed will contain at most a polynomial of degree 4 in x , y , or z , whose roots can be solved analytically. To the best of this author's knowledge, Eq. (4.2) has not been presented elsewhere in existing literature. We notice that the previous expression is indeed separable, since it is possible to express the square of the index as a sum of separate functions, each dependent solely on x , y or z , thus matching the definition for separable GRIN profiles provided in [12].

The implicit midpoint method, a second-order symplectic method is presented as an example of a method whose implementation need not be expensive. Its construction and application to a group of GRIN lenses whose profiles are of the form of Eq. (4.2) follows in the next section.

4.2 The implicit midpoint method

4.2.1 Construction by composition

In Chapter 2, we constructed the symplectic Euler method via the splitting technique [29, 37], where separate symplectic methods are derived from Hamilton's equations in order to solve for the position and optical momentum of a given ray. However, a symplectic numerical method may also be created via composition, where each finite step is divided into a number of sub-steps before alternating between the use of an implicit and explicit method to solve for the position and optical momentum at each subsequent sub-step. In order to derive the implicit midpoint

method, we first recall the forward Euler method

$$\begin{aligned} p_{i_{n+1}} &= p_{i_n} - \frac{\partial}{\partial q_i} [H(q_{i_n})] \Delta t, \\ q_{i_{n+1}} &= q_{i_n} + \frac{\partial}{\partial p_i} [H(p_{i_n})] \Delta t, \end{aligned} \quad (4.3)$$

By itself, the forward Euler method is a cheap yet inaccurate choice for GRIN ray tracing [56]. Its implicit counterpart, the backward Euler method, is usually more stable but not necessarily more accurate [38], given by

$$\begin{aligned} p_{i_{n+1}} &= p_{i_n} - \frac{\partial}{\partial q_i} [H(q_{i_{n+1}})] \Delta t, \\ q_{i_{n+1}} &= q_{i_n} + \frac{\partial}{\partial p_i} [H(p_{i_{n+1}})] \Delta t. \end{aligned} \quad (4.4)$$

A subtle difference between Eqs. (4.4) and (4.3) may be seen; index $n + 1$ now appears on both sides of the equality. In cases where the partial derivative terms are strongly nonlinear, fixed-point iteration schemes such as the Newton-Raphson method are required to solve for terms with index $n + 1$ at each step. Often, this greatly increases the computational cost of using implicit methods, meaning they are typically reserved for instances where their stability justifies any additional computational expense [99]. To construct a second-order method, we first divide Δt by two, with the first half-step being carried out by the backward Euler method and the forward Euler method performing the second half-step, as follows

$$\begin{aligned} p_{i_{n+1/2}} &= p_{i_n} - \frac{\partial}{\partial q_i} [H(q_{i_{n+1/2}})] \frac{\Delta t}{2}, \\ q_{i_{n+1/2}} &= q_{i_n} + \frac{\partial}{\partial p_i} [H(p_{i_{n+1/2}})] \frac{\Delta t}{2}, \end{aligned} \quad (4.5)$$

$$\begin{aligned} p_{i_{n+1}} &= p_{i_{n+1/2}} - \frac{\partial}{\partial q_i} [H(q_{i_{n+1/2}})] \frac{\Delta t}{2}, \\ q_{i_{n+1}} &= q_{i_{n+1/2}} + \frac{\partial}{\partial p_i} [H(p_{i_{n+1/2}})] \frac{\Delta t}{2}, \end{aligned} \quad (4.6)$$

then, substituting Eqs. (4.5) into Eqs. (4.6),

$$\begin{aligned} p_{i_{n+1}} &= p_{i_n} - \frac{\partial}{\partial q_i} [H(q_{i_{n+1/2}})] \frac{\Delta t}{2} - \frac{\partial}{\partial q_i} [H(q_{i_{n+1/2}})] \frac{\Delta t}{2} \\ &= p_{i_n} - \frac{\partial}{\partial q_i} [H(q_{i_{n+1/2}})] \Delta t, \\ q_{i_{n+1}} &= q_{i_n} + \frac{\partial}{\partial p_i} [H(p_{i_{n+1/2}})] \frac{\Delta t}{2} + \frac{\partial}{\partial p_i} [H(p_{i_{n+1/2}})] \frac{\Delta t}{2} \\ &= q_{i_n} + \frac{\partial}{\partial p_i} [H(p_{i_{n+1/2}})] \Delta t. \end{aligned} \quad (4.7)$$

Rearranging Eqs. (4.5) and (4.6) in terms of $q_{i_{n+1/2}}$ and $p_{i_{n+1/2}}$, we find

$$p_{i_{n+1/2}} = p_{i_n} - \frac{\partial}{\partial q_i} \left[H(q_{i_{n+1/2}}) \right] \frac{\Delta t}{2}, \quad (4.8)$$

$$= p_{i_{n+1}} + \frac{\partial}{\partial q_i} \left[H(q_{i_{n+1/2}}) \right] \frac{\Delta t}{2}, \quad (4.9)$$

$$q_{i_{n+1/2}} = q_{i_n} + \frac{\partial}{\partial p_i} \left[H(p_{i_{n+1/2}}) \right] \frac{\Delta t}{2}, \quad (4.10)$$

$$= q_{i_{n+1}} - \frac{\partial}{\partial p_i} \left[H(p_{i_{n+1/2}}) \right] \frac{\Delta t}{2}, \quad (4.11)$$

Since $p_{i_{n+1/2}}$ and $q_{i_{n+1/2}}$ are obtained either by a positive half-step from p_{i_n} and q_{i_n} or a negative half-step from $p_{i_{n+1}}$ and $q_{i_{n+1}}$, the implicit midpoint method is symmetric, making it suitable for modelling the physics of GRIN optics by ensuring Helmholtz reciprocity (i.e. each ray trace is reversible). Finally, adding Eq. (4.10) to Eq. (4.11) and Eq. (4.8) to Eq. (4.9), we obtain

$$\begin{aligned} 2p_{i_{n+1/2}} &= p_{i_{n+1}} + \frac{\partial}{\partial q_i} \left[H(q_{i_{n+1/2}}) \right] \frac{\Delta t}{2} + p_{i_n} - \frac{\partial}{\partial q_i} \left[H(q_{i_{n+1/2}}) \right] \frac{\Delta t}{2}, \\ 2q_{i_{n+1/2}} &= q_{i_{n+1}} - \frac{\partial}{\partial p_i} \left[H(p_{i_{n+1/2}}) \right] \frac{\Delta t}{2} + q_{i_n} + \frac{\partial}{\partial p_i} \left[H(p_{i_{n+1/2}}) \right] \frac{\Delta t}{2}, \end{aligned}$$

and hence, we observe

$$\begin{aligned} p_{i_{n+1/2}} &= \frac{p_{i_{n+1}} + p_{i_n}}{2}, \\ q_{i_{n+1/2}} &= \frac{q_{i_{n+1}} + q_{i_n}}{2}, \end{aligned}$$

which gives us the implicit midpoint method

$$\begin{aligned} p_{i_{n+1}} &= p_{i_n} - \frac{\partial}{\partial q_i} \left[H \left(\frac{q_{i_{n+1}} + q_{i_n}}{2} \right) \right] \Delta t, \\ q_{i_{n+1}} &= q_{i_n} + \frac{\partial}{\partial p_i} \left[H \left(\frac{p_{i_{n+1}} + p_{i_n}}{2} \right) \right] \Delta t. \end{aligned} \quad (4.12)$$

Proving the symplecticity of the implicit midpoint method is demonstrated in Appendix A.2. Composition is thus a highly versatile method for constructing symplectic numerical methods of arbitrarily high order. For example, the implicit midpoint method could itself be composed to create a fourth-order method, which could, in turn, be composed to create an eighth-order method and so on. However, the construction of such high-order methods quickly becomes unwieldy without the assistance of a computer algebra system, making these methods prohibitively expensive for lenses whose indices are not of the form in Eq. (4.2). We may now examine the iteration schemes derived by applying the implicit method to a group of GRIN lenses with refractive indices of the necessary form, allowing it to be used without the need for a root solver.




Figure 4.1: Exact traces available for the lenses given in Table 4.1.

Lens name	Index type	Index profile	General solution?
Lüneburg [52]	Spherical	$n(y, z) = \sqrt{2 - \frac{y^2 + z^2}{R_0^2}}$	Yes
Ilinsky [43]	Spherical	$n(y, z) = \sqrt{\frac{y^2 + (z-f)^2}{(f-r_1)^2}}$	No
Tarkhanov [91]	Axial	$n(z) = \frac{n_0}{\sqrt{1-2n_0 \frac{z}{f}}}$	No

Table 4.1: A selection of lenses whose GRIN profiles of the form identified in Eq. (4.2). R_0 represents the radius of the Lüneburg lens, while r_1 represents the radius of the front surface for the Ilinsky lens. Finally, n_0 is the refractive index value on the optical axis for the Tarkhanov lens while f represents the relevant focal distance for each lens.

4.2.2 Deriving the necessary iteration scheme for each lens

We now apply the implicit midpoint method to a selection of lenses whose index profiles may be written in the requisite form of Eq. (4.2) and derive the necessary iterative schemes in each case. Table 4.1 provides additional information about the index profiles of the lenses examined, as well as the existence of solutions for general ray traces. Here, we notice the ability to use the implicit midpoint method without a root-solver is not strictly limited to axial or spherical index profiles, with the Lüneburg [52] and Ilinsky [43] lenses providing examples of spherical profiles, while the Tarkahnov [91] lens demonstrates an axial instance. The existence of exact traces for axially-collimated rays in the Ilinsky and Tarkhanov lenses is well documented [35,43], as depicted in Figure 4.1. However, closed-form solutions for general rays deviating from these special cases usually exist only within the paraxial region in the vicinity of the optical axis [31]. Thus, the use of numerical methods is necessary in all other instances.

Beginning with the Lüneburg lens, we first derive its optical Hamiltonian by substituting its index profile from Table 4.1 into Eq. (4.1).

$$H_{\text{Lüneburg}} = \frac{1}{2} \left(p_y^2 + p_z^2 + \frac{y^2 + z^2}{R_0} \right) - 1, \quad (4.13)$$

where R_0 is the lens' radius with $0 \leq y^2 + z^2 \leq R_0^2$. Partial differentiation of Eq. (4.13) yields

$$\frac{\partial H}{\partial p_i} = p_i, \quad \frac{\partial H}{\partial q_i} = \frac{q_i}{R_0}. \quad (4.14)$$

Substituting Eqs. (4.14) into Eqs. (4.12), we obtain

$$\begin{aligned} p_{i_{n+1}} &= p_{i_n} - \frac{1}{R_0} \left[q_{i_n} + (p_{i_{n+1}} + p_{i_n}) \frac{\Delta t}{2} + q_{i_n} \right] \frac{\Delta t}{2}, \\ q_{i_{n+1}} &= q_{i_n} + \left[p_{i_n} - \frac{1}{R_0} (q_{i_{n+1}} + q_{i_n}) \frac{\Delta t}{2} + p_{i_n} \right] \frac{\Delta t}{2}, \end{aligned}$$

then, expanding brackets and grouping similar terms

$$\begin{aligned} p_{i_{n+1}} \left[1 + \frac{(\Delta t)^2}{4R_0} \right] &= p_{i_n} \left[1 - \frac{(\Delta t)^2}{4R_0} \right] - \frac{q_{i_n}}{R_0} \Delta t, \\ q_{i_{n+1}} \left[1 + \frac{(\Delta t)^2}{4R_0} \right] &= q_{i_n} \left[1 - \frac{(\Delta t)^2}{4R_0} \right] + p_{i_n} \Delta t. \end{aligned} \quad (4.15)$$

Finally, dividing both sides by $1 + [(\Delta t)^2/4R_0]$

$$\begin{aligned} p_{i_{n+1}} &= p_{i_n} \left[\frac{4R_0 - (\Delta t)^2}{4R_0 + (\Delta t)^2} \right] - q_{i_n} \left[\frac{4\Delta t}{4R_0 + (\Delta t)^2} \right], \\ q_{i_{n+1}} &= q_{i_n} \left[\frac{4R_0 - (\Delta t)^2}{4R_0 + (\Delta t)^2} \right] + p_{i_n} \left[\frac{4R_0\Delta t}{4R_0 + (\Delta t)^2} \right]. \end{aligned} \quad (4.16)$$

Hence, we observe it is possible to isolate $p_{i_{n+1}}$ and $q_{i_{n+1}}$, thereby deriving recurrence relations in terms of p_{i_n} and q_{i_n} for the Lüneburg lens. Repeating this analysis for the Ilinsky lens, we find its optical Hamiltonian may be written as

$$H_{\text{Ilinsky}} = \frac{1}{2} \left[p_y^2 + p_z^2 - \frac{y^2 + (z - f)^2}{(f - r_1)^2} \right], \quad (4.17)$$

with $y^2 + z^2 \leq r_1^2$, where r_1 is the radius of the lens' front surface and f is the focal distance. However, we notice that for the Ilinsky lens, Hamilton's equations with respect to y and z are not identical for each variable, as we saw with the Lüneburg lens. The iteration schemes for y and z must now be considered separately, where the relevant expressions are:

$$\begin{aligned} p_{y_{n+1}} &= p_{y_n} + \frac{1}{(f - r_1)^2} \left[\frac{y_n + y_{n+1}}{2} \right] \Delta t, \\ p_{z_{n+1}} &= p_{z_n} + \frac{1}{(f - r_1)^2} \left[\frac{z_n + z_{n+1}}{2} - f \right] \Delta t, \end{aligned} \quad (4.18)$$

$$\begin{aligned} y_{n+1} &= y_n + \left(\frac{p_{y_{n+1}} + p_{y_n}}{2} \right) \Delta t, \\ z_{n+1} &= z_n + \left(\frac{p_{z_{n+1}} + p_{z_n}}{2} \right) \Delta t. \end{aligned} \quad (4.19)$$

Then, we substitute Eqs. (4.18) into Eqs. (4.19) and vice versa, expanding the resulting expres-

sions and grouping similar terms, as we did before

$$\begin{aligned}
p_{y_{n+1}} \left[1 - \left(\frac{\Delta t}{2(f-r_1)} \right)^2 \right] &= p_{y_n} \left[1 + \left(\frac{\Delta t}{2(f-r_1)} \right)^2 \right] + \frac{y_n}{(f-r_1)^2} \Delta t, \\
p_{z_{n+1}} \left[1 - \left(\frac{\Delta t}{2(f-r_1)} \right)^2 \right] &= p_{z_n} \left[1 + \left(\frac{\Delta t}{2(f-r_1)} \right)^2 \right] + \frac{(z_n-f)}{(f-r_1)^2} \Delta t, \\
y_{n+1} \left[1 - \left(\frac{\Delta t}{2(f-r_1)} \right)^2 \right] &= y_n \left[1 + \left(\frac{\Delta t}{2(f-r_1)} \right)^2 \right] + p_{y_n} \Delta t, \\
z_{n+1} \left[1 - \left(\frac{\Delta t}{2(f-r_1)} \right)^2 \right] &= z_n \left[1 + \left(\frac{\Delta t}{2(f-r_1)} \right)^2 \right] + p_{z_n} \Delta t - \left(\frac{f}{2} \right) \left[\frac{\Delta t}{(f-r_1)} \right]^2.
\end{aligned}$$

Finally, dividing across by $1 - \{\Delta t/[2(f-r_1)]\}^2$

$$\begin{aligned}
p_{y_{n+1}} &= p_{y_n} \left[\frac{4(f-r_1)^2 + (\Delta t)^2}{4(f-r_1)^2 - (\Delta t)^2} \right] + y_n \left[\frac{4\Delta t}{4(f-r_1)^2 - (\Delta t)^2} \right], \\
p_{z_{n+1}} &= p_{z_n} \left[\frac{4(f-r_1)^2 + (\Delta t)^2}{4(f-r_1)^2 - (\Delta t)^2} \right] + (z_n-f) \left[\frac{4\Delta t}{4(f-r_1)^2 - (\Delta t)^2} \right], \\
y_{n+1} &= y_n \left[\frac{4(f-r_1)^2 + (\Delta t)^2}{4(f-r_1)^2 - (\Delta t)^2} \right] + p_{y_n} \left[\frac{4\Delta t(f-r_1)^2}{4(f-r_1)^2 - (\Delta t)^2} \right], \\
z_{n+1} &= z_n \left[\frac{4(f-r_1)^2 + (\Delta t)^2}{4(f-r_1)^2 - (\Delta t)^2} \right] + [p_{z_n} - 2f(\Delta t)^2] \left[\frac{4\Delta t(f-r_1)^2}{4(f-r_1)^2 - (\Delta t)^2} \right],
\end{aligned} \tag{4.20}$$

which gives the necessary iteration schemes for the Ilinsky lens. Moving onto the Tarkhanov lens, since its refractive index profile is axial (i.e. it is z -dependent only), we immediately find $p_{y_n} = p_{y_0} \forall n$ and so $y_{n+1} = y_n + p_{y_0} \Delta t$. This is, in effect, the standard Euler method, which is symplectic since $p_{y_n} = p_{y_0}$. The optical Hamiltonian for the Tarkhanov lens is then given by

$$H_{\text{Tarkhanov}} = \frac{1}{2} \left[p_y^2 + p_z^2 - \frac{n_0^2}{1 - 2n_0 \frac{z}{f}} \right], \tag{4.21}$$

where n_0 is the refractive index at the apex of the curved back surface of the lens. Differentiation with respect to z and substitution into Eq. (4.12) provides us iterative schemes for z and p_z

$$p_{z_{n+1}} = p_{z_n} + \left\{ \frac{fn_0^3}{[f - n_0(z_n + z_{n+1})]^2} \right\} \Delta t, \tag{4.22}$$

$$z_{n+1} = z_n + \left(\frac{p_{z_{n+1}} + p_{z_n}}{2} \right) \Delta t. \tag{4.23}$$

Now, substituting Eq. (4.23) into Eq. (4.22) and rearranging

$$(p_{z_{n+1}} - p_{z_n}) \left\{ f - n_0 \left[2z_n + (p_{z_{n+1}} + p_{z_n}) \frac{\Delta t}{2} \right] \right\}^2 - f n_0^3 \Delta t = 0, \quad (4.24)$$

expanding the previous equation with the assistance of a symbolic algebra system [59] then gives a cubic expression $ap_{z_{n+1}}^3 + bp_{z_{n+1}}^2 + cp_{z_{n+1}} + d$, where

$$\begin{aligned} a &= \left(\frac{n_0 \Delta t}{2} \right)^2, \\ b &= ap_{z_n} - n_0 (f - 2n_0 z_n) \Delta t, \\ c &= -ap_{z_n}^2 - 4z_n n_0 (f - n_0 z_n) + f^2, \\ d &= -ap_{z_n}^3 - (b - ap_{z_n})p_{z_n}^2 - (c + ap_{z_n}^2)p_{z_n} - fn_0^3 \Delta t. \end{aligned}$$

Checking the discriminant of this cubic equation indicates that it has three real roots. However, two of these roots give solutions consistent with axially-collimated rays (i.e. the situation presented in Figure 4.1), which is obviously incorrect in this instance. Thus, the desired root is

$$p_{z_{n+1}} = -\frac{1}{3a} \left(b + C + \frac{\Delta_0}{C} \right), \quad (4.25)$$

where

$$\begin{aligned} C &= \sqrt[3]{\frac{\Delta_1 + \sqrt{\Delta_1^2 - 4\Delta_0^3}}{2}}, \\ \Delta_0 &= b^2 - 3ac, \\ \Delta_1 &= 2b^3 - 9abc + 27a^2d. \end{aligned}$$

However, caution is advised, as $\Delta_1^2 - 4\Delta_0^3$ can become negative due to inaccuracies in the floating point arithmetic and so, additional care is need to prevent rounding errors.

4.3 Numerical ray tracing

The accuracy of the implicit midpoint method is now tested against a number of symplectic and nonsymplectic numerical methods when ray tracing the lenses present given in Table 4.1. Since it exists, we make use of the analytical solution available for the Lüneburg lens [52]. However, to the best of our knowledge, general analytical solutions for the Ilinsky and Tarkhanov lenses have not yet been derived; the nonsymplectic RK4 method [89] is used as an alternative in these instances. Based on their impressive performance in spherical lenses, the velocity Verlet method and Ruth's method are also compared with the implicit midpoint method to further assess its suitability for optical problems. The geometry of the Lüneburg lens is identical to

that of the lenses considered in Chapter 3 (i.e. $R_0 = 1.0$ mm), while the Ilinsky lens is of radius $r_1 = 1.875$ mm with an aperture diameter of 2.0 mm. The Tarkhanov lens is of radius $r_2 = -2.0$ mm (its radius is negative, following the Cartesian sign convention) where, like the Ilinsky lens, its aperture diameter is also 2.0 mm. For the Lüneburg lens, the object was placed an infinite distance to the left of the lens while for both the Ilinsky and Tarkhanov lenses the object distance d_o was chosen to be 100.0 mm to the left of each lens.

Figures 4.3 and 4.5 show numerical ray traces within the Lüneburg, Ilinsky and Tarkhanov lenses respectively. Overall, the trace for the Lüneburg lens, is much the same as that in the previous chapter, with all the methods being effectively indistinguishable when $\Delta t = \lambda$ before becoming more easily identifiable as Δt increases. For the largest step size examined, we notice that the implicit midpoint appears to be slightly more accurate than Ruth’s method, which further emphasises the results of our previous numerical experiments, with a second-order method once again outperforming a fourth-order one. For the Ilinsky lens, however, the RK4 method (which is used in the absence of an analytical trace) deviates somewhat from the trajectories returned by the velocity Verlet, implicit midpoint and Ruth’s methods for all step sizes examined, where each of the symplectic methods are virtually identical for all Δt values. However, in the absence of an analytical solution, attempting to ascertain which of the methods is the most accurate will require further work. A similar situation is present in the Tarkhanov lens, though in this instance Ruth’s method, RK4 and the velocity Verlet method display significant overlap at all step sizes, with the implicit midpoint method exhibiting a noticeably different trajectory, which may suggest that the implicit midpoint method is less effective for axial rather than spherical index profiles. Another possible explanation for this discrepancy could be the significant amount of floating point arithmetic required by the implicit midpoint for the Tarkhanov lens, resulting in rounding errors which may adversely affect its accuracy. Yet again, however, without an analytical ray trace, we are somewhat unsure as to which method is best for GRIN ray tracing.

To further assess the suitability of the implicit midpoint method, we may examine the amount of coma introduced within the Ilinsky and Tarkhanov lenses for each step size when the beam is rotated about the optical axis. Coma is an optical aberration which occurs when a lens fails to focus an off-axis ray bundle to a single point. As a result, the object will appear to possess a tail similar to a comet, and so, the name “coma” was chosen to describe this effect. Figure 4.2 depicts a constant-index meniscus lens exhibiting coma for a group of parallel, off-axis rays and also the wavefront corresponding to the comatic image. Wavefronts are discussed in further detail in Chapter 5. The Lüneburg lens, being aplanatic (i.e. free of spherical aberration and coma), is unfortunately unsuitable for this kind of test. Nevertheless, Figure 4.6 shows the amount of coma recorded for each method when the incoming beam is moved one, five and ten degrees off-axis. For the Ilinsky lens, RK4 exhibits significantly less coma for all rotation angles, with each of the symplectic methods giving roughly the same coma value, particularly when

$\Delta t < 50\lambda$. For the Tarkahnov lens, the implicit midpoint method exhibits substantially less coma than all other methods, though the difference is significantly less than that of the Ilinsky lens. Noticeably, the amount of coma recorded with RK4, the velocity Verlet method and Ruth's method is almost identical in each case and seems to decrease slightly as Δt tends to 100λ . By contrast, the implicit midpoint deviates from this trend entirely, exhibiting much less coma than the other three methods considered. It remains mostly flat when the beam is rotated 5 or 10 degrees off-axis before increasing slightly as Δt approaches 100λ . For a single degree off-axis, the amount of coma appears to be minimised when $\Delta t = 25\lambda$ before increasing once again and following the same general trend. To carry our investigation further, these results could be compared with an analytical calculation of how much coma may be expected for both the Ilinsky and Tarkhanov lenses, providing another means of validating the numerical methods presented.




Figure 4.2: Left: A constant-index meniscus lens with $n = 1.40$ exhibiting coma, where the screen is placed at the intersection point of the marginal rays. Right: The wavefront profile associated with coma. Wavefront profiles are explained in further detail in Chapter 5.

Another method by which we may compare the implicit midpoint with the other numerical methods is by considering the deviation in the optical path calculated for forward and reverse ray traces within each lens. Unlike the previous experiment, this approach remains valid for the Lüneburg lens also. Figure 4.7 depicts the difference between the optical paths computed by each method within each lens for three separate initial ray heights, with the diffraction limit being represented by the dashed horizontal line, as before. For the Lüneburg lens, the difference in optical path follows the same general trend for all ray heights, with the implicit midpoint and Ruth's method both offering a diffraction-limited trace for all ray heights provided that $\Delta t \leq 25\lambda$. RK4 and the velocity Verlet instead perform much less satisfactorily, not offering diffraction limited traces for any step size. Within the Ilinsky lens, however, much more variation is observed. This time, the implicit midpoint method appears to be the least accurate of all the methods, though it still provides a diffraction-limited trace for $\Delta t \leq 10\lambda$. RK4 and the velocity Verlet method follow a similar downward trend, being diffraction-limited for all step sizes. While it is also diffraction limited for all Δt values, Ruth's method is much more erratic, showing a particularly drastic decrease of more than three orders of magnitude when $\Delta t = \lambda$ for




Figure 4.3: Implicit midpoint tracing a Lüneburg lens of radius R_0 . Like in Chapter 3, the region shown in each right-hand column is marked on each trace in the left-hand column by a black rectangle and depicts the focal point in greater detail.




Figure 4.4: Implicit midpoint tracing an Ilinsky lens whose front surface is of radius r_1 . The left hand column represents a detailed view of the focus, as before. The detector, represented by the dashed vertical line, is placed at the ideal focus.




Figure 4.5: Implicit midpoint tracing a Tarkhanov lens whose back surface is of radius r_2 . As before, the detector is represented by a dashed vertical line located at the ideal focus.

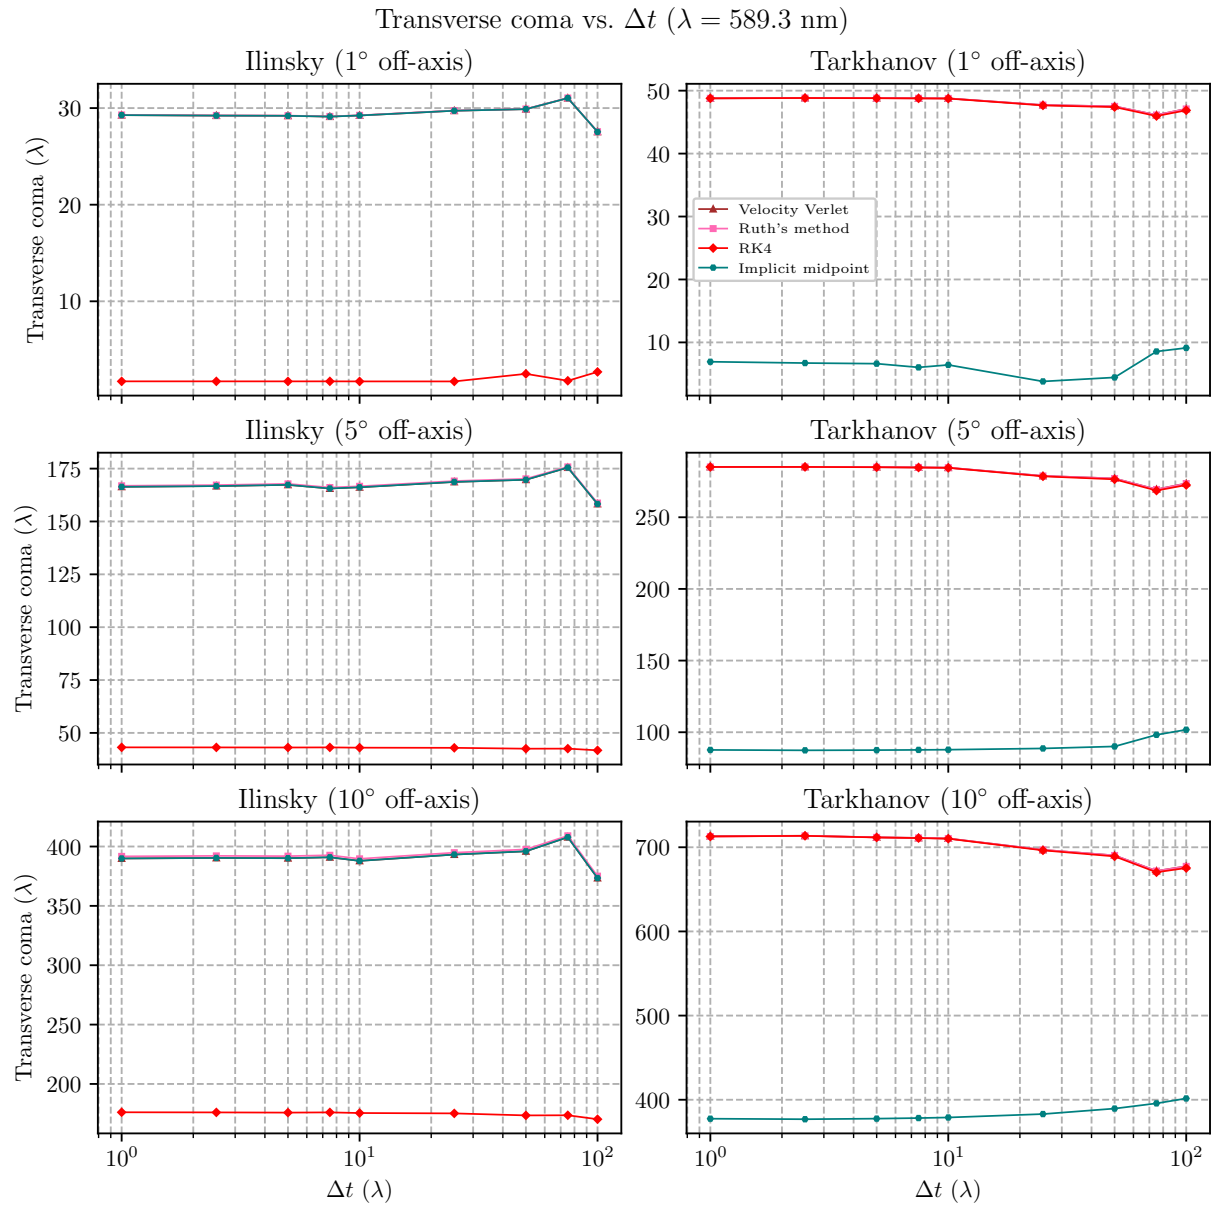


Figure 4.6: The maximum transverse coma value calculated by each method versus Δt for a bundle of rays placed one, five and ten degrees off-axis.

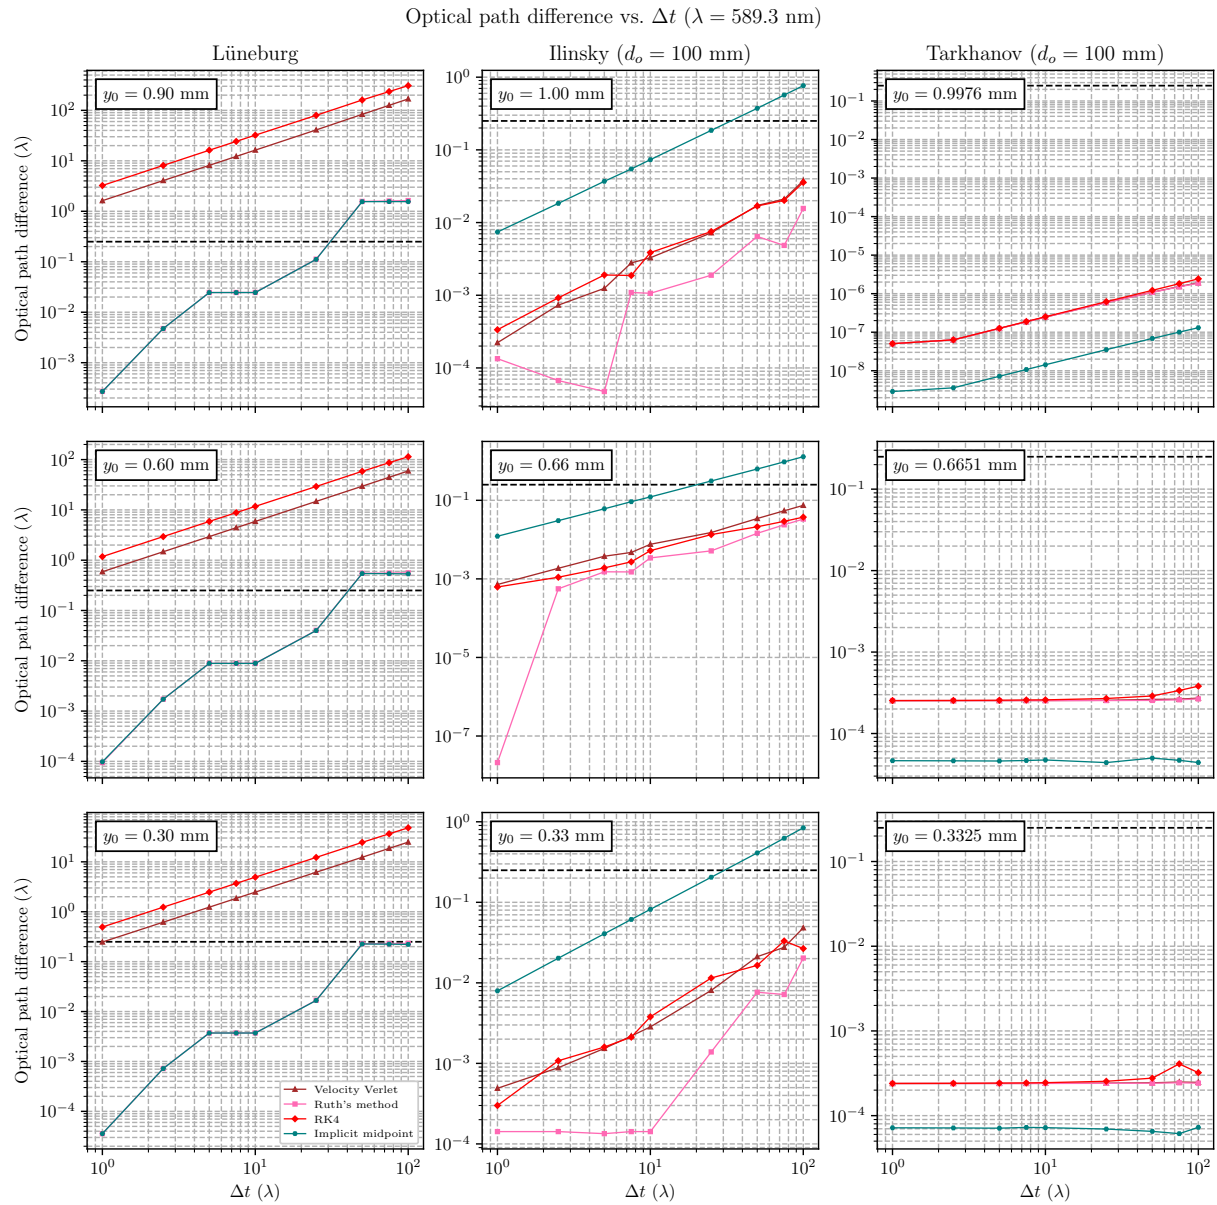


Figure 4.7: Differences in the optical path travelled by forward and reverse traces. The dashed horizontal line represented the diffraction limit, as before, with d_o and y_0 representing the object distance to the left of each lens and the initial ray height, respectively.

a ray height of $y_0 = 0.66$ mm. Finally, the Tarkhanov lens is more similar to the Lüneburg lens, where there are any differences in the difference between the forward and reverse traces appear to be much less significant with all methods well below the diffraction limit. Once again, the implicit midpoint method outperforms the other three methods, while the velocity Verlet method and Ruth's method provide virtually indistinguishable results. RK4 follows the same trend for the most part, particularly for $\Delta t \leq 25\lambda$. It does, however, err more noticeably if the step size is further increased, exhibiting behaviour that we might expect when comparing symplectic and nonsymplectic methods, where symplectic methods are generally more accurate for the same given step size. That said, any discrepancies between RK4 and the three symplectic methods are seen to diminish with an increase in initial ray height.

In summary, we have attempted to numerically trace rays using the implicit midpoint method within the Lüneburg, Ilinsky and Tarkhanov lenses. A bundle of rays emanating from a point source placed on the optical axis located a distance of 100 mm to the left of the Ilinsky and Tarkhanov lenses was chosen to examine their behaviour, while the ray bundles incident on the Lüneburg lens were instead axially collimated. The implicit midpoint method appears to perform reasonably well when compared with a selection of explicit numerical methods, though it seems to favour spherical rather than axial index profiles. However, the use of the implicit midpoint without the need for Newton's method or a similar root-finding scheme is limited to lenses with a separable index profile subject to specific constraints; A number of popular lenses fall outside this category, such as the Maxwell fish-eye [55], Eaton [27] and optical black hole [69] meaning the use of a root-finder cannot be avoided for these lenses. Moreover, root-finding methods would presumably struggle to converge in the vicinity of the singularities present in the Eaton lens and optical black hole. Thus, the overall usefulness of the implicit midpoint method for numerical ray tracing appears to be somewhat limited.

Chapter 5

GRIN aberrometry with a pyramid wavefront sensor

5.1 A lack of testing techniques

While ray tracing simulations have their uses in optical design, a lack of methods (destructive or otherwise) for testing lens prototypes leaves us with no experimental comparison for the ray traces carried out in the previous chapter. Indeed, the means required for the mass production of Ilinsky- and Tarkhanov-type lenses remain to be created. However, another barrier to their implementation lies in the difficulty of performing optical testing on GRIN lenses.

Most techniques available are not suitable for GRIN optics, or do not provide all of the necessary information. For instance, using a Mach-Zehnder interferometer offers only a relative measurement of the refractive index [102]. Another possible metrology technique requires the destruction of the lens in order to perform individual measurements in regions whose refractive index is approximately constant [50]. Furthermore, inferring a material's refractive index by geometric means (such as its position on a detector, say), can easily be solved for media with constant indices, though the same technique is far less useful for GRIN elements; one cannot be sure at which point in the lens an aberration may be introduced. Often, a tomographic method is required, though the lack of scattering in properly fabricated GRIN lenses makes optical coherence tomography unavailable for this purpose.

Nonetheless, we may still gain some information about the quality of each lens simply by ray tracing. A spot diagram, for example, depicts the position of each ray in the detector plane. By comparing the spot size with that of the Airy disk, whose radius is given by $r_{\text{Airy}} = 1.22\lambda\frac{f}{D}$, where D is the lens' aperture diameter, we may investigate whether a lens is capable of producing

a diffraction-limited image of a point source. Figure 5.1 depicts spot diagrams for the Ilinsky and Tarkhanov lenses, where a point source was placed at a distance $d_o = 100$ mm to the left of each lens, as in the previous chapter. The Airy disk is represented by the black circle at the centre of each plot, indicating that for the object distance chosen, both the Ilinsky and Tarkhanov lenses produce defocused images significantly larger than the diffraction limit. However, we also notice that the spot size of 0.0899 mm for the Ilinsky lens is slightly smaller than that of 0.10747 mm in the case of the Tarkhanov lens, which seems to suggest that the Ilinsky lens is slightly more tolerant to misconjugation of the object.

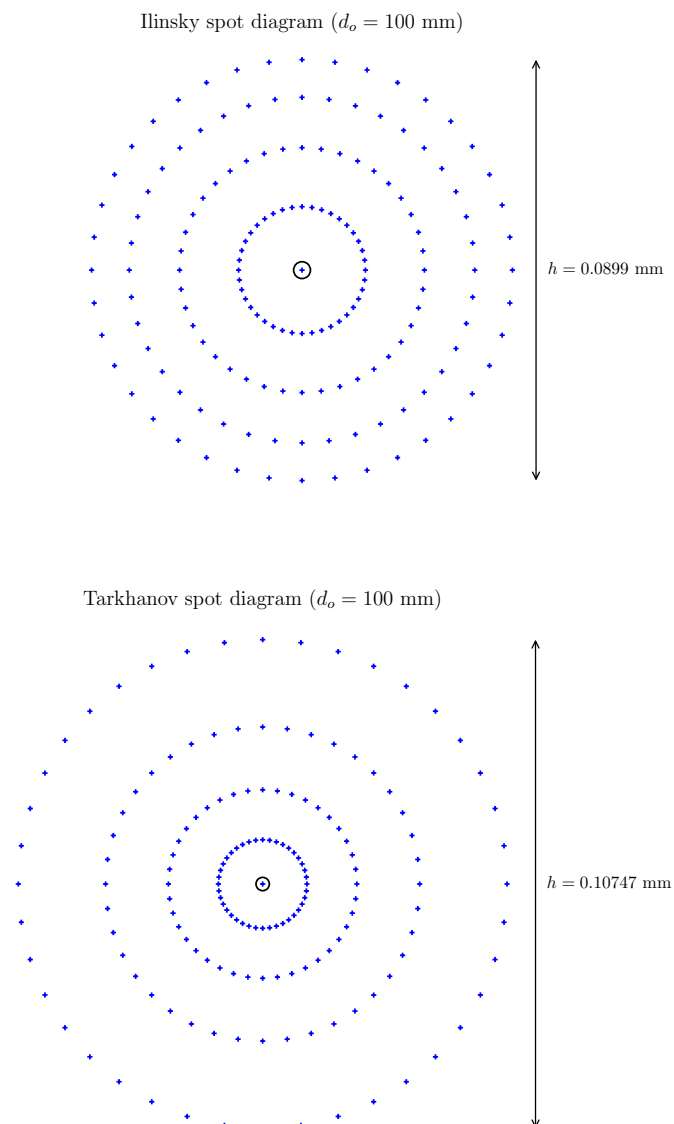


Figure 5.1: Spot diagrams for the Ilinsky (top) and Tarkhanov (bottom) lenses. The Airy disk is represented by a black circle in each case.

Even so, a simple ray trace is incapable of providing us with information about the wavefront error associated with a particular GRIN lens, meaning a device is required to measure the resulting wavefront. While we may not be able to identify which region within a lens is responsible for a particular aberration, we could, in the case of the Illinsky and Tarkhanov lenses, obtain an estimate for the wavefront aberrations introduced when the incident rays are no longer axially collimated. The pyramid wavefront sensor, a component typically employed in adaptive optics systems and well-known for its high sensitivity presents itself as a useful tool for optical testing in such a scenario. The following section provides further information on the pyramid wavefront sensor itself before detailing the mathematical methods employed in wavefront modelling and giving example wavefront profiles for the Illinsky and Tarkhanov lenses. Finally, a polynomial decomposition is performed on the resulting wavefront in each case.

5.2 The pyramid wavefront sensor

The pyramid wavefront sensor was first developed by Ragazzoni in 1996 [72], where an oscillating prism was used to obtain four images of the pupil on a single detector. The operating principle of the pyramid wavefront sensor is based on the Foucault knife-edge test in two orthogonal directions [76]. While the pyramid wavefront sensor is considerably more sensitive than the more mature Shack-Hartmann sensor, it has a much smaller dynamic range. Nevertheless, the dynamic range of the pyramid wavefront sensor may be increased by modulating the sensor's glass pyramid at the expense of a reduction in sensitivity [36].

When simulating a pyramid wavefront sensor, two different approaches may be used. The first treats the glass pyramid as a transmission mask object, where each image of the telescope pupil is created by applying a different Heaviside step filter to the point spread function (PSF) of the telescope entrance pupil [83]. The four pupil images must then be stitched together, creating a single image similar to the detector plane seen in Figure 5.2. The second method of simulating a pyramid wavefront sensor instead treats the pyramid as a phase mask object, adding a phase shift to the wavefront [83]. This method more accurately accounts for the geometry of the glass pyramid, generating a single image of the detector plane without the need to stitch together four individual pupil images [16]. Irrespective of which method may be used to simulate the sensor, the data of interest are the normalised wavefront gradients, obtained from the detector plane intensities via the expressions, first presented in [72]:

$$S_x = \frac{(I_1 + I_4) - (I_2 + I_3)}{I_1 + I_2 + I_3 + I_4}, \quad S_y = \frac{(I_1 + I_2) - (I_3 + I_4)}{I_1 + I_2 + I_3 + I_4}, \quad (5.1)$$

The calculation of these gradients is significantly easier using the transmission mask approach as arithmetic operations are more easily carried out on four individual pupil images. Isolating the four pupils from a single phase mask image is a more tedious task. In fact, for simulation packages




Figure 5.2: Schematic of a pyramid wavefront sensor.

which use the phase mask approach, calculating wavefront gradients is a significant challenge. For example, the open-source library High-Contrast Imaging for Python (HCIPy) employs a phase mask pyramid wavefront sensor, though the function used to estimate the wavefront gradients does not operate as intended and requires further work [71]. Another Python library, AOtools, instead uses a transmission mask pyramid capable of sensing the required wavefront gradients. At the time of writing, AOtools' pyramid does not yet oscillate, placing significant limits on its performance [95].

However, another algorithm is required in order to reconstruct the wavefront from the measured data. To this end, polynomial fitting is a useful tool as it allows any wavefront to be decomposed into a linear combination of mutually orthogonal polynomials, typically obtained via a singular value decomposition [22, 83, 101]. Mathematically, a given wavefront ϕ is then expressed as:

$$\phi = \sum_{n=1}^{\infty} a_n P_n, \quad (5.2)$$

where P_n is the n^{th} term in a series of polynomials chosen to model the wavefront and a_n is its associated coefficient. Obviously, it is impossible to perfectly model any wavefront, since the series in equation (5.2) contains an infinite number of terms. Therefore, any practical attempt

to profile a wavefront will require this series to be truncated eventually. Hence, the competing requirements of accuracy and computation time must be balanced.




Figure 5.3: Detector plane images of transmission mask (left) and phase mask (right).

While any set of orthogonal polynomials is suitable in theory, Zernike polynomials are perhaps the most popular set of polynomials for optical problems, as they represent optimally balanced classical aberrations (i.e. astigmatism, coma, spherical aberration, etc.). Additionally, recursive schemes are available, allowing a large number of Zernike polynomials to be generated quickly [1]. Numerically, Zernike polynomials are given by

$$Z_n^m(r, \theta) = N_n^m R_n^m(r) \Theta(\theta), \quad (5.3)$$

where n is the radial degree and m is the azimuthal order. $R_n^m(r)$ is the radial polynomial, written as:

$$R_n^m(r) = \sum_{s=0}^{\frac{n-|m|}{2}} \frac{(-1)^s (n-s)!}{s! \left(\frac{n+m}{2} - s\right)! \left(\frac{n-m}{2} - s\right)!} r^{n-2s}, \quad (5.4)$$

with $n \in \mathbb{N}_0$ and $n - |m|$ being even. $\Theta(\theta)$ is sometimes referred to as the triangular function [22], described by:

$$\Theta(\theta) = \begin{cases} \sin(m\theta) & m < 0 \\ 1 & m = 0 \\ \cos(m\theta) & m > 0 \end{cases} \quad (5.5)$$

Although other normalisations are equally valid, the one presented here is perhaps the most popular, chosen such that the RMS for each polynomial is unity over the unit circle [8]:

$$N_n^m = \begin{cases} \sqrt{n+1} & m = 0 \\ \sqrt{2(n+1)} & m \neq 0 \end{cases} \quad (5.6)$$

5.3 The test procedure

The premise of testing GRIN elements with a pyramid wavefront sensor is primarily based on replacing the focusing lens in Figure 5.2 with a GRIN lens to be tested. Figure 5.4 shows a simple diagram of the test setup with an Ilinsky lens, where an ideal optic should, in theory, produce a wavefront free of aberrations, or, at least only affected by any aberrations intrinsic to the pyramid element. Thus, we may compare any real lens with a theoretical, ideal one. Additionally, based on the numerical results obtained in the previous chapter, we may also estimate the aberrations introduced when the incident rays are no longer axially collimated, assuming a suitable numerical method has been used for the ray trace.

In our experiments, rather than considering rays emanating from an object at infinity, we instead placed a point source on the optical axis at a distance of $d_o = 100$ mm to the left of the front surface of each lens. A similar analysis of the Lüneburg lens was avoided, since it would have to be brought into contact with the pyramid, resulting in damage to the Lüneburg lens, the pyramid, or both in any practical implementation. 500 rays were traced through each lens, with the detector position chosen in order to remove any misconjugation errors which might affect the incoming wavefront. These values were found to be 5.2219 mm and 3.8265 mm for the Ilinsky and Tarkhanov lenses respectively. The primary wavelength was once again chosen to be $\lambda = 589.3$ nm, assumed to be produced by a monochromatic light source. Based on the results obtained in Chapter 3, Ruth’s method was used to trace both lenses during each simulated test.




Figure 5.4: An Ilinsky lens being tested with the pyramid wavefront sensor.

Figure 5.5 shows the wavefront profiles generated by the Ilinsky and Tarkhanov lenses, with positive spherical aberration being the most obvious aberration in both cases. However, observing the colour bars to the right of each wavefront profile, we notice that the peak-to-valley error for the Ilinsky lens is nearly three times less than that of the Tarkhanov lens. This therefore

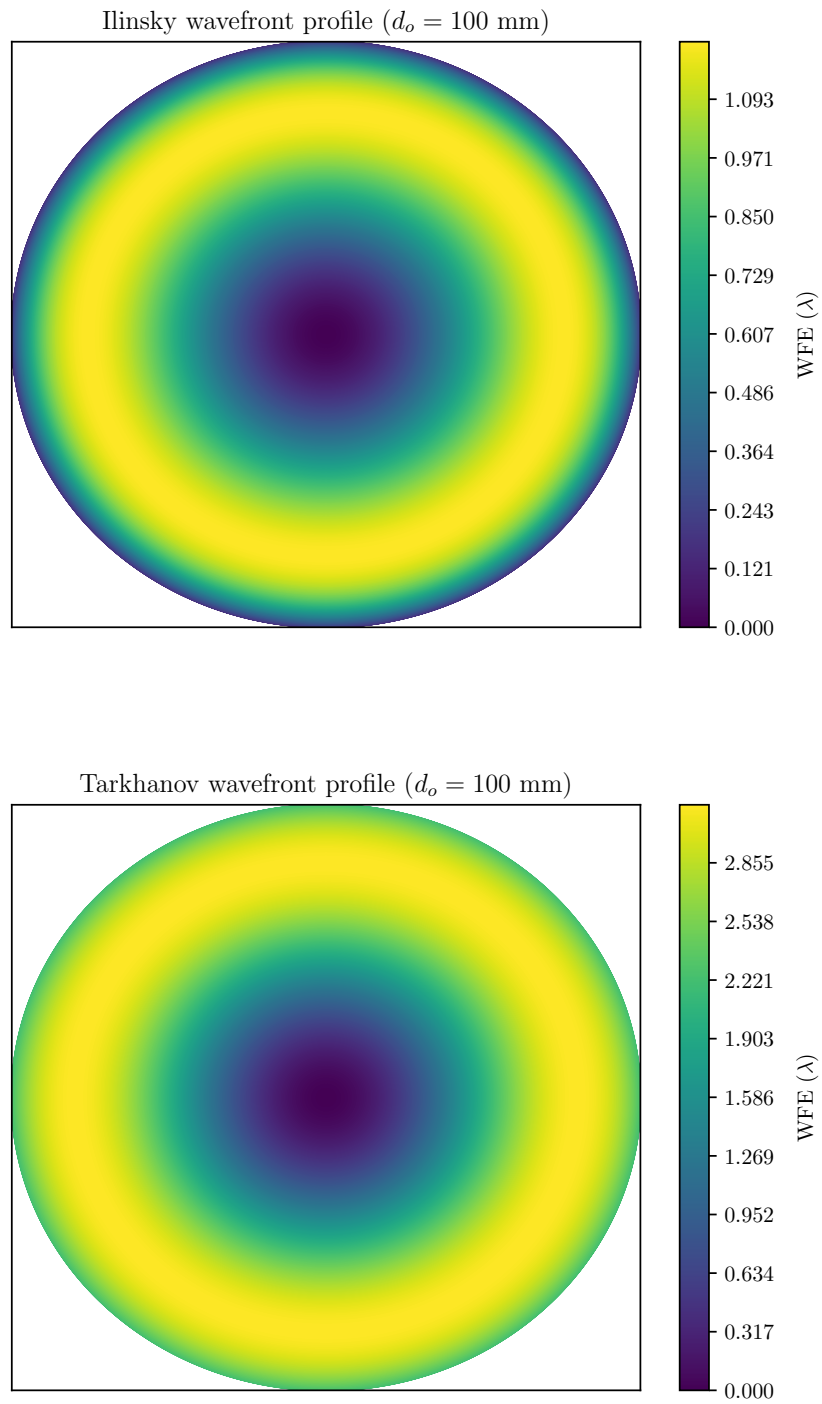


Figure 5.5: Wavefront profiles for Ilinsky (top) and Tarkhanov (bottom) lenses.

suggests that the Ilinsky lens is more tolerant to misconjugation, thereby offering better image quality for objects placed a finite distance from the front surface of the lens.

A decomposition using the first 37 Zernike polynomial coefficients for the wavefronts obtained from both lenses is given in Figure 5.6, where Zernike polynomials are arranged following the Noll indexing scheme [66]. The first three terms (i.e. piston, tip and tilt) have been omitted. As expected, the spherical terms dominate, with spherical aberration (Z_{11}) and defocus (Z_4) being almost balanced in magnitude for both lenses. Other off-axis aberrations (i.e. astigmatism, coma) are virtually negligible. Higher-order spherical terms are also present, with negative fifth-order spherical aberration (Z_{22}) being the next most dominant term for the Ilinsky lens, while negative fifth- and seventh-order spherical aberration (Z_{37}) are nearly identical in magnitude for the Tarkhanov lens. The difference in sign between the two lenses for seventh-order spherical aberration is perhaps due to the fact that their back surfaces are of opposite sign.

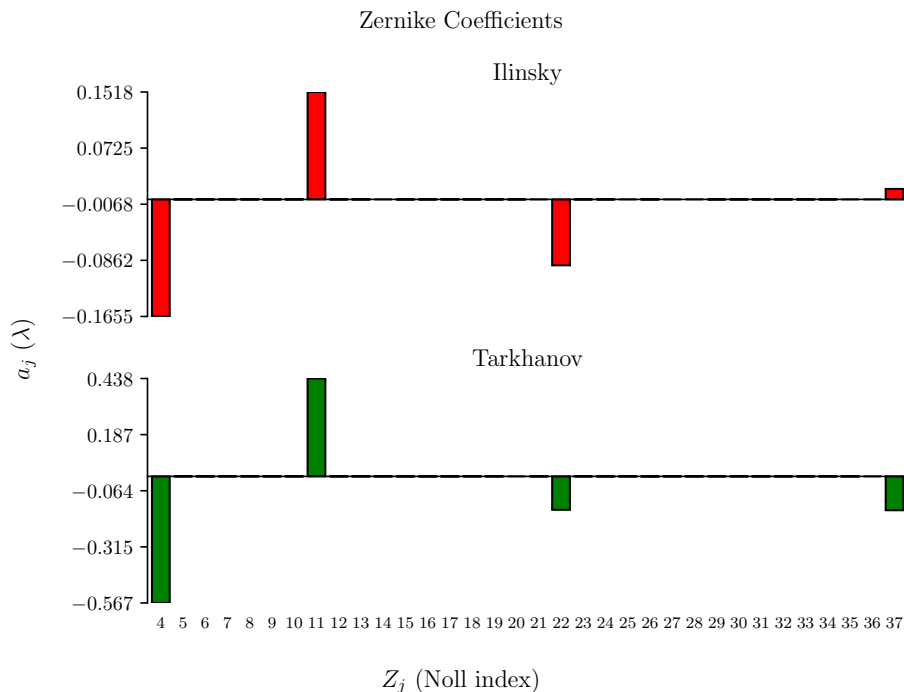


Figure 5.6: A Zernike wavefront decomposition for each lens. Piston, tip and tilt are omitted.

While the simulated experiments here present a strong case for using the pyramid wavefront sensor for aberrometry in general, we have not yet considered how aberrations associated with an off-axis, oblique wavefront may manifest themselves within the Tarkhanov and Ilinsky lenses. Moreover, while the Moore-Penrose pseudoinverse [6] was used to compute a best-fit decomposition to reconstruct the wavefront, other techniques make use of Fourier and Hilbert transforms for wavefront construction, though their accuracy is somewhat limited [10, 83]. A range of iterative methods have also been developed [42]. With this in mind, further work could be done

to see if symplectic algorithms could also be applied to wavefront reconstruction in addition to ray tracing. Since various Hamiltonian formulations of wave-mechanical problems already exist [57, 75], many of these techniques could perhaps be adapted for the purposes of wavefront propagation, reconstruction or solving other problems which may arise when considering GRIN testing from a wave-optical perspective.

Chapter 6

Conclusions and recommendations

Having first provided a brief overview of Hamiltonian optics, this work emphasises the value of symplectic numerical techniques for ray tracing GRIN media, demonstrating their increased accuracy when compared with popular nonsymplectic methods. On the whole, the results presented here make a strong case for the development of specialised symplectic methods for GRIN optics where analytical ray traces may not exist. Furthermore, we have identified the form an index may take in order for the implicit midpoint method to be used without the need for a computationally intensive root-solver and although some practical examples have been identified, several popular GRIN elements including the Maxwell fish-eye and Eaton lens fall outside this category. Higher-order symplectic Runge-Kutta methods may also be tested in the same manner as the implicit midpoint method, though deriving the necessary iterative schemes could prove to be a rather tedious process. Furthermore, our results do not suggest that implicit symplectic methods offer any significant improvement when compared with explicit methods. However, the combination of multiple symplectic methods as part of a predictor-corrector scheme could be a worthwhile area for further investigation.

In addition to future applications in optical design, symplectic methods greatly reduce the computational cost of nonlinear ray tracing, which may also be of some use to those outside the scientific community, particularly for visual effects in video games and entertainment media. Finally, we propose the pyramid wavefront sensor as a means of performing aberrometry on GRIN lenses in order to address the dearth of optical testing methods available for GRIN elements while simultaneously enabling the development of wavefront reconstruction algorithms to be used with the pyramid wavefront sensor or other prism-based wavefront sensors. Finally, pyramid wavefront sensors could also be used to find the necessary conjugation for an axial point source in order to achieve diffraction-limited imaging using the Ilinsky or Tarkhanov lenses.

A.1 Numerical recipes for explicit methods

This appendix provides numerical recipes for the explicit methods used in this thesis so that results may be reproduced. The notation given here will be identical to that in Section 3.2, where Δt represents a small step between iterations, with $P = (p_x^2 + p_y^2 + p_z^2)/2$ and $N = n^2(x, y, z)/2$.

Nonsymplectic methods

The Euler method is given by

$$\begin{aligned} p_{i_{n+1}} &= p_{i_n} - \frac{\partial}{\partial q_i} [N(q_{i_n})] \Delta t, \\ q_{i_{n+1}} &= q_{i_n} + \frac{\partial}{\partial p_i} [P(p_{i_n})] \Delta t. \end{aligned}$$

The RK2 (second-order Runge-Kutta) method, an improvement in accuracy over the standard Euler method, is represented by

$$\begin{aligned} p_{i_{n+1}} &= p_{i_n} - \left(\frac{\hat{k}_{i_1} + \hat{k}_{i_2}}{2} \right), \\ q_{i_{n+1}} &= q_{i_n} + \left(\frac{\tilde{k}_{i_1} + \tilde{k}_{i_2}}{2} \right), \end{aligned}$$

where

$$\begin{aligned} \hat{k}_{i_1} &= \frac{\partial}{\partial q_i} [N(q_{i_n})] \Delta t, & \tilde{k}_{i_1} &= \frac{\partial}{\partial p_i} [P(p_{i_n})] \Delta t, \\ \hat{k}_{i_2} &= (p_{i_n} + \hat{k}_{i_1}) \Delta t, & \tilde{k}_{i_2} &= (q_{i_n} + \tilde{k}_{i_1}) \Delta t. \end{aligned}$$

The RK4 (fourth-order Runge-Kutta) method is itself an improvement in accuracy over RK2 by further subdividing the interval between successive iterations and the use of a weighted average:

$$\begin{aligned} p_{i_{n+1}} &= p_{i_n} - \left(\frac{\hat{k}_{i_1} + 2\hat{k}_{i_2} + 2\hat{k}_{i_3} + \hat{k}_{i_4}}{6} \right), \\ q_{i_{n+1}} &= q_{i_n} + \left(\frac{\tilde{k}_{i_1} + 2\tilde{k}_{i_2} + 2\tilde{k}_{i_3} + \tilde{k}_{i_4}}{6} \right), \end{aligned}$$

where

$$\begin{aligned}\hat{k}_{i_1} &= \frac{\partial}{\partial q_i} [N(q_{i_n})] \Delta t, & \tilde{k}_{i_1} &= \frac{\partial}{\partial p_i} [P(p_{i_n})] \Delta t, \\ \hat{k}_{i_2} &= \left(p_{i_n} + \frac{\hat{k}_{i_1}}{2} \right) \Delta t, & \tilde{k}_{i_2} &= \left(q_{i_n} + \frac{\tilde{k}_{i_1}}{2} \right) \Delta t, \\ \hat{k}_{i_3} &= \left(p_{i_n} + \frac{\hat{k}_{i_2}}{2} \right) \Delta t, & \tilde{k}_{i_3} &= \left(q_{i_n} + \frac{\tilde{k}_{i_2}}{2} \right) \Delta t, \\ \hat{k}_{i_4} &= \left(p_{i_n} + \frac{\hat{k}_{i_3}}{2} \right) \Delta t, & \tilde{k}_{i_4} &= \left(q_{i_n} + \frac{\tilde{k}_{i_3}}{2} \right) \Delta t.\end{aligned}$$

Symplectic methods

The symplectic Euler method differs from the standard Euler method by first calculating a ray's optical momentum before computing the ray's position using this new momentum value.

$$\begin{aligned}p_{i_{n+1}} &= p_{i_n} - \frac{\partial}{\partial q_i} [N(q_{i_n})] \Delta t, \\ q_{i_{n+1}} &= q_{i_n} + \frac{\partial}{\partial p_i} [P(p_{i_{n+1}})] \Delta t.\end{aligned}$$

The velocity Verlet method is created by taking two symplectic Euler method iterations with step size $\Delta t/2$. This method is symmetric, producing nearly-identical trajectories for positive and negative step sizes.

$$\begin{aligned}q_{i_{n+1}} &= q_{i_n} + \frac{\partial}{\partial p_i} [P(p_{i_n})] \Delta t + \frac{\partial}{\partial q_i} [N(q_{i_n})] \frac{(\Delta t)^2}{2}, \\ p_{i_{n+1}} &= p_{i_n} - \frac{\partial}{\partial q_i} [N(q_{i_n}) + N(q_{i_{n+1}})] \frac{\Delta t}{2}.\end{aligned}$$

Ruth's method was initially developed for molecular dynamics simulations. Like the velocity Verlet method, is also symmetric, described by

$$\begin{aligned}
p_{i_{n+1/4}} &= p_{i_n} + \left(\frac{2^{1/3} + 2^{-1/3} + 2}{6} \right) \frac{\partial}{\partial q_i} [N(q_{i_n})] \Delta t, \\
q_{i_{n+1/3}} &= q_{i_n} + \left(\frac{2^{1/3} + 2^{-1/3} + 2}{3} \right) p_{i_{n+1/4}} \Delta t, \\
p_{i_{n+1/2}} &= p_{i_{n+1/4}} - \left(\frac{2^{1/3} + 2^{-1/3} - 1}{6} \right) \frac{\partial}{\partial q_i} [N(q_{i_{n+1/3}})] \Delta t, \\
q_{i_{n+2/3}} &= q_{i_{n+1/3}} - \left(\frac{2^{7/3} + 2^{5/3} + 2}{6} \right) p_{i_{n+1/2}} \Delta t, \\
p_{i_{n+3/4}} &= q_{i_{n+2/3}} - \left(\frac{2^{1/3} + 2^{-1/3} - 1}{6} \right) \frac{\partial}{\partial q_i} [N(q_{i_{n+2/3}})] \Delta t, \\
q_{i_{n+1}} &= q_{i_{n+2/3}} + \left(\frac{2^{1/3} + 2^{-1/3} + 2}{3} \right) p_{i_{n+3/4}} \Delta t, \\
p_{i_{n+1}} &= p_{i_{n+3/4}} + \left(\frac{2^{1/3} + 2^{-1/3} + 2}{6} \right) \frac{\partial}{\partial q_i} [N(q_{i_{n+1}})] \Delta t.
\end{aligned}$$

It is perhaps worth drawing attention to the positive signs before the bracketed coefficients in the expressions for $p_{i_{n+1/4}}$ and $p_{i_{n+1}}$. These are not errors, but rather are results of the coefficients being negative, as mentioned previously in Chapter 3.

Finally, the Chin-Chen 4A method is one of a series of symmetric fourth-order numerical methods created with the aim of solving few-body gravitational problems, which are almost mathematically identical to ray trajectories within the Eaton lens. The method itself is written as

$$\begin{aligned}
p_{i_{n+1/3}} &= p_{i_n} - \frac{\partial}{\partial q_i} [N(q_{i_n})] \frac{\Delta t}{6}, \\
q_{i_{n+1/2}} &= q_{i_n} + p_{i_{n+1/3}} \frac{\Delta t}{2}, \\
p_{i_{n+2/3}} &= p_{i_{n+1/3}} - \frac{\partial}{\partial q_i} [N(q_{i_{n+1/2}})] \frac{2\Delta t}{3} - \frac{\partial^2}{\partial q_i^2} [N(q_{i_{n+1/2}})] \frac{(\Delta t)^3}{72}, \\
q_{i_{n+1}} &= q_{i_{n+1/2}} + p_{i_{n+2/3}} \frac{\Delta t}{2}, \\
p_{i_{n+1}} &= p_{i_{n+2/3}} - \frac{\partial}{\partial q_i} [N(q_{i_{n+1}})] \frac{\Delta t}{6}.
\end{aligned}$$

A.2 Proof that the implicit midpoint is indeed symplectic

In order to prove that the implicit midpoint is symplectic, we must first compute ∇H by differentiating Eq. (4.12)

$$\begin{aligned}\frac{\partial p_{i_{n+1}}}{\partial p_{i_n}} &= \left(1 - \frac{\partial^2 H}{\partial q_i^2} \frac{\Delta t}{2}\right) \frac{\partial q_{i_{n+1}}}{\partial p_{i_n}}, \\ \frac{\partial p_{i_{n+1}}}{\partial q_{i_n}} &= -\left(1 + \frac{\partial q_{i_{n+1}}}{\partial q_{i_n}}\right) \frac{\partial^2 H}{\partial q_i^2} \frac{\Delta t}{2}, \\ \frac{\partial q_{i_{n+1}}}{\partial p_{i_n}} &= \left(1 + \frac{\partial p_{i_{n+1}}}{\partial p_{i_n}}\right) \frac{\partial^2 H}{\partial p_i^2} \frac{\Delta t}{2}, \\ \frac{\partial q_{i_{n+1}}}{\partial q_{i_n}} &= \left(1 + \frac{\partial^2 H}{\partial p_i^2} \frac{\Delta t}{2}\right) \frac{\partial p_{i_{n+1}}}{\partial q_{i_n}}.\end{aligned}$$

Now, substituting each of the expressions into one another, as necessary.

$$\begin{aligned}\frac{\partial p_{i_{n+1}}}{\partial p_{i_n}} &= \left(1 - \frac{\partial^2 H}{\partial q_i^2} \frac{\Delta t}{2}\right) \left(1 + \frac{\partial p_{i_{n+1}}}{\partial p_{i_n}}\right) \frac{\partial^2 H}{\partial p_i^2} \frac{\Delta t}{2}, \\ \frac{\partial p_{i_{n+1}}}{\partial q_{i_n}} &= -\left\{1 + \left[\left(1 + \frac{\partial^2 H}{\partial p_i^2} \frac{\Delta t}{2}\right) \frac{\partial p_{i_{n+1}}}{\partial q_{i_n}}\right]\right\} \frac{\partial^2 H}{\partial q_i^2} \frac{\Delta t}{2}, \\ \frac{\partial q_{i_{n+1}}}{\partial p_{i_n}} &= \left\{1 + \left[\left(1 - \frac{\partial^2 H}{\partial q_i^2} \frac{\Delta t}{2}\right) \frac{\partial q_{i_{n+1}}}{\partial p_{i_n}}\right]\right\} \frac{\partial^2 H}{\partial p_i^2} \frac{\Delta t}{2}, \\ \frac{\partial q_{i_{n+1}}}{\partial q_{i_n}} &= -\left(1 + \frac{\partial^2 H}{\partial p_i^2} \frac{\Delta t}{2}\right) \left(1 + \frac{\partial q_{i_{n+1}}}{\partial q_{i_n}}\right) \frac{\partial^2 H}{\partial q_i^2} \frac{\Delta t}{2}.\end{aligned}$$

Then, isolating each of the first-order partial derivatives

$$\begin{aligned}\frac{\partial p_{i_{n+1}}}{\partial p_{i_n}} &= \left[1 + \frac{\partial^2 H}{\partial q^2} \frac{\partial^2 H}{\partial p^2} \left(\frac{\Delta t}{2}\right)^2\right]^{-1} \left[1 - \frac{\partial^2 H}{\partial q^2} \frac{\partial^2 H}{\partial p^2} \left(\frac{\Delta t}{2}\right)^2\right], \\ \frac{\partial p_{i_{n+1}}}{\partial q_i} &= -\left[1 + \frac{\partial^2 H}{\partial q^2} \frac{\partial^2 H}{\partial p^2} \left(\frac{\Delta t}{2}\right)^2\right]^{-1} \frac{\partial^2 H}{\partial q^2} (\Delta t)^2, \\ \frac{\partial q_{i_{n+1}}}{\partial p_{i_n}} &= \left[1 + \frac{\partial^2 H}{\partial q^2} \frac{\partial^2 H}{\partial p^2} \left(\frac{\Delta t}{2}\right)^2\right]^{-1} \frac{\partial^2 H}{\partial p^2} (\Delta t)^2, \\ \frac{\partial q_{i_{n+1}}}{\partial q_{i_n}} &= \left[1 + \frac{\partial^2 H}{\partial q^2} \frac{\partial^2 H}{\partial p^2} \left(\frac{\Delta t}{2}\right)^2\right]^{-1} \left[1 - \frac{\partial^2 H}{\partial q^2} \frac{\partial^2 H}{\partial p^2} \left(\frac{\Delta t}{2}\right)^2\right].\end{aligned}$$

and so,

$$\nabla H = \left[1 + \frac{\partial^2 H}{\partial q^2} \frac{\partial^2 H}{\partial p^2} \left(\frac{\Delta t}{2}\right)^2\right]^{-1} \begin{pmatrix} 1 - \frac{\partial^2 H}{\partial q^2} \frac{\partial^2 H}{\partial p^2} \left(\frac{\Delta t}{2}\right)^2 & -\frac{\partial^2 H}{\partial q^2} (\Delta t)^2 \\ \frac{\partial^2 H}{\partial p^2} (\Delta t)^2 & 1 - \frac{\partial^2 H}{\partial q^2} \frac{\partial^2 H}{\partial p^2} \left(\frac{\Delta t}{2}\right)^2 \end{pmatrix}.$$

For notational simplicity, let

$$\alpha = \left[1 - \frac{\partial^2 H}{\partial q^2} \frac{\partial^2 H}{\partial p^2} \left(\frac{\Delta t}{2} \right)^2 \right], \quad \beta = -\frac{\partial^2 H}{\partial q^2} (\Delta t)^2,$$

$$\gamma = \left[1 + \frac{\partial^2 H}{\partial q^2} \frac{\partial^2 H}{\partial p^2} \left(\frac{\Delta t}{2} \right)^2 \right]^{-1}, \quad \delta = \frac{\partial^2 H}{\partial p^2} (\Delta t)^2.$$

Then, substituting into Eq. (2.20)

$$\begin{aligned} (\nabla H)^T \mathbf{J} (\nabla H) &= \gamma^2 \begin{pmatrix} \alpha & \beta \\ \delta & \alpha \end{pmatrix}^T \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \begin{pmatrix} \alpha & \beta \\ \delta & \alpha \end{pmatrix}, \\ &= \gamma^2 \begin{pmatrix} \alpha & \delta \\ \beta & \alpha \end{pmatrix} \begin{pmatrix} \delta & \alpha \\ -\alpha & -\beta \end{pmatrix}, \\ &= \gamma^2 \begin{pmatrix} 0 & \alpha^2 - \beta\delta \\ \beta\delta - \alpha^2 & 0 \end{pmatrix}, \end{aligned}$$

algebraic manipulation reveals $\gamma^2(\alpha^2 - \beta\delta) = 1$. Hence

$$\begin{aligned} (\nabla H)^T \mathbf{J} (\nabla H) &= \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}, \\ &= \mathbf{J}, \end{aligned}$$

Thus, the implicit midpoint method is a symplectic numerical method. \square

A.3 Proof of the nonexistence of explicit symplectic Runge-Kutta methods

In general, a Runge-Kutta method may be written as follows

$$y_{n+1} = y_n + \Delta t \sum_{i=1}^m b_i k_i,$$

with

$$\begin{aligned} k_1 &= f[t_n, y_n], \\ k_2 &= f[t_n + c_2 \Delta t, y_n + (a_{21} k_1) \Delta t], \\ k_3 &= f[t_n + c_3 \Delta t, y_n + (a_{31} k_1 + a_{32} k_2) \Delta t], \\ &\vdots \\ k_m &= f \left[t_n + c_m \Delta t, y_n + \Delta t \sum_{j=1}^m a_{mj} k_j \right], \end{aligned}$$

where each a_{ij} is a coefficient in the Runge-Kutta matrix, b_i is the associated weight and c_i is known as a node. As in the rest of the text, m represents the method's order. Alternatively, we may represent a Runge-Kutta method in the form of a Butcher tableau [15], shown below

$$\begin{array}{c|cccc} c_1 & a_{11} & a_{12} & \dots & a_{1m} \\ c_2 & a_{21} & a_{22} & \dots & a_{2m} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ c_m & a_{m1} & a_{m2} & \dots & a_{mm} \\ \hline & b_1 & b_2 & \dots & b_m \end{array} \quad \text{or, more succinctly} \quad \begin{array}{c|c} \mathbf{c} & \mathbf{a} \\ \hline & \mathbf{b}^T \end{array}$$

For a Runge-Kutta method to be explicit, its Runge-Kutta matrix \mathbf{a} must be strictly lower triangular (i.e. $a_{ij} = 0$ for $j \geq i$). Furthermore, for a Runge-Kutta method to be useful, it must be consistent, where a Runge-Kutta method is classified as consistent if and only if $\sum_{i=1}^m b_i = 1$. However, in order for a Runge-Kutta method to be symplectic, it must also satisfy the relation $\mathbf{b}^T \mathbf{a} + \mathbf{b} \mathbf{a}^T - \mathbf{b}^T \mathbf{b} = \mathbf{0} \forall i, j = 1, \dots, m$.

We first assume there exists a method whose Runge-Kutta matrix is strictly lower triangular while also satisfying $\mathbf{b}^T \mathbf{a} + \mathbf{b} \mathbf{a}^T - \mathbf{b}^T \mathbf{b} = \mathbf{0} \forall i, j = 1, \dots, m$. Yet, if we examine the case for terms on the main diagonal (i.e. where $i = j$), we notice $b_i \cdot 0 + b_j \cdot 0 - b_i b_j = 0$, since for explicit methods $a_{ij} = a_{ji} = 0$. Hence, we require $b_i b_j = b_i^2 = 0 \forall i = 1, \dots, m$, which contradicts $\sum_{i=1}^m b_i = 1$, making it impossible for a Runge-Kutta method to be both explicit and symplectic. In other words, symplectic Runge-Kutta methods are necessarily implicit. \square

Bibliography

- [1] T. B. Andersen. Efficient and robust recurrence relations for the Zernike circle polynomials and their derivatives in Cartesian coordinates. *Optics Express*, 26(15):18878–18896, 2018.
- [2] R. Antonelli. starless. <https://github.com/rantonels/starless>, 2015. [GitHub repository; accessed 14-June-2023].
- [3] V. I. Arnol'd. *Mathematical Methods of Classical Mechanics*, volume 60. Springer Science & Business Media, 2013.
- [4] J. Babington. Freeform aberrations in phase space: an example. *JOSA A*, 34(6):1045–1053, 2017.
- [5] J. Babington. Phase space aberrations in freeform and gradient index imaging systems. *Optical Engineering*, 57(10):105106–105106, 2018.
- [6] J. C. A. Barata and M. S. Hussein. The Moore–Penrose pseudoinverse: A tutorial review of the theory. *Brazilian Journal of Physics*, 42:146–165, 2012.
- [7] S. Behnel, R. Bradshaw, C. Citro, L. Dalcin, D. S. Seljebotn, and K. Smith. Cython: The best of both worlds. *Computing in Science & Engineering*, 13(2):31–39, 2010.
- [8] J. Bentley and C. Olson. *Field Guide to Lens Design*. SPIE, 2012.
- [9] A. Bertrou-Cantou, E. Gendron, G. Rousset, V. Deo, F. Ferreira, A. Sevin, and F. Vidal. Confusion in differential piston measurement with the pyramid wavefront sensor. *Astronomy & Astrophysics*, 658:A49, 2022.
- [10] C. Z. Bond, C. M. Correia, J.-F. Sauvage, K. El Hadi, B. Neichel, and T. Fusco. Fourier wavefront reconstruction with a pyramid wavefront sensor. In *Adaptive Optics Systems VI*, volume 10703, pages 1226–1233. SPIE, 2018.

- [11] M. Born and E. Wolf. *Principles of optics: electromagnetic theory of propagation, interference and diffraction of light*. Elsevier, 1997.
- [12] H. A. Buchdahl. Rays in gradient-index media: separable systems. *JOSA*, 63(1):46–49, 1973.
- [13] H. A. Buchdahl. Kepler problem and Maxwell fish-eye. *American Journal of Physics*, 46(8):840–843, 1978.
- [14] H. A. Buchdahl. *An Introduction to Hamiltonian Optics*. Courier Corporation, 1993.
- [15] J. C. Butcher. A stability property of implicit Runge-Kutta methods. *BIT Numerical Mathematics*, 15(4):358–361, 1975.
- [16] M. Carillet, C. Vérinaud, B. Femenía, A. Riccardi, and L. Fini. Modelling astronomical adaptive optics-I. The software package CAOS. *Monthly Notices of the Royal Astronomical Society*, 356(4):1263–1275, 2005.
- [17] A. Chacon and A. Vladimírsky. Fast two-scale methods for eikonal equations. *SIAM Journal on Scientific Computing*, 34(2):A547–A578, 2012.
- [18] J. Chaves. *Introduction to Nonimaging Optics*. CRC press, 2008.
- [19] S. A. Chin and C. R. Chen. Forward symplectic integrators for solving gravitational few-body problems. *Celestial Mechanics and Dynamical Astronomy*, 91:301–322, 2005.
- [20] L. Clarke, I. Glendinning, and R. Hempel. The MPI message passing interface standard. In *Programming Environments for Massively Parallel Distributed Systems: Working Conference of the IFIP WG 10.3, April 25–29, 1994*, pages 213–218. Springer, 1994.
- [21] L. Dagum and R. Menon. OpenMP: an industry standard API for shared-memory programming. *IEEE computational science and engineering*, 5(1):46–55, 1998.
- [22] G. M. Dai. Modal wave-front reconstruction with Zernike polynomials and Karhunen–Loève functions. *JOSA A*, 13(6):1218–1225, 1996.
- [23] M. A. DeGosson. *Symplectic Methods in Harmonic Analysis and in Mathematical Physics*. Springer Science & Business Media, 2011.
- [24] A. Demetriadou and Y. Hao. Slim Luneburg lens for antenna applications. *Optics express*, 19(21):19925–19934, 2011.

- [25] D. Donnelly and E. Rogers. Symplectic integrators: an introduction. *American Journal of Physics*, 73(10):938–945, 2005.
- [26] G. Du, M. Liang, R. A. Sabory-Garcia, C. Liu, and H. Xin. 3-D printing implementation of an X-band Eaton lens for beam deflection. *IEEE Antennas and Wireless Propagation Letters*, 15:1487–1490, 2016.
- [27] J. E. Eaton. An Extension of the Luneberg-Type Lenses. Technical report, Naval Research Lab, Washington D.C., 1953.
- [28] E. Fadeev, V. Kosyakov, and A. Tikhvatulin. Mathematical simulation of the technology and properties of gradient spherical lenses. *Technical Physics*, 43:1199–1202, 1998.
- [29] K. Feng and M. Qin. *Symplectic Geometric Algorithms for Hamiltonian Systems*. Springer, 2010.
- [30] K. Feng and D. Wang. Variations on a theme by Euler. *Journal of Computational Mathematics*, pages 97–106, 1998.
- [31] C. Flynn and A. V. Goncharov. Semi-Analytical Finite Ray-Tracing Through the Quadratic Symmetric GRIN Lens. *Optica Open preprint*, 8 2023.
- [32] E. Forest and R. D. Ruth. Fourth-order symplectic integration. *Physica D: Nonlinear Phenomena*, 43(1):105–117, 1990.
- [33] J. Franco and I. Gómez. Symplectic explicit methods of Runge–Kutta–Nyström type for solving perturbed oscillators. *Journal of Computational and Applied Mathematics*, 260:482–493, 2014.
- [34] S. K. Gray, D. W. Noid, and B. G. Sumpter. Symplectic integrators for large scale molecular dynamics simulations: a comparison of several explicit methods. *The Journal of Chemical Physics*, 101(5):4062–4072, 1994.
- [35] H. Gross. *Handbook of Optical Systems, Volume 1, Fundamentals of Technical Optics*, volume 1. Wiley Online Library, 2005.
- [36] S. Haffert. Generalised optical differentiation wavefront sensor: a sensitive high dynamic range wavefront sensor. *Optics Express*, 24(17):18986–19007, 2016.
- [37] E. Hairer, C. Lubich, and G. Wanner. *Geometric Numerical Integration: structure-preserving algorithms for ordinary differential equations*. Springer, 2006.

- [38] E. Hairer and G. Wanner. Euler methods, explicit, implicit, symplectic. *Encyclopedia of Applied and Computational Mathematics*, pages 451–455, 2015.
- [39] W. R. Hamilton. Theory of systems of rays. *The Transactions of the Royal Irish Academy*, pages 69–174, 1828.
- [40] W. R. Hamilton. *On a General Method in Dynamics*. Richard Taylor, 1834.
- [41] L. N. Hand and J. D. Finch. *Analytical Mechanics*. Cambridge University Press, 1998.
- [42] V. Hutterer and R. Ramlau. Nonlinear wavefront reconstruction methods for pyramid sensors using Landweber and Landweber–Kaczmarz iterations. *Applied Optics*, 57(30):8790–8804, 2018.
- [43] R. Ilinsky. Gradient-index meniscus lens free of spherical aberration. *Journal of Optics A: Pure and Applied Optics*, 2(5):449, 2000.
- [44] M. Kang, H. Huang, H. Xu, and F. Liu. Optical black-hole analog in inhomogeneous photonic lattice. In *APS March Meeting Abstracts*, volume 2019, pages G70–219, 2019.
- [45] H. Kinoshita, H. Yoshida, and H. Nakai. Symplectic integrators and their application to dynamical astronomy. *Celestial Mechanics and Dynamical Astronomy*, 50:59–71, 1991.
- [46] V. Lakshminarayanan, A. Ghatak, and K. Thyagarajan. *Lagrangian Optics*. Springer Science & Business Media, 2002.
- [47] S. K. Lam, A. Pitrou, and S. Seibert. Numba: A LLVM-based Python JIT compiler. In *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*, pages 1–6, 2015.
- [48] R. B. Leighton and M. Sands. *The Feynman Lectures on Physics*. Addison-Wesley Boston, MA, USA, 1965.
- [49] B. Leimkuhler and S. Reich. *Simulating Hamiltonian Dynamics*. Cambridge University Press, 2004.
- [50] D. H. Lippman, N. S. Kochan, T. Yang, G. R. Schmidt, J. L. Bentley, and D. T. Moore. Freeform gradient-index media: a new frontier in freeform optics. *Optics Express*, 29(22):36997–37012, 2021.
- [51] W. Liu, H. Hu, F. Liu, and H. Zhao. Manipulating light trace in a gradient-refractive-index medium: a Lagrangian optics method. *Optics Express*, 27(4):4714–4726, 2019.

- [52] R. K. Luneburg. *Mathematical Theory of Optics*. University of California Press, 1966.
- [53] M. Malisoff. Bounded-from-below solutions of the Hamilton-Jacobi equation for optimal control problems with exit times: vanishing Lagrangians, eikonal equations, and shape-from-shading. *Nonlinear Differential Equations and Applications NoDEA*, 11:95–122, 2004.
- [54] C. Mateo-Segura, A. Dyke, H. Dyke, S. Haq, and Y. Hao. Flat Luneburg lens via transformation optics for directive antenna applications. *IEEE Transactions on Antennas and Propagation*, 62(4):1945–1953, 2014.
- [55] J. C. Maxwell. *The Scientific Papers of James Clerk Maxwell*, volume 2. University Press, 1890.
- [56] B. McKeon and A. V. Goncharov. Symplectic numerical methods in optics and imaging: ray tracing in spherical gradient-index lenses and computer-generated image rendering. *Optica Open preprint*, 7 2023.
- [57] R. McLachlan. Symplectic integration of Hamiltonian wave equations. *Numerische Mathematik*, 66:465–492, 1993.
- [58] E. Merchand. *Gradient Index Optics*. Elsevier, 2012.
- [59] A. Meurer, C. P. Smith, M. Paprocki, O. Čertík, S. B. Kirpichev, M. Rocklin, et al. SymPy: symbolic computing in Python. *PeerJ Computer Science*, 3:e103, Jan. 2017.
- [60] A. L. Mikaelian. SELFOC dielectric waveguides. *Soviet Journal of Quantum Electronics*, 7(2):266, 1977.
- [61] A. L. Mikaelian and A. M. Prokhorov. V. self-focusing media with variable index of refraction. In *Progress in Optics*, volume 17, pages 279–345. Elsevier, 1980.
- [62] S. Mikkola and K. Tanikawa. Explicit Symplectic Algorithms For Time-Transformed Hamiltonians. *Celestial Mechanics and Dynamical Astronomy*, 74:287–295, 1999.
- [63] J. C. Miñano, P. Benítez, and A. Santamaría. Hamilton-Jacobi equation in momentum space. *Optics Express*, 14(20):9083–9092, 2006.
- [64] D. T. Moore. Gradient-index optics: a review. *Applied Optics*, 19(7):1035–1038, 1980.
- [65] J. Nichols. Approximate solutions for propagating light beams in a material with quadratic, exponential, and quartic transverse refractive index profiles. *Journal of Optics*, 22(4):045601, 2020.

- [66] R. J. Noll. Zernike polynomials and atmospheric turbulence. *JOSA*, 66(3):207–211, 1976.
- [67] R. L. Nowack. Wavefronts and solutions of the eikonal equation. *Geophysical Journal International*, 110(1):55–62, 1992.
- [68] H. Ohno. Symplectic ray tracing based on Hamiltonian optics in gradient-index media. *JOSA A*, 37(3):411–416, 2020.
- [69] H. Ohno and T. Usui. Gradient-index dark hole based on conformal mapping with etendue conservation. *Optics Express*, 27(13):18493–18507, 2019.
- [70] R. Pegis. I. The modern development of Hamiltonian optics. In *Progress in Optics*, volume 1, pages 1–29. Elsevier, 1961.
- [71] E. H. Por, S. Y. Haffert, V. M. Radhakrishnan, D. S. Doelman, M. van Kooten, and S. P. Bos. High Contrast Imaging for Python (HCIPy): an open-source adaptive optics and coronagraph simulator. In *Adaptive Optics Systems VI*, volume 10703, page 1070342. International Society for Optics and Photonics, 2018.
- [72] R. Ragazzoni. Pupil plane wavefront sensing with an oscillating prism. *Journal of Modern Optics*, 43(2):289–293, 1996.
- [73] A. Rangwala, A. Ghodgaonkar, and R. Tewari. Paraxial ray tracing in inhomogeneous optical media. *Optical Engineering*, 37(3):1025–1032, 1998.
- [74] J. Rebordao and M. Grosmann. Exact Solutions Of Eikonal Equation And Their Relation To Hamilton Characteristic Functions. In *Optical System Design, Analysis, and Production*, volume 399, pages 178–185. SPIE, 1983.
- [75] S. Reich. Multi-symplectic Runge–Kutta collocation methods for Hamiltonian wave equations. *Journal of Computational Physics*, 157(2):473–499, 2000.
- [76] G. Rodríguez, J. Villa, R. Ivanov, E. González, and G. Martínez. Foucault test: a quantitative evaluation method. *JOSA A*, 33(8):1604–1611, 2016.
- [77] P. Saha and S. Tremaine. Symplectic integrators for solar system dynamics. *The Astronomical Journal*, 104:1633–1640, 1992.
- [78] J. M. Sanz-Serna. Symplectic integrators for Hamiltonian problems: an overview. *Acta numerica*, 1:243–286, 1992.
- [79] J. M. Sanz-Serna. Symplectic Runge–Kutta and related methods: recent results. *Physica*

- D: Nonlinear Phenomena*, 60(1-4):293–302, 1992.
- [80] T. R. Satoh. Symplectic ray tracing: a new approach to nonlinear ray tracing by using Hamiltonian dynamics. In *Visualization and Data Analysis 2003*, volume 5009, pages 277–285. SPIE, 2003.
- [81] J. A. Sethian. A fast marching level set method for monotonically advancing fronts. *Proceedings of the National Academy of Sciences*, 93(4):1591–1595, 1996.
- [82] A. Sharma, D. V. Kumar, and A. K. Ghatak. Tracing rays through graded-index media: a new method. *Applied Optics*, 21(6):984–987, 1982.
- [83] I. Shatokhina, V. Hutterer, and R. Ramlau. Review on methods for wavefront reconstruction from pyramid wavefront sensor data. *Journal of Astronomical Telescopes, Instruments, and Systems*, 6(1):010901, 2020.
- [84] C. J. Sheil and A. V. Goncharov. Accommodating volume-constant age-dependent optical (AVOCADO) model of the crystalline GRIN lens. *Biomedical optics express*, 7(5):1985–1999, 2016.
- [85] P. Shirley and R. K. Morley. *Realistic ray tracing*. AK Peters, Ltd., 2008.
- [86] G. F. Simmons. *Differential Equations with Applications and Historical Notes*. CRC Press, 2016.
- [87] R. D. Skeel, G. Zhang, and T. Schlick. A family of symplectic integrators: stability, accuracy, and molecular dynamics applications. *SIAM Journal on Scientific Computing*, 18(1):203–222, 1997.
- [88] M. Soskin and M. Vasnetsov. Singular optics. *Progress in Optics*, 42(4):219–276, 2001.
- [89] E. Süli and D. F. Mayers. *An Introduction to Numerical Analysis*. Cambridge University Press, 2003.
- [90] J. L. Synge. *Geometrical Optics: an introduction to Hamilton’s method*. Cambridge University Press, 1937.
- [91] V. I. Tarkhanov. Aplanatic gradient lens. USSR Author Certificate (analog of patent) No. 1337861, G02 B 6/00.
- [92] M. Testorf, B. Hennelly, and J. Ojeda-Castañeda. *Phase-Space Optics: fundamentals and applications*. McGraw-Hill Education, 2010.

- [93] V. A. Toponogov. *Differential Geometry of Curves and Surfaces*. Springer, 2006.
- [94] A. Torre. *Linear Ray and Wave Optics in Phase Space: bridging ray and wave optics via the Wigner phase-space picture*. Elsevier, 2005.
- [95] M. Townson, O. Farley, G. O. de Xivry, J. Osborn, and A. Reeves. AOtools: a Python package for adaptive optics modelling and analysis. *Optics Express*, 27(22):31316–31329, 2019.
- [96] S. G. Turyshev and V. T. Toth. Resolved imaging of exoplanets with the solar gravitational lens. *Monthly Notices of the Royal Astronomical Society*, 515(4):6122–6132, 2022.
- [97] T. Tyc and A. Danner. Frequency spectra of absolute optical instruments. *New Journal of Physics*, 14(8):085023, 2012.
- [98] L. Verlet. Computer ”experiments” on classical fluids. I. Thermodynamical properties of Lennard-Jones molecules. *Physical review*, 159(1):98, 1967.
- [99] G. Wanner and E. Hairer. *Solving Ordinary Differential Equations II: stiff and differential-algebraic problems*. Springer Berlin, Heidelberg, New York, 1996.
- [100] K. B. Wolf. *Geometric Optics on Phase Space*. Springer Science & Business Media, 2004.
- [101] J. Ye, W. Wang, Z. Gao, Z. Liu, S. Wang, P. Benítez, J. C. Miñano, and Q. Yuan. Modal wavefront estimation from its slopes by numerical orthogonal transformation method over general shaped aperture. *Optics Express*, 23(20):26208–26220, 2015.
- [102] A. J. Yee and D. T. Moore. Free-space infrared Mach–Zehnder interferometer for relative index of refraction measurement of gradient index optics. *Optical Engineering*, 56(11):111707–111707, 2017.
- [103] H. Yoshida. Construction of higher order symplectic integrators. *Physics Letters A*, 150(5-7):262–268, 1990.
- [104] H.-K. Zhao. A fast sweeping method for eikonal equations. *Mathematics of computation*, 74(250):603–627, 2005.
- [105] Y. Y. Zhao, Y. L. Zhang, M. L. Zheng, X. Z. Dong, X. M. Duan, and Z. S. Zhao. Three-dimensional Luneburg lens at optical frequencies. *Laser & Photonics Reviews*, 10(4):665–672, 2016.