Metarefraction

Alasdair Craig Hamilton
MSci

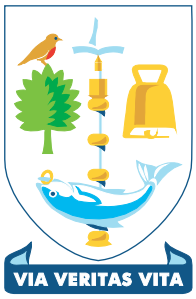

Submitted in fulfilment of the requirements for
the Degree of PhD

Department of Physics & Astronomy
Faculty of Physical Sciences
University of Glasgow

September 1, 2010

Abstract

Imagine a thin sheet that performs optical illusions on the scene behind it. For example, a window that appears to reverse depth and to image objects in front of the sheet, or alternatively swimming goggles that cancel the refraction of surrounding water. This thesis will explore how such sheets may be realized.

With the refinement of optical fabrication technologies, it is now possible to mass-produce miniaturized optical components. Repeating them over the surface of a sheet, their combined effect may realize optical effects from the structure, rather than the substance, of the sheet. Specifically, such components may realize arbitrary ray-direction mappings at each point on the sheet. Here such mappings, *metarefractions*, are explored from a range of perspectives.

This thesis will explore the inception, theoretical development and ultimately the experimental realization of metarefraction. At its core, this work is primarily mathematical in nature but draws upon both experimental and computational techniques in order to test and visualize the concepts that will be discussed. Examples of such ray-direction mappings will be explored as will their ray- and wave-optical implications.

This thesis is structured as follows: Initially, the definition of metarefraction, along with some existing examples, is presented. Then, ray mappings are related to negative refraction, a subject that metarefraction has a surprising number of parallels to. New forms of metarefraction are then introduced, before being incorporated into imaging systems. Later, ray-optical transformations, such as metarefraction, are shown to be limited by implicit wave-optical restrictions. In some cases, these vastly reduce the number of light fields that may be exactly transformed. After this, the most general possible metarefraction is sought, and a simple case is realized experimentally. Further restrictions are then determined, before finishing with a discussion and summary, and by considering possible directions that future work could develop in.

Preface

Over the course of my PhD, I have had the good fortune to explore a number of areas of physics. My research began by exploring atom optics before completely changing direction. After this, I worked on the incorporation of transformation optics and integral imaging into a new form of ray-optical transformation.

In my first year, I focused on BEC phase imprinting — work that was abandoned when it transpired that another group had already published the same idea. In the December of that year, the focus then changed during a visit to Aimé Cotton laboratory, in Orsay. There I modelled a novel optical setup to reduce laser speckle in an optical 'lens' in order to shape a cold atom cloud. It was on these simulations that I spent much of the first year of my PhD.

During the Orsay trip, my supervisor, Johannes Courtial, gave a talk on his ideas on Dove-prism sheets. I found the optical properties of these structures to be both fascinating and counter intuitive. It is from that talk that all of the work in this thesis since developed. Along the way, I have visited micro fabrication facilities at Strathclyde University and in Durham, simulated and visualized a number of such sheets, explored their ray- and wave-optical properties and ultimately built metarefracting sheets.

Discussions of such sheet structures has also lead to additional research. For example, at the EOS'08 conference, we met Carl Paterson who had previously studied limitations of related transformations. It became apparent that some of our sheets were examples of optical transformations that lacked an exact wave-optical analogue. Later, in collaboration with Mark Dennis, we explored how ray direction was limited by wave-optical restrictions.

(Signature) _____

Acknowledgements

This work was performed in collaboration with Johannes Courtial to whom I owe considerable thanks for his creativity, patience and enthusiasm throughout. Aspects of the computational and theoretical work was also in collaboration with Bhuvanesh Sundar, Sharvil Talati, John Nelson, Martin Šarbort and Tomáš Tyc. The experimental work was in collaboration with Michael Blair, Leo Clark, Alasdair Houston and Gary Smith and Jonathan Leach. Further mention will be made of each person's contribution where appropriate throughout this text. Finally, I would like to thank everyone who has studied in Professor Miles Padgett's optics group for their contributions to the fun, friendly and interesting environment they have each contributed to.

Publications

- Alasdair C. Hamilton and Johannes Courtial. Imaging with parallel ray-rotation sheets. *Optics Express*, 16:20826–20833, 2008.
- Alasdair C. Hamilton and Johannes Courtial. Optical properties of a dove-prism sheet. *J. Opt. A: Pure Appl. Opt.*, 10:125302–125302, 2008.
- Alasdair C. Hamilton and Johannes Courtial. Metamaterials for light rays: ray optics without wave-optical analog in the ray-optics limit. *NJP*, 11:013042–013042, 2009.
- Bhuvanesh Sundar, Alasdair C. Hamilton, and Johannes Courtial. Fermats principle and the formal equivalence of local light-ray rotation and refraction at the interface between homogeneous media with a complex refractive index ratio. *Optics Lett.*, 34:374–376, 2009.
- Mark R. Dennis, Alasdair C. Hamilton, and Johannes Courtial. Superoscillation in speckle patterns. *Opt. Lett.*, 33:2976–2978, 2008.
- Alasdair C. Hamilton and Johannes Courtial. Generalized refraction using lenslet arrays. *J. Opt. A: Pure Appl. Opt.*, 11:065502, 2009.
- Michael Blair, Leo Clark, E. Alasdair Houston, Gary Smith, Jonathan Leach, Alasdair C. Hamilton, and Johannes Courtial. Experimental demonstration of a light-ray-direction-flipping METATOY based on confocal lenticular arrays. *Opt. Comm.*, 282:4299–4302, 2009.

Presentations

- **23$^{\text{rd}}$ September 07** A talk at the European Graduate College workshop, Büsum, Germany;
- **24$^{\text{th}}$ June 08** Poster presentation at the Scottish University Physics Alliance annual meeting, Glasgow University, Scotland;
- **26$^{\text{th}}$ August 08** Two talks and a poster presentation at Photon 08, Herriot Watt University, Scotland;
- **1$^{\text{st}}$ October 08** A talk and a poster were presented at the European Optical Society annual meeting, Paris, France;
- **9$^{\text{th}}$ June 09** A talk at the European Optical Society, Jena, Germany;
- **17$^{\text{th}}$ September 09** A talk at the International Conference on Electromagnetics in Advanced Applications, Torino, Italy.

Contents

CHAPTER 1

Introduction

"I set out to find one thing and found something better."

— *James Lovelock*

1.1 Metarefraction

Metarefraction consists of a local and pixelated ray-direction transformation of light rays as they pass through a metarefracting sheet. It is local in the sense that ray position is only slightly offset, transversely, during propagation through the sheet, and is pixelated in that most parts of the wavefront, locally, are well approximated by a single, well-defined light ray. For visual applications, offsets of a few hundreds of wavelengths should be barely visible. A metarefracting sheet then acts approximately like a homogeneous, light-ray-direction changing window (Figure 1.1(a)); the light-ray-direction change can be seen as a generalization of refraction at the interface,

$$\mathbf{r}' = \mathbf{T}(\mathbf{r}), \tag{1.1}$$

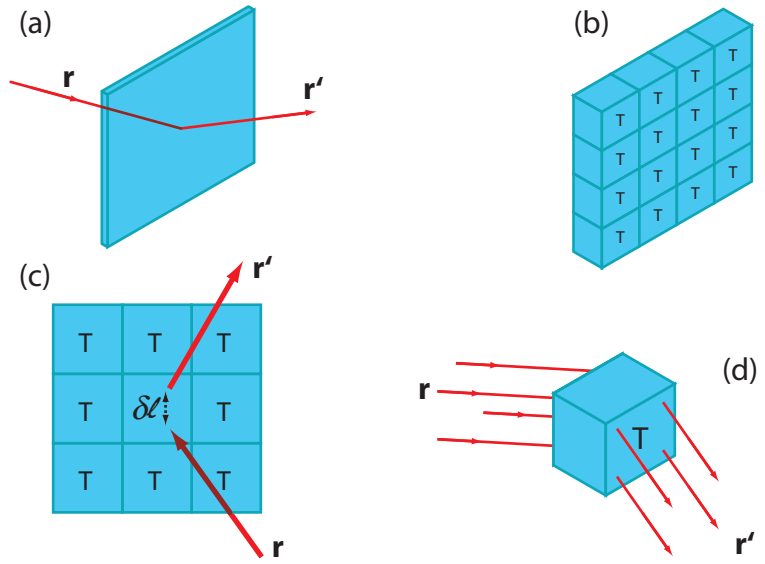

Figure 1.1: (a) A metarefracting sheet is analogous to a refracting interface which maps an incoming light-ray direction into an outgoing light-ray direction, T : $\mathbf{r} \rightarrow \mathbf{r}'$. (b) Local light-ray direction-transformation is defined in terms of the combined effect of constituent metarefracting components (each represented by a square containing the letter **T**) that each act separately on those incident light rays. (c) Now only a section of the sheet is viewed front on, the transformation may introduce only offsets, δl, that cannot be larger than the component size. (d) The ray direction transformations should also be telescopic, i.e. parallel rays incident on a component are all mapped into the same new direction.

where **T** is the law of metarefraction, **r** is the incident light-ray direction and $\mathbf{r}'$ is the transformed light-ray direction.

More precisely, there are three requisite properties that a metarefracting sheet must satisfy. Firstly, the sheet must be describable as a collection of independent optical components that separately transform the direction of incident light as is shown in Figure 1.1(b). Secondly, the components must not introduce a significant offset, δl, in the position of the ray as it propagates within the sheet. That is, lateral offsets must be smaller than of the size of the optical components as shown in Figure 1.1(c). This is so as to ensure

a good visual quality of the mapping for a sufficient density of these optical components. Finally, when mapping ray direction, every sheet component must preserve collimation of a pencil of rays. This is defined here as a *telescopic* ray mapping and is shown in Figure 1.1(d). This is in order to ensure that a component such as a ground glass plate would not satisfy the requirements as light rays will be inconsistently redirected.

Allowing non-zero offsets allows a far greater variety of metarefractions than would be possible were only zero-offset metarefractions considered. Ray optics is the limiting case of wave optics in the limit of small wavelength. That is, phase front curvature also tends to infinity, and so locally all waves act as plane waves. Light is then well described by rays that travel through the scene along curvilinear paths. However, optical wavelengths are finite, albeit very small, and so in practice each wave has a finite curvature and is subject to additional wave-optical restrictions (see Chapter 6). In order to avoid such strict requirements, a transformation need only be approximated as described above. Allowing offsets then frees metarefraction to approximate a far greater range of possible mappings than could be exactly realized while preserving the simplicity of its ray-optical description.

This chapter will discuss the visual distortions that metarefracting sheets may produce. After introducing existing examples of metarefraction, comments will be made regarding the format and notations used throughout this document.

1.2 Imaging and non-imaging visual distortions

There are two key attributes with which to describe the view through a metarefracting sheet. Firstly, all metarefracting sheets must realize a *visual distortion* i.e. when viewed through a pinhole, rays that reach an observer from a specific

direction must all originate from a small number of localized regions from the far side of the sheet. Where this is not true, rays from many locations blur together, giving rise to significant blurring across the sheet.

Secondly, some such sheets are also able to perform *imaging* in the sense that in addition to producing a visual distortion, they will continue to produce a sharp image when viewed through an extended aperture. This then requires the sheet to satisfy parallax, i.e. to relate, for any given point light source, the outgoing direction of light rays (from the sheet) to the point on the sheet through which they exit. Consistency of parallax will be used as a test of imaging later in this document.

The reason to distinguish between visual distortions and imaging is that, for visual applications, the imaging requirement may often be unnecessarily restrictive. This is of particular relevance in that the human eye has a relatively small aperture so subtle aberrations remain undetected. An example of a visual distortion in which, strictly, imaging does not take place is viewing a stick through an air-water interface. The same is true for several metarefracting sheets.

Additional care must be taken when using such sheets in more precise optical systems. Light rays that intersect at different points on the sheet are each transformed independently. Each metarefracting component may then affect polarization and coherence separately and where such rays intersect to form image points or at caustics, the resulting polarisation and coherence may be elaborate and potentially disruptive. Wave-optical effects will be ignored here, but may be optimized to better approximate what from now on will simply to be assumed to be bright regions.

When constructed, sheet behaviour will be complicated by both physical and engineering constraints. In order to transform light-ray direction, a

component will often employ fine structures such as a refracting surface or an aperture which will work imperfectly. As a result, most components can only approximate idealized metarefraction to some limiting angular resolution. Similarly, they will have a limited field of view over which they will transform telescopically.

The behaviour of such components will, when shrunk to a small enough scale, be diffraction limited. This of course applies only to non-trivial components that consist of some inhomogeneous part. A sufficiently homogeneous mapping may work perfectly which will neither redirect, nor appreciably offset, incident light rays. In contrast, however, the effect of miniaturising even trivial components will be to causes diffraction effects to become significant. When the size of components are of the order of magnitude where diffraction effects become significant, interaction between components through evanescent coupling [1] can also become important.

1.3 Existing examples of metarefraction

While the definition of metarefraction as a paradigm is a new concept, there are scattered examples of optical systems already in existence that satisfy the above requirements. A real world example of metarefraction is the refraction performed by a Fresnel lens.

Geometrical imaging may be understood in terms of Fermat's principle [3] which states that the optical path length between the object and image point is stationary with respect to perturbations of the ray path. A consequence of this is that the optical path length along each route that each ray follows through a lens is the same. The lens is thicker at the centre to increase the optical path length there. This is in order to compensate for the physically shorter path at the centre of the lens. This then satisfies the criteria for

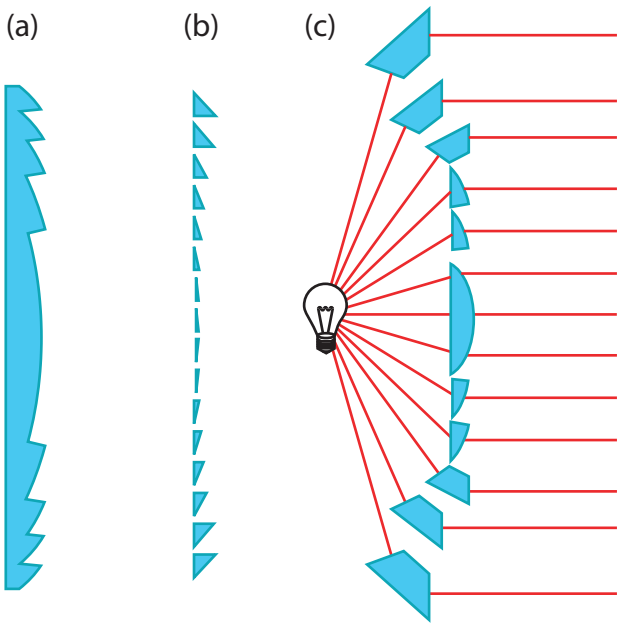

Figure 1.2: (Looking from the side). An example of a metarefracting sheet is a (a) Fresnel lens which consists of (b) a series of concentric annuli with triangular cross sections, such that locally they each act as a prism, i.e. each of which acts as a metarefracting component. This is an example of an inhomogeneous sheet in that the parameter that describe the constituent prisms, the prism angle, changes across the sheet. A consequence of this inhomogeneity is that the sheet performs imaging. (c) The Fresnel lens has long been generalized by adding Dove-prisms at the edges of the sheet [2]. The resulting catadioptric optical system offers a larger field of view and as a result has been widely used in lighthouses.

constructive interference which results in a bright spot. A Fresnel lens builds on the original lens design by taking the original optical path profile modulo the sheet thickness so as to produce a piecewise continuous wavefront.

The behaviour of a Fresnel lens may alternatively be reinterpreted as a metarefraction. A Fresnel lens, shown in Figure 1.2(a), may be constructed from a collection of small prisms, Figure 1.2(b). In this sense, a Fresnel lens satisfies the required criteria of locality of transformation, insignificance in

offset and the telescopic property.

This is a complicated example of a metarefraction in that the transformation is inhomogeneous, i.e. the law of refraction varies with position. This is however not the most complicated example of a metarefraction. It can, for example be further extended in its field of view through the incorporation of Dove-prisms at the edges of the sheet as is shown in Figure 1.2(c). Such systems have found use in lighthouses in order to collimate as much light as possible. In contrast to the Fresnel lens, the transformations that will be introduced in the following chapters are all assumed to be homogeneous.

1.4 Notation and thesis structure

A consistent notation will, where possible, be used throughout this document. Furthermore abstract index notation will be interspersed with explicit vector notation where appropriate [4].

Ray directions in two and three dimensions will be written in a number of coordinate systems, but the most useful will prove to be as the two dimensional angle mapping, $\alpha \mapsto \alpha'$. In three dimensions, metarefractions will be described in terms of Cartesian coordinates

$$(r_x, r_y, r_z) \mapsto (r'_x, r'_y, r'_z). \tag{1.2}$$

where the sheet is orientated perpendicularly to the z axis. Additionally, use will also be made of (unit length) spherical polar coordinates

$$(r_\theta, r_\phi) \mapsto (r'_\theta, r'_\phi) \tag{1.3}$$

which consist of polar and azimuthal angles respectively. Both Cartesian and spherical polar coordinates will be interspersed where appropriate.

This thesis will be divided into several parts. Firstly, an overview was provided above which defined and introduced metarefraction. In the next

chapter, background material will be reviewed that explores some of the inspiration of metarefraction. Perhaps surprisingly, an inherently ray-optical transformation possesses a number of connections with a wave-optical subject called negative refraction. This and related subjects will be explored.

There will then be a pair of example transformations that demonstrate the effect of a metarefracting sheet. These include the Dove-prism sheet in Chapter 3 and the ray-rotation sheet in Chapter 4. These simple metarefractions will demonstrate some effects that may be realized, and that a ray-direction transformation must be recast in terms of the apparent position of a point light source in order to understand the visual distortion. Simulated views of these examples will also be presented. Several examples of metarefractions outlined in the following chapters have been demonstrated with rendering ray-tracing software, POV-Ray [5], which enables a full ray-optical simulation through a collection of prisms and lenses.

After presenting example metarefractions, conceptually and practically interesting properties will be discussed. These will focus on the use of metarefracting sheets to perform imaging; on the ability to (approximately) realize wave-optically-forbidden light fields, and hence to recognize the existence of a new and unexplored kind of transformation. Also, there is a surprisingly natural representation of one such transformation in terms of an extended form of Snell's law and which, despite the lack of wave-optical analogue, may be derived from Fermat's principle — a wave-optical concept!

After finishing discussing the first two kinds of metarefracting components, another transformation will be explored in Chapter 7: metarefraction by confocal lenslet arrays. In addition to reproducing and generalizing previous ray transformations, the commercial availability of components required to construct a confocal lenslet array means that a new form of metarefraction may

be experimentally demonstrated as is outlined in Chapter 8.

The penultimate chapter will explore the combination of confocal lenslet arrays into the most general possible metarefracting sheet which can map ray direction without an offset. This document will then conclude with a summary chapter which will provide an overview of this work as well as a discussion of possible future generalization of metarefraction.

CHAPTER 2

Background

> "Any sufficiently advanced technology is indistinguishable from magic."
>
> — *Arthur C. Clarke*

2.1 Overview

Before expanding on my own work, the work of others that has influenced metarefraction will be outlined. In particular this chapter will explore negative refraction which will first be discussed in a historical context. Then, the manner by which it may be realized will be explored and finally, one of its applications will be discussed.

It is through approximate models of light that scientists and engineers are able to concisely imagine and express the myriad of optical effects and technologies that are available today. The most familiar theories, the ray- and wave-optical models, have proven enormously useful. These models are however simplifications and so inherently skew our view of nature.

That nature behaves very similarly to such models can, through a lack of counterexamples, lead to falsely assumed restrictions. An example of which is the belief that refractive index could only be positive. However, this was not immediately clear from the defining expressions, and one must look very hard to find examples of negative refraction.

As well as by exploring the subtleties and extremities of existing theories, progress may be made through the reinterpretation of existing concepts in terms of new paradigms. One example of which is transformation optics which utilizes the equivalence between spatial curvature and medium-property transformation to control the propagation of light through an optical system. Simply having changed the way in which refractive index is defined, negative refractive index appears quite naturally.

The structure of this chapter will be as follows. In the next section, negative refraction will be explored. Then, transformation optics will be explored, particularly in terms of a simple special case. After this, metamaterials will be introduced which will give a concrete mechanism by which negative refraction may be realized and transformation optics implemented. The penultimate section will discuss the 'perfect lens' which uses negative refraction to perform a form of imaging that will prove relevant in a number of later chapters. This chapter will conclude with a brief summary, discussing the merits of metamaterials.

2.2 Negative refraction

The study of negative refractive index began in 1967 when Veselago explored what the implications were of a material having a negative refractive index [6]. Refractive index [7] represents the ratio between the speed of light in a vacuum

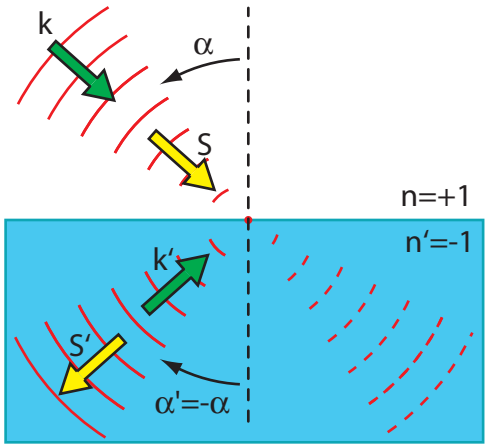

Figure 2.1: Within a negative refractive index medium, a wave will propagate with a negative phase velocity. This means the phase fronts and hence the wave vectors travel backwards while the group velocity, energy and information that the wave carries continue to travel forwards. At the interface, Snell's law implies that the refracted ray is reflected, but that the ray direction is reflected along the sheet plane.

and the phase velocity of light in a particular medium,

$$n = \frac{c}{v_p}, \tag{2.1}$$

and its sign reflects the fact that the velocity of light in a vacuum, c, and in the medium, v_p, may be in the same, or in opposite directions. So, intuitively, negative refraction corresponds to a negative phase velocity, and hence to backwards propagating phase fronts.

Refractive index may also be defined in terms of the medium with which it is associated:

$$\boldsymbol{n} = \pm\sqrt{\boldsymbol{\epsilon}_r \boldsymbol{\mu}_r}. \tag{2.2}$$

Here relative permittivity is $\boldsymbol{\epsilon}_r$ and the relative permeability is $\boldsymbol{\mu}_r$ respec-

tively [8]. The key realization that Veselago had was that, although in conventional materials the sign of the square root is positive, this need not always be true. The sign of the refractive index follows from that of permeability and permittivity.

These medium properties in turn follow from the constitutive equations [9] which are

$$\mathbf{D} = \epsilon_0 \boldsymbol{\epsilon}_r \mathbf{E}, \tag{2.3a}$$

$$\mathbf{B} = \mu_0 \boldsymbol{\mu}_r \mathbf{H}. \tag{2.3b}$$

These equations relate resultant fields (the electric displacement field, $\mathbf{D}$, and the magnetic flux density, $\mathbf{B}$) to their respective causative fields (i.e. the electric field, $\mathbf{E}$, and the magnetic field, $\mathbf{H}$). Although permittivity and permeability can be scalars for some materials, they are in general rank two tensor quantities.

For most materials, both permittivity and permeability are positive. Permittivity of metals can be negative above their plasma frequency and unusual materials such as magnesium fluoride have a negative permeability. In contrast when only one of the pair is negative, the material is opaque. Under these conditions (i.e. when $\epsilon_r \mu_r < 0$), Maxwell's equations no longer have a travelling wave solution. Now working with scalar permittivity and permeability, this may be demonstrated by combining Equation (2.1) and Equation (2.2) to get

$$v_p = \frac{c}{\sqrt{\epsilon_r \mu_r}} \tag{2.4}$$

and on taking the square root of a negative, the phase velocity is imaginary.

Veselago could find no material in which both properties were negative. Despite the lack of materials demonstrating negative refraction, Veselago calculated a number of interesting properties, for example modified Čerenkov

radiation and the Doppler effect, for what was then only a hypothetical material.

In negative refractive index materials, the sign of both permittivity and permeability are negative. In this case, the phasefronts do travel, but do so backwards, in the opposite direction from the group velocity, against the progressing wavefront. This may be seen from the relationship between the electromagnetic field and the Poynting vector,

$$\mathbf{S} = \mathbf{E} \times \mathbf{H}. \tag{2.5}$$

Conventionally, materials are 'right handed' in the sense that the above relation obeys the right hand rule. In the case of negative permeability and permittivity, the Poynting vector changes sign. This follows from the constitutive relation, Equation (2.3b). The equation, and hence the material, is 'left handed'. As a result of this, for the same direction of energy flow, the causative electromagnetic field is oppositely orientated from normal as shown in Figure 2.1.

2.3 Transformation optics

Transformation optics is a recent merging of medium and coordinate-system concepts. Initially it was conceived of in order to generalize (finite-difference time-domain) electromagnetic field simulations [12] to arbitrary coordinate system grids [13, 14, 15]. Subsequently, it has been applied as a very general optical design technique that naturally incorporates smoothly varying positive and negative refractive indices.

It is natural to talk about a visual distortion as being 'stretched' or 'twisted', using language that describes spatial distortions. This is in con-

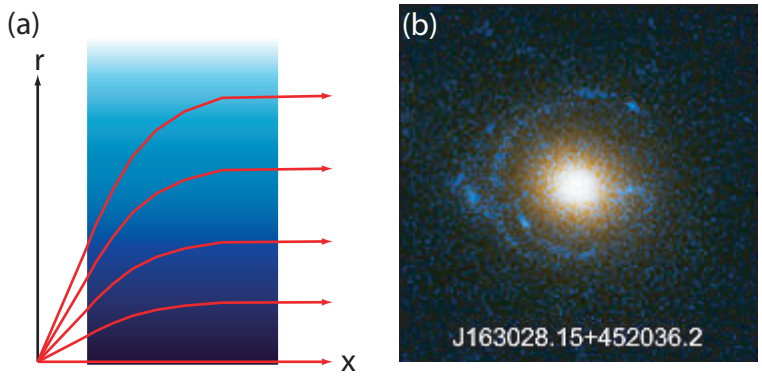

Figure 2.2: Light rays can curve as they propagate in space as a result of varying medium properties or from spatial curvature. An example of a varying medium is (a) a graded refractive index lens [10] (shown here as viewed in cross section). Such a lens can collimate light using a parabolic radial refractive index gradient. (b) Alternatively, curved space may bend light rays. For example, the mass of an intermediate galaxy (yellow) distorts the light from a more distant galaxy (blue) into Einstein rings (image credit: A. Bolton [11]).

trast to optical design which shapes the optical path of light by controlling medium properties.

Physically, there are two reasons why the path of a light ray will bend as it propagates through space. Either the surrounding medium properties, permittivity or permeability, vary in space. This is described by Maxwell's equations [16], an example in which they vary smoothly is in the case of a graded refractive index lens which is shown in Figure 2.2(a)).

Alternatively, the space that the ray propagates through may bend [17], for example in the case of gravitational lensing [18]. Here, the curvature of space changes as a result of an intermediate mass such as a galaxy (Figure 2.2(b)). From general relativity, space may be curved by the presence of mass or energy. Different spatial curvatures correspond to different ray paths through the save volume. Transformation optics uses the mathematical equivalence of medium properties and spatial curvature to change the language in which curved opti-

cal paths may be described. Through this analogy, the visual distortion that would result from a curved space may be realized by the equivalent material properties.

From the work of Pendry [13, 14] and Shyroki [15], the two interpretations of optical path may be interchanged. They allowed permeability and permittivity to be written as

$$\epsilon^{\gamma\delta} = \epsilon g^{\gamma\delta} \sqrt{\det g_{\alpha\beta}}, \tag{2.6a}$$

$$\mu^{\gamma\delta} = \frac{\mu g^{\gamma\delta}}{\sqrt{\det g_{\alpha\beta}}}. \tag{2.6b}$$

That is, they may be written as functions of an equivalent metric tensor, $g_{\alpha\beta}$, of the equivalent curved space. For a fuller and accessible treatment of the subject, consult Leonhardt and Philbin [19]. This reinterpretation of material properties in terms of a new coordinate system has subsequently provided a new paradygm with which to design optical systems.

A wide range of applications of transformation optics exist [20]. These vary from optical novelties to effective extensions of fundamental theories [21] and they even extend to acoustics [22, 23]. Also, objects may be hidden by shrinking their visual distortion to a negligible size as has been demonstrated in ray tracing simulations [24, 25] and experimentally at microwave frequencies [26]. Also, apparent topology changes may be realized by hiding a connecting tunnel of space between two parts of the system. The resulting visual distortion then suggests that the ends of the tunnel are unconnected [21].

Transformation optics is formulated without specifying the method by which a particular permeability and permittivity may be realized. As a result, it is independent of the source of the medium properties. While existing optical fabrication techniques allow fine control over refractive index, for example in Reference [27], conventional materials are limited to positive refractive in-

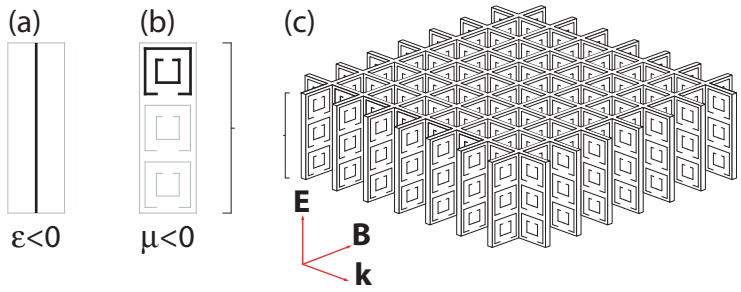

Figure 2.3: Metamaterials achieve their optical properties through microscopic structure and not atomic structure. (a) By embedding wires within a region smaller than the wavelength of incident light, its permittivity will be negative across a frequency range. (b) More elaborately, a magnetic 'meta-atom' may be built which has a negative permeability as realized by a split ring resonator. (c) By repeating both structures within a volume, the entire volume may be described by these effective medium properties.

dices. A wave-optical method by which to implement negative refraction and the smoothly varying positive and negative refractive index profiles produced by transformational optics is metamaterials. These will be discussed in the following section.

2.4 Metamaterials

Veselago's exploration of negative refraction ended when he could not find any materials with simultaneously negative permittivity and permeability. However, what could not be found in nature was later fabricated.

Metamaterials gain their material properties through structure and not substance. This relies on subwavelength structures to interact with the electric and magnetic fields. The first such structure was originally used to perform plasma simulations [28] and consisted of metal rods that give a negative frequency response to an electric field (as is shown in Figure 2.3(a)). This is

an example of a material in which bulk permittivity was entirely different to that of its constituent parts.

Another structure was required to also artificially alter the permeability. This was complicated by the relatively small magnetic response of natural materials. Progress was made when researchers sought to replace atoms — which are in effect circulating currents that give rise to a magnetic dipole moment — with a much better source, the split-ring resonator. As shown in Figure 2.3(b), split-ring resonators each consist of a tiny circuit that produces a significant magnetic moment when at resonance with the external field. At microwave frequencies metamaterials have been experimentally demonstrated [29]. However, progress in the fabrication of optical equivalents has been slower due to the lack of equivalent, pre-existing fabrication techniques at these smaller size scales.

These two ideas have since been combined into a single, effective medium in which both permittivity and permeability are simultaneously negative (see Figure 2.3(c)). A metamaterial is constructed by repeating such subwavelength components throughout its volume. The key property is that, due to their size, each circuit interacts not with light as a whole, but with the electric and magnetic fields that constitute it. Therefore, as long as each circuit is smaller than half a wavelength (from the Nyquist limit [30]), each meta-atom in a material may separately interact with the local electric and magnetic fields.

2.5 The perfect lens

Shown in Figure 2.4, one truly remarkable application of metamaterials is the perfect lens, which was first suggested by Pendry [31]. This is able to image finer details than that of an ordinary lens. This is because high spatial

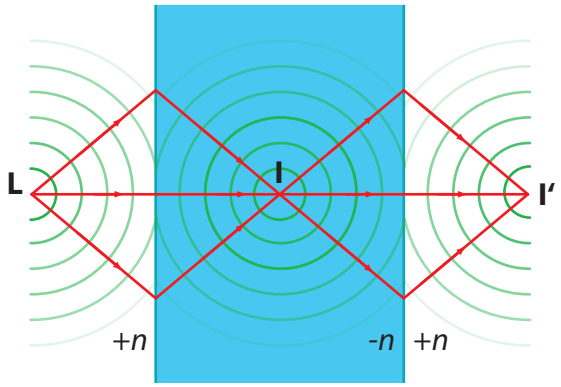

Figure 2.4: Shown in blue and viewed from the side, a 'perfect lens' consists of a block with an equal magnitude, but opposite sign of refractive index, $-n$, to that of its surroundings, n. The lens images a point light source, **L**, twice: first at **I** inside the block and then at **I′**. Rays (shown in red) are reflected at each interface such that diverging rays converge. In addition to travelling waves, evanescent components — which carry high spatial frequency information — also propagate. These decay exponentially in positive index media so are normally lost. However, in negative media they are amplified and so the near field information may be recovered by the perfect lens.

frequency information is transported by evanescent waves which ordinarily decay exponentially as they propagate until the 'near field' information is lost [1]. Metamaterials are able to amplify these waves and so preserve the information that they contain.

The amplification of evanescent components by a negative refractive index medium is a consequence of waves having a negative optical path length. The phase fronts literally propagate backwards and decay as they do so. Equivalently, in the direction of information flow, their amplitudes grow exponentially.

By ensuring that each wave propagates through the same path length of negative index medium as it does through positive space on either side of the lens, the total optical path length is zero and the evanescent components

should regain their original amplitudes. Furthermore, from Fermat's principle [3] the optical path is stationary. Alternatively, in terms of Huygens' principle [32], wavefronts propagate through the system through equal path lengths (all of length zero). They will then combine coherently to interfere constructively at the image.

A variant of this has been demonstrated using a curved lens in a nanoscale experiment by coating an object in layers of silver which has a negative permeability [33]. Such a setup has obvious applications in subwavelength microscopy.

Ray-optical interpretation

It is more relevant to the following chapters to discuss the perfect lens within the context of ray optics. Each interface between positive and negative index media separately images point light sources into the other medium. Ray optically, this is a consequence of the reflection of the components of ray direction that lie parallel to the interface. Snell's law continues to work, and implies that the angle of refraction is the negative of the incident angle.

This transverse ray-direction reflection follows from the refraction of rays at the interface between positive and negative media of equal magnitude. This is shown wave optically in Figure 2.1. The figure shows wavefronts (solid lines) travelling to one side of the refracting interface and rather than continue in the same direction (dashed lines), the component of the ray direction that is parallel to the interface is reflected backwards (again shown as solid lines). Then, from Snell's law the refracted angle is the negative of the incident angle.

Diverging rays then converge after each interface to form an image. From equality in magnitude between incident and refracted angle, before and after refraction, the rates of divergence and convergence are equal. As a conse-

quence, the object and image distances are equal to one another. A consequence of this is that the depth of further away objects correspond to closer image distances. Depth has therefore been reversed. This is known as pseudoscopic imaging [34].

The behaviour of the perfect lens is then to combine two such interfaces (on either side of the negative refractive index block) to pseudoscopically imaging the scene twice (Figure 2.4). The advantage of this is that neither the object plane nor the image plane need be immersed in a negative index medium and that the overall optical path length may be set to zero for the correct lens thickness, returning evanescent components to their original magnitude.

2.6 Summary

Clearly metamaterial volumes and metarefracting sheets operate on a different length scale, to a different optical approximation and are designed quite differently. However, both gain optical properties from structure and not substance. They both transcend the conventional limitations placed upon them, be it the effective medium properties or the angle of refraction. And as will be shown in the following chapters, both may realize negative refraction and pseudoscopic imaging.

CHAPTER 3

The Dove-prism sheet

> "It can scarcely be denied that the supreme goal of all theory is to make the irreducible basic elements as simple and as few as possible without having to surrender the adequate representation of a single datum of experience."
>
> — *Albert Einstein*

3.1 Introduction

Having already briefly introduced the motivation behind metarefraction, and placed it within the context of existing work of other researchers, the purpose of this chapter is to explore what is perhaps the simplest example of metarefraction — local light-ray direction reflection along one axis. That is, a single component of the direction of each light ray (taken here to be the x-direction) is reflected about the point at which the ray intersects the sheet:

$$r'_x = -r_x \tag{3.1a}$$

$$r'_y = +r_y \tag{3.1b}$$

$$r'_z = +r_z \tag{3.1c}$$

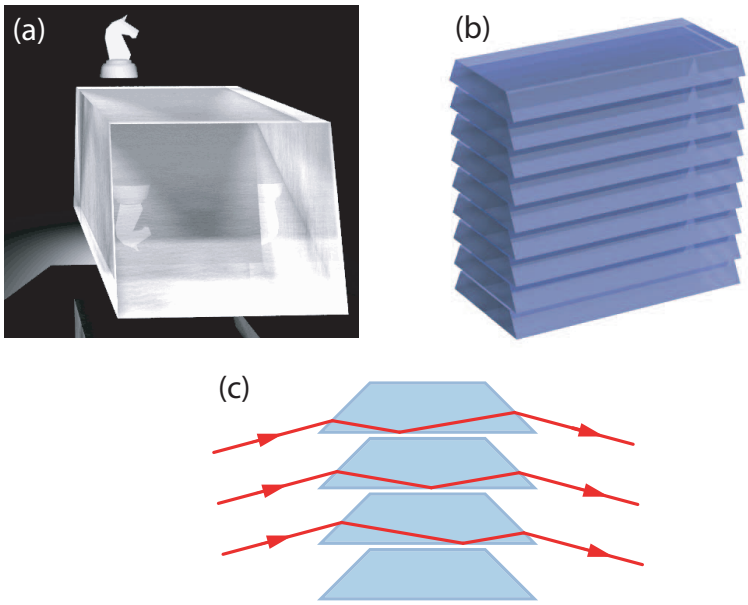

Figure 3.1: (a) In order to perform ray-direction flipping, a Dove prism may be used. In the case of a single prism, each reflection takes place about some common centre rather than the point at which each ray intersects the sheet. (b) In order to avoid this, many short but wide-apertured prisms may be stacked to produce a "Dove-prism sheet". (c) On tracing a series of light rays through the sheet, their ray direction may be seen to be reflected about one axis with only minor position offset.

where the sheet normal will be taken to lie in the z direction. While reflection may be realized by a number of well known components, what is different about this transformation is that the plane of reflection lies on the point where the ray and sheet intersect.

This chapter will explore this ray-flipping metarefraction from a number of perspectives. Firstly, by what mechanism may ray flipping be approximated? Additionally, the visual qualities are of interest: What does the visual distortion through the sheet look like? Can a simple expression be derived to describe it and if so what form does it take? And lastly, why is such a mapping interesting?

How may ray-direction reflection be realized? A simple reflection of ray direction may be realized by a single Dove prism. As an example, rays observed propagating upwards (Figure 3.1(a)) were initially downwards propagating and vice versa. However, the sought transformation should perform the reflection without such a large offset.

The offset cannot be eliminated, but it can be reduced by several orders of magnitude (i.e. below the angular resolution of the observer who's field of view includes the entire sheet). In order to do so, multiple Dove prisms are stacked as shown in Figure 3.1(b). Choosing a stack of short and wide-apertured prisms of sufficient proximity to one another reduces the offset size. Offsets are of the order of the height of each prism, so by reducing the prism height by several orders of magnitude, the offset is similarly reduced (see Figure 3.1(c)).

This structure will be referred to as a Dove-prism sheet, and was first proposed in a paper by Lian and Chang [35] for beam steering applications before being patented by Watkins [36]. Vastly increasing the number of Dove-prisms, Nelson and Courtial realized that such a structure could produce a sharp visual distortion [37]. Later, I was involved in the investigation of the optical properties of such a structure with Dr Courtial [38].

When viewed from the side, a Dove prism is a trapezium, but transverse to this plane, the Dove prism is homogeneous and so has no net effect on that component of the light ray direction. A Dove prism therefore only transforms the light ray direction two dimensionally. When projecting the scene onto this plane, this projection will be referred to as the 'prism projection'. It is on this plane that the relevant behaviour of a Dove-prism sheet may be captured.

A Dove-prism sheet does not perform imaging because it only reflects one transverse component of a light ray. Diverging rays are made to reconverge after passing through the sheet only along that axis. In the orthogonal direction,

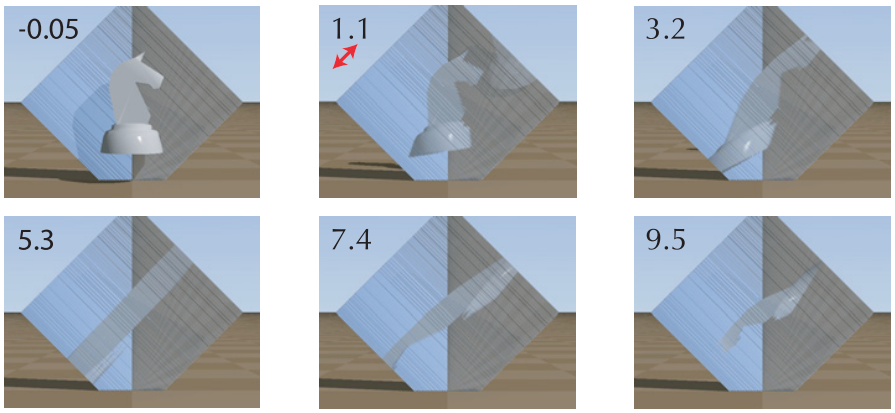

Figure 3.2: Chess piece seen through a Dove-prism sheet for various distances of the piece behind the sheet. The reflection direction is marked as a red arrow and is 45° to the horizontal. The chess piece's distance behind the sheet is given in the corner of each frame (in units of the side length of the floor tiles). The sheet is a distance $z_2 = 5.5$ in front of the camera and contains 100 Dove prisms.

normal to the prism projection, the ray continues to diverge.

The lack of imaging does not mean that the view through a Dove-prism sheet is blurry. So long as it is viewed through a small aperture, the resulting visual distortion will appear sharp because each prism 'looks' in a slightly different direction to that of its neighbour with a small field of view. This sharpness of the visual distortion requires each prism to be smaller than the angular resolution of the observer.

While reduction of prism size is desirable to achieve good visual quality, a compromise must be reached to avoid diffraction. Additionally, prisms must be sufficiently separated to avoid coupling of evanescent waves between neighbouring prisms, and must be significantly larger in height than this gap.

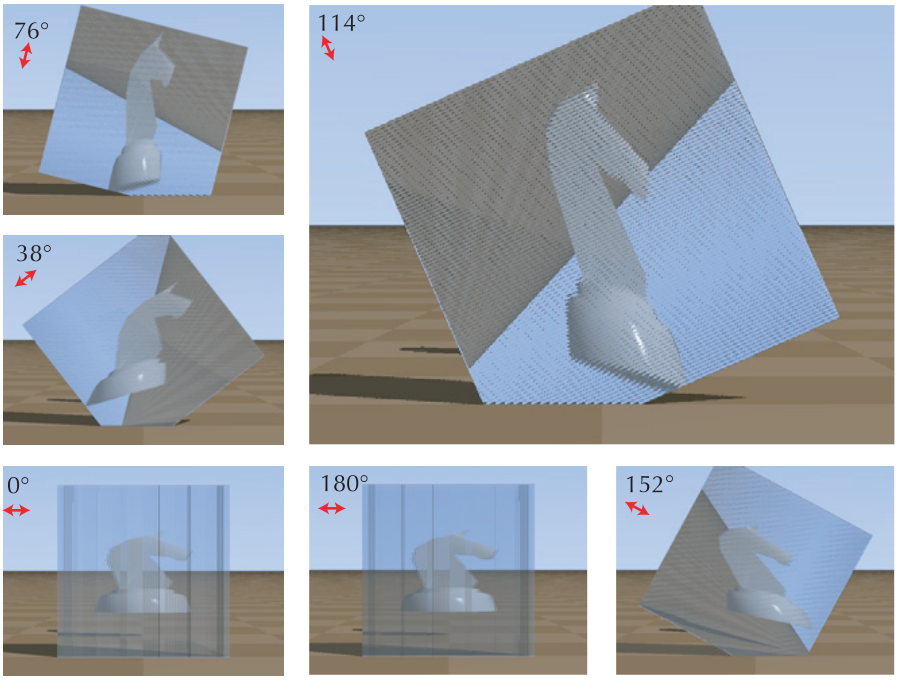

Figure 3.3: A chess knight seen through a Dove-prism sheet for various ray-inversion directions. The arrows in the top left corner of each frame indicate the direction along which the ray direction is inverted in the different frames; its angle with respect to the horizontal is also stated. The chess piece is positioned a distance $z_1 = 2$ floor tiles behind the Dove-prism sheet, which is a distance $z_2 = 5.5$ floor tiles in behind of the camera.

3.2 Visual qualities

While such a structured sheet may be concisely expressed through its law of metarefraction, it is more directly useful to instead consider the resulting visual distortion. That is, given the location of the eye and point light source, where on the sheet does the transformed light that intersects the eye appear to come from? As might be expected, the resulting view does not bear an obvious connection to a reflection. What then does the view through a Dove-prism sheet look like?

There are two particularly informative perspectives from which to view

the transformed scene. In order to explore such optical properties, sheets may be simulated by the ray-tracing software POV-Ray [5]. As is shown in Figure 3.2, the view through a Dove-prism sheet distorts the horizontal coordinate of each object. This in effect skews the object. The key length scale is the depth at which the apparent position of objects appear to diverge towards infinity from the perspective of the current observer. This occurs when the object-sheet separation is equal to that of the camera-sheet separation. Another parameter of interest is the orientation of the flip axis, as is shown in Figure 3.3. This extends the law of refraction from flipping along the x axis to an arbitrary direction on the x-y plane. As the sheet is rotated through 180°, the axis along which the object appears to stretch changes and the horizon spins through 360°.

Ray-direction flipping, involving the transformation of only a single ray-direction component, and transforms the ray only along the prism projection, nonetheless takes place in three dimensions. Despite this, the problem can easily be reduced in dimension by noting that while each ray tranvels further, the distance that it travels along the z axis remains unchanged. The interesting behaviour is then entirely captured by projecting the rays onto the x-y plane.

When looking through the sheets, an observer sees the location of point light sources according to some transformation. However, what the observer sees is not intuitively connected to the underlying metarefraction transformation. Specifically, the apparent ray direction may be described by an implicit equation that expresses the apparent point light source direction (i.e. the ray-sheet intersection point) in terms of the initial ray direction and the metarefraction. This implicit expression must be solved.

An expression for the apparent location of a point light source will now be

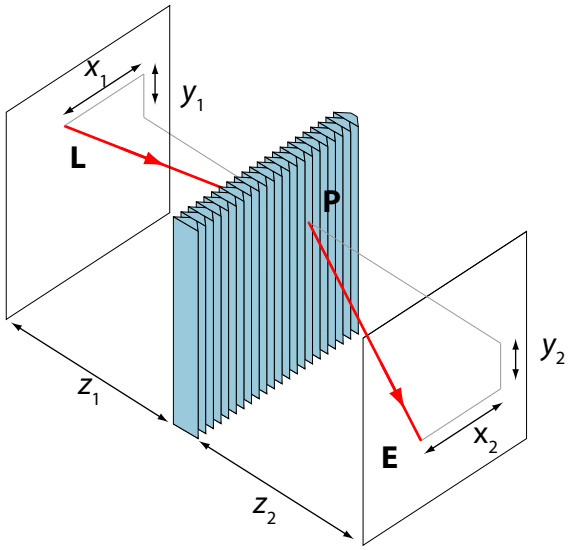

Figure 3.4: A side view of a ray path through a Dove-prism sheet from a point light source, **L**, to an eye point, **E**, through a point, **P** on the sheet. The first light ray travels a displacement (x_1, y_1, z_1) and the mapped, second ray travels a displacement, (x_2, y_2, z_2).

presented. Figure 3.4 shows the location of a point, **P**, on a Dove-prism sheet where a light ray intersects as the ray propagates from the point light source, **L**, to the observer, **E**. As before, the coordinate system is chosen such that the sheet is orientated perpendicularly to the z-axis, and the ray-direction reflection is along the x-axis. The horizontal ray-direction flip may therefore be expressed as

$$\frac{x_1}{z_1} = -\frac{x_2}{z_2} \tag{3.2}$$

where the initial ray vector is (x_1, y_1, z_1) and the metarefracted ray vector is (x_2, y_2, z_2). Meanwhile, the vertical ray direction is unaffected, so

$$\frac{y_1}{z_1} = \frac{y_2}{z_2}. \tag{3.3}$$

The (known) displacement between light source and observer, (x_{LE}, y_{LE}, z_{LE}),

is the sum of the two rays. In the x direction this is

$$x_{LE} = x_1 + x_2 \tag{3.4}$$

and the y direction is

$$y_{LE} = y_1 + y_2. \tag{3.5}$$

Eliminating the transformed light ray vector, it is now possible to solve for the initial light ray

$$(x_1, y_1) = \left(\frac{x_{LE} z_1}{z_1 - z_2}, \frac{y_{LE} z_1}{z_1 + z_2} \right). \tag{3.6}$$

As stated above, only the x and y coordinates are sought. The z-component is fixed as the depth between the point light source and the sheet, z_1, is known. So too is the distance from the sheet to the camera, z_2.

This equation can be more concisely written in dimensionless coordinates, which are defined as

$$\xi = \frac{x_1}{x_{LE}}, \tag{3.7a}$$

$$\psi = \frac{y_1}{y_{LE}}, \tag{3.7b}$$

$$\zeta = \frac{z_1}{z_2}. \tag{3.7c}$$

On substituting into Equation (3.6)

$$(\xi, \psi) = \left(\frac{\zeta}{\zeta - 1}, \frac{\zeta}{\zeta + 1} \right). \tag{3.8}$$

Now consider the quantitative behaviour of the dimensionless equation. As the depth ratio, ζ, is increased from zero the actual point and apparent point start in the same place, but move apart. As the depth ratio tends to unity, the x-coordinate of the apparent point tends to infinity. This takes place when the camera is as far in front of the sheet as the light source is behind. Camera-sheet distance provides, as was observed in Figure 3.2, the

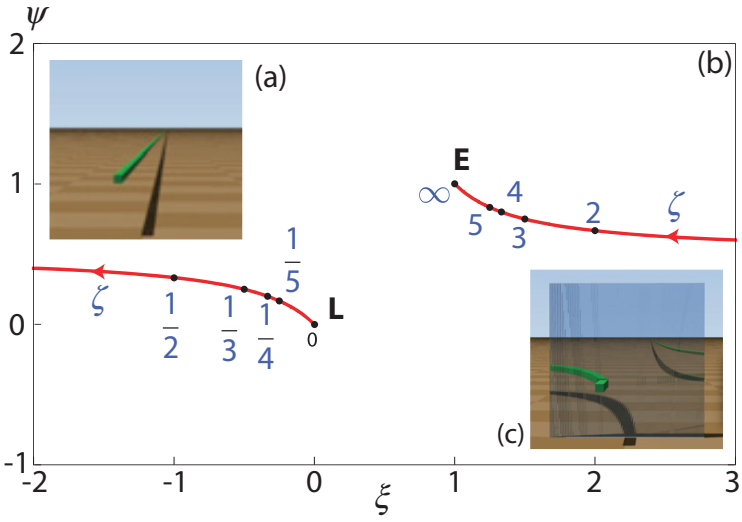

Figure 3.5: (a) The simulated view shows a straight rod, parallel to the viewing direction of a pinhole camera and extending towards infinity. A Dove-prism sheet orientated to flip the horizontal component of the light ray direction may be described by Equation (3.7a). (b) The visual distortion of the rod should therefore display a hyperbola as plotted in the graph. This hyperbola tends to the object and eye x-y projection as the object-sheet separation tends to zero and infinity respectively. The rod is at $(0, 0)$ and the observer is at $(1, 1)$. (c) The same scene as in Figure (a) has been rendered, but now it includes a Dove-prism sheet which is five and a half floor units in front of the camera. The visual distortion of the rod seen through the sheet reproduces the quantitative result shown in the parametric plot.

natural length scale of the system. As the object moves further from the sheet, its appears moves towards the projection of the eye on the sheet.

By eliminating ζ, Equation (3.8) may alternatively be written as

$$2\xi\psi - \xi - \psi = 0. \tag{3.9}$$

This equation implies that a straight rod that is normal to the sheet will appear as a hyperbola when viewed through a Dove-prism sheet. This is confirmed in the simulation shown in Figure 3.5.

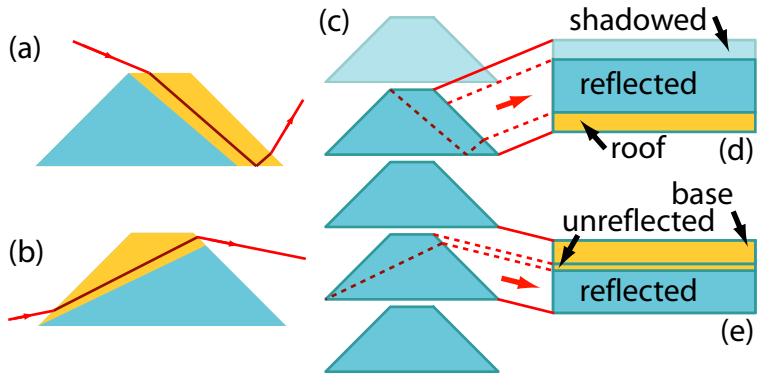

Figure 3.6: (a) In addition to an inverted view, when an observer looks from above at the bottom of the front face of a Dove prism, an observer will see the view through the roof of the prism. (b) Alternatively, when looking from below, an observer will see an unreflected scene. (c) When stacked to form a Dove-prism sheet, each prism displays a series of strips of the visual distortion and any 'parasitic images' that are visible at that viewing angle. Across a single Dove prism, the fractional area of each view may be found by projecting collimated light rays onto a plane. As the strips, like the prisms, are periodic, the same relative brightness for a single prism corresponds to that of a region of the sheet. The relative brightnesses of each varies across the sheet while the types of errors depend on whether the sheet is viewed from (d) above or (e) below.

3.3 Dove-prism sheet quality

There are a number of ways in which the quality of a Dove-prism sheet is limited. These include a limited field of view, 'pixelation' of the visual distortion, a decreased transmittance at larger viewing angles and the presence of 'parasitic images'. Here parasitic images do not imply imaging but rather alternative, undesired visual distortions due to alternative light-ray paths through the sheet.

For low densities of Dove prisms, the view will appear 'pixelated' or 'noisy'. This is because the field of view of each prism is sufficiently large for an appreciable portion of the scene to be visible through it. Each strip of the

scene is individually reflected, and the border between neighbouring strips then appears as a sharp discontinuity.

The presence of dimming and parasitic images are due to errors over portions of the front face of each prism. The view through an individual Dove prism is divided into several strips, each of which displays either a visual distortion or a parasitic image. If, when looking through a Dove-prism sheet, individual prisms are thinner than the observer's minimum resolvable distance, strips of the view blur together across that local region. The resulting visual distortion superimposes the ray-flipping visual distortion with the various parasitic images with angle-specific relative brightnesses.

Each strip that appears across the face of a Dove prism arises due to a particular route that rays may take through the prism. When viewing a single Dove prism from above, the bottom of its front face displays the view through the roof of the prism (Figure 3.6(a)). Alternatively, when viewing the front face from below, another parasitic image, an unreflected view of the scene, appears in the top corner (Figure 3.6(b)).

Variations in the brightness across the Dove-prism sheet are due to each strip of the visual distortion or parasitic image appearing over a different area of the front aperture of each prism. The visual distortion is brightest at the centre, while parasitic images may be seen through a significant fraction of the aperture at larger angles.

On combining prisms to form a Dove-prism sheet, there are additional scattering mechanisms due to interaction between neighbouring prisms (Figure 3.6(c)). Each prism displays several strips which include the visual distortion and the various parasitic images. The sheet then displays a series of alternatively coloured stripes which, locally, are approximately periodic. For a particular visual distortion or parasitic image, the resulting fraction of

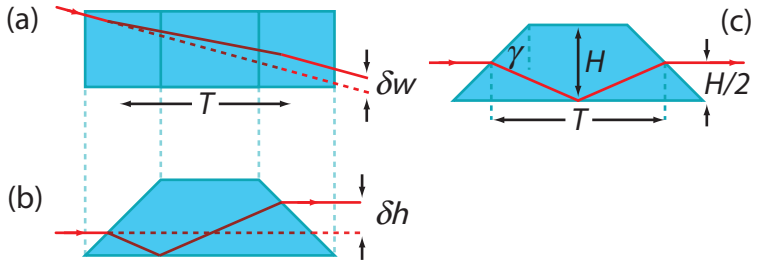

Figure 3.7: The reflection of a light ray as viewed through the top and side of a Dove prism. While propagating internally, the ray is offset in width and height. (a) The offset in width, δw, may be minimized by reducing the thickness of the prism, T, (defined here as the prism thickness at half its height). (b) The vertical position of the ray is reflected on passing through the prism. The offset, δh, is twice the distance between the ray height and the axis of reflection. (c) This is minimal when the axis of reflection is at the centre of the prism.

transmitted light is therefore proportional to the fractional area of its corresponding strip. Projecting all strips from the local area onto a plane, and selecting a periodic section of the resulting pattern, the fractional area of the visual distortion or any of the parasitic images may be determined.

Because front faces of each prism are inclined, their projections overlap when viewed from above (Figure 3.6(d)). Prisms can therefore occlude part of the aperture of their lower neighbour. The part of the front aperture that is in partial shadow displays the visual distortion. Prism overlap therefore decreases the fractional brightness of the visual distortion in favour of the other parasitic-images.

Alternatively, looking upwards through the front face of a Dove in the sheet, the top of the prism displays both the unreflected view through the prism and the base of the neighbouring prism (Figure 3.6(e)).

The Dove-prism offsets light ray position along both axes of the sheet plane. Offsets take place perpendicularly to the flip axis (see Figure 3.7(a)), but are relatively insignificant relative to those along the flip axis. A side effect

of ray-direction reflection with a Dove-prism sheet is that the view through each prism is also reflected. This leads to a small offset in position of the order of the height of each prism which may be minimized by correct prism design. The reflection takes place along a mirror plane parallel to the prism base (see Figure 3.7(b)). Its position depends on a number of parameters, including the prism angle, γ, (measured between the left face and the vertical), the refractive index, n, and the thickness to height ratio of the prism. This ratio minimizes the offset when the mirror plane is halfway between the roof and base of the prism (see Figure 3.7(c)). It may be shown that this occurs when

$$\frac{D}{H} = \cot\left[\gamma - \sin^{-1}\left(\frac{\sin\gamma}{n}\right)\right] \tag{3.10}$$

where the prism height is H and the prism thickness, T, which is measured at half the height of the prism.

The sheet design may be optimized in order to reduce the effect of the angular and positional errors discussed above. Offsets are smallest for small enough prism sizes and so are simple to remove, but of the angular errors, only the 'pixelation' may be reduced.

3.4 Connection to negative refraction

A Dove-prism sheet, when considered only in terms of the prism projection (i.e. the x-z plane) is simply a reflection of a light ray direction perpendicular to the sheet. Analogously, as was discussed in Chapter 2, the same law of refraction is observed across an interface between positive and negative media of equal magnitude (i.e. between n and $-n$). The later law follows from Snell's law, but applies along both transverse axes. In order to generalize a Dove-prism sheet to this three dimensional form, a second Dove-prism sheet may be placed behind the first in order to reflect along the other transverse axis.

The law of metarefraction for this crossed Dove-prism sheet is

$$r'_x = -r_x \tag{3.11a}$$

$$r'_y = -r_y \tag{3.11b}$$

$$r'_z = +r_z \tag{3.11c}$$

and corresponds to pseudoscopic imaging (see Section 2.5). This will be discussed further in Chapter 5.

3.5 Summary

The work explored in this chapter was begun by Johannes Courtial and John Nelson. I have been involved in the exploration of Dove-prism parameters and much of the simulation work. I also took part in the mathematical analysis of the Dove-prism sheet. This chapter explored work that has been published in References [37] and [38].

In this chapter, a Dove-prism sheet has been introduced. Specifically, it was shown that the reflection of a single light-ray direction, with negligible offset, may be achieved by a stack of wide-apertured Dove prisms. The apparent direction of a point light source then appears to move away from its actual location of the light when it is just behind the sheet to negative infinity (along the flip direction) as the object moves as far behind the sheet as the camera is in front. As it continues to move further back, the object then reappears from positive infinity (again along the flip direction) and appears to move towards the eye position's projection onto the sheet. Equation (3.9) was derived which quantifies the behaviour of the sheet making a straight line appear as a hyperbola.

While a Dove-prism sheet's ray-direction mapping is extremely simple and so too is the resulting view, the connection between the two is not. In addition,

the Dove-prism sheet acts as a fundamental component in realizing other more elaborate metarefractions (such as the crossed Dove-prism sheet). Such extensions will be the subject of Chapter 4 and 5. In particular, the latter chapter will return to the idea of pseudoscopic imaging and will explore the potential of using Dove-prism sheets to perform more general forms of imaging.

CHAPTER 4

The ray-rotation sheet

> I was like a boy playing on the sea-shore, and diverting myself now and then finding a smoother pebble or a prettier shell than ordinary, whilst the great ocean of truth lay all undiscovered before me.
>
> — *Isaac Newton*

4.1 Introduction

Here the crossed Dove-prism sheet will be generalized. In Section 3.4 the crossed Dove-prism sheet was described as a pair of ray-direction reflections along both transverse axes of the sheet, and was shown to be analogous to a refractive index interface between media of opposite sign. This metarefraction is also equivalent to a single light-ray direction rotation around the sheet normal by 180°. Here it will be shown that this ray-rotating metarefraction may be generalized to rotate by an arbitrary angle.

Now consider how such a metarefraction may be realized. As before, the sheet may be constructed from a pair of parallel Dove-prism sheets. Now however the flip axes of the two sheets are at an arbitrary angle to one another.

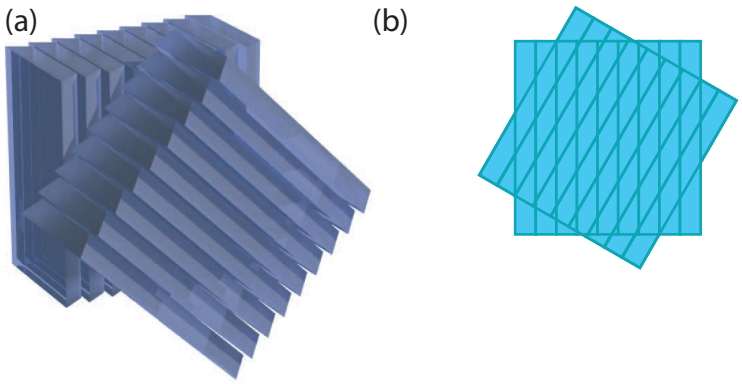

Figure 4.1: (a) The figure shows a ray-rotation sheet from one side. Such a sheet consists of a pair of parallel Dove-prism sheets in series. These sheets are rotated relative to one another, but remain coplanar. Figure (b) schematically shows the boundaries of a pair of Dove-prisms. Here, the sheet is divided into diamond 'pixels' which separately transform the local light ray direction by performing a pair of reflections.

The constituent metarefracting component is a pair of Dove-prisms then share a common optical axis, but are rotated relative to one another. An example of such a sheet is shown in Figure 4.1(a). The resulting metarefraction is for the light-ray direction to rotate by a fixed angle, α. The law of metarefraction for local light ray direction rotation may be expressed as

$$r'_\phi = r_\phi + \alpha, \tag{4.1}$$

where, as before, the z axis is taken to be the sheet normal around which the rotation takes place. The rotation therefore only changes the azimuthal component of the ray direction, r'_ϕ, while orthogonal ray directions remain unchanged. These components include the ray direction component normal to the sheet, r_z, and the polar ray direction angle, r_θ, between the ray and the sheet normal.

As stated in the previous chapter, Dove-prism sheets introduce a small offset of the order of the height of a prism. This offset will be introduced by each

sheet in turn about a different axis which leads to diamond shaped pixelation of any visual distortion. This is shown schematically in Figure 4.1(b).

Whereas an idealized (i.e. an offset free) crossed Dove-prism sheet has an unambiguous wave-optical interpretation, the same is not true for an idealized ray-rotation sheet. With the exception of two angles, it is strictly speaking impossible as will be explored in Chapter 6. Such an approximation of ray-rotation using a metarefraction is perhaps as close to an idealized transformation as is physically realizable. The lack of exact wave-optical implementation may well explain why (to the best of the author's knowledge) ray-rotation has never previously been explored.

Following the same approach as was applied in the previous chapter, the visual distortions of ray rotation will be explored. In order to present the qualitative features, ray-traced simulations have been included. A more quantitative approach will then explore patterns in terms of a simple expression for the apparent position of a point light source.

This chapter will be structured as follows. Firstly, the mechanism by which two Dove-prism sheets perform ray rotation will be discussed. Then, the transformation's visual qualities will be explored in some detail. Next a quantitative analysis will fully describe the visual distortion and finally this chapter will conclude with a short summary.

4.2 Transformation mechanism

It was stated above that a metarefracting component comprising a pair of Dove-prisms is able to azimuthally rotate the local light ray direction. That a pair of reflections is equivalent to a rotation is a curious, but trivial mathematical identity.

Now consider the effect of a single constituent pair of Dove prisms, i.e. one

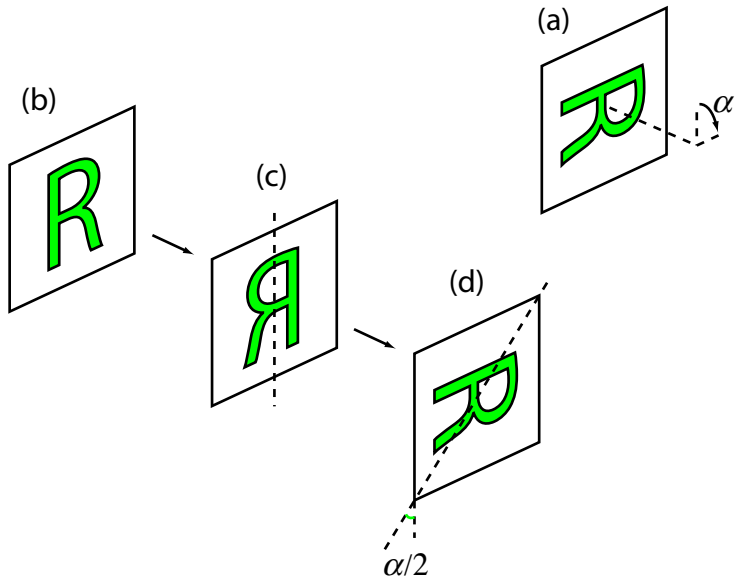

Figure 4.2: A ray-rotation sheet may be realized by combining two Dove-prism sheets. In order to do so, consider the behaviour of a single pair of Dove prisms that constitute the resulting metarefracting component. (a) The effect of such components is to rotate the local light ray direction. Locally, this is equivalent to rotating the incident scene. (b) Here the intensity pattern incident on the first Dove prism is chosen to be the letter 'R'. (c) The view through the Dove-prism is the reflection of the original pattern. (d) Viewing this reflected pattern through the second Dove prism, a second reflection takes place. The resulting transformations are equivalent.

of the diamonds in Figure 4.1(b). As long as both Dove-prisms reflect along the sheet plane, the ray direction component normal to the sheet remains unchanged by either Dove-prism. This discussion therefore reduces to the projection onto the sheet plane. The identity may then be demonstrated pictorially by rotating the letter 'R' as shown in Figure 4.2(a). This may be broken down into the various steps that take place. The view incident on the far prism's far face is shown in Figure 4.2(b). Now looking through the first prism, the visual distortion is a reflected version of the scene, as is shown in Figure 4.2(c). The second Dove prism then reflects again, this time at an angle

$\alpha/2$ to the first. The resulting visual distortion of the two is now a rotated version of the original scene. In Figure 4.2(d) the letter now appears to have been rotated by α.

In what way does the relative prism orientation relate to the resulting ray-rotation angle? Consider the effect of a pair of Dove-prism sheets on a ray direction vector. The ray passes through each sheet in turn, flipping it twice. A flip along an arbitrary axis on the x-y plane may be written as a change of basis by a rotation, say $\alpha/2$, from the x-axis, flipping along the new x-axis, then changing basis back by the reverse rotation:

$$\begin{bmatrix} \cos(\alpha/2) & \sin(\alpha/2) \\ -\sin(\alpha/2) & \cos(\alpha/2) \end{bmatrix} \begin{bmatrix} -1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \cos(\alpha/2) & -\sin(\alpha/2) \\ \sin(\alpha/2) & \cos(\alpha/2) \end{bmatrix}, \tag{4.2}$$

which may more concisely be written as

$$\begin{bmatrix} -\cos\alpha & \sin\alpha \\ \sin\alpha & \cos\alpha \end{bmatrix}. \tag{4.3}$$

Applying a second flip, this time along the x-axis, the resulting matrix is

$$\begin{bmatrix} \cos\alpha & -\sin\alpha \\ \sin\alpha & \cos\alpha \end{bmatrix} \tag{4.4}$$

which is the matrix describing rotation through α. That is, a pair of flips at an angle $\alpha/2$ to one another is equivalent to a rotation by α.

4.3 Optical properties of a ray-rotation sheet

As in the case of a Dove-prism sheet, a ray-rotation sheet may be best understood in terms of ray-traced images. The quantitative behaviour of the sheet is quite unexpected as demonstrated in the two ray-traced images that are shown in Figures 4.3 and 4.4.

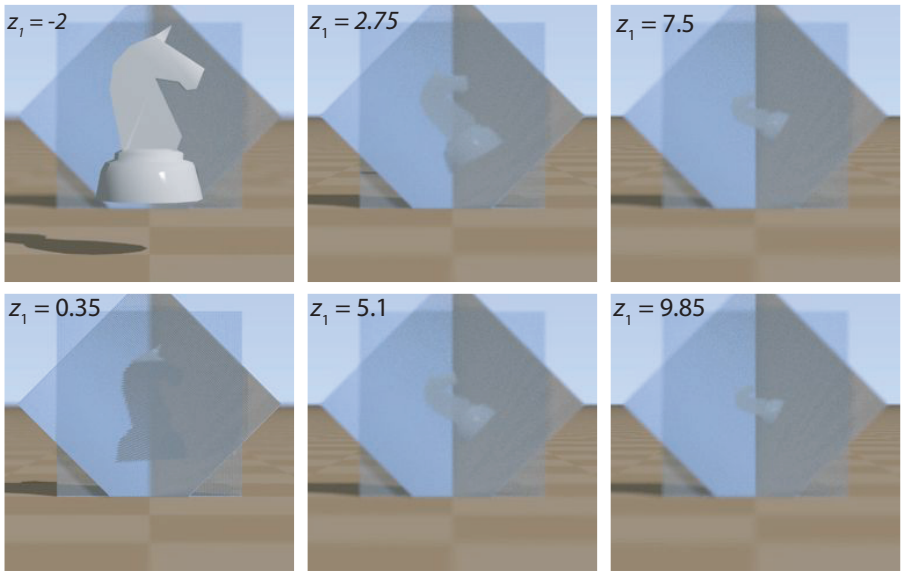

Figure 4.3: The apparent orientation of the chess knight is dependant on the particular depth. Chess piece seen through a pair of Dove-prism sheets rotating the direction of transmitted light rays through $\alpha = 90°$ around the sheet normal. Different frames are calculated for various distances of the chess piece behind the sheets (2×100 Dove prisms). The distance between the sheets and the camera is $z_2 = 6$ (in units of the floor-tile side length).

First considering Figure 4.3 where depth of a chess knight has been varied, the orientation, despite appearing to rotate, has not been changed throughout. In moving the object then it appears to travel along different axes and as a result will move around some centre. While subtle, the knight is actually distorted as well as being reorientated and so does not simply undergo a rotation. This said, the knight does not appear to be significantly offset. The horizon, however, remains consistently distorted by some fixed angle. Therefore at infinity (at least for the horizon) the line is transformed into a second line. When viewing any point on the knight, it actually appears to rotate in a separate circle, independently of the whole knight.

In Figure 4.4, the ray-rotation angle has been varied. In order to make

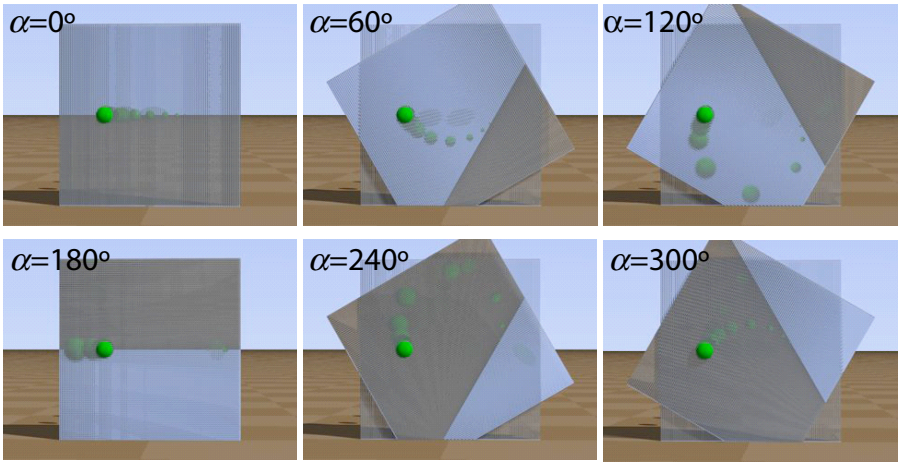

Figure 4.4: The frames show the apparent positions of several green spheres, all positioned along a sheet normal through the point at various distances from the sheet. Different frames are calculated for different ray-rotation angles, α. The green spheres play the role of a light source located at different distances behind the sheet. In terms of the distances between the camera and the sheet, these distances are $z = 2^n z'$, for $n = -5, \ldots, 5$. The camera is directed normal to the sheet at a distance $z' = 6$ floor-tile side lengths in front.

clearer the circular paths traced as the object depth is varied, the chess piece has been replaced by a series of spheres. Each sphere is located immediately behind the previous one, but when viewed through a ray-rotation sheet, it instead appears to trace the circumference of a circles. By varying the rotation angle, the underlying shape is made clearer and is indeed always circular. The particular path does however appear to move, changing both centre and radius with rotation angle. Additionally, looking at the horizon in Figure 4.4 demonstrates that the horizon in not distorted by ray-rotation sheet so while it is rotated, the horizon remains a line. In fact, the angle by which the horizon rotates is equal to the ray-rotation angle, α.

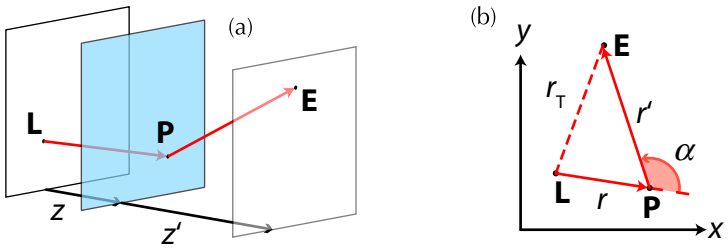

Figure 4.5: (a) The side view of a ray-rotation sheet as light-rays propagate between a point light source, **L**, and an eye position, **E**, through a point on the sheet, **P**. (b) As the z components of both light rays are known, it is simpler to work with the projection of the light rays onto the x-y plane. The second light ray, **PE** , which has a projected length, r', is rotated by an angle, α, relative to the first ray, **LP**, with corresponding projection length r.

4.4 Quantitative analysis

This chapter has so far explored the qualitative aspects of local light ray-rotation, but in order to describe the transformation more precisely, a quantitative description of what an observer sees when looking through such a sheet will now be investigated.

The constituent transformation on each light ray is a rotation and is therefore rotationally symmetric. Clearly the overall effect is very different to that of the constituent mapping, but objects do appear to rotate as the depth of an object is changed which implies that underlying metarefraction is relatively straightforward.

For a given observer's eye position, the distorted point light source is restricted to a line by the observed ray direction. Light rays propagate through the sheet from a known point light source, **L**, to a known eye position, **E**, through an unknown point on the sheet, **P**. As the point lies on the sheet, which is a known distance from the light source, z, and from the observer, z', it then remains only to solve for its x and y coordinates.

As a ray-rotation sheet only turns the ray direction, $\mathbf{LP} = (x, y, z)$, azimuthally, the polar angle, r_θ, remains unchanged during ray rotation into $\mathbf{PE} = (x', y', z')$. The tangent of polar angle is the ratio between the x-y ray projection length and the z component of the light ray. This ratio is therefore also preserved:

$$\tan r_\theta = \frac{\rho}{z} = \frac{\rho'}{z'}, \tag{4.5}$$

where $\rho = \sqrt{x^2 + y^2}$ and $\rho' = \sqrt{x'^2 + y'^2}$ denote the *projected* ray lengths before and after the ray mapping. Combining this with the definition of ray rotation, Equation (4.1), an expression for the projected light rays may be written as

$$\begin{pmatrix} x' \\ y' \end{pmatrix} = \frac{z'}{z} \begin{bmatrix} \cos\alpha & -\sin\alpha \\ \sin\alpha & \cos\alpha \end{bmatrix} \begin{pmatrix} x \\ y \end{pmatrix}. \tag{4.6}$$

As is summarized in Figure 4.5, in order to simplify the following treatment, the sheet plane will now be interpreted as being the complex plane. Now, the x-axis becomes the real axis and the y-axis becomes the imaginary axis. Correspondingly, the ray direction projections will similarly be described as complex numbers. Now (x, y) becomes $r = x + iy$, (x', y') is now $r' = x' + iy'$ and similarly $\mathbf{LE}$ becomes r_T. This representation is simpler because by Equation (4.6) may be rewritten as

$$r' = \frac{z'}{z} \exp(i\alpha) r. \tag{4.7}$$

The vector sum of the two rays is the (known) projected distance between the light source and eye,

$$r_T = r + r'. \tag{4.8}$$

Upon substituting for r' from Equation (4.7), this becomes

$$r_T = r \left[1 + \frac{z'}{z} \exp(i\alpha) \right], \tag{4.9}$$

which may be rearranged for the (unknown) ray direction projection

$$r = \frac{1}{1 + z'/z \exp(i\alpha)} r_T. \tag{4.10}$$

What is of interest is the visual distortion seen by an observer looking through the ray-rotation sheet. This distorted location of each point light source, in turn, may be inferred from the apparent ray direction projection

$$r' = r_T - r \tag{4.11}$$

where r is given by Equation (4.11) and hence

$$r' = r_T \left[1 - \frac{1}{1 + z'/z \exp(i\alpha)} \right] \tag{4.12}$$

which on simplifying may be written as

$$r' = r_T \left[\frac{z}{z'} \exp(-i\alpha) + 1 \right]^{-1}. \tag{4.13}$$

This equation may be rewritten as a Möbius transformation [39] of the ray projection:

$$f(\zeta) = \frac{a\zeta + b}{c\zeta + d}, \tag{4.14}$$

where

$$\zeta = z/z' \exp(-i\alpha) \tag{4.15}$$

and the coefficients are $a = 0$, $b = r_T$, $c = 1$ and $d = 1$. Equation (4.15) describes a point on a circle: varying the object or image distance varies radius and varying rotation angle rotates varies the point on the circle's circumference. It is true that for any Möbius transform that lines and circles are mapped into other lines or circles. Therefore varying object distance, image distance or rotation angle moves the viewing direction in a line or circle as was observed in Figure 4.4. The Möbius transformation is only a consequence of the artificial choice of working on the complex plane. The real vector-space

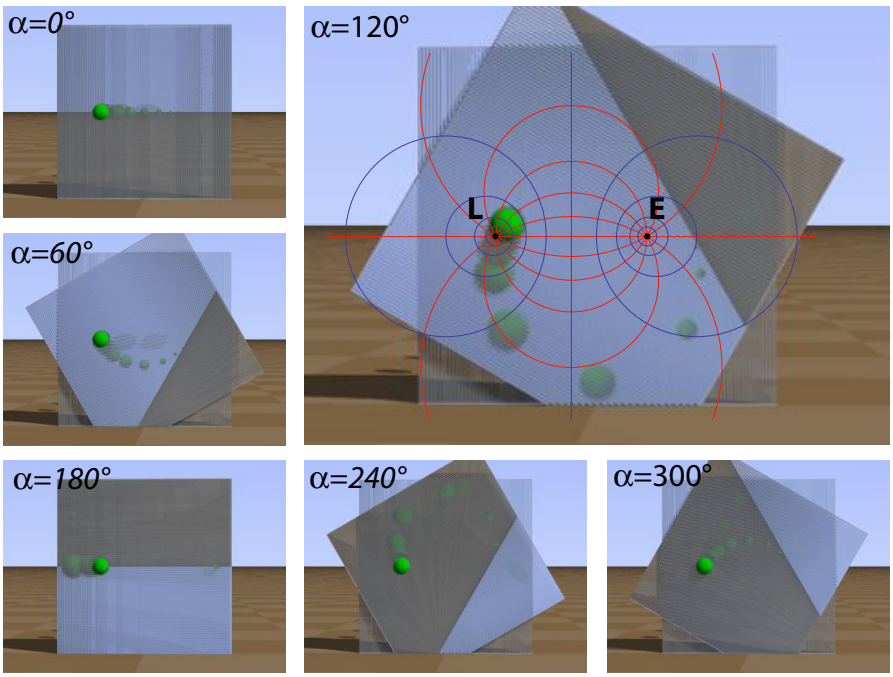

Figure 4.6: The transformation of a point light source by a ray-rotation sheet is to make appear the point according to a bipolar coordinate transformation. Changing the object's depth moves it around circles as does changing the ray-rotation angle.

equivalent is a bipolar coordinate transformation as shown in Figure 4.6, but the same argument applies. In this case the circles are known as Apollonian circles.

4.5 Summary

This work was performed in collaboration with Bhuvanesh Sundar, John Nelson and Johannes Courtial and closely follows the content of Reference [40].

Local light-ray direction rotation has been shown to be constructable from a pair of Dove-prism sheets. By varying either the ratio of sheet separation of the sheet and observer or the rotation angle, the apparent location of a point light source was observed to rotate around a circle.

There is a rotational symmetry that is introduced into the system by the constituent ray-direction mapping. Because of this, it is not surprising that there was a corresponding spherical symmetry in observed ray location. The resulting ray mapping was a bipolar transformation and was expressed in complex numbers as Equation (4.13).

Now that a ray-rotation transformation has been described and its visual properties outlined, what will follow in the next chapter is an exploration of imaging as may be performed by the metarefracting components that have been introduced so far.

CHAPTER 5

Imaging

For long you live and high you fly
And smiles you'll give and tears you'll cry
And all you touch and all you see
Is all your life will ever be.

— *Breathe, by Pink Floyd*

5.1 Introduction

This chapter is concerned with perhaps the most practical application of metarefraction, imaging. Metarefraction is able to realize several forms of imaging. The most obvious, and perhaps the only example already in common use, is the Fresnel lens. It is shown in Figure 1.2 and uses radially orientated prisms as constituent components in order to realize the metarefraction. This is able to reproduce the effect of a parabolic lens and has already been discussed in Section 1.3.

The second example to be explored was that of a Dove-prism sheet, which was discussed in Section 3.4. The direction of light rays that diverge from a point light source will be transversely reversed upon passing through the sheet

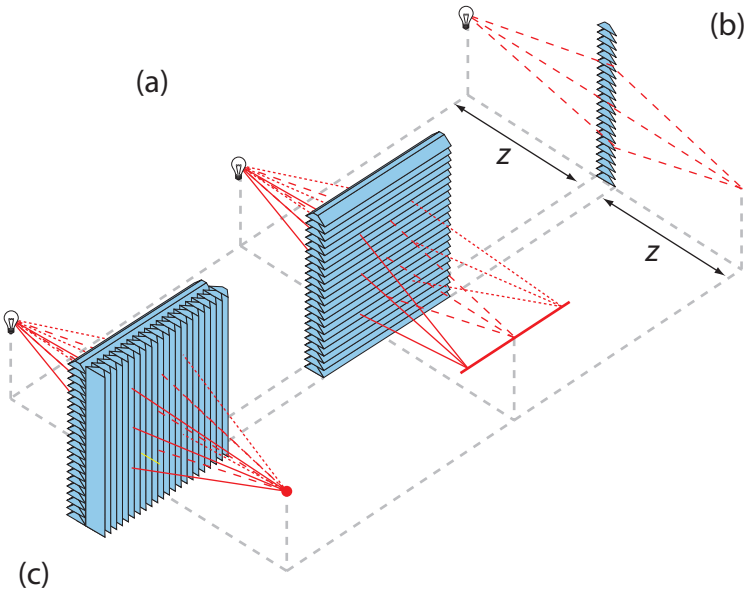

Figure 5.1: Imaging may be performed in a restrictive sense with (a) a Dove-prism sheet in two dimensions along the (b) prism projection and (c) with a Ray-rotation sheet in three dimensions (where the rotation angle is 180°). In either case, the image and object distances must be equal in order for imaging to take place.

and so converge to a point as is shown in Figure 5.1(a). This then images from a point anywhere on a line to a corresponding line a distance as far in front of the sheet as the original line was behind. A Dove-prism sheet is therefore able to image in the sense shown in Figure 5.1(b), along the prism projection.

Such imaging was then generalized to three dimensions using crossed Dove-prism sheets as shown in Figure 5.1(c). In either case, as the direction of the light rays are only reflected, rays travel as far a depth after the sheet to converge to a point as they travelled while diverging before the sheet. As a consequence of this, depth appears to be reversed: closer objects appearing behind further ones and vice versa. Such a sheet is equivalent as far as ray-optical imaging is concerned to an interface between positive and negative refractive indices of equal magnitude. This form of imaging is called pseudo-

scopic imaging as discussed in Section 2.5 and is discussed further by Nelson and Courtial in Reference [37].

Generalizing further by considering a ray-rotation sheet, imaging is observed only in the two special cases where the ray-rotation angle, α, is 0° which immediately reduce to direct transmission (the identity ray mapping); and to 180° which corresponds to a crossed Dove-prism sheet. For any other rotation angle, the sheet produces a non-imaging visual distortion as was discussed in Chapter 4.

The imaging performed by a Dove-prism sheet, sign-changing refractive-index interface and crossed Dove-prism sheets is unusual. The peculiarity to all three imaging systems is that they are all homogeneous. That is, there is no optical-axis 'centre' as there is in a Fresnel lens. As a consequence, such systems may have applications in the construction of large apertured systems without significant aberration. In practice, however, the quality of the imaging performed by metarefracting sheets is significantly undermined by the pixelated and ray-optical nature of the transformation.

Additionally, the flexibility in the forms of imaging that any homogeneous metarefracting sheets may perform is limited. There is no preferred displacement direction in which to transform. As a result of the azimuthal symmetry of the ray-rotation transformation and the lack of a central optical axis position, there can be no scaling nor rotation of the object plane when mapped onto the image plane and so imaging must therefore be an identity mapping.

Despite imaging being quite restricted when performed by a ray-rotation sheet, this is no longer true of two or more such sheets. In this case imaging is far more unusual; it is more specialized in that imaging is only possible between two specific planes. However, these planes may be chosen by correct selection of appropriate sheet parameters which may be changed as required

(i.e. they need not be selected during manufacture). While these two planes may be sharply resolved by an observer, the visual distortion from and to all others appears to be both blurred and rotated.

5.2 Multiple ray-rotation sheet imaging

This form of imaging may be generalized to a type of imaging which takes place when multiple ray-rotation sheets are combined. Such series of sheets offer the possibility of imaging between planes at arbitrary distances as well as the ability to control the location of the object and image plane by varying the sheets' ray-rotation angles and the inter-sheet distance(s). A simulation of such a system of two such sheets is shown in Figure 5.2(a). Here the image distance and chess-knight distance were selected to complement the sheet separation and ray-rotation angles such that, despite appreciable focal blur being simulated in the scene, the chess knight appears in sharp focus. Here, focal blur is simulated by averaging the sampled intensity of multiple rays that intersect the camera through a slightly different point on its aperture, all of which are detected at the same depth plane and from the same direction. This form of imaging will be discussed further throughout the rest of this chapter and follows the work performed in Reference [41] which was performed with Dr Courtial.

Imaging with two ray-rotation sheets may be shown to take place. This can be demonstrated by looking for consistency in parallax between different viewing directions, this time shown in Figure 5.2(b). The same scene has added to it additional coloured chess knights at one half, twice and four times the separation between the original knight and the far sheet. Here the camera is moved horizontally while the chess knights are kept stationary and upright behind the two separated ray-rotation sheets. As the camera moves, all the

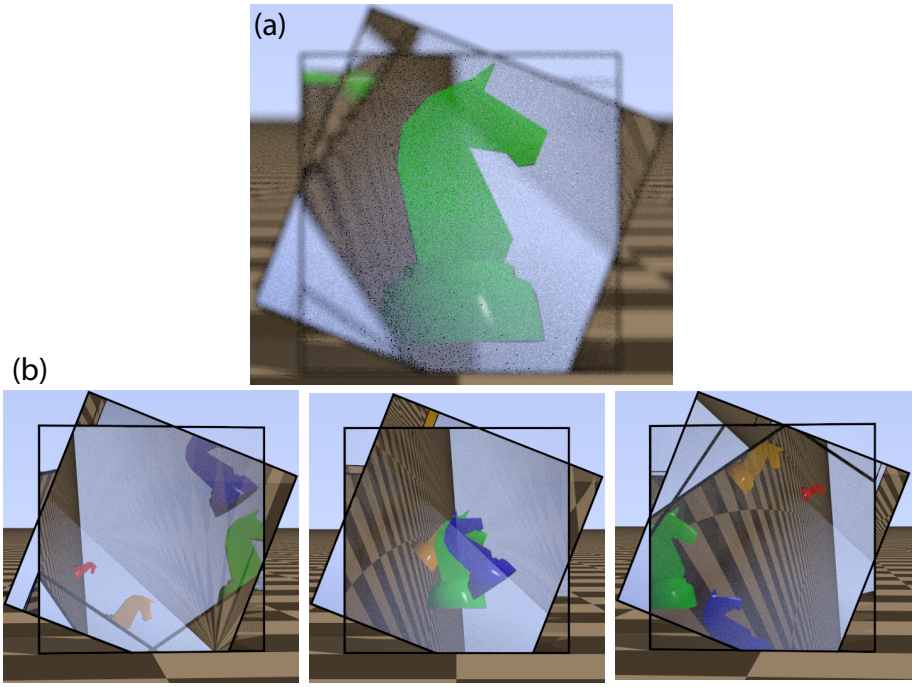

Figure 5.2: Two simulations each looking through a pair of ray-rotation sheets, where each sheet has a ray-rotation angle of $\alpha = 137°$, each consisting of two Dove-prism sheets (each of 1000 Dove prisms). The sheets are separated by a distance $s = 5$ and the object and image distance are $o = i = 3.42$ (in units of the floor-tile length). (a) A green chess knight in the object plane behind the two sheets appears as a sharp image despite rendering the scene with a finite aperture. (b) Now without focal blur and halving the size of the knight, additional (upright) chess knights have been added to the previous scene. The new chess knights are of equal size and are placed in different planes behind the ray-rotation sheets: at distances $o/2$ (blue), $2o$ (orange), and $4o$ (red). Three frames were calculated for different horizontal camera positions. This makes those chess pieces that are not in the object plane appear to move vertically which cannot be reconciled with parallax.

newly added chess knights appear to rotate and to move along lines that are inclined at some angle to the horizon. In these cases, the position of a point light source cannot be inferred using parallax and hence imaging, in the sense discussed in Chapter 1, does not take place. In contrast the original (green)

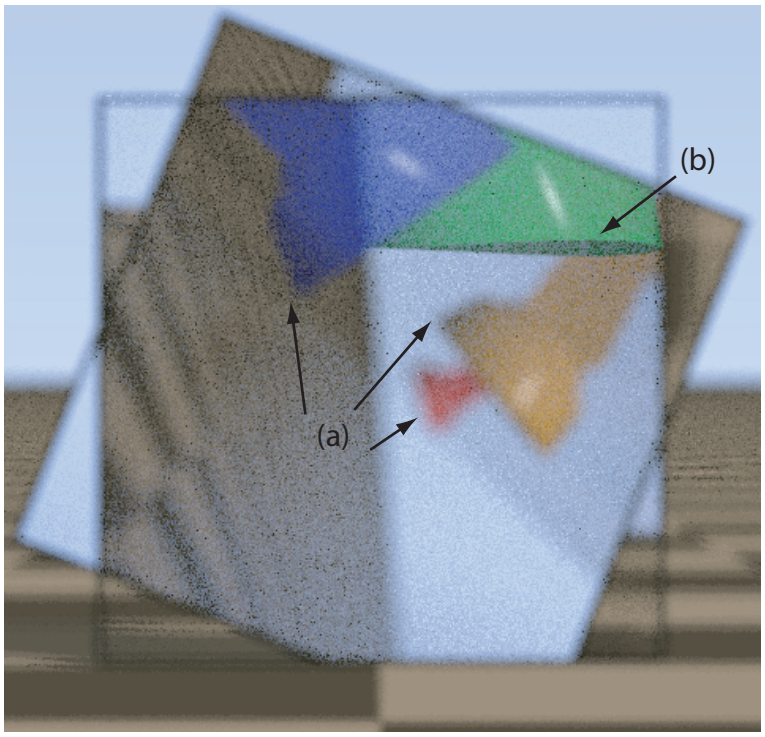

Figure 5.3: A reproduction of the central image in Figure 5.2(b) with focal blur and where the knights have been moved from directly behind the sheet to the upper right corner. While there is always a single object plane, there are no other planes from which imaging takes place. All other planes appear both blurred as well as appearing rotated, at a constant rate, relative to the imaging plane. This leads to both (a) objects at different depths appearing to be rotated by different amounts, and (b) giving objects the appearance of being twisted in depth.

knight does obey parallax, moving horizontally without being rotated.

Clearly imaging does not always take place and an exhaustive search would show only one object plane and one image plane exists for a given set of sheet rotation-angles and separation. A proof of this will be given in the next section.

Now the imaged chess knight appears in sharp focus. As before, the other chess knights are observed to be rotated, but this is also true to a lesser extent

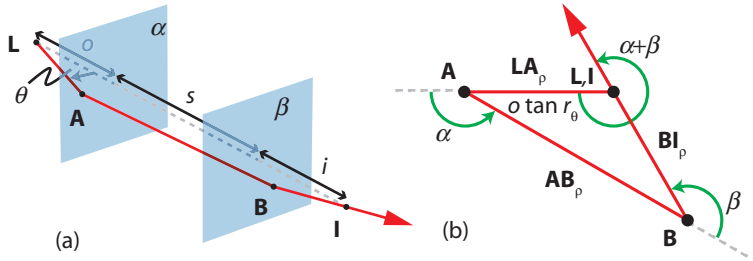

Figure 5.4: (a) The light path of a single ray is shown travelling between a point light source, **L** and image point, **I**. The ray is seen propagating through a pair of ray-rotation sheets, separated by a finite distance as viewed from the side. (b) The same ray path has now been projected onto the sheet plane. From this perspective, the ray path is a polygon of three lines, each rotated from the previous line by the rotation angle of the sheet that lies between them. The rotation angles are α and β respectively. The axis of ray rotation is around the sheet plane and is therefore transverse to the sheet normal. This direction may then be ignored and as a consequence, the sheet plane contains all of the light-ray dynamics.

of the individual depth planes on which the imaged chess knight lies. Each of these is also slightly rotated which leads to an additional visual distortion of any extended object. This becomes apparent on closer inspection of the centre image in Figure 5.2(b) and is more clearly reproduced in Figure 5.3 with focal blur (in the sense described above), and where the knights have been moved to the top right corner of the sheet.

5.3 Object and imaging plane selection

This section will explore two-ray-rotation-sheet imaging from a quantitative perspective. The system that will perform this imaging consists of a pair of ray-rotation sheets that an observer looks through in order to observe an image of the scene. As shown in Figure 5.4(a), both sheets lie parallel to the x-y plane.

This section will derive expressions for the required object distance, o, and image distance, i, (which are both measured relative to their nearest sheet). These quantities will be shown to be functions of the rotation angles of the two sheets, α and β respectively, and of the sheet separation, s.

Each light ray travels along a polygonal path through the system. The ray leaves the light-source, **L**, and intersects the first sheet at **A**. After being rotated by an angle, α, the ray leaves and travels in a straight line until it intersects the second sheet, this time at point **B**. After a second rotation by β, the ray then propagates to the image point, **I**.

Due to the similarity in problem to determining the visual distortion due to a single ray-rotation sheet, a similar treatment will be applied here. In Chapter 4, the law of metarefraction for a single ray-rotation sheet is given by Equation (4.1). The equation only alters the azimuthal angle, r_ϕ, and so does not affect orthogonal ray-direction components. Therefore the polar angle of the light ray, r_θ, and hence

$$\tan r_\theta = \frac{r_\rho}{r_z}, \tag{5.1}$$

remain unchanged. Here $r_\rho = (r_x^2 + r_y^2)^{1/2}$ is the component of the light ray direction in the sheet plane (i.e. the x-y plane). Two ray-rotation-sheet imaging is composed of two such azimuthal transformations, but neither affect the polar ray direction component. The three stages of the light ray path then all share the same polar angle

$$\tan r_\theta = \frac{\mathrm{LA}_\rho}{o} = \frac{\mathrm{AB}_\rho}{s} = \frac{\mathrm{BI}_\rho}{i}. \tag{5.2}$$

This allows the problem of solving for quantities in the z direction (the object and image distances) to be reinterpreted in terms of the projection of rays onto the sheet plane. This projection is shown in Figure 5.4(b). Solving

Equation (5.2) for the ray-projection lengths gives

$$\mathrm{LA}_\rho = o \tan r_\theta, \tag{5.3a}$$

$$\mathrm{AB}_\rho = s \tan r_\theta, \tag{5.3b}$$

$$\mathrm{BI}_\rho = i \tan r_\theta. \tag{5.3c}$$

As the ray-rotation sheets are both homogeneous, they are incapable of displacing or globally rotating an image. This is because there is no preferred offset direction in which to translate or rotate around. Imaging can therefore only ever be an identity mapping from the object plane to the image plane. Then, in order for imaging to take place, light rays that diverge from a point on the object plane must be rotated by the two sheets through appropriate angles in order to ensure that the resulting image point is not offset. That is, the initial and final light-ray projections into the sheet plane coincide:

$$(L_x, L_y) = (I_x, I_y). \tag{5.4}$$

In terms of the projection of the light ray, the polygonal ray path closes to form a triangle as shown in Figure 5.4(b).

The rotation angles of the sheet can be related to the interior angles of this triangle. The first two interior angles of the triangle are complementary to their ray-rotation angle and the third is inferred from all angles summing to 180°. The interior angles are

$$\angle\mathrm{LAB} = \pi - \alpha, \tag{5.5a}$$

$$\angle\mathrm{ABI} = \pi - \beta, \tag{5.5b}$$

$$\angle\mathrm{BIA} = \alpha + \beta - \pi. \tag{5.5c}$$

In order for the line segments to form a triangle, they must satisfy the law of sines. This may be written as

$$\frac{\mathrm{LA}_\rho}{\sin\angle\mathrm{ABI}} = \frac{\mathrm{AB}_\rho}{\sin\angle\mathrm{BIA}} = \frac{\mathrm{BI}_\rho}{\sin\angle\mathrm{LAB}}. \tag{5.6}$$

Combining Equation (5.3), Equation (5.5), and Equation (5.6) and on eliminating the common factor (i.e. $\tan r_\theta$) and simplifying with $\sin(\pi - x) = \sin x$ and $\sin(x - \pi) = -\sin x$, the object and image distances are

$$o = -\frac{s \sin \beta}{\sin(\alpha + \beta)}, \tag{5.7a}$$

$$i = -\frac{s \sin \alpha}{\sin(\alpha + \beta)}. \tag{5.7b}$$

The case of imaging (or otherwise) between a pair of sheets has been discussed. This may be generalized to an arbitrary number of ray-rotation sheets. General ray-rotation sheet imaging takes place between two principal planes [42]. For example, a single sheet only performs trivial imaging between two sides of the sheet: it maps only the sheet plane to itself. More generally, $n - 1$ sheets image when the closed path on the projection plane forms an n-sided polygon (generalizing a three sided triangle). The condition for imaging with a series of ray-rotation sheets then generalizes to

$$o = -\frac{1}{\sin \alpha} \sum_j s_j \sin \left(\sum_{k=j+1}^{N} \alpha_k \right), \tag{5.8a}$$

$$i = -\frac{1}{\sin \alpha} \sum_j s_j \sin \left(\sum_{k=1}^{j} \alpha_k \right) \tag{5.8b}$$

where s_1, s_2, ..., s_{N-1} is the separation between the sheets and the effective ray-rotation angle is given by $\alpha = \sum_j \alpha_j$.

Note that there is in general only one object plane and one image plane. There will, however, be special cases in which imaging does take place between intermediates, which corresponds to the ray-projection polygon closing at multiple points.

5.4 Summary

In this chapter, imaging by metarefracting elements and systems thereof has been explored. This is perhaps the most immediately useful application of metarefraction. A Dove-prism sheet images in two dimensions and crossed Dove-prism sheet image pseudoscopically. A series of ray-rotation sheets only image between principal planes. Here the view is an identity map while other planes appear both blurred and rotated. The apparent position of a point light source may be recovered with Equation (5.7) in the case of a pair of ray-rotation sheets and more generally with Equation (5.8).

CHAPTER 6

Wave-optical limitations of ray rotation

> "The fact that we live at the bottom of a deep gravity well, on the surface of a gas covered planet going around a nuclear fireball 90 million miles away and think this to be normal is obviously some indication of how skewed our perspective tends to be."
>
> — *Douglas Adams*

6.1 Introduction

The ray-optical consequences of metarefraction have already been briefly discussed in Section 1.1. There the discussion focussed on the size limitations of individual components. In this chapter, the wave-optical behaviour of metarefraction will be explored from a different perspective. While metarefraction was defined to allow minor offsets, this chapter will consider whether an offset-free equivalent transformation could be constructed, even in principle.

It will be shown that some offset-free ray mappings can only be realized for specific incident light fields. This is a consequence of implicit wave-optical requirements applying equally when working in the ray optical limit in terms of rays. Firstly, this requirement will be explored in terms of wave optics before the topic of translating from ray-optics back to wave-optics will be discussed. This requirement will then be applied to a light-ray field transformed by a ray-rotation sheet.

In order to simplify this discussion, only coherent, monochromatic light will be considered, and then only in terms of scalar wave optics. The following discussion will explore the wave-optical requirement in two dimensions, as a mapping between two planes, for example between either side of a metarefracting sheet. However, this discussion also applies to a wider range of optical systems.

6.2 Implicit wave-optical restrictions

Ultimately, any optical system must map between physically realizable, 'allowed', fields. That is all ray-optical mappings must map and incident light field into a second field that satisfies the ray optical approximation. Ray optics is an incomplete description of light and by translating into wave optics the light field is described more comprehensively. The conservation of underlying wave-optical concepts can be expressed mathematically as a constraint on the incident and outgoing light rays which must also be satisfied.

A transformation from an allowed field into a forbidden one is clearly impossible, and its real-world implementation must approximate such a transformation by an allowed equivalent. There therefore exists a class of apparently ray-optically valid mappings that, despite appearing to be self consistent, cannot actually transform most incident light fields.

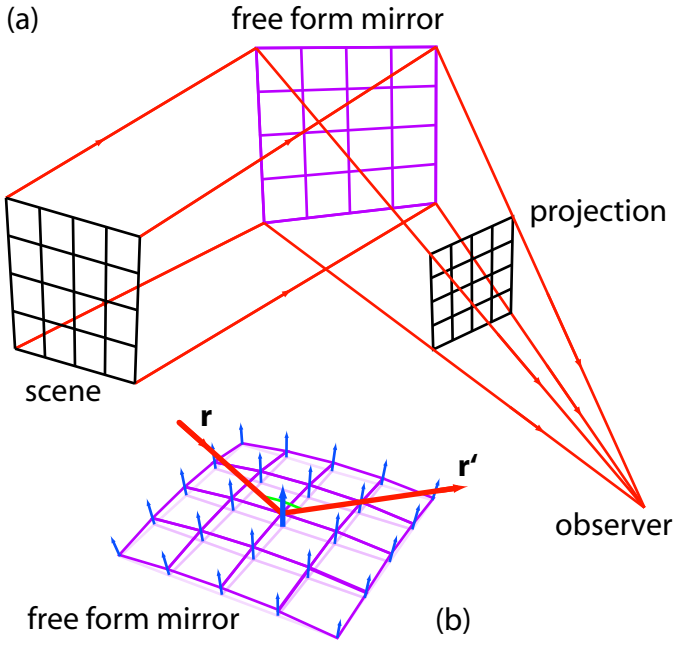

Figure 6.1: (a) In attempting to construct a free-form mirror that magnified a scene without distortion, Hicks [43] realized that the mirror surface could be inferred from the manner by which it reflected known incident rays towards a known eye position. (b) In integrating from a point at the centre of the mirror, he sought to calculate the surface from its gradient. However, the resulting surface was dependent on the path of integration, implying that the surface was inconsistently defined.

The form of wave-optical consistency that will prove relevant to offset-free metarefraction has already been discussed by previous authors. The work is most closely related to the work of Paterson [44, 45] who's work is discussed in detail below. However, it is instructive to begin by considering the work of Hicks [43] who sought to redirect an incident intensity pattern of initially collimated light rays towards an observer.

Specifically, Hicks studied the design of a distortion-free magnifying mirror surface. His aim was to redirect collimated light rays towards an observer's eye without distorting the intensity cross section. The experimental setup is

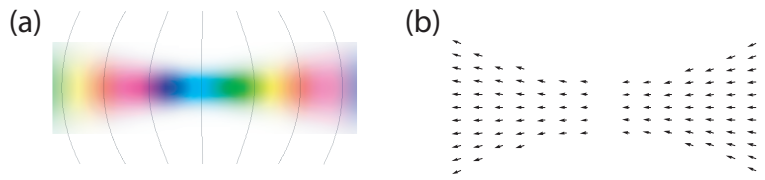

Figure 6.2: Ray-optical transformations, such as metarefraction, can be translated back into wave optics; for example, for the cross section of a Gaussian beam. (a) The intensity is common to both descriptions and is here shown as saturation. However, there is no single ray-optical replacement to phase (shown as hue and with solid grey lines). There are a number of alternatives that include the Poynting vector, the momentum flow and the phase-front-normal. (b) This last quantity is concisely represented by the direction in which the phase gradient points.

shown in Figure 6.1(a). Despite well approximating such a transformation, he realized that it was wave-optically impossible to do so using his approach.

The method that Hicks tried was to infer the shape of the mirror from its surface normal, and in turn to determine the normal at each point from the direction that correctly reflected the incident and reflected light ray pair that intersected there (see Figure 6.1(b)). However, the surface normals could not be translated into a single, consistent mirror surface. Although the reflection was defined in terms of apparently 'allowed' ray optics, the equivalent wave-optical description was 'forbidden'.

The same consistency requirement applies to offset-free metarefraction. In order to explore this further, a wave-optical description of metarefraction is required.

6.3 Generalizing from ray to wave optics

When relating ray optics to wave optics, the former may be derived from the latter. Crudely, this involves working with a simplified version of Maxwell's

equations [16] in the limit of infinitely small wavelength [46]. In contrast, translating back to the more general wave optics is more ambiguous. Ray optically, a light-ray field may be viewed as a real vector field of ray directions, $\mathbf{r}(x, y)$, and a scalar intensity field at every point in space, $I(x, y)$. Note that we are concerned here only with the slice through the light field on either side of a transforming sheet (which is taken to lie on the x-y plane).

The two slices through the field may be more completely described in terms of wave optics. For the sake of simplicity, this discussion will be restricted to scalar wave theory. Within this context, a coherent light field may be described wave-optically by a single, complex scalar field

$$u(x, y) = \sqrt{I(x, y)} \exp(i\phi(x, y)), \tag{6.1}$$

or in terms of the two constituent real scalar fields: intensity $I(x, y)$ and phase $\phi(x, y)$. An example of this is shown in Figure 6.2(a). This intensity is common to both ray- and wave-optical descriptions, so only the ray direction needs to be translated into phase in order to convert between these two field representations.

There are at least three reasonable choices for a wave-optical description of light-ray direction: These include the direction of energy flow, i.e. the Poynting vector; the direction of momentum flow, which can flow around the former, for example near the centre of an optical vortex; and the phase gradient direction. It is this last interpretation that is used here, where rays are seen as vectors normal to surfaces of constant phase and may be written as

$$\boldsymbol{\nabla}\phi(x, y) \propto \mathbf{r}(x, y). \tag{6.2}$$

An example of this (in which the ray direction has been normalized) is shown in Figure 6.2(b). This equation may be seen as being closely related to the eikonal equation (see Reference [46] and in particular Equation (24) in Reference [47]).

This requirement has made no assumptions on the nature of the scale factor in Equation (6.2). While quantitatively it is unimportant, it is needed to express a requirement on any outgoing ray mapping. In order to better understand this proportionality, a scale factor, which will be referred to here as the phase gradient length, $f(x, y)$, may be included explicitly. The ray direction is then taken to be of unit length, $\hat{\mathbf{r}}$. The more general version of Equation (6.2) becomes

$$\boldsymbol{\nabla}\phi(x, y) = f(x, y)\hat{\mathbf{r}}(x, y). \tag{6.3}$$

The ray direction and phase gradient must be parallel. This is however not always true for some values of the phase gradient length. For example, when it is negative then rays point in the opposite direction to the phase gradient and describes a different light field. Alternatively, when the phase gradient is zero or infinite then the ray direction is undefined. Phase gradient length is therefore restricted to values within the range

$$0 < f(x, y) < \infty. \tag{6.4}$$

Where phase gradient length is inferred from a given ray direction to take values outwith this range, the field is inconsistent and therefore physically cannot exist. The interpretation of this factor is discussed in more detail in References [48, 49, 50].

6.4 Wave-optical requirements

Above, the work by Hicks was introduced as an example in which ray-optical transformations suffered from implicit wave-optical restrictions. His transformation inconsistently described a mirror surface, and hence the phase of reflected light in terms of their gradients.

Another example of work in which the phase was inconsistently described was that of Paterson [44, 45]. His work investigated the rotation of an arbitrary intensity pattern about its centre. The experiment used a phase hologram to add a spatially varying phase to an incident light field pattern and observed the far field through a lens.

The aim of Paterson's experiment was to find a single hologram that produced a second intensity pattern in the far field that was an identical copy of the incident intensity, except that the entire pattern was rotated around the optical axis of the system.

In order to perform this transformation, Paterson sought an expression for the required phase hologram within the paraxial limit. In this context, the resulting intensity is related to the phase hologram in terms of the hologram's gradient. See Section 2.1.1 in Reference [44] for more details. As in the case of Hicks, the required phase pattern was defined in terms of its gradient vector field. And as before, while the vector field was well defined, it described an inconsistent phase.

Exactly this issue renders a number of offset-free forms of metarefraction unphysical. Apparently well defined phase gradients, i.e. ray directions, can imply an inconsistent phase. That is, metarefraction describes a visual distortion in which the phase cannot exist, even in principle.

So in a number of contexts, phase, a scalar field [51], may be described (to within an additive constant) in terms of its gradient. The gradient is a vector field, but a vector field may be the gradient to more elaborate, less consistent structures than a scalar field. So what limitations should be placed on such a vector field in order for it to qualify as the gradient of a scalar field?

A point on a scalar field may be determined by integrating from some initial point. For the vector field to consistently describe a scalar field, the

height of a point must be independent of the path of integration. In order to test for this consistency, the net change in phase around a closed path should be zero:

$$\Delta\phi = 0. \tag{6.5}$$

The differential form of this expression may be written as

$$\left.\frac{\partial^2\phi}{\partial y\partial x}\right|_{(x_0,y_0)} = \left.\frac{\partial^2\phi}{\partial x\partial y}\right|_{(x_0,y_0)}. \tag{6.6}$$

which tests whether a vector field, $\nabla\phi$, is a valid gradient to the scalar field ϕ. Mathematically, it is simply the z component of the familiar vector identity [52]

$$\boldsymbol{\nabla}\times(\boldsymbol{\nabla}\phi) = \mathbf{0}. \tag{6.7}$$

Rather than transform into a wave-optical description before applying the requirement, it is simpler to translate the phase-gradient requirement explicitly in terms of ray direction. Replacing the first order derivatives of phase in Equation (6.6) with the ray direction, as defined in Equation (6.2), the constant of proportionality cancels to become

$$\frac{\partial r_y}{\partial x} = \frac{\partial r_x}{\partial y}. \tag{6.8}$$

Or using the more general relationship between phase gradient and ray direction, Equation (6.3), the phase-gradient requirement may instead be written as

$$f\frac{\partial\hat{r}_y}{\partial x} + \frac{\partial f}{\partial x}\hat{r}_y = f\frac{\partial\hat{r}_x}{\partial y} + \frac{\partial f}{\partial y}\hat{r}_x. \tag{6.9}$$

6.5 Allowed and forbidden metarefractions

The wave-optical inconsistency of an apparently ray-optically valid metarefraction will now be demonstrated. Ray-rotation is chosen as it may be par-

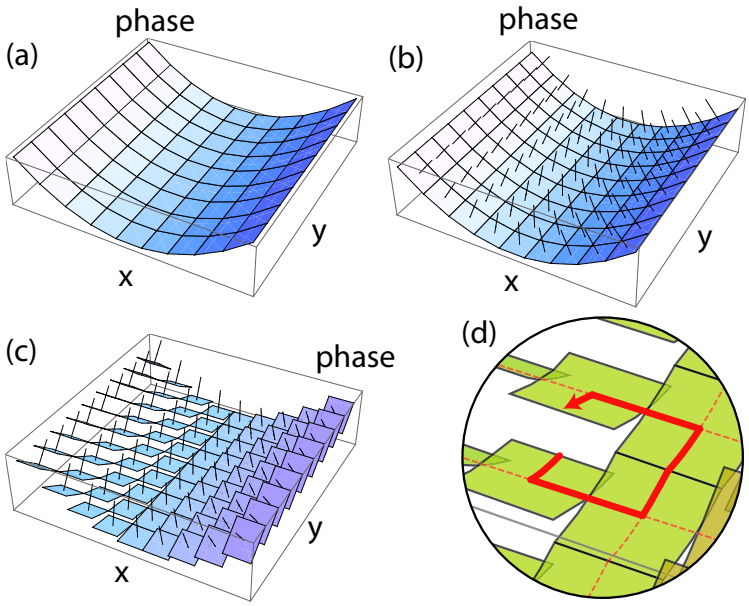

Figure 6.3: Ray rotation involves an azimuthal rotation of each light ray about the point at which the ray intersects the sheet. Wave optically, this may be described by how the phase gradient is transformed. (a) A plot of a parabolic phase field that is incident on the front of the sheet. (b) The light-ray direction is given by the corresponding phase gradient direction. (c) On passing through the sheet, the phase gradient direction is rotated by 90° at every point. The resulting phase field may be inferred from these gradients. This would be achieved, in principle, by moving each square up and down. (d) However, as can be seen in the phase plot, there can be no combination of offsets that joins the pieces of the surface into a piecewise-smooth whole.

ticularly clearly demonstrated. A more concise mathematical equivalent will then be presented for offset-free Dove-prism and crossed Dove-prism sheets.

Consider an arbitrary but non-zero light field with a phase that is constant along the y direction of the sheet plane, but parabolic along the x direction:

$$\phi(x, y) = x^2. \tag{6.10}$$

This is shown in Figure 6.3(a). Initially, this discussion will be restricted to the simpler of the two phase gradient definitions. The ray direction can be calculated by differentiating according to Equation (6.2) (with unity as a

proportionality factor) to get

$$\mathbf{r}(x, y) = (2x, 0), \tag{6.11}$$

and is shown in Figure 6.3(b). This is an allowed ray field according to Equation (6.8). After locally rotating this direction vector through 90°, the resulting vector field, shown in Figure 6.3(c), is

$$\mathbf{r}'(x, y) = (0, 2x) \tag{6.12}$$

and nolonger satisfies the requirements as the gradient to a scalar field.

This can also be seen visually in Figure 6.3(d). Regardless of how each polygon is moved up or down, the sides cannot smoothly join their neighbours at more than a finite number of points. Without being able to smoothly join the *infinitesimal* components, no smooth phase surface can be constructed, even in principle. Ray-rotation of an incident plane wave is therefore impossible.

The above discussion used the simple expression for the ray direction and ignored the effect of ray rotation on phase gradient length, $f(x, y)$. While it was sufficient in most practically interesting cases, the more specific description, Equation (6.9), is also informative. Explicitly including the phase gradient length in the description of the field, the phase gradient requirement is then a partial differential equation of phase gradient length as well as ray direction, $\mathbf{r}(x, y)$.

Applying this requirement to the rotated light-ray field, Equation (6.12), and cancelling ray direction, the expression reduces to

$$\frac{\partial f(x, y)}{\partial x} = -\frac{f(x, y)}{x}. \tag{6.13}$$

This partial differential equation may be solved to get

$$f(x, y) = \frac{g(y)}{x} \tag{6.14}$$

where $g(y)$ is an arbitrary function.

Those phase gradient lengths that do describe a consistent phase front must also propagate forwards in a well defined direction. The length of the phase gradient must remain positive, non-zero and finite (i.e. it is restricted by Equation (6.4)) which on substituting Equation (6.14), requires that

$$0 < \frac{g(y)}{x} < \infty. \tag{6.15}$$

This expression can only ever be true for at most half of the x axis. Offset-free ray-rotation may therefore only be implemented in a special case.

Rather than search for incident light fields for which a metarefraction does not work correctly, here a more systematic method will be introduced. For an optical mapping to be allowed, it must only produce allowed light-ray-direction fields that satisfy Equation (6.8).

The resultant light-ray-direction field is a function of the incident light-ray-direction field. As a consequence of this, it is possible to express the requirement of the output field in terms of the metarefraction of the incident field. On expanding Equation (6.8) using the chain rule, the simple test may be written as

$$\sum_{i=x,y,z} \frac{\partial r'_y}{\partial r_i} \frac{\partial r_i}{\partial x} = \sum_{i=x,y,z} \frac{\partial r'_x}{\partial r_i} \frac{\partial r_i}{\partial y} \tag{6.16}$$

where transformed rays are primed and the untransformed rays are unprimed. The $\partial \mathbf{r}'/\partial \mathbf{r}$ factors are derivatives of the metarefraction. On substituting for this law, this equation tests whether a metarefraction is allowed.

This may be used, for example, to test whether an offset-free Dove-prism sheet is allowed. That is, whether an alternative optical device, regardless of how it was implemented, could perform ray-direction reflection without any transverse change in the ray position on passing through such a sheet. The partial derivatives of Equation (3.1) with respect to the incident direction

components are

$$\frac{\partial r'_h}{\partial r_i} = \begin{cases} -1 & h = i = x \\ +1 & h = i = y \text{ or } h = i = z \\ 0 & \text{otherwise.} \end{cases} \tag{6.17}$$

Now substituting these into Equation (6.16), it reduces to

$$\frac{\partial r_y}{\partial x} + \frac{\partial r_x}{\partial y} = 0. \tag{6.18}$$

On substituting for the phase from Equation (6.2), this is

$$\left(\frac{\partial^2}{\partial x \partial y} + \frac{\partial^2}{\partial y \partial x} \right) \phi = 0. \tag{6.19}$$

In combination with Equation (6.6), this places extreme restrictions on the form of incident light fields that could be correctly transformed, even in principle. There therefore cannot, even in principle, exist a metarefraction that reflects one component of the ray direction without also transversely offsetting the position of the ray.

The Crossed Dove-prism is however allowed. This may be demonstrated by taking the partial derivatives of the law of metarefraction for a crossed Dove-prism sheet (i.e. Equation (3.11)):

$$\frac{\partial r'_h}{\partial r_i} = \begin{cases} -1 & h = i = x \text{ or } h = i = y \\ +1 & h = z \\ 0 & \text{otherwise.} \end{cases} \tag{6.20}$$

Now substituting these into Equation (6.16), it becomes

$$\frac{\partial r_y}{\partial x} = \frac{\partial r_x}{\partial y}. \tag{6.21}$$

This expression, in terms of phase (on substituting Equation (6.2)), becomes

$$\frac{\partial^2 \phi}{\partial x \partial y} = \frac{\partial^2 \phi}{\partial y \partial x} \tag{6.22}$$

which simply states that the incident field must satisfy the same requirement that was applied to the refracted field. This is simply a restatement of Equation (6.6) for the incident field. Offset-free crossed Dove-prism sheets are therefore a valid ray-optical transformation.

Note that these analyses are restricted to near the ray-optical limit. Yet while there may be unusual structures in which such requirements do not apply, they would be sufficiently specialized to be beyond the scope of this document.

6.6 Summary

This chapter gave an example of a wave-optical restriction and discussed how it might be applied to a mapping that is expressed ray-optically. Much of the subject matter contained within this document was originally published in Reference [48] and is the combined work of Mr Sundar, myself and Dr Courtial. The systematic test has been developed as part of a collaboration with Martin Šarbort and Tomáš Tyc.

It has been shown by previous authors that any ray-optical transformation must result in a phase that satisfies Equation (6.6) in order for it to describe a consistent phase surfaces. This requirement was restated, in terms of ray optics and then in terms of a metarefracted field. These requirements on the ray direction demonstrated the (in)consistency of several offset-free metarefractions.

This discussion was restricted to scalar waves in the ray-optical limit. Where this is true, no ray flipping or ray rotation is possible, and in more unusual situations, they remain limited to mapping between very specific light-ray fields. A similar, but more elaborate, analysis in terms of the Poynting vector should apply to a wider range of electromagnetic fields. This is however

dependent on matters such as polarization which is beyond the scope of this document.

CHAPTER 7

The confocal lenslet array

O wad some Power the giftie gie us
To see oursels as ithers see us!

— *To a Louse, by Robert Burns*

7.1 Introduction

So far only Dove-prism sheets have been considered. This chapter will seek a more general law of metarefraction than has been explored so far, i.e. the most flexible function that simultaneously maps each incident light ray direction into an arbitrary outgoing direction. In practice, such a sheet can only have a finite domain of parameter values, and hence can simultaneously map ray directions only within a range of outgoing angles. Nonetheless, the flexibility of a metarefracting sheet may be extended beyond ray-direction flipping and ray rotation.

In order to do so, a new metarefracting component will be explored. Shown in Figure 7.1(a), an alternative to a Dove prism is a confocal pair of cylindrical lenses. The resulting sheet that may be constructed from these will be referred

to as a confocal cylindrical-lenslet array. In terms of the metarefraction that it performs, it is equivalent to a Dove-prism sheet. That is, either form of sheet is able to realize local light ray flipping, rotating and imaging just as it may be performed by one or more Dove-prism sheets. Additionally however, the design of an individual confocal lenslet array may also be generalized beyond these forms of metarefraction.

There are a number of useful design parameters that control the performance of a confocal lenslet array. While varying a single parameter is sufficient to control one ray direction, multiple parameters are required in order to simultaneously, but separately control the redirection of multiple rays. That is, the law of metarefraction should be as flexible as possible in possessing the greatest possible number of independent parameters. Here the term 'independent' implies that changing each parameter leads to a qualitatively different behaviour to any arbitrary combination of all others.

The effect of a simple confocal cylindrical-lens pair, Figure 7.1(b), is to reflect one component of the light ray direction. This is achieved by focussing an incident, collimated beam through the shared focal line. From the focus, rays diverge towards the second lens whereupon they are again collimated. The reflection in direction and equality between incident and transformed angle follows from the horizontal symmetry of the ray paths about the focal line. Transverse to this, i.e. along the axis of the lens cylinder, the ray component remains unchanged throughout.

One way in which a confocal cylindrical-lenslet array may be extended is to replace the cylindrical lenses with corresponding spherical lenses in order to form a confocal spherical-lenslet array (Figure 7.1(c)). These are otherwise equivalent in behaviour except that the ray-direction reflection takes place simultaneously along both transverse axes of the lenses. Such a sheet performs

a metarefraction that is equivalent to that of a crossed Dove-prism sheet. Later, the sheet will be further generalized to use confocal elliptical lenses.

After exploring the historical development of various types of confocal lenslet array, this chapter will consider their construction and constituent parts. In the next section, the related subject of integral imaging will be discussed, in terms of which metarefraction may be reinterpreted. Having introduced this, the laws of metarefraction for both generalized cylindrical- and elliptical-confocal lenslet arrays will then be derived.

Pairs of spherical-lenslet arrays have been discussed by a number of authors. This literature may be grouped into two categories: The first of which seeks to image between nearby planes. In addition to neutron/x-ray imaging [53], these have found applications within photocopiers and oscilloscope screens [53, 54, 55]. The second category, derived from integral imaging, performs visual distortions and satisfy the definition as metarefracting components. These build on the work of Lippmann [56] and include the Gabor superlens [1] [57] and the Moiré magnifier [58]. The use of such components has been in the formation and magnification of various visual distortions. As well as being examples of integral imaging, this latter category may also be considered as inhomogeneous confocal lenslet arrays.

Alternative homogeneous generalizations have since been considered, first in terms of an analogy with negative refractive index which was explored by Johannes Courtial [59], before other parameters were explored by Sharvil Talati, myself and Courtial [60].

Confocal cylindrical-lenslet arrays may be constructed as shown in Figure 7.1(a), but it is impractical to align so many small lenslets individually. It is simpler to instead use a pair of prefabricated lenticular arrays. These

[1] The term 'superlens' in this context derives from the notion of the superlattice of a crystal and has no connection with the metamaterial superlens of Pendry.

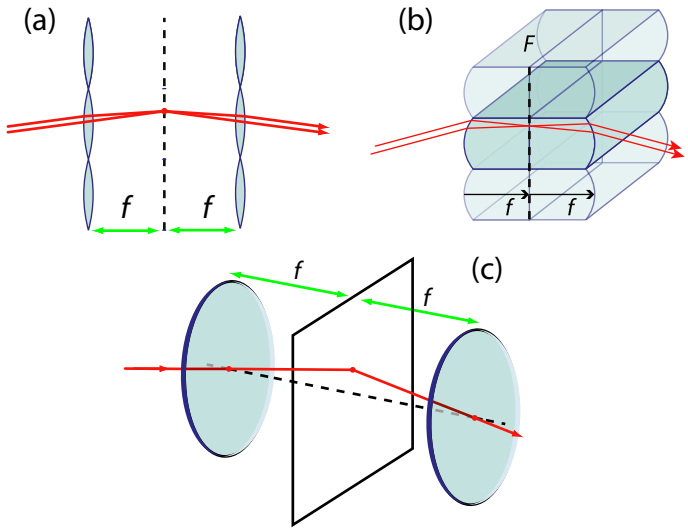

Figure 7.1: (a) The (schematic) side view of a confocal cylindrical-lenslet array. Two representative light rays are shown passing through the sheet. The system is confocal so transforms incident collimated light rays into a second beam of collimated rays. The effect of the lens pair is to reflect the transverse ray component. (b) In practice, such a sheet structure may be realized using a pair of lenticular arrays placed back-to-back and with a combined optical thickness of twice the focal length of the constituent lenses. (c) The sheet may be extended by replacing cylindrical lenses with spherical lenses.

are plastic sheets with curved lens-shaped indentations on one side. Lenticular arrays are commercially available as they are often used in tasks such as Shack-Hartmann wavefront sensing [61]. Confocal lenslet arrays (cylindrical or spherical) may then be constructed by placing two lenticular arrays back-to-back. The cylindrical case is shown in Figure 7.1(b). This constitutes perhaps the simplest method of confocal lenslet array construction. Correctly selecting their thickness ensures that the lenses on their surface are separated by two focal lengths of optical path, ensuring the resulting lens pairs are confocal.

Constituent lenticular arrays may be produced through a number of processes. Standard lenticular-array fabrication includes melting micro spheres [62, 63] onto a plastic block or moulding the requisite surface structure

from a metal template. An alternative approach is to replace curved lenses with graded refracting index lenses (GRIN lenses). These may be designed as discussed in Reference [64] and fabricated as described in Reference [10], most likely through diffusion [65].

Due to the homogeneity in depth of both lenses in a confocal GRIN lens pair, the resulting structure is also homogeneous in depth. Such a structure may therefore be constructed in exactly the same way as a single lens, except for lens thickness being doubled. This is also true for a lenticular GRIN array. The fabrication of GRIN lenticular arrays have already been thoroughly investigated [64]. Simply by doubling the thickness of the array should produce a confocal spherical-lenslet array. Its fabrication, although more specialized, would require only one such sheet to be fabricated and would therefore avoid all alignment issues.

7.2 Integral imaging

A confocal spherical-lenslet array may be interpreted in terms of integral imaging. Originally proposed by Lippmann [56], integral imaging was then referred to as integral photography, but the core concept has since applied in a wider range of optical systems. The device used to capture an integral image is known as a polydioptric camera or as a plenoptic camera [66] and has given rise to a number of components [57, 58, 67, 68].

A polydioptric camera is simply a lenticular array a focal length away from the integral image that it creates. The resulting integral image consists of a lattice of images, each formed by a separate lens that images from a different location. For all lenses, multiple images of the same scene, all taken from neighbouring perspectives, are imaged side by side onto a common plane — the integral image.

A conventional camera measures only the direction of light rays. All rays incident upon a single lens with a common direction are mapped onto a corresponding point on the image plane, regardless of where on the lens the ray intersected. However, on replacing the single lens with a lenticular array, the image that each constituent lens captures is specific to the location of that lens. In so doing, both the position and direction of a light field is measured. Measurement precision is limited by both the angular resolution of the (now smaller) lenses and spatially by the lens diameter. Angular and spatial precision are optimized at the expense of one another. Integral image quality is discussed further in Reference [69].

In order to view the scene that an integral image records, the same polydioptric camera may be used, but where light paths are reversed. Light rays travel from a point light source on the integral image back through the polydioptric camera, again a focal length away. The part of the light-ray field that was captured from where an observer is now positioned is retransmitted to the observer. In this way, an integral image acts like a ray-optical hologram, but the analogy is not valid wave optically in that a polydioptric camera does not capture phase.

While the integral image may be used as a three-dimensional photograph, it can also be interpreted as an information transformation. That is, various visual distortions may be performed by the manipulation of the integral image representation of a scene. An integral image may therefore be post-processed in order to select information either computationally or optically. This may, for example, reproduce focal blur from an arbitrary, hypothetical observation point [66, 70] or be used to infer subtle detail in the scene than could be captured from a single lens [67]. Alternatively, there are several examples in which a second lenticular array transforms back after scaling, offsetting or

selectively absorbing the integral image [57, 58, 68].

This form of such an information processing configuration, i.e. using a second lenticular array in order to view the first, generalizes a confocal lenslet array. The discussion will now be restricted to such a setup in which corresponding lens pairs in both arrays are arranged confocally, but in between which is some additional optical transformation. Only inhomogeneous optical transformations have been explored by the aforementioned authors, but the following discussion will be limited to homogeneous transformations.

Despite having reflected the polydioptric camera position in depth along the optical axis of the lenses, the lens that a ray leaves from remains unchanged as is position in the visual distortion of the integral image. The ray direction is however altered in that light rays were twice Fourier transformed by the recording and the viewing lenticular arrays rather than previously where the transformation was undone on passing back through the original lens.

The Gabor superlens and Moiré magnifier are implemented using inhomogeneous variations of the confocal spherical-lenslet array. Because of their inhomogeneity, they are therefore beyond the scope of this discussion. In contrast, a number of homogeneous lens properties have not elsewhere been explored (so far as the author is aware). What will now follow will be the generalization of the confocal lenslet array in a number of different ways.

7.3 Generalized confocal cylindrical-lenslet arrays

Regardless of how the design of a confocal cylindrical-lenslet array is generalized, each lens pair must remain confocal. Each lens pair must share a common focal line (while in the spherical lens case, each lens pair must share a common focal plane and more elaborately still, confocal elliptical lens pairs must share common focal lines along both transverse axes). Several parameters do not

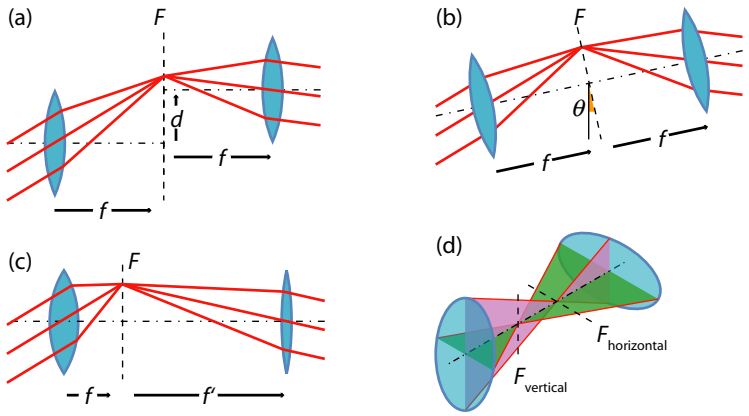

Figure 7.2: Confocal cylindrical lenses may be generalized in a number of different ways. While the telescopic property must be preserved, i.e. that the lenses must still share a focal line, they can be (a) shifted along the focal line relatively (by some transverse distance, d), (b) rotated (by some angle θ) and (c) have their relative focal lengths changed (and so have their focal line, F, moved along the optic axis). (d) A further extension is to work with elliptical lenses in order to simultaneously transform in both directions. The resulting astigmatism present in the latter case is shown in green and purple and so the lenses have two sets of focal lines along orthogonal axes (marked $F_{\text{horizontal}}$ and F_{vertical}).

affect the resulting metarefraction and so are ignored here.

Nonetheless, there remains some flexibility in how a confocal lenslet array is constructed, as is summarized in Figure 7.2. For example, the lenses shape may be varied, as can the relative orientation of confocal lenses to the sheet plane. Additionally, the offset between the shared focal lines and their position in depth offers an additional degree of freedom. Here such generalizations as well as their implications will be explored in more detail.

As is shown in Figure 7.2(a), the first parameter that may be varied is the offset between the two lenses. This offset must be entirely transverse to the optical axes of the lenses or else the focal planes of the two lenses will no longer coincide. This breaks the symmetry of the simple confocal lens pair and so incident rays will be transformed asymmetrically. The offset may be

described by a scalar displacement, d, which is a measure of the displacement of the second lens relative to the first, transverse to the optical axis.

A second generalization, as is shown in Figure 7.2(b), consists of rotating both lenses around a fixed point. The rotation is parametrised by an anti-clockwise rotation, θ, of the lens pair relative to the sheet. The rotation of the lenses corresponds to an opposite rotation of the rays. This change of basis from the input ray direction, α, and output ray direction, α', may be written in terms of their corresponding angles in the rotated basis,

$$\tilde{\alpha} = \alpha - \theta, \tag{7.1a}$$

$$\tilde{\alpha}' = \alpha' - \theta. \tag{7.1b}$$

This lens orientation parameter, θ, may be interpreted as approximately equivalent the previous offset parameter, δ. Observing that the focal planes are inclined with respect to one another implies that one is simply an aberrated version of the other. However, their effect is qualitatively different, affecting different incident angles differently and while they are approximately equivalent, as is shown below, they enable the independent control of multiple ray directions and the inclusion of both as separate parameters contributes to the flexibility of the ray-direction mapping.

The last parameter is the distance of the focal lines along the optic axis, between the two lenses. Additionally, the focal lines may even be moved in front or behind both lenses. Changing the position of the focal plane is equivalent to varying either focal length and is shown in Figure 7.2(c). This last parameter is related to the refractive index analogy as will be discussed below.

The above effects may be incorporated into a single description of the metarefraction performed by a generalized confocal cylindrical-lenslet array.

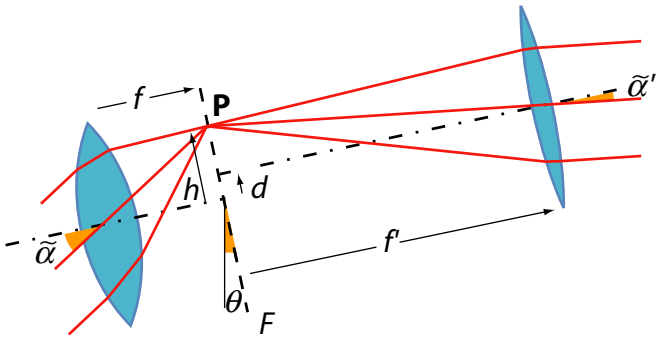

Figure 7.3: The most general cylindrical confocal lens pair setup. It is described by a number of parameters which include the focal lengths of the left and right lenses, f and f' respectively; in terms of the lens pair's orientation relative to the sheet normal, θ; and in terms of the offset of the optical axis of the right lens relative to the that of the left lens (shown with dashed lines). As before, the key property of this component is that it retains the telescopic property.

Shown in Figure 7.2, it will be sufficient here to consider the behaviour along a single plane on which the focal lines and optical axes lie. This will be referred to as the prism projection in order to remain consistent with previous chapters. Perpendicular to this plane, the ray direction component remains almost unchanged throughout.

Described by three parameters, the law of metarefraction for a generalized cylindrical case is significantly more elaborate than for ray flipping. However, an expression with which to describe it may be derived. This is done by relating the incident ray direction to the common point on the focal plane, which will be referred to here as the confocal point, **P** (as shown in Figure 7.3). This in turn may be related to the outgoing light ray. Both connections are made in terms of each light ray that passes, without deviation, through the centre of one of the two lens.

The tangent of the incident angle (relative to the shared direction of the lens' optical axis), $\tan\tilde{\alpha}$, is therefore equal to the ratio of the focal point height,

h, to the focal length, f, that separates the lens and focal line. Changing basis according to Equation (7.1)(a), the input angle becomes $\alpha - \theta$, viz.

$$\tan(\alpha - \theta) = \frac{h}{f}. \tag{7.2}$$

The optical axes of each lens are transversely shifted, relative to one another, by a distance d. Again in terms of the sheet basis, Equation (7.1)(b), the equivalent relation for the second lens is

$$\tan(\alpha' - \theta) = \frac{d - h}{f'}. \tag{7.3}$$

Eliminating h from Equations (7.2) and (7.3) gives

$$f \tan(\alpha - \theta) = d - f' \tan(\alpha' - \theta). \tag{7.4}$$

Simplifying by defining and substituting the dimensionless displacement,

$$\delta = \frac{d}{f}, \tag{7.5}$$

and the effective refractive index ratio,

$$\eta = -\frac{f'}{f}, \tag{7.6}$$

(the latter of which will be explained below). The resulting equation is a three-parameter model that describes the metarefraction that is performed by a two-dimensional generalized confocal lenslet array:

$$\tan(\alpha' - \theta) = \frac{\tan(\alpha - \theta) - \delta}{\eta}. \tag{7.7}$$

Note that, in the small angle approximation, δ, and θ, are equivalent, but for larger angles of incident, a pencil of rays will fan out differently on changing the lens orientation, θ, to changing the offset, δ. Again, multiple rays must be transformed simultaneously in order to distinguish between δ and θ.

For zero offset and rotation, but arbitrary effective refractive index ratio, the expression reduces to

$$\tan\alpha' = \frac{\tan\alpha}{\eta}. \tag{7.8}$$

This equation is very similar to Snell's law

$$\sin\alpha' = \frac{\sin\alpha}{n_2/n_1}, \tag{7.9}$$

particularly when using the small angle approximation where the two tangent functions may then be replaced with sine functions. It is in this paraxial limit where lower quality, aberrated confocal lenslet arrays work best. The effective refractive index ratio, η, is approximately equivalent to the ratio of the indices of refraction of an equivalent refracting interface, n_2/n_1.

The minus sign in Equation (7.6) ensures that the effective refractive index ratio is of opposite sign to that of the focal-length ratio as required. (A convex-convex lens pair with positive focal-length ratio reflects rays and so is analogous to a negative effective refractive index ratio.) The exploration of the effective refractive index ratio is discussed in more detail in Reference [59].

7.4 Three-dimensional model

The law of metarefraction that describes a generalized confocal spherical-lenslet array may be interpreted in terms of light-ray projections onto two orthogonal planes that lie transversely to the sheet and map light rays according to the equation derived in the previous section. Alternatively, both equations may be combined into a single model as will now be outlined.

The constituent component, a pair of confocal spherical-lenses, is shown in Figure 7.4. Lens-pair orientation is now described by a Euler rotation [71]. This is a three dimensional rotation matrix, $\mathbf{R}(\theta, \phi, \psi)$, parametrised by a z-axis rotation by θ then a x-axis rotation by ϕ and finally a second z-axis

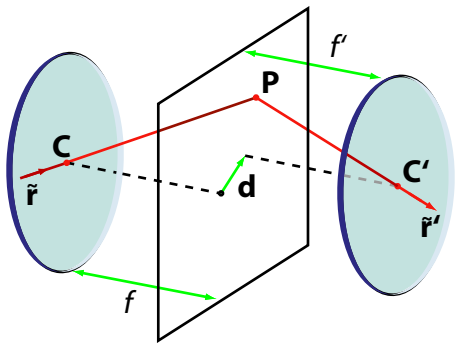

Figure 7.4: The passage of two separate light rays through a pair of lenses, **C** and **C**′. As in Figure 7.3, two light rays are considered (here they are not continued beyond the shared focal point). An input light ray direction $\tilde{\mathbf{r}}$ is mapped into an outgoing light ray direction $\tilde{\mathbf{r}}'$. As before, the left and right sheet are separated from the focal plane by distances f and f' respectively. The vector quantity **d** describes the displacement between the optical axes of the left and right lenses.

rotation by ψ. Rotation about the axis of the lens is unimportant, so ϕ is unimportant. Therefore only two parameters are required, but the system will be later generalized whereupon θ will prove useful. This change in bases of the input ray direction, $\mathbf{r}$, and output ray direction, $\mathbf{r}'$, will therefore be written as

$$\tilde{\mathbf{r}} = \mathbf{R}(\theta, \phi, \psi)\, \mathbf{r}, \tag{7.10a}$$

$$\tilde{\mathbf{r}}' = \mathbf{R}(\theta, \phi, \psi)\, \mathbf{r}'. \tag{7.10b}$$

By working in this local coordinate system, denoted by a tilde, the three Euler angles are implicitly included in the ray mapping. To the best of the author's knowledge, there is no method of more concisely incorporating the Euler angles into the metarefraction expression and so will not be discussed further. There are however other parameters. Shown in Figure 7.4, these are analogous to the transverse offset and the effective refractive index discussed above, but are now vector quantities.

As in two dimensions, the direction of the outgoing light ray that passes through the centre of the second lens is undeviated. For collimated outgoing light rays, all have passed through a single point on the focal plane, again referred to as the confocal point, $\mathbf{P}$. As a result, the new ray direction can be found by subtracting the centre of the second lens, $\mathbf{C}'$, from the point, $\mathbf{P}$, on the focal plane:

$$\mathbf{PC}' = \frac{f'}{\tilde{r}'_z}\tilde{\mathbf{r}}'. \tag{7.11}$$

Similarly, the incident light ray direction, $\tilde{\mathbf{r}}$, may be related to the left lens centre, $\mathbf{C}$, and to the shared confocal point, $\mathbf{P}$:

$$\mathbf{CP} = \frac{f}{\tilde{r}_z}\tilde{\mathbf{r}}. \tag{7.12}$$

These quantities may be related via the separation between lens centres,

$$\mathbf{CC}' = \mathbf{CP} + \mathbf{PC}'. \tag{7.13}$$

In order for the lenses to remain confocal, their separation in depth is the sum of focal lengths, $f + f'$, while their transverse displacement is arbitrary, d_x along the x axis and d_y along the y axis. Summing, the lens centres must therefore be separated by

$$\mathbf{CC}' = \begin{pmatrix} d_x \\ d_y \\ f + f' \end{pmatrix}. \tag{7.14}$$

Substituting Equations (7.11), (7.12) and (7.14) in Equation (7.13) gives

$$\begin{pmatrix} d_x \\ d_y \\ f + f' \end{pmatrix} = \frac{f}{\tilde{r}_z}\tilde{\mathbf{r}} + \frac{f'}{\tilde{r}'_z}\tilde{\mathbf{r}}' \tag{7.15}$$

solving for the outgoing ray direction to within an arbitrary scaling factor, this becomes

$$\tilde{\mathbf{r}}' \propto \begin{pmatrix} d_x \\ d_y \\ f + f' \end{pmatrix} - \frac{f}{\tilde{r}_z}\tilde{\mathbf{r}}. \tag{7.16}$$

This simplifies to

$$\tilde{\mathbf{r}}' \propto \begin{pmatrix} d_x - f\tilde{r}_x/\tilde{r}_z \\ d_y - f\tilde{r}_y/\tilde{r}_z \\ f' \end{pmatrix}. \tag{7.17}$$

In order to separate the effects of refractive index into the relevant ray direction component, Equation (7.17) may be divided by the z component:

$$\tilde{\mathbf{r}}' \propto \begin{pmatrix} d_x/f' - f\tilde{r}_x/f'\tilde{r}_z \\ d_y/f' - f\tilde{r}_y/f'\tilde{r}_z \\ 1 \end{pmatrix}. \tag{7.18}$$

Factorizing the x and y elements converts the equation,

$$\tilde{\mathbf{r}}' \propto \begin{pmatrix} [d_x/f - \tilde{r}_x/\tilde{r}_z]\, f/f' \\ [d_y/f - \tilde{r}_y/\tilde{r}_z]\, f/f' \\ 1 \end{pmatrix}, \tag{7.19}$$

into a form into which the dimensionless parameters may be substituted for

$$\tilde{\mathbf{r}}' \propto \begin{pmatrix} (\tilde{r}_x/\tilde{r}_z - \delta_x)/\eta \\ (\tilde{r}_y/\tilde{r}_z - \delta_y)/\eta \\ 1 \end{pmatrix}. \tag{7.20}$$

The minus sign appears due to the definition of effective refractive index and the dimensionless displacement is defined as

$$(\delta_x, \delta_y) = \left(\frac{d_x}{f}, \frac{d_y}{f}\right). \tag{7.21}$$

Effective refractive index may be similarly generalized, but in order to do so the focal lengths along the two axes of the lens, f_x and f_y, must first be separately varied. Such changes in focal lengh describe the metarefraction performed by a confocal elliptical-lenslet array along both axes. As a result, the dimensionless displacement may more generally be written as

$$(\delta_x, \delta_y) = \left(\frac{d_x}{f_x}, \frac{d_y}{f_y}\right). \tag{7.22}$$

Similarly, separate effective refractive indices may be introduced along the x and y axes:

$$(\eta_x, \eta_y) = \left(\frac{f'_x}{f_x}, \frac{f'_y}{f_y}\right). \tag{7.23}$$

Generalizing Equation (7.19) results in

$$\tilde{\mathbf{r}}' \propto \begin{pmatrix} (\tilde{r}_x/\tilde{r}_z - \delta_x)/\eta_x \\ (\tilde{r}_y/\tilde{r}_z - \delta_y)/\eta_y \\ 1 \end{pmatrix}. \tag{7.24}$$

The analogy with reflection across a planar interface may now be extended to include anisotropic refraction. This is exact only for small angles while aberrations will cause problems at larger angles.

Changing basis of the rays using the aforementioned Euler rotation,

$$\mathbf{R}(\theta, \phi, \psi)\,\mathbf{r} = (\tilde{r}_x, \tilde{r}_y, \tilde{r}_z) \tag{7.25}$$

the metarefraction is expressed in its most general form. The law of metarefraction for a generalized elliptical-confocal lenslet array may therefore be written as

$$\mathbf{r}'(\mathbf{r}; \eta_x, \eta_y, \delta_x, \delta_y, \mathbf{R}) \propto \mathbf{R}^T \left(\frac{\tilde{r}_x/\tilde{r}_z - \delta_x}{\eta_x}, \frac{\tilde{r}_y/\tilde{r}_z - \delta_y}{\eta_y}, 1\right) \tag{7.26}$$

where normalization of ray length is ignored.

7.5 Summary

In this chapter, confocal lenslet arrays and their generalizations have been explored. In particular, a historical introduction to the subject in terms of the related field of integral imaging was presented, as were potential fabrication methods and the derivation of laws of metarefraction for confocal cylindrical-lenslet arrays and for confocal elliptical-lenslet arrays.

There are a number of parallels between integral imaging and metarefraction. Both work ray optically, both transform between either sides of planar sheets and both use structure and not substance to realize a transformation analogous to refraction. Additionally, a metarefracting sheet may be thought of as transforming an incident light field in terms of its integral image as a confocal spherical-lenslet array consists of a pair of polydioptric cameras.

While confocal lenslet array shaping was previously considered in terms of a number of parameters, these were all inhomogeneous. Here, homogeneous parameters have been explored. Specifically, a confocal lenslet array may reproduce some of the visual qualities of negative refraction, either in two dimensions, as can the Dove-prism sheet; or in three dimensions, as a pair of crossed Dove-prism sheets can (see Figure 5.1). Additionally, the confocal lenslet array may be generalized to such an extent that its description includes transverse offsets, Euler rotations and a vector (i.e. anisotropic) effective refractive index ratio.

Transformation quality is dependent on well optimizing relevant engineering parameters that were previously ignored. But quality also depends on the parameters discussed above that affect the metarefraction. This has been summarized in Reference [72].

CHAPTER 8

Experimental ray flipping

"Exact science, Mr Angier, is not an exact science."

— *The Prestige*

8.1 Introduction

Until now the discussion of metarefraction has almost entirely been theoretical. In this chapter, experimental conformation of ray-flipping will be presented. This will demonstrate key qualitative aspects of the transformation. The demonstration will use a confocal cylindrical-lenslet array to perform the ray flipping. In order to do so, experimental images will be compared to equivalent ray-traced scenes.

In order to confirm ray-flipping, three experiments are performed. These will explore visual distortions associated with the two quantities that parametrize the metarefraction. The metarefraction is defined as reflecting one component of the ray direction parallel to the sheet. This is orientated at some angle to the horizon, β. The second parameter is the ratio between

the sheet-object and sheet-camera distances, ζ. The characteristic effect of ray flipping is the hyperbolic visual distortion of a line perpendicular to the sheet and this will be demonstrated in the first experiment. The other two experiments will explore the visual distortion along the sheet plane of a relatively flat object. In their respective experiments, the distance ratio and then the axis of reflection will be varied and compared to equivalent ray-traced simulations.

The next section will describe how the experiment was assembled and outline the key aspects in its design. Particular attention will be directed towards the choice of the camera lens, sheet assembly and shielding that significantly improved the quality of the results. After this, experimental results will be presented as a series of photographs and associated comments. In the last section, the quality of the results will be discussed and conclusions drawn regarding sources of blurring in the visual distortions.

8.2 Method

In order to demonstrate ray flipping, three versions of the same experiment were performed. In all cases, the experimental setup consisted of an object which was viewed through confocal cylindrical-lenslet array. The experimental setup is shown in Figure 8.1(a). The most sensitive part of the experiment was the construction and alignment of the confocal cylindrical lenslet array.

A confocal cylindrical-lenslet array can be constructed from a pair of plastic cylindrical lenslet arrays ("lenticular arrays") placed back-to-back as discussed in the previous chapter and shown schematically in Figure 7.1(a). These were commercially sourced from Edmund Optics (B43–028, $f = 0.085\,\text{in}$) which avoided the expense and complexity of having them custom built. The constituent lenses on the cylindrical-lenslet arrays will only be

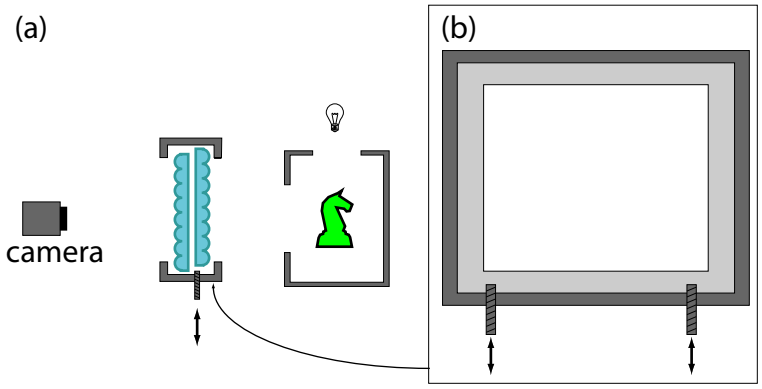

Figure 8.1: The ray-flipping metarefraction discussed in previous chapters was experimentally demonstrated. (a) The experiment consisted of a camera that viewed an object through a confocal cylindrical-lenslet array. The camera, of aperture $f/22$, was a distance 93.6 cm infront of the sheet. In order to resolve objects more closely, a macro lens was used which gave rise to a smaller field of view. Two objects were used: The first experiment used a metal rod which was perpendicular to the sheet surface. In the other two experiments, the rod was replaced with a chess knight (shown). Unless otherwise stated, the distance between the chess knight and the sheet was 43.7 cm. The object was illuminated while in a light box. Such shielding was used in order to avoid directly illuminating the sheets (which scattered a lot). (b) The confocal cylindrical-lenslet array consisted of a pair of cylindrical-lenslet arrays arranged back-to-back and held in place with a metal frame. One sheet would be aligned relative to the other using two grub screws, each separately pushing on a side of the latter sheet.

confocal if the optical path length between each lens pair is equal to the sum of their focal lengths. By selecting correct lenslet array thicknesses, and by placing them in a frame that held them together, the lenslet arrays remained approximately confocal (see Figure 8.1(a)).

In addition, the sheets had to be correctly orientated in order for the focal lines of the corresponding lenses to coincide. By adding a pair of grub screws to the sides of the frame that holds the lenslets, one cylindrical-lenslet array could be moved slightly along two axes relative to the other. The lenslets could then be shifted while remaining confocal (see Figure 8.1(b)).

Light rays that pass through one lens may not necessarily pass through the corresponding lens on the other lenslet array. Instead, it may travel through a neighbour of the latter lens giving rise to a significant loss mechanism. This may alternatively be described in terms of how this light appears. That is, the visual distortion gains additional "parasitic images" (here imaging does not actually take place, and the term is used to clearly distinguish undesired, alternative visual distortions). As the lens-pair offset is smoothly increased from zero, i.e the focal lines of the two confocal cylindrical-lenslet arrays are displaced, the loss fraction (for normal incidence) increases until all of the light passes through the wrong lens and the parasitic images are increasingly brightly illuminated. This topic is discussed more thoroughly in Reference [72].

In order to minimize this offset and hence the brightness of the parasitic images, a laser was shone perpendicularly through the sheet. This resulted in a comb of spots that were incident on a screen. These corresponded to parts of the beam passing through various lens pairs, each with a different offset, and so each beam left in a different direction. By having previously marked the destination of the beam without the sheet, the offset could be minimized by moving the brightest spot of the comb to this position.

The sheets scattered a significant amount of stray light from other paths towards the camera. In order to minimize this, shielding was placed around the illuminated object. While allowing the object to be illuminated and viewed, it stopped the light source from directly illuminating the sheets and scattering significant amounts of light into the camera.

The Canon EOS 450D camera with a Canon EF 100 mm $f/2.8$ macro lens was used to photograph the scene. The scene was viewed through the smallest available aperture setting ($f/22$) in order to minimize blurring. Focal distance was empirically chosen in order to further reduce blurring. The macro lens

was used in order to enable focussing on the sheets from shorter distances than would have otherwise been possible. This had a noticeable effect on blurring, but was not necessary. Additionally, the lens offered a narrow field of view.

The three experiments performed using this setup are as follows. The first object that was used was a metal rod that was arranged perpendicularly to the sheet. This was in order to demonstrate that the visual distortion of a line, which was normal to the sheet, would appear as a hyperbola.

Next, the object was replaced with a chess knight (shown in Figure 8.1). The sheet was then rotated azimuthally over a range of angles such that the sheet flipped light rays along a range of different axes. These were then compared with an equivalent ray-traced simulation to verify the visual distortion was as expected.

This experiment was also repeated for various camera distances in order to test variations in the sheet-camera to sheet-object distance ratio. This parameter was again compared with equivalent simulations in order to check the validity of the performed transformation.

8.3 Results

The experiments were carried out as stated above, and are summarized in Figures 8.2, 8.3 and 8.4. Overall the experiments worked as expected, the later two corresponding well to their simulation. However, all experimental photographs were subject to significant blurring.

The first experimental test was that the visual distortion of a metal rod, orientated approximately perpendicularly to the ray-flipping sheet, appeared as a hyperbola. Looking through the confocal cylindrical-lenslet array and through the circular aperture of the light box, Figure 8.2 confirms this prediction. The metal bar touches the sheet and diverges to the left beyond the

Figure 8.2: Experimental test of the predication straight line perpendicular to a ray-flipping sheet, when viewed through the sheet, appears bent into a hyperbola. The photo shows a straight metal bar that is approximately perpendicular to the ray-flipping sheet, seen through a confocal cylindrical-lenslet array, orientated such that it inverts the horizontal light-ray-direction component. The end of the bar (small back disk) touches the ray-flipping sheet.

field of view of the sheet. As the rod is of finite length, it does not extend far enough to reappear from the right hand side. As well as the rod, there is a wooden block on which it rests and on which there is a shadow cast by the rod. Despite the camera focal distance having been chosen by the experimenter to minimize blur, a noticeable degree of blurring remains. Close inspection of the image also reveals subtle 'pixelation' due to the constituent lenses.

The next experiment explored the effect of varying the orientation of the sheet, parametrized by the angle β with which the x axis is inclined with respect to the horizon. Meanwhile everything else remained constant. Here the experimental setup was the same as that of the previous experiment except that the metal rod was replaced with a wooden chess knight. Figure 8.3 compares the results of five angles against equivalent ray-traced scenes. The ray-traced scenes only include a chess knight surrounded by a black background. Clearly there is a good match between theory and experiment. However, there is significant additional blur in the experimental images.

The last experiment to be performed was similar to the previous one ex-

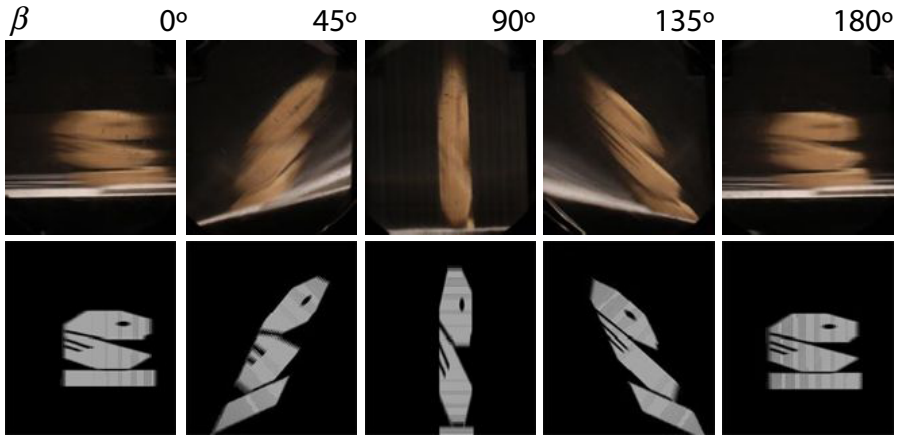

Figure 8.3: The view through a ray-flipping sheet for various sheet orientations, described by the angle, β, of the flip axis, relative to the horizontal. The top images are experimental photographs for a range of orientations. The simulated view through a corresponding Dove-prism sheet is shown in the images below. The experiment was performed for a camera distance of 96.3 cm (measured to the detector plane) and the object (a chess piece) at a distance of 43.7 cm. The scene includes a chequerboard floor and a chess piece viewed through a confocal cylindrical-lenslet array. The simulation was performed using the ray-tracing software POV-Ray and contained an equivalent chess knight in front of a black background.

cept that the sheet orientation was fixed while the distance ratio was varied. As before, the results (shown in Figure 8.4) were compared with equivalent simulations only of the chess knight. Again, the experimental and simulated visual distortions are in good agreement. Notably, the amount of focal blur increases for larger camera distances.

8.4 Discussion and conclusions

The three experiments have demonstrated that a confocal cylindrical-lenslet array performs ray-flipping as expected. These experiments tested both sheet parameters as well as the resultant visual distortion both normal-to and along

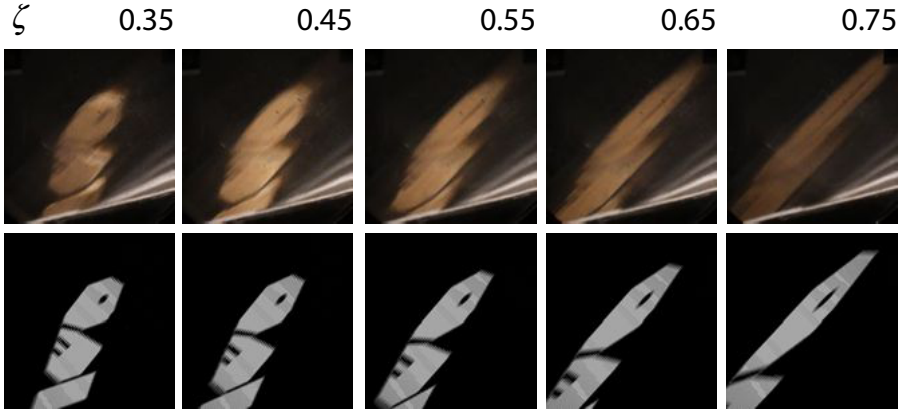

Figure 8.4: The view through a ray-flipping sheet for various object distances. Different columns correspond to different ratios of sheet-object distance to sheet-camera distance, which is labelled as ζ. The angle between the ray-flipping sheet and the horizontal was $\beta = 45°$ and the camera distance was 96.3 cm. The top images are photos of the scene, including a chequerboard floor and a chess knight, taken through a confocal cylindrical-lenslet array while the bottom image shows the simulated view through a corresponding Dove-prism sheet. In the simulation, an equivalent chess piece was rendered in front of a black background.

the plane of the sheet.

Light ray flipping is, however, subject to significant performance issues. The most serious problem was the presence of parasitic images which were caused by the misalignment of focal lines of corresponding lenslets. As a consequence, light passing through non-corresponding lens pairs. These parasitic images overlapped with the visual distortion, limiting the field of view of the sheets. In order to eliminate these, absorbing baffles could be placed between individual confocal lens pairs. Alternatively, the width of the lenses could be increased, or their focal lengths reduced. This would move the parasitic images to greater angles.

Another limitation was blurring over the visual distortion. This may be attributed to variation in sheet thickness which meant that the sheets were

nolonger confocal. An additional cause of blurring could have been because the sheets did not image and a finite camera aperture was used.

For further information on the experimental work, consult Reference [73]. The photography and alignment was carried out by Michael Blair, Leo Clark, Alasdair Houston and Gary Smith as part of a undergraduate project. I was involved in the design and assembly of the confocal cylindrical-lenslet array frame, the light box and some of the initial alignment. In addition, I supervised the student project along with Dr Jonathan Leach and Dr Johannes Courtial.

CHAPTER 9

Systems of confocal lenslet arrays

"Plurality ought never be posited without necessity."

— *William Ockham*

9.1 Introduction

The most general possible metarefracting sheet presented so far is the generalized confocal lenslet array. Confocal lenses are, to the best of the author's knowledge, the most flexible optical component to preserve the telescopic property. It is then reasonable to ask whether a (potentially infinite) series of generalized confocal lenslet arrays is able to perform a completely arbitrary metarefraction. Here it will be shown that this is not the case, and consider what restrictions apply.

A metarefraction may be limited in two ways: Firstly, as restrictions on the range of ray mappings that the metarefraction may perform. This could,

for example, be due to wave-optical restrictions on the light field as discussed in Chapter 6, or as engineering limitations of the constituent devices such as field of view limitations of confocal lenses.

Secondly, only a finite number of light-ray paths may be independently controlled. That is, the ray mapping is parametrized by a finite number of independent parameters. Varying each parameter simultaneously affects many ray paths, so selecting the direction into which an incident angle is refracted will, as a side effect, also offset other ray paths. In order to separately perturb each ray path, the affect on others must be cancelled out by other parameter perturbations.

In order to quantify how many degrees of freedom such a sheet can possess, consider that a confocal lenslet array is able to transform the local light ray direction two dimensionally according to Equation (7.7), and three dimensionally according to Equations (7.25) and (7.26). As was discussed in Chapter 7, there are a large number of parameters that describe individual confocal lens pairs, but many could be eliminated. Such a cylindrical structure possesses 3 independent parameters while an elliptical structure possesses 7 independent parameters in three dimensions.

Adding additional sheets contributes additional parameters, but many of these are redundant. Here, it will be shown that the combination of multiple sheets would offer at most only three additional parameters, regardless of how many sheets are added, and then only for generalized confocal elliptical-lenslet arrays. For the cylindrical-lenslet case, all additional parameters are fully dependent on those of a single sheet.

This chapter will be structured as follows: The next sections will explore the parameter and angle spaces associated with various series of generalized confocal lenslet arrays. After which, a numerical method will be introduced

which estimated the number of degrees of freedom that a series of generalized confocal lenslet arrays possess. Then, in two dimensions, the previous result will be confirmed analytically. This chapter will conclude with a brief summary and a discussion regarding the domain of series of such sheets.

9.2 Parameter space

A particular metarefraction, i.e. a function that maps an incident direction into a refracted direction, may be described by some collection of parameters, p_i. These then constitute a point, $\mathbf{p}$, in parameter space. These parameters include the effective refractive index, η, the offset, δ, and the lens orientation, θ, all of which were summarized in Chapter 7. The last two parameters are so similar to one another, that they are equivalent for small angles of refraction, however, they do transform the ray direction differently for larger angles.

The domain of a generalized confocal lenslet array is bound by various physical restrictions. The effective refractive index ratio, η, is minus unity times the ratio of focal lengths of the two lenses (i.e. Equation (7.6)). Both individual focal lengths are limited and therefore so too is the effective refractive index ratio. Similarly, the lens shift, δ, is limited to a range $[-1/2, +1/2]$ before less light passes through each lens into its corresponding lens than through the lens and the corresponding lens's neighbour (at least for small incident angles). Finally, rotation angle is the most restrictive of the parameters as the transmittance rapidly falls for larger rotations. Transmittance corresponds to the fraction of the aperture that will correctly transform incident light rays.

In the case of a two-dimensional generalized confocal lenslet array, param-

eter space is three-dimensional:

$$\mathbf{p} = (\eta, \delta, \theta). \tag{9.1}$$

The resulting metarefraction is then a function of input angle and this parameter vector:

$$\alpha' = \mathrm{T}(\mathbf{p}, \alpha). \tag{9.2}$$

Similarly in three dimensions parameter space is seven-dimensional:

$$\mathbf{p} = (\eta_x, \eta_y, \delta_x, \delta_y, \theta, \phi, \psi) \tag{9.3}$$

where (η_x, η_y) are the (possibly anisotropic) effective refractive indices along the x and y axes of the sheet while (δ_x, δ_y) are the dimensionless vector offset and θ, ϕ, ψ are Euler angles that describe the change of basis from the sheet coordinate system to that of each confocal lens pair. The parameter θ is significant only where $\eta_x \neq \eta_y$ as otherwise the lenses are rotationally symmetric about their normal.

The corresponding metarefraction may be seen as the mapping

$$\hat{\mathbf{r}}' = \mathbf{T}(\mathbf{p}, \hat{\mathbf{r}}). \tag{9.4}$$

More general components may also be realized by appending each sheet's parameters in turn. In multiple 2D sheets, this is

$$\mathbf{p} = (\eta_1, \delta_1, \theta_1, \eta_2, \delta_2, \theta_2, \ldots). \tag{9.5}$$

In order to implement multiple generalized confocal lenslet array metarefractions, one after the other, the resulting law of refraction, Equation (9.2) in 2D (or Equation (9.4) in 3D) is produced by the iterative application of the law of refraction for a single generalized confocal lenslet array, i.e. Equation (7.7) (or Equations (7.25) and (7.26) in 3D). For example, in 2D, and for two sheets, the law of metarefraction is

$$\mathrm{T}\left((\eta_1, \delta_1, \theta_1, \eta_2, \delta_2, \theta_2), \alpha\right) = \mathrm{T}\left((\eta_2, \delta_2, \theta_2), \mathrm{T}((\eta_1, \delta_1, \theta_1), \alpha)\right). \tag{9.6}$$

9.3 Angle space

The independence of metarefracting parameters will be discussed in terms of their effects on the metarefraction. While there are an infinite number of input and output angle pairs, it will be sufficient here to work with a finite set of angles. Each input angle has a corresponding output angle

$$\alpha_i' = \mathrm{T}(\mathbf{p}, \alpha_i). \tag{9.7}$$

Here it will be sufficient to consider a sample size of the order of tens of input angles to capture the most significant behaviour of the ray mapping. For a fixed set of representative input angles, the metarefraction may be approximately described by a corresponding set of output angles. It will prove to be convenient to consider this set as a vector of output angles, $\boldsymbol{\alpha}'$. Therefore, the output angle coordinate may be interpreted as being a point in angle space. (Input angles, α_i, are evenly sampled across the range $[-90°, +90°]$).

As was the case of the parameters, all angles are bound due to physical constraints. The angles are restricted to the working angles of the sheet. Here they will be taken to be any forward direction, but in practice will be further restricted. What is true in the discussion here for the entire hemisphere is no less restrictive for smaller angle ranges. In fact, the opposite is true: for example in the paraxial limit it may be shown that in two dimensions, a single generalized confocal lenslet array is well described by only two parameters.

9.4 Numerical analysis

This section will determine the number of ways in which different collections of ray paths may be independently varied. In order to do so, consider that a particular metarefraction may be approximately described by the output angle

vector, which is a function of the parameter vector. Consider a small perturbation in parameter space, δp_j. The resultant perturbation in angle space may be determined by expanding the metarefraction described by Equation (9.7) using the chain rule:

$$\delta\alpha_i' = \sum_{j=x,y,z} \frac{\partial \mathrm{T}(\mathbf{p}, \alpha_i)}{\partial p_j}\, \delta p_j. \tag{9.8}$$

This may be written as a matrix multiplication

$$\boldsymbol{\delta\alpha'} = \mathbf{J}\, \boldsymbol{\delta p} \tag{9.9}$$

where $\mathbf{J} = \boldsymbol{\partial T}/\boldsymbol{\partial p}$ is the Jacobian matrix, and the i^{th} element of $\mathbf{T}$ is $\mathrm{T}(\mathbf{p}, \alpha_i)$.

The effect of varying a single parameter on the output angle vector is quantified by the corresponding column vector of the Jacobian. The elements of this vector describe how each component in angle space varies as the parameter is perturbed. The column vector, multiplied by the corresponding perturbation, may therefore be interpreted as a displacement vector in angle space.

The perturbation of each parameter corresponds to a different displacement in angle space. Independence of parameters corresponds to independence of their corresponding displacement in angle space. The number of independent parameters then corresponds to the number of linearly independent angle-space vectors. The dimensionality of the column space of the Jacobian therefore corresponds to the number of linearly independent parameters. Given a large enough row space, the rank of the Jacobian may be used to count this.

To reiterate, when the rank of the Jacobian is calculated for a sufficiently large number of sample angles, it equals the number of dimensions that span the column space of the Jacobian. This in turn gives the minimum number of parameters that must be retained, without loosing flexibility in perturbing the

angle mapping. Perturbations are chosen randomly, but for a large enough number of random angle vectors, they should prove representative.

The Jacobian is calculated for a particular law of metarefraction. Here we are interested in comparing the various numbers of combinations of generalized confocal lenslet arrays. By comparing the Jacobian rank for different numbers of sheets, a relationship between number of parameters and number of constituent sheets may be determined.

The Jacobian was calculated analytically and evaluated for fifty evenly distributed incident angles within $[-90°, +90°]$ and for randomly chosen parameters. In order for the following numerical work to be comprehensive, larger parameter ranges were sampled over those physical ranges stated above. Both η and δ were taken to be within $[-5, +5]$ and rotation angles were selected such that no ray deviation was ever greater than $90°$ from its original orientation.

A Jacobian rank was calculated for each of up to five different nested generalized confocal lenslet array models. In each case, the metarefraction performed by a generalized confocal cylindrical-lenslet array, Equation (7.7), was recursively evaluated with the first three parameters, then the resulting angle was refracted again with the next three parameters. This was repeated until no more parameters were left.

The rank of the Jacobian was always found to be three. That is, there are three linearly independent parameters for a generalized confocal cylindrical-lenslet array regardless of how many subsequent sheets were added. A single sheet offered as many degrees of freedom as it is in principle possible to achieve.

The same analysis was then generalized to three dimensions, i.e. to the generalized confocal elliptical-lenslet array, where the space of output angles was reinterpreted as the space of ray direction components. Each output angle component is described by a function of the input angle vector and a

parameter vector. Angle perturbations may again be described in terms of corresponding perturbations in parameter space. Again fifty evenly spaced sample rays were considered (i.e. 150 ray direction components). The rank of the Jacobian was again calculated for a range of nested generalized confocal lenslet arrays.

This time, adding a second sheet did increase the seven parameters of a single sheet, but only to ten parameters. Furthermore, no amount of additional sheets further increased the number of independent parameters. A pair of sheets therefore offers as many degrees of freedom as can be realized with confocal lenslet arrays.

9.5 2D analytic proof

The number of independent parameters that describe a series of confocal lenslet arrays has already been measured numerically. In two dimensions, this may be confirmed analytically. Specifically, it is possible to analytically equate the parameters that describe the metarefraction performed by one and two generalized confocal lenslet arrays. First consider Equation (7.7), and in particular expand the tangents using the identity

$$\tan(\alpha - \beta) = \frac{\tan\alpha - \tan\beta}{1 + \tan\alpha\tan\beta}, \tag{9.10}$$

to get

$$\frac{\tan\alpha - \tan\zeta}{1 + \tan\alpha\tan\zeta} = \delta + \eta\frac{\tan\alpha' - \tan\zeta}{1 + \tan\alpha'\tan\zeta}. \tag{9.11}$$

From this, we may solve for $\tan\alpha'$ to get

$$\tan\alpha' = \frac{(-1 + \delta\tan\zeta - \eta\tan^2\zeta)\tan\alpha + (\tan\zeta + \delta - \eta\tan\zeta)}{(\tan\zeta + \delta\tan^2\zeta + \eta\tan\zeta)\tan\alpha + (\tan^2\zeta - \delta\tan\zeta - \eta)} \tag{9.12}$$

The pattern

$$\tan\alpha' = \frac{A\tan\alpha + B}{C\tan\alpha + D} \tag{9.13}$$

emerges, where the empirical parameters are

$$A(\eta, \delta, \theta) = -1 + \delta \tan\zeta - \eta \tan^2\zeta, \tag{9.14a}$$

$$B(\eta, \delta, \theta) = (1 - \eta) \tan\zeta + \delta, \tag{9.14b}$$

$$C(\eta, \delta, \theta) = (1 + \eta) \tan\zeta + \delta \tan^2\zeta, \tag{9.14c}$$

$$D(\eta, \delta, \theta) = -\eta - \delta \tan\zeta + \tan^2\zeta. \tag{9.14d}$$

The equivalence between one and two 2D generalized confocal lenslet arrays reduces to the equality

$$\alpha''(\eta', \delta', \theta'; \alpha'(\eta, \delta, \theta; \alpha)) = \alpha''(\eta'', \delta'', \theta''; \alpha). \tag{9.15}$$

The pair of transformations may be rewritten as

$$\tan\alpha' = \frac{A \tan\alpha + B}{C \tan\alpha + D}, \tag{9.16a}$$

$$\tan\alpha'' = \frac{A' \tan\alpha' + B'}{C' \tan\alpha' + D'}, \tag{9.16b}$$

in terms of primed and unprimed empirical parameters while the single transformation may be written as

$$\tan\alpha'' = \frac{A'' \tan\alpha + B''}{C'' \tan\alpha + D''}, \tag{9.17}$$

in terms of a third set of doubly primed empirical parameters.

All that then remains is to equate parameters that describe the two step transformation with those for the one step transformation,

$$A'' = A'\,A + B'\,C, \tag{9.18a}$$

$$B'' = A'\,B + B'\,D, \tag{9.18b}$$

$$C'' = C'\,A + D'\,C, \tag{9.18c}$$

$$D'' = C'\,B + D'\,D. \tag{9.18d}$$

This relation underlines the equivalence between one and two generalized confocal lenslet arrays. This implies that there are the same number of control parameters in either case, confirming the 2D numerical results.

This of course ignores the difference between the domains of Equations (9.16) and (9.18). That is, is there a value of A, B, C and D for every choice of A', B', C' and D' and vice versa in Equation (9.18).

9.6 Summary

Here it has been shown that an infinite series of confocal lenslet arrays cannot perform an arbitrary metarefraction. Furthermore, an upper limit on the number of degrees of freedom has been determined for any number of sheets. Only in three dimensions did adding a single additional sheet offer any extra independent parameters, and even then only three independent parameters were gained. In two dimensions, no additional independent parameters were gained.

These results have explored the number of independent parameters, and hence the number of light rays that may be separately redirected simultaneously, but the range of reachable combinations of outgoing rays was not explored. That is, physical limitations of sheet parameters restrict the possible values of the outgoing ray directions, but by adding additional parameters by combining multiple sheets, parameter space increases and as a result there may be combinations of outgoing ray directions that could not otherwise be achieved.

While not necessarily the most general possible means by which to realize a metarefraction, a pair of confocal elliptical-lenslet array contains the greatest (known) number of degrees of freedom of any thin sheet. Additionally, it achieves this with the minimum number of required components.

As an aside, were offsets to be included, multiple confocal lenslet arrays would offer a clear advantage. A consequence of allowing offsets would be that, within such a volume, a new form of transformation optics (discussed in Section 2.3) could be implemented. However, due to significant compound loss and complexity in fabrication and alignment, such a structure is currently impractical.

CHAPTER 10

Conclusion

I shall be telling this with a sigh
Somewhere ages and ages hence:
Two roads diverged in a wood, and I
I took the one less travelled by,
And that has made all the difference.

— *The Road Not Taken, by Robert Frost*

This chapter will gather together the various threads that have been explored in this thesis. It will be divided into three parts, each exploring a different aspect of the subject. The next section will summarize the exploration of metarefraction that has been presented above. After which, the progress made will be discussed in a wider context. Finally, a brief discussion regarding possible future directions of this work will be considered.

10.1 Summary

Initially, metarefraction was introduced in Chapter 1. This began with a definition of metarefraction and presented a number of examples. Throughout,

metarefracting sheets were interpreted as being analogous to refraction. The background chapter explored recent advances in the description of refraction, in particular the history of negative refraction as well as its realization and experimental demonstration with metamaterials.

Chapters 3 and 4 then explored two new forms of metarefraction: local-light-ray-direction flipping using Dove-prism sheets and local-light-ray-direction rotation using rotated Dove-prism sheets. Here, the analogy between a refracting interface and a metarefracting sheet was explored. The law of metarefraction at crossed Dove-prism sheets was equivalent to a sign change in refractive index across the interface. Furthermore, ray-rotation gave rise to an entirely new law of refraction.

A useful optical transformation that a number of metarefracting systems performed was geometrical imaging. In particular, pseudoscopic imaging was reviewed in two and three dimensions before the non-trivial imaging between two or more ray-rotation sheets was discussed in Chapter 5.

Notably, several idealized metarefractions lack a wave-optical description as they implicitly describe the gradients of inconsistent phase surfaces. While 'pixelated' equivalent transformations were entirely possible, the same is not true of some metarefractions that were everywhere-continuous. Chapter 6 firstly considered the translation of a ray-optical mapping into the language of wave optics. In the ray-optical limit, the ray-direction was shown to be irrotational so those ray-mappings that altered this aspect of an incident light field were wave-optically impossible.

The ideas explored in Chapter 6 implied that offset-free metarefraction is only superficially correct, i.e. that in many cases it may only ever be approximated by piecewise-smooth mappings. There is therefore an entire class of optical transformations that may be well described and approximated yet

have never been systematically explored until now.

In Chapter 7 another metarefracting component, the confocal lenslet array, was discussed. This is a structure consisting of series of telescopic lens pairs that may be extended into a great many variants. One parameter, for example, was interpreted as being analogous to a refractive index ratio between two media. This extended the metarefracting-sheet/refracting-interface analogy from a sign changing interface to any possible refractive index interface. A simple confocal cylindrical-lenslet array was then demonstrated experimentally in Chapter 8.

The penultimate chapter explored the limitations on transformations that a series of generalised confocal lenslet arrays could perform. Ultimately, their useful parameter space only grew so far which limited the range of transformations that could be performed by an arbitrary, offset-free series of confocal lenslet arrays.

10.2 Comments

There had only been scattered examples of metarefraction until the systematic exploration that I was involved in. The subject developed quite naturally and offered a range of topics worth investigating [37, 38, 40, 41, 59, 60, 72, 73] . Despite having lead to few applications, it has been a precursor to a range of other topics [48, 49, 50, 74]. Perhaps some of these will provide applications (which are currently missing) in order to motivate further work.

Of particular note are the range of visual distortions that may be realized by metarefractions. This subject is such that while being simple to describe, it produces quite unintuitive visual distortions, even for the simplest metarefractions. This has made the subject an aesthetically as well as an intellectually interesting topic to study.

10.3 Future work

There are any number of directions in which metarefracting sheets may be developed. The variety of unexplored metarefractions alone may describe any number of interesting visual distortions. Additionally, few inhomogeneous sheets have been explored. Constructing imaging systems from inhomogeneous sheets is significantly more flexible; resultant images, for example, are no longer restricted to 1:1 mappings.

Of particular interest would be the development of Dove-prism sheets, and to construct custom-built confocal lenslet arrays as optical novelties. Proof of principle demonstrations are necessary to justify their development. Lenslet array technology already exists that can realize ray-flipping, but the experiments performed so far were proof-of-principle demonstrations that used far from optimal designs. Principal aims would then be to improve the field of view and reduce the visibility and frequency of "parasitic images".

Metarefracting sheets were initially conceived as equivalent to generalized refraction across an interface. This analogy may be developed by considering not only ray direction, but also attenuation being a function of incident angle. An existing, trivial example of such an angle-specific attenuation is that of a light control film [75]. Such sheets consist of miniature Venetian blinds which are used to reduce glare on computer screens, and for privacy on bank machine displays. The effect of the sheets is to absorb all light from any angle greater than a cut-off angle.

The greatest generalization possible, while continuing to work in terms of light rays, is to allow significant ray-position offsets during transformation. Such a general ray mapping could for example be achieved with a pair of confocal lenslet arrays that focus on either side of a fibre bundle. Such a system is shown in Figure 10.1. By interconnecting fibres, a many-to-many

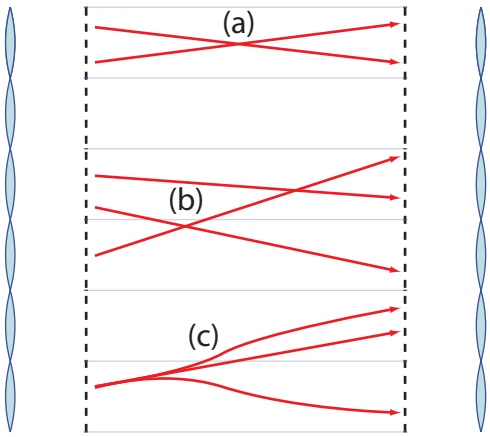

Figure 10.1: The most general form of ray position and direction mapping consists of rearranging the intensity pattern of an integral image using a fibre bundle. The side view of such a structure is shown schematically. Lenses focus the local plenoptic function onto one side of a fibre bundle. This is then mapped (shown as as red arrows) to the focal plane of the other lenslet array. (a) A metarefraction may be implemented by mapping incident angle to an outgoing angle. (b) Such a mapping may be extended to include large offsets, or (c) to include a many-to-many mapping by joining fibres.

mapping is possible over integral image spaces. A fibre bundle is, however, 'pixelated' both in position and in angle space. Additionally, the bundle would be too thick to be considered as a sheet.

Metarefraction may be developed in a number of ways, and there are any number of directions in which to direct future research. Examples include exploring the wave-optical behaviour of such components including how the polarization is affected, the study of optical quality that may be achieved by better understanding alternative light-ray paths and the search for ways to coherently realize metarefractions, i.e. offset free equivalents. While a number of the metarefractions have been thoroughly investigated, there are likely to be many more forms yet to be discovered. Only time will tell what visual distortions may be realized.

Bibliography

[1] Hecht. *Optics*, chapter 4, page 107. Adison-Wesely, 2nd edition, 1974.

[2] B. K. Johnson. *Optics and Optical Instruments*, chapter Practical optics, page 157. Dover, 1960.

[3] M. Born and E. Wolf. *Principles of Optics*, chapter App. I, pages 719–720. Pergamon Press, 6th edition, 1986.

[4] G. Stephenson and P. M. Radmore. *Advanced Mathematical Methods for engineering and science students*, chapter 1, pages 1–4. Cambridge University Press, 1990.

[5] C. Cason, T. Froehlich, N. Kopp, and R. Parker. POV-Ray – The Persistence of Vision Raytracer. http://www.povray.org/, 2003.

[6] V. G. Veselago. The electrodynamics of substances with simultaneously negative values of ϵ and μ. *Sov. Phys. Uspekhi*, 10:509–514, 1968.

[7] H. Benson. *University Physics*, chapter 35, page 713. John Wiley and sons inc., 1996.

[8] E. Hecht. *Optics*, chapter 3, pages 36–37. Addison-Wesley, 2nd edition, 1974.

[9] M. Born and E. Wolf. *Principles of Optics*, chapter 1, page 3. Pergamon Press, 1986.

[10] D. T. Moore. Gradient-index optics: a review. *Appl. Opt.*, 19:1035–1038, 1980.

[11] A. Bolton. The Sloan Lens ACS Survey, 2010.

[12] A. Taflove and S. C. Hagness. *Computational electrodynamics.* Artech House, 3rd edition, 2005.

[13] A. J. Ward and J. B. Pendry. Refraction and geometry in Maxwells equations. *J. Modern Opt.*, 43:773–793, 1996.

[14] A. J. Ward and J. B. Pendry. Calculating photonic greens functions using a nonorthogonal finite-difference time-domain method. *Phys. Rev. B*, 58:7252 – 7259, 1998.

[15] D. M. Shyrok. Note on transformation to general curvilinear coordinates for Maxwells curl equations. arXiv.physics/0307029v1 [physics.optics], 2003.

[16] E. Hecht. *Optics*, chapter 1, pages 1–3. Adison-Wesely, 2nd edition, 1974.

[17] L. D. Landau and E. M. Lifshitz. *The classical theory of fields*, chapter 10, pages 275–277. Elsevier, 1975.

[18] A. Einstein. On the influence of gravity on the propagation of light. *Annalen der Physik*, 35:898, 1911.

[19] U. Leonhardt and T. G. Philbin. General relativity in electrical engineering. *New J. Phys.*, 8:247, 2006.

[20] H. Ma, S. Qu, Z. Xu, and J. Wang. Wave-shape-keeping media. *Opt. Lett.*, 34:127–129, 2009.

[21] A. Greenleaf, Y. Kurylev, M. Lassas, and G. Uhlmann. Electromagnetic wormholes and virtual magnetic monopoles from metamaterials. *Phys. Rev. Lett.*, 99:183901, 2007.

[22] H. Chena and C. T. Chan. Acoustic cloaking in three dimensions using acoustic metamaterials. *Appl. Phys. Lett.*, 91:183518, 2007.

[23] S. A. Cummer, B. Popa, D. Schurig, D. R. Smith, J. Pendry, M. Rahm, and A. Starr. Scattering theory derivation of a 3D acoustic cloaking shell. *Phys. Rev. Lett.*, 100:024301, 2008.

[24] D. Schurig, J. Pendry, and D. Smith. Calculation of material properties and ray tracing in transformation media. *Opt. Express*, 14:9794–9804, 2006.

[25] U. Leonhardt. Optical conformal mapping. *Science*, 312:1777–1780, 2006.

[26] D. Schurig, J. J. Mock, B. J. Justice, S. A. Cummer, J. B. Pendry, A. F. Starr, and D. R. Smith. Metamaterial electromagnetic cloak at microwave frequencies. *Science*, 314:977, 2006.

[27] C. Ye and R. R. McLeod. Grin lens and lens array fabrication with diffusion-driven photopolymer. *Opt. Lett.*, 33:2575–2577, 2008.

[28] W. Rotman. Plasma simulation by artificial dielectrics and parallel-plate media. *IRE Trans. Antennas Propag.*, 10:82–95, 1962.

[29] R. Shelby, D. Smith, and S. Schultz. Experimental verification of a negative index of refraction. *Science*, 292:77–79, 2001.

[30] Ş. Selçuk Bayin. *Essentials of mathematical methods in science and engineering*. Wiley, 2008.

[31] J. B. Pendry. Negative refraction makes a perfect lens. *Phys. Rev. Lett.*, 85:3966–3969, 2000.

[32] Hecht. *Optics*. Adison-Wesely, 2nd edition, 1974.

[33] N. Fang and X. Zhang H. Lee, C. Sun. Subdiffraction-limited optical imging with a silver superlens. *Science*, 308:534–537, 2005.

[34] N. J. Wade. Philosophical instruments and toys — optical devices extending the art of seeing. *J. Hist. Neurosci.*, 13:102–124, 2004.

[35] T. Lian and M. Chang. New types of reflecting prisms and reflecting prism assembly. *Opt. Eng.*, 35:3427–3431, 1996.

[36] Robert A. Watkins. Multiple Dove prism assembly. US patent number 6,097,554, 2000.

[37] J. Courtial and J. Nelson. Ray-optical negative refraction and pseudoscopic imaging with Dove-prism arrays. *New J. Phys.*, 10:023028, 2008.

[38] A. C. Hamilton and J. Courtial. Optical properties of a Dove-prism sheet. *J. Opt. A: Pure Appl. Opt.*, 10:125302–125302, 2008.

[39] R. Penrose. *The road to reality*, chapter 8.3, pages 142–145. Vintage books, 2005.

[40] A. C. Hamilton, J. Nelson B. Sundar, and J. Courtial. Local light-ray rotation. *J. Opt. A: Pure Appl. Opt.*, 11:085705, 2009.

[41] A. C. Hamilton and J. Courtial. Imaging with parallel ray-rotation sheets. *Opt. Express*, 16:20826–20833, 2008.

[42] J. Morgan. *Introduction to Geometrical and Physical Optics*. McGraw-Hill, 1953.

[43] A. Hicks. Controlling a ray bundle with a free-form reflector. *Opt. Lett.*, 33:1672–1674, 2008.

[44] C. Paterson. *Computer-generated difractive optics with spiral phase dislocations*. PhD thesis, Imperial College, 1996.

[45] C . Paterson. Diffractive optical elements with spiral phase dislocations. *J. Mo*, 41:757–765, 1994.

[46] L. D. Landau and E. M. Lifshitz. *The classical theory of fields*, chapter 53, pages 140–146. Elsevier, 4th revised english edition edition, 1975.

[47] M. Born and E. Wolf. *Principles of Optics*, chapter 3, pages 110–114. Pergamon Press, 1986.

[48] A. C. Hamilton and J. Courtial. Metamaterials for light rays: ray optics without wave-optical analog in the ray-optics limit. *New J. Phys.*, 11:013042–013042, 2009.

[49] M. R. Dennis, A. C. Hamilton, and J. Courtial. Superoscillation in speckle patterns. *Opt. Lett.*, 33:2976–2978, 2008.

[50] M. V. Berry and M. R. Dennis. Natural superoscillations in monochromatic waves in D dimensions. *J. Phys. A*, 44:022003, 2009.

[51] L. D. Landau and E. M. Lifshitz. *The classical theory of fields*, chapter 48, page 125. Elsevier, 4th revised english edition edition, 1975.

[52] P. Lorrain, D. R. Corson, and F. Lorrain. *Electromagnetic fields and waves*. Freeman, 3rd edition, 1996.

[53] M. A. Piestrup. Large area x-ray and neutron imaging using three-dimensional arrays of microlenses. *Rev. Sci. Instrum.*, 75:4769–4774, 2004.

[54] R. H. Anderson. Close-up imaging of documents and displays with lens arrays. *Appl. Opt.*, 18:477–484, 1979.

[55] Vesselin Shaoulov, Ricardo Martins, and Jannick P. Rolland. Compact microlenslet-array-based magnifier. *Opt. Lett.*, 29:709–711, 2004.

[56] G. Lippmann. La photographie intégrale. *C. R. Séances Acad. Sci.*, 146:446–51, 1908.

[57] C. Hembd-Solner, R. F. Stevens, and M. C. Hutley. Imaging properties of the gabor superlens. *J. Opt. A: Pure Appl. Opt.*, 1:94–102, 1999.

[58] M. C. Hutley, R. Hunt, R. F. Stevens, and P. Savander. The moiré magnifier. *Pure Appl. Opt.*, 3:133–142, 1994.

[59] J. Courtial. Ray-optical refraction with confocal lenslet arrays. *New J. Phys.*, 10:083033, 2008.

[60] A. C. Hamilton and J. Courtial. Generalized refraction using lenslet arrays. *J. Opt. A: Pure Appl. Opt.*, 11:065502, 2009.

[61] B. C. Platt and R. Shack. History and principles of Shack-Hartmann wavefront sensing. *J. Refract. Surg.*, 17:S573–S577, 2001.

[62] Ming-Hsien Wu and G. M. Whitesides. Fabrication of two-dimensional arrays of microlenses and their applications in photolithography. *J. Micromech. Microeng.*, 12:747–758, 2002.

[63] Ph Nussbaum, R. Volkely, H. P. Herzig, M. Eisner, and S. Haselbeck. Design, fabrication and testing of microlens arrays for sensors and microsystems. *Pure Appl. Opt.*, 6:617636, 1997.

[64] Keisuke Araki. Compound eye systems for nonunity magnification projection. *Appl. Opt.*, 29:4098–4104, 1990.

[65] K. Hermans, S. Z. Hamidi, A. B. S., C. W. M. Bastiaansen, and D. J. Broer. Rapid, direct fabrication of antireflection-coated microlens arrays by photoembossing. *Appl. Opt.*, 47:6512–6517, 2008.

[66] R. Ng, M. Levoy M. Brédif, G. Duval, M. Horowitz, and P. Hanrahan. Light field photography with a hand-held plenoptic camera. *Standford Tech Report*, 02:1–11, 2005.

[67] A. V. Kanaev, D. A. Scribner, J. R. Ackerman, and E. F. Fleet. Analysis and application of multiframe superresolution processing for conventional imaging systems and lenslet arrays. *Appl. Opt.*, 46:4328, 2007.

[68] M. Fuchs, R. Raskar, H. Seidel, and H. P. A. Lensch. Towards passive 6D reflectance field displays. In *SIGGRAPH '08: ACM SIGGRAPH 2008 papers*, pages 1–8, New York, NY, USA, 2008. ACM.

[69] C. B. Burckhardt. Optimum parameters and resolution limitation of integral photography. *JOSA*, 58:71–76, 1968.

[70] Ren Ng. Fourier slice photography. *International Conference on Computer Graphics and Interactive Techniques*, 24:735 – 744, 2005.

[71] H. Goldstein, C. Poole, and J. Safko. *Classical mechanics*, chapter 4.4, pages 150–154. Addison Wesley, 1992.

[72] J. Courtial. Standard and non-standard metarefraction with confocal lenslet arrays. *Opt. Comm.*, 282:2634–2641, 2009.

[73] M. Blair, L. Clark, A. Houston, G. Smith, J. Leach, A. C. Hamilton, and J. Courtial. Experimental demonstration of a light-ray-direction-flipping METATOY based on confocal lenticular arrays. *Opt. Comm.*, 282:4299–4302, 2009.

[74] B. Sundar, A. C. Hamilton, and J. Courtial. Fermats principle and the formal equivalence of local light-ray rotation and refraction at the interface between homogeneous media with a complex refractive index ratio. *Opt. Lett.*, 34:374–376, 2009.

[75] J. Morimoto, S. J. Lee, and H. Y. Choi. Light control film. US patent number 2006/0261318, 2006.