TUDelft

Delft University of Technology

Ray-based methods for simulating aberrations and cascaded diffraction in imaging systems

Mout, Marco

DOI

Publication date
2019

Document Version
Final published version

Important note
To cite this publication, please use the final published version (if applicable).
Please check the document version above.

RAY-BASED METHODS FOR SIMULATING ABERRATIONS AND CASCADED DIFFRACTION IN IMAGING SYSTEMS

Ray-based methods for simulating aberrations and cascaded diffraction in imaging systems

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology
by the authority of the Rector Magnificus, prof. dr. ir. T.H.J.J. van der Hagen,
Chair of the Board for Doctorates
to be defended publicly on
Friday 12 July 2019 at 10:00 o'clock

by

Beat Marco Mout

Master of Science in Applied Physics, Delft University of Technology, the Netherlands
born in Hengelo, the Netherlands

This dissertation has been approved by the promotors

Composition of the doctoral committee:
Rector Magnificus   chairperson
Prof. dr. H. P. Urbach   Technische Universiteit Delft, promotor
Dr. F. Bociort   Technische Universiteit Delft, copromotor

Independent members:
Prof. dr. H. Gross   Friedrich-Schiller-Universität Jena, Germany
Prof. dr. ir. W.L. IJzerman Technische Universiteit Eindhoven
Prof. dr. ir. C. Vuik   Technische Universiteit Delft
Prof. dr. F. Wyrowski   Friedrich-Schiller-Universität Jena, Germany
Prof. dr. B. Rieger   Technische Universiteit Delft, reserve member

Other members:
Dr. A. Flesch   Robert Bosch GmbH, Germany

The beauty of this is its simplicity.
If the plan gets too complex something always goes wrong.

Joel & Ethan Coen

CONTENTS

SUMMARY

In designing an optical instrument, computer simulations can help to better understand and improve the system. A commonly used simulation method is the exit pupil diffraction integral (EPDI). It uses ray tracing to propagate the light from the object until the exit pupil in the final medium, thereby accounting for the full geometry and all aberrations of the optical system. The propagation from the exit pupil to the focal area is performed using a wave optical propagator. The EPDI accurately describes the effects of aberrations, but assumes all the diffraction to occur at the exit pupil. This exit pupil approximation fails for systems that experience a cascade of diffraction effects. In this thesis, we are primarily interested in simulating such systems without neglecting the effects of aberrations. This requires a simulation method that goes beyond the exit pupil approximation.

To this purpose we investigate two Monte Carlo ray-based diffraction methods. The advantage of ray-based methods is that they accurately account for the aberrations of the optical system. Ray-based diffraction simulations start by initiating rays at the input plane. For Monte Carlo methods the set of all possible rays is sampled randomly. These rays are then traced through part of the optical system until they reach a diffracting surface like an aperture. At this surface secondary rays are initiated based on a physical principle that is specific to the ray-based diffraction method at hand. These secondary rays, which in a Monte Carlo method are again sampled randomly, are then traced through the next part of the system. At every diffracting surface new rays are initiated and then traced further through the system until the rays reach the plane of interest, where they are used to calculate the electric field. Although the procedures for ray-based diffraction simulations are very similar, the theory behind the methods and the quantity carried by the (secondary) rays differ between the methods. In this thesis we focus on one method that transports the Wigner distribution function using rays and one method that combines the Huygens-Fresnel principle with ray tracing.

The Wigner distribution function (WDF) can be used as a mathematical representation of the scalar electric field in spatial and frequency coordinates. Upon propagation through a paraxial optical system (i.e., a system where all rays remain close to the axis of symmetry) the WDF undergoes a coordinate transformation. If one relates a ray to every point in the four dimensional coordinate system of the WDF (two position and two frequencies or direction coordinates), this ray undergoes the same coordinate transformation as the WDF. In other words, the WDF along a paraxial ray remains constant.

In Wigner-based ray tracing (WBRT) this property is used to transport the WDF using rays. The intensity can be calculated by summing the WDF of all the rays arriving in a pixel. These WDF values are real but not necessarily positive. The idea of WBRT is to capture the diffraction effects by use of the WDF while capturing the full geometry of the system by using ray tracing. We show that WBRT can accurately model lenses with low numerical apertures, but that it can lead to inaccurate and even unphysical

results in the presence of aberrations. The latter is explained using an analytical model of WBRT. Although our results do not allow drawing conclusions about Wigner-based methods in general, they do demonstrate that WBRT is unsuited for the simulation of both diffraction and aberrations.

Huygens-Fresnel path integration (HFPI), the second method presented in this thesis, combines the Huygens-Fresnel principle with ray tracing. The Huygens-Fresnel principle states that the propagation of light can be described by placing point sources on a wavefront and summing the fields of all these point sources. HFPI applies this principle at surfaces that introduce diffraction (e.g., apertures) and uses Monte Carlo ray tracing to propagate the fields of the point sources. The rays in an HFPI simulation have a phase determined by their optical path length. The field in a pixel is the sum of the complex amplitudes of all the rays passing through the pixel. By using vectorial amplitudes, HFPI can also account for polarization.

The Huygens-Fresnel principle gives an intuitive understanding of HFPI, but for a more rigorous theory we turn to a physical interpretation of diffraction integrals. Scalar HFPI is based on the Rayleigh-Sommerfeld diffraction integral, vectorial HFPI on the m-theory integrals. We show that these integrals can be interpreted as a decomposition of the field in a plane into (magnetic) dipole fields. These secondary fields are of such nature that they can be propagated using geometrical optics.

We simulate the propagation of plane waves, spherical waves, and a Gaussian beam through a range of (aberrated) systems with one or more diffracting surfaces. The HFPI results are shown to be in good agreement with those of reference methods. We explain that the use of a Monte Carlo ray tracer causes a small systematic error and a statistical error whose variance is inversely proportional to the number of rays. A big advantage of the Monte Carlo procedure is that it eliminates the need for an explicit (re)sampling.

HFPI is a robust method that can account for aberrations and cascaded diffraction. It is relatively easy to implement if one has access to a Monte Carlo ray tracer that can keep track of optical path lengths. A disadvantage of the method is that it is computationally demanding, but this drawback is strongly reduced by the perfect scalability of the method on modern parallel architectures (e.g., cloud computing). We also present several techniques to reduce the computational demand of the method. One of these associates a plane wave to every ray in the focal region, which reduces the simulation times by about two orders of magnitude.

With HFPI we introduce a method that can be used to simulate the propagation of coherent fields (e.g., laser beams) through optical systems for which the aberrations cannot be neglected and where diffraction effects are introduced at one or more surfaces. We believe that the method can aid in developing new optical instruments.

SAMENVATTING

Bij het ontwikkelen van optische systemen kunnen computersimulaties helpen om de systemen te begrijpen en te verbeteren. Een daarvoor veel gebruikte simulatiemethode is de uittreepupil diffractie-integraal (EPDI). Deze methode transporteert het licht, vanaf het object naar de uittreepupil in het laatste medium van het optische systeem, met behulp van stralen. De paden van de stralen worden door de geometrie van het optische systeem bepaald waardoor de aberraties van het systeem correct worden gesimuleerd. Vanaf de uittreepupil wordt het licht verder naar het brandpunt getransporteerd met behulp van een fysisch-optische methode. De EPDI neemt de effecten van aberraties correct mee, maar de methode veronderstelt dat alle diffractie zijn oorsprong heeft in de uittreepupil. Deze uittreepupilbenadering is onjuist voor systemen met meervoudige diffractie. In dit proefschrift concentreren wij ons op het simuleren van systemen met meervoudige diffractie zonder daarbij de effecten van aberraties te verwaarlozen. Daarvoor zijn simulatiemethodes nodig die geen gebruik maken van de uittreepupilbenadering.

Hiertoe onderzoeken we twee straaloptische Monte-Carlo-diffractiemethoden. Het voordeel van straaloptische methoden is dat zij de aberraties van optische systemen correct modelleren. Bij straaloptische diffractiesimulaties worden allereerst in het startvlak van het systeem stralen geïnitieerd. Monte-Carlo-methoden doen dit door willekeurig stralen te kiezen uit de set van alle mogelijke stralen. Daarna worden deze stralen tot het eerste diffractieve oppervlak (bijvoorbeeld een apertuur) van het systeem getransporteerd, waar vervolgens secundaire stralen worden gestart. Dit gebeurt op basis van een fysisch model dat per methode verschilt. Monte-Carlo-simulaties genereren deze stralen met een toevalsgenerator. De stralen vervolgen daarna hun weg door het systeem, waarbij in elk diffractief vlak de inkomende stralen stoppen en nieuwe stralen starten. Met behulp van de stralen die het laatste oppervlak van het systeem (bijvoorbeeld de sensor) bereiken, wordt het lokale elektrische veld berekend. Deze afloop van een diffractie-simulatie is weliswaar hetzelfde voor de verschillende straaloptische methoden, maar de achterliggende theorie en de waarde die door de (secundaire) stralen wordt getransporteerd verschillen van methode tot methode. In dit proefschrift onderzoeken we een methode die de Wignerdistributie met behulp van stralen transporteert en een methode die het principe van Huygens-Fresnel combineert met een straaloptische simulatie.

De Wignerdistributie (WDF) is een reële functie waarmee we een scalair elektrisch veld kunnen beschrijven middels coördinaten in het ruimtelijk en het frequentie domein. Een bijzondere eigenschap van deze functie is dat het effect van de voortplanting door een paraxiaal optisch systeem (oftewel een systeem waarin alle stralen dicht bij de symmetrieas blijven) middels een coördinatentransformatie kan worden beschreven. De vierdimensionale coördinaten van de WDF kunnen we ook gebruiken om de positie en de richting van een straal te beschrijven. Bij de voortplanting door een paraxiaal systeem ondergaat deze straal dezelfde coördinatentransformatie als de WDF. Met andere

woorden, langs de baan van een paraxiale straal blijft de WDF constant.

De Wigner straalsimulatie (WBRT) gebruikt deze eigenschap om de WDF met behulp van stralen te transporteren. De intensiteit kunnen we berekenen door de WDF van alle stralen die op een pixel vallen bij elkaar op te tellen. Deze WDF-waarden zijn reëel, maar niet altijd positief. Het concept achter WBRT is om de diffractie middels de WDF en de geometrie van het optische systeem middels de straalsimulatie te modelleren. In dit proefschrift tonen wij aan dat men met WBRT systemen met lage numerieke aperturen correct kan simuleren, maar dat de methode onjuiste en zelfs onfysische resultaten kan genereren wanneer het systeem aberraties heeft. Dit laatste verklaren we nader met behulp van een analytisch model van WBRT. Op basis van onze resultaten zijn weliswaar geen algemene conclusies te trekken over methoden die de Wignerdistributie gebruiken, maar de resultaten tonen wel eenduidig aan dat WBRT niet geschikt is voor het simuleren van zowel diffractie als aberraties.

Huygens-Fresnel pad integratie (HFPI), de tweede methode die we in dit proefschrift behandelen, combineert een straaloptische voortplanting met het principe van Huygens-Fresnel. Volgens dit principe kan de propagatie van licht gemodelleerd worden door puntbronnen op het golffront te plaatsen en de velden van deze puntbronnen bij elkaar op te tellen. HFPI past dit principe toe op locaties waar diffractie ontstaat (bijvoorbeeld in aperturen). De velden van de puntbronnen worden met behulp van een straaloptische Monte-Carlo-methode getransporteerd, waarbij de fase van elke straal door zijn optische weglengte wordt bepaald. We kunnen het elektrische veld berekenen door per pixel de complexe amplitudes van de stralen bij elkaar op te tellen. Wanneer we de stralen een vectoriële amplitude geven houden we ook rekening met polarisatie.

Het principe van Huygens-Fresnel helpt om het concept van HFPI te begrijpen. De theoretische onderbouwing van de methode baseert echter op een fysische interpretatie van diffractie integralen. Scalaire HFPI is gebaseerd op het Rayleigh-Sommerfeld diffractie integraal en vectoriële HFPI op de integralen van de m-theorie. We laten zien dat men deze integralen kan interpreteren als een decompositie van het elektromagnetische veld in een vlak in (magnetische) dipoolvelden. Deze secundaire velden kunnen vervolgens met behulp van geometrische optica worden getransporteerd.

In dit proefschrift simuleren we de voortplanting van vlakke golven, bolgolven en een Gauss-straal door verschillende systemen met en zonder aberraties en met één of meerdere diffractieve vlakken. De resultaten van HFPI komen daarbij goed overeen met die van referentiemethoden. Het gebruik van een straaloptische Monte-Carlo-methode resulteert in een kleine systematische fout en een statische fout waarvan de variantie omgekeerd evenredig is met het aantal stralen. Een groot voordeel van de Monte-Carlo-methode is dat er geen expliciete bemonstering nodig is.

HFPI is een robuuste methode waarmee men de effecten van aberraties en meervoudige diffractie correct kan simuleren. De methode is relatief eenvoudig te implementeren mits men over een straaloptische Monte-Carlo-simulator beschikt die de optische weglengte van stralen kan berekenen. Een nadeel van de methode is dat zij veel rekentijd vraagt. Daar staat tegenover dat HFPI bij uitstek geschikt is voor parallellisatie op moderne computerarchitectuur (bijvoorbeeld cloud computing). Bovendien presenteren we meerdere methoden om de rekentijd voor HFPI te verminderen. Bij een van deze methoden wordt elke straal van een vlakke golf voorzien. Dit reduceert de rekentijd met

ongeveer twee ordegroottes.

Met HFPI, de methode ontwikkelt in dit proefschrift, kan men de voortplanting van coherente velden (bijvoorbeeld een laserstraal) door optische systemen simuleren, ook wanneer meerdere vlakken in het systeem diffractie veroorzaken. Vanwege deze toegevoegde waarde denken we dat de methode een rol kan spelen bij het ontwikkelen van nieuwe optische systemen.

1

INTRODUCTION

This thesis addresses simulation methods for the propagation of light through an optical system (e.g., a microscope). Such simulations have to describe two types of effects: aberrations and diffraction. Section 1.2 introduces these effects and the simulation methods that are used to describe them. The heart of this thesis, as explained in Section 1.3, is the introduction and analysis of two methods that aim to describe both aberrations and diffraction within a single ray-based simulation. This chapter only provides a brief introduction to this thesis. An outline of the rest of the thesis is given in Section 1.4. Before starting the technical discussion, Section 1.1 briefly places the research in the context of the mutually reinforcing advances in science and optical instruments.

1.1. SHORT HISTORIC PERSPECTIVE

Optical instruments have enabled many scientific breakthroughs, and science has enabled many breakthroughs in optical instruments. Perhaps one of the most prominent examples of this bidirectional influence is the work by Johannes Kepler. Using his observations through a telescope he determined the laws of planetary motion, and using his observation of the telescope he wrote Dioptrice, a book with great influence on the theory and practice of telescope making [1]. Many of the other great scientists of the seventeenth century (e.g., Isaac Newton and Christiaan Huygens) used optical instruments for their research and contributed to the development of optics and optical instruments. Until the nineteenth century the production of optical instruments still mainly relied on trial-and-error methods [2]. During the eighteen hundreds, scientists like Gauss, Seidel, Fraunhofer, Maxwell, and Abbe contributed to a better understanding of optical instruments and devised ways to build them according to mathematical calculations. Nowadays, the development of optical instruments is guided by such calculations and numerical simulations. Much like advances in science and in optical instruments mutually reinforced each other, the improvement in optical lithography (a crucial step in the manufacturing of computer chips) led to more computational power which was used to improve lithographic machines. These developments enabled the optical simulation

1

methods presented in this thesis, which in turn may enable the development of new optical instruments and machines.

1.2. SIMULATING ABERRATIONS AND DIFFRACTION IN OPTICAL SYSTEMS

When designing a technical product, computer simulations help to better understand and improve the design without having to build a prototype. For optical systems one can distinguish between two phenomena that affect the optical performance of the system: aberrations and diffraction. In aberration-free imaging systems, all light rays coming from an object point will be concentrated in an image point. This is called perfect imaging within geometrical optics. In real optical systems, which use lenses and mirrors to concentrate rays, perfect imaging is usually not achieved and the rays are spread out over a small region. The way the rays are distributed is described by the aberrations of the system. In reality, light is a wave and rays are just a simplified model to understand how light propagates. As a result, light can never be concentrated in an infinitely small point, even if all the rays pass through a single point as depicted in Fig. 1.1a. This deviation from geometrical optics is what is known as diffraction.

(a) Ray tracing through a singlet.

(b) A wave optical simulation of a singlet.

Figure 1.1: A comparison between a geometrical and a wave optical simulation of the propagation of light through a singlet. According to geometrical optics, when there are no aberrations all rays intersect at a single point. This causes a very local energy concentration in this perfect image point. In reality, the light is spread out over a small region due to diffraction. The wave optical simulation shows this diffraction effect, but, compared to geometrical optics, most wave optical methods require a stronger simplification of the optical system. In the depicted wave-optical model the lens is replaced by a thin plate that approximates its effect.

Aberrations and diffraction are usually treated separately in optical simulations. The aberrations are simulated by tracing rays and diffraction is described using a wave optical model. Ray tracing can be used to propagate light through a series of lenses using a simplified model of light propagation [3] which neglects part of the wave nature of light, i.e., diffraction. Wave optical models (e.g., the Collins integral [4]) on the other hand need a simplified model of the optical system (as illustrated in Fig. 1.1b), because the wavelength of light is extremely small compared to the size of optical components.

Most simulation methods can therefore only accurately describe either aberrations or diffraction. Some methods can, to some extent, switch between the ray and the wave model. The exit pupil diffraction integral [5] for instance uses ray tracing to propagate light through all the lenses of an optical system until the final uniform medium. It then uses a wave optical method to propagate the light through this medium to the image plane. This method assumes that all diffraction effects are only important for this final propagation step. For some systems, this assumption is not valid.

It was already mentioned that due to diffraction, light cannot be concentrated in an image point. Another diffraction effect is that light to some extent bends around corners. If you illuminate a small hole, the light will not only continue straight through the hole (as the ray model predicts), but will radiate in all directions. Therefore, if we simulate a system with multiple small holes (i.e., apertures) behind each other we observe a cascade of diffraction effects. This is not compatible with the assumptions of the exit pupil diffraction integral. Therefore cascaded diffraction has often been analyzed for systems that were strongly simplified or had very specific optical properties (e.g., in [6, 7]). It is the aim of this thesis to develop a simulation method for cascaded diffraction that also accurately describes the aberrations of the system.

1.3. Two ray-based diffraction methods

This thesis is an investigation into two methods for simulating aberrations and cascaded diffraction. Both methods are Monte Carlo ray-based diffraction methods. Ray-based diffraction methods use ray tracing and can therefore propagate light through a series of lenses or mirrors. In order to describe diffraction, these methods initiate secondary rays at the surfaces that introduce diffraction (e.g., an aperture), as shown Fig. 1.2. Some

Figure 1.2: A graphical representation of a ray-based diffraction simulation. Rays are traced until the aperture where they initiate secondary sources to describe the wave optical effect of diffraction.

prior ray-based diffraction methods, used as reference methods in this thesis, are the Gaussian beam decomposition [8], Beam Synthesis Propagation [9] in the commercial software package CODE V, and the vectorial ray-based diffraction integral [10]. Contrary to those methods, the methods developed in this thesis are Monte-Carlo methods. This indicates that the sampling is random, which means that the starting position and/or direction of every ray is chosen randomly. The common simulation procedure for Monte Carlo ray-based diffraction simulations is depicted in the flow chart shown in Fig. 1.3.

The two methods investigated in this thesis are Wigner-based ray tracing (WBRT) and Huygens-Fresnel path integration (HFPI). Although the procedures for these simulation tools are very similar, the theory behind the methods and the quantity carried by

Figure 1.3: A flow chart illustrating the steps of a Monte Carlo ray-based diffraction method like Huygens-Fresnel path integration or Wigner-based ray tracing. The methods differ in the quantities carried by the rays.

the (secondary) rays differ between the methods.

In WBRT, rays transport the Wigner distribution function. This distribution function [11] is a transformation that can be applied to the electric field (i.e., to light). It has the special property, that for many systems it follows the same propagation rules as geometrical optical rays, yet it still accounts for diffraction effects [12]. The idea of Wigner-based ray tracing (WBRT) is therefore to use ray tracing to transport the Wigner distribution function.

The second method, HFPI, relies on the Huygens-Fresnel principle. This principle states that the propagation of light can be described, by placing point sources on a wavefront of the light and summing the fields of all these point sources [13, p. xxvii-xxviii]. HFPI applies this principle at surfaces that introduce diffraction and uses ray tracing to propagate the point sources.

The research presented in this thesis is the first extensive study of Wigner-based ray tracing and Huygens-Fresnel path integration for the simulation of optical systems. WBRT had been applied in the field of computer graphics, where it was used to render images that included diffraction effects (e.g., of a light reflecting of a CD [14]). We implement the method, apply it to the simulation of optical systems, and compare the results to those of other physical models. Our numerical results and analytical model show that WBRT is unable to simultaneously describe the effects of aberrations and diffraction.

The domain of validity of HFPI turns out to be much broader than that of WBRT. This thesis describes the numerical procedure and provides the theoretical foundation for HFPI. The method is developed for both scalar and vectorial optics, and is validated against multiple reference models for both cases. It is a robust method that can accurately describe the effects of aberrations and diffraction beyond the exit pupil approximation.

1.4. OUTLINE OF THIS THESIS

Since the topic of this thesis is the simulation of light propagation, Chapter 2 gives an overview of the most important propagation methods in the context of this thesis. It starts with the description of diffraction integrals (i.e., the m-theory [15], Rayleigh-Sommerfeld diffraction integral, and the Collins integral [4]), a commonly used wave optical model for the propagation of light through a uniform medium or a strongly idealised optical system. We then present geometrical optics and ray tracing, its numerical implementation. Ray tracing is the main method for simulating the effect of the shape

of lenses and mirrors in optical systems. Since the theoretical basis for Huygens-Fresnel path integration is the combination of a diffraction integral and geometrical optics, these models are explained in detail. The chapter continues with methods that are used for reference simulations in later chapters: the exit pupil diffraction integral [5] and some ray-based diffraction methods [8–10]. The final topic of Chapter 2 is the Wigner distribution function [11], which is the basis of Wigner-based ray tracing.

The three following chapters treat the methods which are the core topic of this thesis: Wigner-based ray tracing in Chapter 3, the scalar version of Huygens-Fresnel path integration in Chapter 4, and its extension to vectorial optics in Chapter 5.

Wigner-based ray tracing relies heavily on the theory presented in Section 2.7, so the method can be concisely presented in Chapter 3. The chapter continues by comparing the results of WBRT to those of reference simulations for the propagation through free-space, a lens, and an idealised so called thin lens. An analytical model of WBRT is used to support the findings of the numerical model. The chapter finishes with a discussion of the finding that WBRT produces unphysical results in the presence of aberrations.

The topic of Huygens-Fresnel path integration is split over two chapters: Chapter 4 treats the scalar case and Chapter 5 presents the method for vectorial optics. Both chapters follow roughly the same outline as the chapter on WBRT: treating first the theory behind the method and its numerical implementation, then comparing its performance to that of reference methods, and concluding with a discussion on the method and its applications. Furthermore, Chapter 4 also contains a section on strategies to reduce the computational load of HFPI. The most useful of these strategies is also applied to vectorial HFPI in Chapter 5.

Huygens-Fresnel path integration is probably most easily understood if one reads Chapter 4. On the other hand, Chapter 5 on vectorial HFPI offers a more fundamental derivation of the method. Since scalar and vectorial HFPI share most of their properties, these properties are only listed once in the discussion in Section 4.6. The discussion of the chapter on vectorial HFPI is used to compare the concept of the method to that of other simulation tools. If the reader is going over this outline to find how to get the most out of this thesis with the smallest effort, the discussion in Section 4.6 and the method description in Section 4.2 are strongly suggested.

The main findings, a discussion on their merits, and some suggestions for future research are presented in Chapter 6, the final chapter of this thesis.

The reader that has overlooked the introduction of an abbreviation, notation, or mathematical or physical quantity, can find them listed in Appendix A. A very compact overview of the main simulation methods treated in this thesis is provided in Appendix B.

2

SIMULATION METHODS FOR OPTICAL SYSTEMS

This chapter gives an overview of methods used for simulating optical systems. Many of the methods treated in this chapter will reappear in later chapters, either as the basis for Wigner-based ray tracing and Huygens-Fresnel path integration or as a reference method.

2.1. Introduction

The challenge in simulating macroscopic optical systems is the large discrepancy between the wavelength of light and the size of the systems. Since these differ by many orders of magnitude, it is difficult to capture both the wave nature of light and the full geometry of the optical system in a single simulation method. The classical methods have therefore treated the two phenomena separately, either simplifying the optical system as a series of thin elements to capture the wave optical effects, or simplifying the propagator to capture the full geometry of the system and thereby losing some of the wave optical effects. The first is done by diffraction integrals, which is the topic of Section 2.3, and the second by geometrical optics, treated in Section 2.4. These two sections will be referred to frequently in later chapters, since they form the basis for the methods developed in this thesis.

The problem in simulating imaging systems is that both the geometry of the optical system and the wave nature of light affect the image quality of the optical system. In the exit pupil diffraction integral, the topic of Section 2.5, this is addressed by first propagating the light through the optical system using geometrical optics and then using a diffraction integral for the final propagation step in free space. This way the geometry of the system is captured fully, and the wave optical effects related to the focussing of the light are also simulated correctly.

For some systems diffraction already significantly affects the field before the final propagation step. This can for instance happen if the input field is diffracting (e.g., a relatively strongly diffracting laser beam) or if the field passes through a pinhole. Cascaded diffraction is the phenomenon where such a diffracted field is diffracted yet again (e.g., by an aperture). Several ray-based diffraction methods have been developed to simulate these types of systems. The methods developed in this thesis also belong to this category. Section 2.6 gives an overview of some of the more prominent ray-based diffraction methods.

The division between wave and geometrical optics is an inelegant feature of optics. The Wigner distribution function (WDF) is a phase-space distribution that has been used to bridge these two theories. In Section 2.7 it will be shown how the WDF captures the wave nature of light, while adhering to the simple propagation laws of geometrical optics.

A large number of simulation methods are touched upon in this chapter and the chapters to come. Appendix B provides a very compact overview of the main methods, including references to the relevant sections in this thesis.

Before proceeding to any of these optical propagation methods, the fundamental laws of electro-magnetism will first be derived in Section 2.2.

2.2. FUNDAMENTALS OF ELECTROMAGNETISM

Light is an electromagnetic wave and its behaviour is fully described by the Maxwell equations

$$\nabla \times \tilde{E} = -\frac{\partial \tilde{B}}{\partial t}, \tag{2.1}$$

$$\nabla \times \tilde{B} = \mu_0 \epsilon_0 \frac{\partial \tilde{E}}{\partial t} + \mu_0 \tilde{J}, \tag{2.2}$$

$$\nabla \cdot \tilde{E} = \frac{q}{\epsilon_0}, \tag{2.3}$$

$$\nabla \cdot \tilde{B} = 0, \tag{2.4}$$

where $\tilde{E}$ is the electric field, $\tilde{B}$ the magnetic flux density, ϵ_0 the permittivity of vacuum, μ_0 the permeability of vacuum, $\tilde{J}$ the current density and q the charge density. Note that throughout this thesis, all vectors are indicated by bold symbols (e.g., A) and the tildes are used for time-varying quantities[1]. These relations are not very convenient when dealing with electromagnetic fields in matter, if we do not know the distribution of currents and charges within the material. It is therefore common practice to treat the charges and currents that are due to the interaction of the electromagnetic field with the material, the bound charges (q_b) and currents (J_b), separate from the free charges (q_f) and currents (J_f). This is done by defining two auxiliary fields, the electric displacement field $\tilde{D}$ and the magnetic field $\tilde{H}$, using the polarization density $\tilde{P}$ and the magnetization $\tilde{M}$

$$\tilde{D} = \epsilon_0 \tilde{E} + \tilde{P}, \tag{2.5}$$

$$\tilde{H} = \frac{1}{\mu_0} \tilde{B} - \tilde{M}, \tag{2.6}$$

and the following relations

$$q_b = -\nabla \tilde{P}, \tag{2.7}$$

$$\tilde{J_b} = \nabla \times \tilde{M} + \frac{\partial \tilde{P}}{\partial t}, \tag{2.8}$$

$$q = q_f + q_b, \tag{2.9}$$

$$\tilde{J} = \tilde{J_f} + \tilde{J_b}. \tag{2.10}$$

This results in the following form of the Maxwell equations:

$$\nabla \times \tilde{E} = -\frac{\partial \tilde{B}}{\partial t}, \tag{2.11}$$

$$\nabla \times \tilde{H} = \frac{\partial \tilde{D}}{\partial t} + \tilde{J_f}, \tag{2.12}$$

$$\nabla \cdot \tilde{D} = q_f, \tag{2.13}$$

$$\nabla \cdot \tilde{B} = 0. \tag{2.14}$$

---

[1] All conventions, abbreviations, and symbols used in this thesis are listed in Appendix A.

So far, this only introduces an extra set of definitions, but these new equations become convenient when combined with the following model for the material

$$\tilde{D} = \epsilon \tilde{E}, \tag{2.15}$$

$$\tilde{H} = \frac{1}{\mu} \tilde{B}, \tag{2.16}$$

where ϵ and μ are the permittivity and permeability of the medium. These constitutive relations assume the material to be isotropic, linear, and without dispersion. If we further assume the medium to be free of external sources and losses

$$J_f = \mathbf{0}, \tag{2.17}$$

$$q_f = 0, \tag{2.18}$$

and the solution to be time harmonic

$$\tilde{H} = \Re\left\{H \exp\left(-i\omega t\right)\right\}, \tag{2.19}$$

where $\Re$ takes the real part and ω is the real and positive angular frequency of the light, we obtain

$$\nabla \times E = i\omega\mu H, \tag{2.20}$$

$$\nabla \times H = -i\omega\epsilon E, \tag{2.21}$$

$$\nabla \cdot (\epsilon E) = \mathbf{0}, \tag{2.22}$$

$$\nabla \cdot \left(\mu H\right) = \mathbf{0}. \tag{2.23}$$

These are the Maxwell equations for time harmonic fields in source-free, linear, isotropic, and dispersion-free media. Note that E and H are complex quantities and that the physical field can be obtained by multiplying with the time harmonic term and taking the real part (as in Eq. (2.19)). The vectorial wave equations can be obtained by combining Eqs. (2.20) and (2.21)

$$\nabla \times (\nabla \times E) - \nabla \ln(\mu) \times (\nabla \times E) - k^2 E = \mathbf{0}, \tag{2.24}$$

$$\nabla \times (\nabla \times H) - \nabla \ln(\epsilon) \times (\nabla \times H) - k^2 H = \mathbf{0}, \tag{2.25}$$

with k the wavenumber in matter defined by

$$k = \omega\sqrt{\mu\epsilon}. \tag{2.26}$$

For uniform media, in which the permittivity ϵ and the permeability μ are constant, this reduces to

$$\nabla \times (\nabla \times E) - k^2 E = 0, \tag{2.27}$$

which is equivalent to

$$\nabla^2 E + k^2 E = 0, \tag{2.28}$$

since the field is divergence free. Equivalent equations for the magnetic field can be obtained by substituting E for H in Eqs. (2.27) and (2.28). The scalar version of Eq. (2.28)

is known as the Helmholtz equation. The wave equations, Eq. (2.24) and the Helmholtz equations, are at the basis of the diffraction integrals that describe the propagation of light through uniform media.

Before presenting these diffraction integrals we will derive the boundary conditions for the electromagnetic field at the interface between two media. First we will derive the boundary conditions for the tangential components and then for the normal components of the field.

At the boundary between two media one can analyse a small surface with two edges tangential to the interface and two edges perpendicular to the interface (see Fig. 2.1a). Integrating Eq. (2.20), which assumes the absence of sources, over this surface S gives

(a) The surface used for the tangential components.

(b) The box used for the normal components.

Figure 2.1: The surface and box used to derive the boundary conditions of the electromagnetic field. The interface is drawn in light blue, the surface and box are shown in dark blue. The arrows indicate the direction of integration, i.e. dl.

$$\iint (\nabla \times \boldsymbol{E}) \cdot \hat{\boldsymbol{n}} \, dS = \iint i\omega\mu \boldsymbol{H} \cdot \hat{\boldsymbol{n}} \, dS, \tag{2.29}$$

where $\hat{\boldsymbol{n}}$ is the outward pointing surface normal. The hat of $\hat{\boldsymbol{n}}$ indicates that the vector is normalized (a convention used throughout this thesis). On the left side of this equation we will apply Stokes' theorem, which states that for any vector $\boldsymbol{A}$

$$\iint (\nabla \times \boldsymbol{A}) \cdot \hat{\boldsymbol{n}} \, dS = \oint \boldsymbol{A} \cdot d\boldsymbol{l}, \tag{2.30}$$

where $\boldsymbol{l}$ is the boundary of the surface S. The right side of Eq. (2.29) can be reduced to zero if we take the edges perpendicular to the interface to be infinitesimally small. We thus obtain

$$\oint \boldsymbol{E} \cdot d\boldsymbol{l} = 0. \tag{2.31}$$

Since the edges perpendicular to the interface are arbitrarily small, only the edges tangential to the interface contribute to the line integral of Eq. (2.31). Therefore the tangential components of the electric field must be continuous at the boundary between two media. By substituting Eq. (2.20) for Eq. (2.21) in Eq. (2.29) one can derive that the tangential components of the magnetic field are also continuous across interfaces.

To obtain the boundary conditions for the normal components of the electric field we use a box with in each medium one side tangential to the interface, as illustrated in Fig. 2.1b. Using the divergence theorem,

$$\iiint \nabla \cdot \boldsymbol{A} \, dV = \oiint \boldsymbol{A} \cdot \hat{\boldsymbol{n}} \, dS, \tag{2.32}$$

where A is an arbitrary vector, V is a volume with surface S, and $\hat{n}$ the outward pointing surface normal, we rewrite a volume integral over Eq. (2.22) into a surface integral

$$\iiint \nabla \cdot (\epsilon E) \, dV = 0, \tag{2.33}$$

$$\oiint (\epsilon E) \cdot \hat{n} \, dS = 0. \tag{2.34}$$

By taking the thickness of the box to be infinitesimally small, the only contributions to the integral of Eq. (2.34) come from the two sides tangential to the interface. Considering that the normals of these two sides point in opposite directions, we find that the component of ϵE normal to the interface is continuous across the interface. By performing the same derivation using Eq. (2.23) instead of Eq. (2.22), one can show that the component of μH normal to the interface is also continuous across the interface.

We thus find that at the interface between two media, the tangential components of the electric and magnetic field, and the normal components of ϵE and μH are continuous.

The relations derived in this section assume the media to be isotropic, linear, and source free. They also neglect dispersion and losses, i.e., the materials are assumed to be fully transparent. From here on we will also assume all materials to be non-magnetic, i.e., $\mu = \mu_0$, since this is essentially true for all media that are transparent in the visible spectrum [16, p. 66].

2.3. DIFFRACTION INTEGRALS

The central problem of this thesis is the simulation of the propagation of light. Although the Maxwell equations fully describe the behaviour of electromagnetic fields, they do not in themselves provide a method to solve this propagation problem. Diffraction integrals, which can be derived using the Maxwell equations, do provide a method for field propagation. First the Rayleigh-Sommerfeld diffraction integral and its vectorial equivalent, the m-theory diffraction integrals, are introduced in Section 2.3.1. As the name 'diffraction integral' implies, these integrals describe the wave optical (or diffraction) effects of field propagation. The disadvantage of these methods is that they can only describe the propagation through a uniform medium, and therefore cannot be used to simulate the propagation through an optical system consisting of lenses. The Collins integral, treated in Section 2.3.2, is a diffraction integral that can deal with optical systems, but only by greatly simplifying their effects. All diffraction integrals assume fully coherent light, a concept explained in Section 2.3.3. Diffraction integrals are frequently used to calculate the effects of apertures. Section 2.3.4 describes the model for apertures used throughout this thesis.

2.3.1. RAYLEIGH-SOMMERFELD AND M-THEORY DIFFRACTION INTEGRALS

In this thesis we will extensively use the Rayleigh-Sommerfeld diffraction integral (RSDI) and the m-theory diffraction integrals (or simply m-theory), derived by, e.g., Severin [15]. Both express the field in a point in space as an integral over the field in a plane, the former for the scalar and the later for the vectorial field. They thus offer a way to propagate the

field through a homogeneous medium. Before presenting the integrals, we will give an outline of their derivation.

The starting point for the derivation of these integrals is the divergence theorem of Eq. (2.32). For the RSDI we take

$$A = E\nabla G - G\nabla E, \tag{2.35}$$

and for the m-theory

$$A = E \times (\nabla \times G) - G \times (\nabla \times E), \tag{2.36}$$

with E or $\boldsymbol{E}$ the electric field, and G or $\boldsymbol{G}$ a Green's function. A Green's function is any solution of the differential equation

$$L\boldsymbol{G}(\boldsymbol{r}, \boldsymbol{r_1}) = \boldsymbol{a}\delta(\boldsymbol{r_1} - \boldsymbol{r}), \tag{2.37}$$

where L is a linear differential operator, δ is the delta function, $\boldsymbol{r}$ and $\boldsymbol{r_1}$ are positions in space, and $\boldsymbol{a}$ is an arbitrary vector. For the scalar case, $\boldsymbol{G}$ reduces to G and the vector $\boldsymbol{a}$ is omitted. For the RSDI one needs a Green's function of the Helmholtz equation

$$L_H = \nabla^2 + k^2, \tag{2.38}$$

which operates on the independent variable $\boldsymbol{r}$ and not on the fixed position $\boldsymbol{r_1}$. One solution, which will be frequently used in this thesis, is[2]

$$G_H(\boldsymbol{r}, \boldsymbol{r_1}) = \frac{\exp(ik|\boldsymbol{r_1} - \boldsymbol{r}|)}{4\pi|\boldsymbol{r_1} - \boldsymbol{r}|}. \tag{2.39}$$

The differential operator of the m-theory is

$$L_m = k^2 - \nabla \times \nabla \times, \tag{2.40}$$

for which a Green's function is [17]

$$\boldsymbol{G_m}(\boldsymbol{r}, \boldsymbol{r_1}) = \left(1 + \frac{1}{k^2}\nabla\nabla\cdot\right) G_H(\boldsymbol{r}, \boldsymbol{r_1})\boldsymbol{a}. \tag{2.41}$$

Using these Green's functions in Eq. (2.32), with A given by Eq. (2.35) or (2.36), reduces the volume integral on the left side of Eq. (2.32) to $E(\boldsymbol{r_1})$ or $\boldsymbol{a} \cdot \boldsymbol{E}(\boldsymbol{r_1})$ respectively. Since the surface integral remains, the field in the point $\boldsymbol{r_1}$ can be determined by the integral over a closed surface S.

Both the RSDI and m-theory are derived by splitting S into two parts, a plane S_0 and a half sphere with an infinite radius S_∞. The point $\boldsymbol{r_1}$ is within this volume at the centre of the half sphere, and by convention is usually taken to be to the right of S_0. The integral over S_∞ reduces to zero under conditions known as the radiation conditions. Both in the scalar and vectorial case, these conditions state that no (superposition of) plane waves can radiate through S_∞ into the volume. Under that condition, the only

---

[2]The same function but with a negative exponential is also a valid Green's function, as are certain linear combinations of the function with the negative and the positive exponential. The desired diffraction integral cannot be derived using the Green's function with a negative exponential.

2

remaining integral is that over S_0, which has two terms. One is proportional to E (or $\boldsymbol{E}$), and one to ∇E (or $\nabla \times \boldsymbol{E}$). The second term can be eliminated with the aid of a mirror point $\boldsymbol{r_1'}$, which is the mirror image of the point $\boldsymbol{r_1}$ in the plane S_0. Using this mirror point, one can define the new Green's functions

$$G_H^- = G_H(\boldsymbol{r_0}, \boldsymbol{r_1}) - G_H(\boldsymbol{r_0}, \boldsymbol{r_1'}), \tag{2.42}$$

$$\boldsymbol{G_m^-} = \left(1 + \frac{1}{k^2}\nabla\nabla\cdot\right)\left(G_H(\boldsymbol{r_0},\boldsymbol{r_1})\boldsymbol{a} - G_H(\boldsymbol{r_0},\boldsymbol{r_1'})\boldsymbol{a'}\right), \tag{2.43}$$

with $\boldsymbol{r_0}$ a point on S_0 and $\boldsymbol{a'}$ the mirrored version of $\boldsymbol{a}$

$$\boldsymbol{a'} = \boldsymbol{a} - 2(\boldsymbol{a}\cdot\hat{\boldsymbol{n}})\hat{\boldsymbol{n}}. \tag{2.44}$$

These Green's functions reduce the integrals to our final results, which depend on the vector from the point of integration to the point of interest

$$\boldsymbol{\rho} = \boldsymbol{r_1} - \boldsymbol{r_0}, \tag{2.45}$$

whose length will be denoted by ρ and its unit vector by $\hat{\boldsymbol{\rho}}$. For the scalar case we obtain the Rayleigh-Sommerfeld diffraction integral of the first kind

$$E(\boldsymbol{r_1}) = -\frac{\mathrm{i}k}{2\pi}\iint \frac{\exp(\mathrm{i}k\rho)}{\rho}\hat{\boldsymbol{\rho}}\cdot\hat{\boldsymbol{n}}_0\, E\, \mathrm{d}S_0, \tag{2.46}$$

and for the vectorial case the diffraction integral of the m-theory

$$\boldsymbol{E}(\boldsymbol{r_1}) = -\frac{\mathrm{i}k}{2\pi}\iint \frac{\exp(\mathrm{i}k\rho)}{\rho}\hat{\boldsymbol{\rho}}\times(\hat{\boldsymbol{n}}_0\times\boldsymbol{E})\,\mathrm{d}S_0, \tag{2.47}$$

where $\hat{\boldsymbol{n}}_0$ is the surface normal of S_0 that points towards the half space containing $\boldsymbol{r_1}$. In both Eq. (2.46) and (2.47) we have neglected a term proportional to $1/\rho$, which is sufficiently accurate if the point of interest is many wavelengths away from S_0 ($1/\rho \ll k$). By taking the curl with respect to $\boldsymbol{r_1}$, one can obtain the magnetic field

$$\boldsymbol{H}(\boldsymbol{r_1}) = \frac{-\mathrm{i}k^2}{2\pi\omega\mu}\iint \frac{\exp(\mathrm{i}k\rho)}{\rho}\hat{\boldsymbol{\rho}}\times\left[\hat{\boldsymbol{\rho}}\times(\hat{\boldsymbol{n}}_0\times\boldsymbol{E})\right]\,\mathrm{d}S_0, \tag{2.48}$$

with μ the magnetic permeability of the medium μ (which is assumed to be equal to the permeability of free space throughout this thesis).

In the derivation of Eq. (2.46), (2.47), and (2.48) we have assumed that only the light passing through S_0 affects the field in $\boldsymbol{r_1}$ and that S_0 is a plane. Furthermore, it is assumed that the volume to the right of S_0, which by convention contains $\boldsymbol{r_1}$, extends to infinity; is source free; and contains a homogeneous, linear, and isotropic medium (e.g., air).

Under these conditions the Rayleigh-Sommerfeld diffraction integral of Eq. (2.46) and the m-theory diffraction integral of Eq. (2.47) can be used to calculate the scalar or vectorial electric field.

2.3.2. Collins integral

A significant disadvantage of the diffraction integrals described in Section 2.3.1 is that they can only propagate the field through a homogeneous medium. The Collins integral [4] overcomes this issue by using a simplified model of the optical system.

The Collins integral is a generalisation of the Fresnel integral, which itself is an approximation of the Rayleigh-Sommerfeld diffraction integral, Eq. (2.46). To obtain the Fresnel integral from Eq. (2.46), the distance ρ is approximated by the Taylor expansion

$$\rho = \sqrt{(x - x_0)^2 + (y - y_0)^2 + z^2} = z\sqrt{\frac{(x - x_0)^2}{z^2} + \frac{(y - y_0)^2}{z^2} + 1},$$

$$\approx z + \frac{(x - x_0)^2}{2z} + \frac{(y - y_0)^2}{2z} + ..., \tag{2.49}$$

where it is assumed that z is much larger than $(x - x_0)$ and $(y - y_0)$. By using the second order Taylor expansion in the exponential of Eq. (2.46) and the first order expansion otherwise, we obtain

$$E(x, y, z) = \frac{-ik \exp(ikz)}{2\pi z} \iint E(x_0, y_0, 0) \exp\left[\frac{ik}{2z}\left((x - x_0)^2 + (y - y_0)^2\right)\right] dx_0 dy_0. \tag{2.50}$$

This assumes that the surface normal and the main direction of propagation are along the z-axis. With a quadratic phase function one can also approximate the propagation through a system of lenses and homogeneous media. Such a description results from the paraxial approximation, which assumes the light to be close to the optical axis and the direction of propagation to be nearly parallel to the optical axis. Using a notation that may at first appear a bit arbitrary, a second order diffraction integral that describes an optical system can have the following form

$$E_1(\boldsymbol{x_1}) = \frac{-ik}{2\pi\sqrt{|\mathbb{B}|}} \iint E_0(\boldsymbol{x_0})$$
$$\exp\left[ik\left(\frac{1}{2}\boldsymbol{x_0}^T\mathbb{B}^{-1}\mathbb{A}\boldsymbol{x_0} - \boldsymbol{x_0}^T\mathbb{B}^{-1}\boldsymbol{x_1} + \frac{1}{2}\boldsymbol{x_1}^T\mathbb{D}\mathbb{B}^{-1}\boldsymbol{x_1}\right)\right] d^2\boldsymbol{x_0}, \tag{2.51}$$

where $\boldsymbol{x} = (x, y)$, the integral is over all the components of the vector ($d^2\boldsymbol{x_0} = dx_0 dy_0$), capital letters in blackboard bold (e.g., $\mathbb{A}$) are matrices, and a constant phase term has been neglected. It turns out that the matrices $\mathbb{A}$, $\mathbb{B}$ and $\mathbb{D}$ are part of the matrix that describes the passage of rays through a general paraxial system without obscurations [4]

$$\begin{pmatrix} \boldsymbol{x_1} \\ \boldsymbol{p_1} \end{pmatrix} = \begin{bmatrix} \mathbb{A} & \mathbb{B} \\ -\left(\mathbb{B}^{-1} - \mathbb{A}^T\mathbb{D}\mathbb{B}^{-1}\right)^T & \mathbb{D} \end{bmatrix} \begin{pmatrix} \boldsymbol{x_0} \\ \boldsymbol{p_0} \end{pmatrix} = \begin{bmatrix} \mathbb{A} & \mathbb{B} \\ \mathbb{C} & \mathbb{D} \end{bmatrix} \begin{pmatrix} \boldsymbol{x_0} \\ \boldsymbol{p_0} \end{pmatrix}, \tag{2.52}$$

where T indicates the transpose, $\mathbb{B}^{-1}$ is the inverse of $\mathbb{B}$, $\boldsymbol{x_0}$ and $\boldsymbol{p_0}$ are the ray positions and the ray directions multiplied with the refractive index in the input plane before the system and $\boldsymbol{x_1}$ and $\boldsymbol{p_1}$ those in the output plane after the system. The Collins integral, Eq. (2.51), and the ray transfer matrix, Eq. (2.52), are valid for paraxial systems. Furthermore, $\mathbb{B}$ should have full rank, so the final position of the rays are a function of their initial directions (i.e., the input and output planes cannot be conjugate planes). Finally,

the system should transmit all light. Apertures can be accounted for in a multi-step process, first calculating the field from the input plane until the aperture using Eq. (2.51) or Eq. (2.52), then implementing the effects of the aperture (see Section 2.3.4), and finally applying the Collins integral or ray transfer matrix for the second part of the system. Some further properties of the matrices in Eq. (2.52) are that the determinant of the complete matrix is one, $\mathbb{D}\mathbb{B}^{-1}$ and $\mathbb{B}^{-1}\mathbb{A}$ are symmetric and $\mathbb{D}\mathbb{A}^T - \mathbb{C}\mathbb{B}^T = \mathbb{I}$ with $\mathbb{I}$ the identity matrix [18].

The wave optical propagation through any paraxial system that meets the before mentioned criteria can thus be described by a diffraction integral, Eq. (2.51), whose parameters are closely linked to the parameters describing the propagation of rays through the same paraxial system, Eq. (2.52).

2.3.3. SPATIAL COHERENCE

Suppose there are two radiating point sources $E_{i,1}$ and $E_{i,2}$ located at two points in the input plane that generate the fields $E_{o,1}$ and $E_{o,2}$ at the output plane. If the time harmonic part of the field, which we have omitted since Eq. 2.19, is the same for both point sources, we call the sources fully coherent. The total squared amplitude of the field at the output plane is then

$$|E_o|^2 = |E_{o,1} + E_{o,2}|^2. \tag{2.53}$$

The diffraction integrals presented earlier assume a more general field distribution at the input plane but the principle is the same: the total field is a coherent sum, analogue to Eq. 2.53, of the contributions from all points in the input plane. For this to be true the light at the input plane must be fully coherent.

One can also assume that the phase of the two sources is completely uncorrelated. We call this (spatially)[3] incoherent light, and in that case the total squared amplitude of the field at the output plane is

$$|E_o|^2 = |E_{o,1}|^2 + |E_{o,2}|^2. \tag{2.54}$$

Suppose we have a more general spatially incoherent field at the input plane and we want to calculate the squared amplitude of the field at the output plane. We can no longer apply the diffraction integrals since they assume fully coherent light. Instead we can calculate the field contribution from a single point (r_1) in the input plane by applying a diffraction integral to the field $E_i \delta(r - r_1)$, where E_i is the complex field in the input plane and r the spatial coordinates in the input plane. The total squared amplitude at the output plane can be obtained by repeating this procedure for all points in the input plane and summing the squared amplitudes.

Finally, the electric field can also be partially (spatially) coherent. We can often model this by starting our simulation at the actual coherent or incoherent source of the light (instead of at the plane where the field is partially coherent).

---

[3]Light also has a temporal coherence. This expresses the ability of light coming from a point to interfere with light coming from the same point at a later point in time. Our focus in this thesis is on simulation methods for optical systems. Imaging systems are designed such that the light emitted from one point in the object plane will arrive at a single point in the image plane with a constant time delay, independent of the path through the optical system. We will therefore ignore the topic of temporal coherence.

Throughout this thesis we will assume the light to be fully coherent. However, the simulation methods presented in this thesis can be extended to spatially incoherent light analogue to the way we applied the diffraction integrals to incoherent light: use the simulation method to calculate the field contribution from single points in the input plane and obtain the total squared amplitude of the field by incoherently summing all these field contributions.

2.3.4. KIRCHHOFF BOUNDARY CONDITIONS

Diffraction integrals are frequently used to calculate the electric field a certain distance behind an aperture (for example a perfectly conducting screen with a hole in it). These integrals can propagate the light through the free space behind an aperture, but they do not provide a model for the aperture itself.

In this thesis we will treat apertures as binary thin elements. The thin element approximation means that the field behind the aperture (E) is the product of a so-called mask function (B) and the field in front of the aperture (E_0)

$$E(x) = B(x)E_0(x). \tag{2.55}$$

Note that this approximation means that the effect of the element is concentrated in a plane (of zero thickness) and is local. In general B can be a complex function (thereby changing the amplitude and the phase of the field) but we treat apertures using a binary function. Inside the aperture the field is unaffected ($B = 1$), and outside the aperture (where the perfectly conducting screen is) the light is completely blocked ($B = 0$).

These approximations are known as the Kirchhoff boundary conditions[4]. Although they do not correctly describe the field near the edge of an aperture, they have proven to be sufficiently accurate for most diffraction problems dealing with optical instruments [13, p. 422].

2.4. GEOMETRICAL OPTICS

Geometrical optics is one of the most commonly used methods to simulate optical systems consisting of lenses and mirrors, because it can deal with the general geometry of the optical elements (i.e., the lenses and mirrors). Since it does not rely on the paraxial approximation (although paraxial geometrical optical models exist, e.g., Eq. (2.52)), geometrical optics can describe the propagation of light further away from the optical axis, at large angles with the optical axis, and through non-spherical optical surfaces. Geometrical optics assumes light to propagate along rays, paths that are straight in uniform media and that can change direction at optical interfaces. Such a model is consistent with the particle model of light, but it can also be derived from the Maxwell equations (see, e.g., Born and Wolf [13, p. 117–127]). From this derivation it follows that rays carry a phase, which is in correspondence with the wave nature of light. However, not all wave effects of light are captured by geometrical optics. It fails whenever the electric field changes significantly over a small distance. In particular it wrongly predicts infinitesimally sharp transitions from shadow to light.

---

[4]Technically the Kirchhoff boundary conditions also state that the derivative of the field outside of the aperture is zero.

Depending on whether rays carry a phase term or not, geometrical optics describes the propagation of light as a particle, or approximates the propagation of light waves. Since this can be the source of some confusion, and because geometrical optics (as a wave optical approximation) is essential for the methods presented in later chapters, a derivation of the propagation laws of geometrical optics is presented in Section 2.4.1. The implementation of these propagation laws in two types of simulation methods, ray tracing and Monte Carlo ray tracing, are then treated in Sections 2.4.2 and 2.4.3.

2.4.1. DERIVATION OF GEOMETRICAL OPTICS

Sommerfeld and Runge [3, p. 289-293] derived the laws and the limits of geometrical optics by inserting an ansatz for the electric field into the wave equation. The ansatz they used separated a phase-only term, $\exp[ik_0 s(r)]$ (where k_0 is the wavenumber in free space), from the rest of the field. Although Sommerfeld and Runge used a scalar field, the same procedure can be followed for the vectorial field (see, e.g., Born and Wolf [13, p. 117–127]). We will treat the scalar and vectorial case simultaneously, roughly following Sommerfeld and Runge [3] and Born and Wolf [13].

The vectorial and scalar ansatz read

$$E(r) = \bar{E}(r)\exp(ik_0 s(r)), \tag{2.56}$$

$$E(r) = \bar{E}(r)\exp(ik_0 s(r)), \tag{2.57}$$

with $\bar{E}$ real and where $\bar{E}$ can be complex to represent certain polarizations (i.e., non-linear polarized light). The wave equations can be applied to these fields. The vectorial wave equation is given by Eq. (2.24) and is

$$-\nabla\ln(\mu)\times(\nabla\times E)+\nabla(\nabla\cdot E)-\Delta E-\omega^2\epsilon\mu E=0.$$

The scalar theory uses the Helmholtz equations for the electric field,

$$\nabla^2 E + k^2 E = 0, \tag{2.58}$$

which is the scalar version of Eq. (2.28). After substituting Eq. (2.56) and (2.57) into Eq. (2.24) and (2.58) respectively, one obtains terms proportional to $k_0{}^2$, $k_0{}^1$, and $k_0{}^0$. Since in optics $k_0 \gg 1$, it makes sense to treat these different orders separately and require that the three groups equate to zero individually.

By setting the $k_0{}^2$ terms to zero one obtains, in both the scalar and vectorial case, the eikonal equation

$$|\nabla s|^2 - n^2 = 0, \tag{2.59}$$

where

$$n = \sqrt{\frac{\epsilon\mu}{\epsilon_0\mu_0}}, \tag{2.60}$$

is the refractive index of the medium. Note that we assume all materials to be non-absorbing, which means that the imaginary part of the refractive index is assumed to be zero[5]. Since the phases of Eq. (2.56) and (2.57) are determined by s, surfaces with a

---

[5]Laws for geometrical optics in absorbing media have been derived, e.g., by Chang *et al.* [19]. They showed that when a plane wave enters such a medium, the normals to the plane of constant phase and to the plane of constant amplitude no longer coincide. This substantially complicates a geometrical optical propagation of the field.

constant value of s are wavefronts of equal phase[6]. The gradient ∇s will be perpendicular to the wavefront and in that direction s will change proportional to the refractive index n. In geometrical optics, this direction

$$\hat{\boldsymbol{\rho}} \equiv \frac{\nabla s}{n}, \tag{2.61}$$

is the direction of the rays.

The second group of terms, which is proportional to k_0, will determine the amplitude change of the field (see, e.g., Born and Wolf [13, p. 126]). For uniform media, these terms reduce to

$$\bar{\boldsymbol{E}} \, \nabla^2 s + 2 \left(\nabla s \cdot \nabla \right) \bar{\boldsymbol{E}} = 0, \tag{2.62}$$

$$\bar{E} \, \nabla^2 s + 2 \nabla s \cdot \nabla \bar{E} = 0, \tag{2.63}$$

which are differential equations that tell us that the change in the amplitude of the field along ∇s is proportional to the curvature of the phase function (s). We will later derive a more practical rule governing the amplitude of the field, using the conservation of energy.

The remaining terms, proportional to $k_0{}^0$, are (for materials with uniform magnetic properties)

$$\text{error} = \nabla^2 \bar{\boldsymbol{E}} + \nabla \left(\bar{\boldsymbol{E}} \cdot \nabla \ln(\epsilon) \right), \tag{2.64}$$

and should remain small compared to k_0 for geometrical optics to be valid. This assumption breaks down when the amplitude of the field or ϵ changes abruptly, as is for instance the case at the edge of geometrical optical shadows, in a focal region, or at optical interfaces. For the scalar case we already implicitly assumed the medium to be homogeneous by using the Helmholtz equation. The scalar term proportional to $k_0{}^0$ is

$$\text{error} = \nabla^2 \bar{E}, \tag{2.65}$$

which requires the amplitude changes to be gradual (with respect to the wavelength).

The electric field is thus given by Eq. (2.56) or (2.57) and its propagation is governed by Eqs. (2.59) and (2.62) or (2.63) provided Eq. (2.64) or (2.65) remains small. We will now derive the corresponding magnetic field for the vectorial case so that we can obtain the Poynting vector, which describes the energy flow of the electromagnetic field. We will assume that for the scalar case the same rules will govern the flow of energy.

The magnetic field can be obtained by taking the curl of the electric field

$$\boldsymbol{H} = \frac{-\mathrm{i}}{\omega \mu} \nabla \times \boldsymbol{E},$$

$$\approx \frac{\sqrt{\epsilon_0 \mu_0}}{\mu} e^{\mathrm{i} k_0 s(\boldsymbol{r})} \nabla s \times \bar{\boldsymbol{E}},$$

$$= \sqrt{\frac{\epsilon}{\mu}} e^{\mathrm{i} k_0 s(\boldsymbol{r})} \hat{\boldsymbol{\rho}} \times \bar{\boldsymbol{E}}, \tag{2.66}$$

---

[6]For elliptically polarized light the concept of a wavefront is more difficult to define. One can however write the field as a combination of two linearly polarized fields which do have wavefronts at surfaces of constant s.

where we have used Eqs. (2.59, 2.60) and assumed

$$|k_0 \nabla s \times \bar{E}| \gg |\nabla \times \bar{E}|. \tag{2.67}$$

This assumption states that the amplitude of the field $\bar{E}$ orthogonal to ∇s does not change significantly on the length scale of a wavelength.

The average flow of electromagnetic energy is given by the time-averaged Poynting vector (see, e.g., [16, p.48–50])

$$S = \frac{1}{2} E \times H^*, \tag{2.68}$$

where the asterisk denotes the complex conjugate and we have used that the material is non-magnetic. Applying this to Eqs. (2.56, 2.66) gives

$$
\begin{aligned}
S &= \frac{1}{2} E \times H^*, \\
&= \frac{1}{2} \sqrt{\frac{\epsilon}{\mu}} \left[\left(\bar{E} \cdot \bar{E}^* \right) \hat{\rho} + \left(\bar{E} \cdot \hat{\rho} \right) \bar{E}^* \right], \\
&\approx \frac{1}{2} \sqrt{\frac{\epsilon}{\mu}} |\bar{E}|^2 \frac{\nabla s}{n},
\end{aligned}
\tag{2.69}
$$

under the assumption of a source-free, homogeneous medium. Using that $\nabla \cdot E = 0$, the term neglected in Eq. (2.69) can be rewritten as

$$\left(\bar{E} \cdot \hat{\rho} \right) \bar{E}^* = \frac{1}{n} \left(\bar{E} \cdot \nabla s \right) \bar{E}^* = \frac{\mathrm{i}}{n \, k_0} \left(\nabla \cdot \bar{E} \right) \bar{E}^*, \tag{2.70}$$

which is very small compared to $|\bar{E}|^2 \hat{\rho}$ if $\bar{E}$ does not change significantly on the length scale of a wavelength. We thus obtain a Poynting vector with a magnitude of

$$|S| = \sqrt{\frac{\epsilon}{4\mu}} |\bar{E}|^2, \tag{2.71}$$

and a direction of propagation along $\hat{\rho}$.

Note that the relation between the electric and magnetic field in geometrical optics (expressed by Eq. (2.66)) is also valid for plane waves (see, e.g., [13, p. 24]). Furthermore, Eq. (2.69) is the Poynting vector of a plane wave [13, p. 24-25]. For this reason, geometrical optical rays are often said to carry local plane waves. This local planar approximation of the wavefront, i.e., geometrical optics, is valid if the radius of curvature of the wavefront is large compared to the wavelength.

The concept of local plane waves is useful at optical interfaces, where the second term in Eq. (2.64) cannot be neglected, and we thus require a modified approach. Because the field locally has all the properties of a plane wave, the transmission through the optical surface is treated as the transmission of a plane wave passing through a plane interface, where the plane interface is defined by the local surface normal. The validity of this approximation was studied for sinusoidal gratings, where it was shown accurate for radii of curvature as small as a few wavelengths [20]. Since the curvatures in lenses are

usually several orders of magnitude larger, it seems reasonable to assume that this local plane approximation will give accurate results. The transmission coefficients for plane waves passing through plane interfaces are given by the Fresnel equations (see, e.g., [13, p. 42]). The corresponding change in the wave vector of the plane wave is equivalent to that of Snell's law

$$n_1 \sin(\theta_1) = n_2 \sin(\theta_2), \tag{2.72}$$

where θ is the angle between the surface normal and the wave vector, and the index 1 is used for the quantities in front of the interface and 2 for the quantities after refraction by the interface. Part of the field is reflected at the surface, which causes a decrease in the amplitude of the transmitted wave. The attenuation is different for the component of the electric field in the plane spanned by the surface normal and the wave vector (the p-component) and for the component in the direction perpendicular to this plane (the s-component). The amplitude transmissions through the interface are given by [13, p. 42]

$$T_s = \frac{2n_1 \cos(\theta_1)}{n_1 \cos(\theta_1) + n_2 \cos(\theta_2)}, \tag{2.73}$$

$$T_p = \frac{2n_1 \cos(\theta_1)}{n_2 \cos(\theta_1) + n_1 \cos(\theta_2)}, \tag{2.74}$$

for the s- and p-polarised light respectively. The phase of the electric field is the same before and after the interface. As a result, the eikonal equation (Eq. (2.59)) is also valid across optical surfaces. If the amplitude and phase are slowly varying functions before the interface, they will be slowly varying functions after the interface, provided the interface is slowly varying. One can thus continue propagating the field using geometrical optics.

Until now, the conservation of energy has only been briefly touched upon by Eq. (2.62). In the following paragraphs we derive a more generally applicable rule for energy conservation in geometrical optics, using the Poynting vector.

The Poynting vector is a relevant quantity because its closed surface integral describes the average rate of energy flow into a volume

$$\frac{\mathrm{d}U}{\mathrm{d}t} = -\oiint \hat{\boldsymbol{n}} \cdot \boldsymbol{S} \, \mathrm{d}S, \tag{2.75}$$

where U is the enclosed electromagnetic energy, t the time, $\hat{\boldsymbol{n}}$ the outwards pointing surface normal and S the surface area (see, e.g., [13, p. 10]). If the medium does not absorb electromagnetic energy and there are no sources inside the volume enclosed by S, the integral reduces to

$$\oiint \hat{\boldsymbol{n}} \cdot \boldsymbol{S} \, \mathrm{d}S = 0. \tag{2.76}$$

We take the closed surface S along a tube of rays, so that we have one surface (S_1) where all the rays enter the volume, one surface (S_2) where all the rays exit the volume and the other sides of the surfaces are parallel to the rays (Fig. 2.2). The closed surface integral Eq. (2.76) reduces to

$$\iint \hat{\boldsymbol{n}} \cdot \boldsymbol{S} \, \mathrm{d}S_2 = \iint -\hat{\boldsymbol{n}} \cdot \boldsymbol{S} \, \mathrm{d}S_1. \tag{2.77}$$

Figure 2.2: A ray tube

Since this is valid for an arbitrary small tube of rays, we can remove the integral sign and, using the following definition for the intensity

$$I \equiv |\hat{\boldsymbol{n}} \cdot \boldsymbol{S}|, \tag{2.78}$$

the intensity law of geometrical optics

$$I_1 \, \mathrm{d}S_1 = I_2 \, \mathrm{d}S_2, \tag{2.79}$$

is obtained. The intensity of the field is therefore inversely proportional to the area spanned by the tube of rays. By combining Eq. (2.78) with Eq. (2.69), one obtains (for non-magnetic materials)

$$I = \hat{\boldsymbol{n}} \cdot \hat{\boldsymbol{\rho}} \, n \sqrt{\frac{\epsilon_0}{4\mu_0}} |\boldsymbol{E}|^2. \tag{2.80}$$

The amplitude of the electric field is therefore

$$|\boldsymbol{E}| = \sqrt{\frac{2I\sqrt{\mu_0}}{\hat{\boldsymbol{n}} \cdot \hat{\boldsymbol{\rho}} \, n\sqrt{\epsilon_0}}}, \tag{2.81}$$

and through its dependence on I, it is a function of the area spanned by a tube of rays.

We have derived the propagation laws of geometrical optics, both for the phase and the amplitude of the vectorial field. The next step is the implementation of these propagation laws into a numerical simulation tool.

2.4.2. RAY TRACING

Since rays propagate along straight lines in homogeneous media and change their direction according to a simple relation (Eq. (2.72)) at interfaces between two media, it is straightforward to compute their trajectories using a computer[7]. The phase of the ray is computed by keeping track of the optical path length (OPL). The OPL is zero at the object and increases at each propagation step by the geometrical length of the ray segment multiplied by the local refractive index. The value of the OPL corresponds to that of the eikonal (introduced in Section 2.4.1) if the eikonal is also defined to be zero at the object. The phase is obtained by multiplying the OPL with the wavenumber of free space k_0.

Finding the amplitude of the field is more elaborate and involves keeping track of the area spanned by a tube of rays. One method is to trace an entire grid of rays simultaneously and calculate the areas between the grid points. A second method is to accompany

---

[7]A description of the related algorithms is beyond the scope of this thesis, so the interested reader is referred to a paper by Spencer and Murty [21].

every ray by a set of neighbouring rays with directions or positions that differ from the main ray by an infinitesimally small amount. These neighbouring rays will span the surface area needed for the amplitude adjustment. For a final method one needs to keep track of the local curvature of the phase front and use (analytical) formula for the amplitude change of the field upon propagation. This is straightforward when propagating from a point source through a uniform medium, but gets much more involved when the field passes through non-spherical interfaces between two media.

The areas of the ray tubes change both upon propagation through homogeneous media and upon refraction. The amplitude of the field of a point source in free space is for instance proportional to the inverse of the distance to the source, since the ray tube area increases with the square of the distance. At refracting interfaces, the area spanned

Figure 2.3: The surface area of a ray tube changes at refraction.

by a ray tube will change (see Fig. 2.3) by a factor

$$\frac{S_2}{S_1} = \frac{\cos(\theta_2)}{\cos(\theta_1)},\tag{2.82}$$

which causes an amplitude change of

$$\frac{|\boldsymbol{E_2}|}{|\boldsymbol{E_1}|} = \sqrt{\frac{\cos(\theta_1)}{\cos(\theta_2)}}.\tag{2.83}$$

Since the Fresnel equations are energy preserving, this factor is already accounted for in the Fresnel coefficients given by Eqs. (2.73, 2.74). In order to prevent counting these factors double, it is convenient to use an adapted version of the Fresnel equations for ray tracing

$$T'_s = \sqrt{\frac{\cos(\theta_2)}{\cos(\theta_1)}} \frac{2n_1 \cos(\theta_1)}{n_1 \cos(\theta_1) + n_2 \cos(\theta_2)},\tag{2.84}$$

$$T'_p = \sqrt{\frac{\cos(\theta_2)}{\cos(\theta_1)}} \frac{2n_1 \cos(\theta_1)}{n_2 \cos(\theta_1) + n_1 \cos(\theta_2)}.\tag{2.85}$$

To perform geometrical optical simulations using ray tracing, one defines the source by specifying the ray distribution at the initial surface and the starting power, phase, and polarization of the rays. The rays are then transported through the optical system, where their paths are determined by Snell's law, Eq. (2.72), their phases by the optical path lengths, and their amplitudes by the (adjusted) Fresnel equations, Eqs. (2.84) and (2.85), and the intensity law of geometrical optics, Eq. (2.79).

2.4.3. MONTE CARLO RAY TRACING

Monte Carlo ray tracing is another numerical implementation of geometrical optics. The main difference with standard ray tracing is that a Monte Carlo ray tracer initiates random rays from a predefined source set and traces each ray independently through the optical system. Upon passing through a pixel, the ray weight is added to the pixel value. The ray paths and phases are determined in the same way as in standard ray tracing, but we will see that the intensity of the field is calculated differently.

Mathematically, a Monte Carlo ray tracer relies on the central limit theorem (see, e.g., Dekking [22, p. 197]). This theorem states that the averaged sum of the function values $w(\boldsymbol{x}_i)$ evaluated at (two-dimensional) positions $\boldsymbol{x}_i$ sampled from a random variable $\boldsymbol{X}$ with probability distribution function $p(\boldsymbol{x})$,

$$w_N = \frac{1}{N} \sum_{i=1}^{N} w(\boldsymbol{x}_i), \tag{2.86}$$

converges in the limit of $N \to \infty$, to the analytical expectation value

$$\langle w(\boldsymbol{X}) \rangle = \iint w(\boldsymbol{x}) p(\boldsymbol{x}) d^2 \boldsymbol{x}, \tag{2.87}$$

where the square brackets denote the expectation value. The central limit theorem also predicts the rate of convergence, expressed by the standard deviation (σ) which is the square root of the variance. The variance is a measure of the spread of random variables, and is defined by

$$V[w(\boldsymbol{X})] = \langle (w(\boldsymbol{X}) - \langle w(\boldsymbol{X}) \rangle)^2 \rangle . \tag{2.88}$$

For the estimate w_N, the standard deviation is

$$\sigma[w_N] = \sqrt{V[w_N]} = \sqrt{\frac{V[w(\boldsymbol{X})]}{N}}, \tag{2.89}$$

which implies a convergence rate of order $\mathcal{O}(N^{-1/2})$. In other words, one can approximate the integral of Eq. (2.87) by averaging a large enough number of random samples of $w(\boldsymbol{x}_i)$.

For Monte Carlo ray tracing, the sample points $\boldsymbol{x}_i$ are the intersection points of rays at a final or intermediate plane where the light distribution is to be determined. The probability of sampling a certain ray is proportional to the local density of rays

$$p(\boldsymbol{x}) = \lim_{\Delta S \to 0, \ N \to \infty} \frac{\Delta N(\boldsymbol{x}, N)}{\Delta S(\boldsymbol{x}) N}, \tag{2.90}$$

with ΔN the number of rays on a surface area ΔS, and N the total number of rays. By choosing

$$w_I(\boldsymbol{x}) = \frac{I(\boldsymbol{x})}{p(\boldsymbol{x})}, \tag{2.91}$$

and limiting the domain of integration of Eq. (2.87) to the area of a pixel S_p, Eq. (2.87) gives the power incident on that pixel. Since generally the intensity distribution in the

pixel plane $I(\boldsymbol{x})$ is not known, one can use the ray tracer to calculate it from the intensity distribution in another plane $I_0(\boldsymbol{x_0})$. Rewriting Eq. (2.79)

$$I(\boldsymbol{x}) = I_0(\boldsymbol{x_0}) \frac{\Delta S_0(\boldsymbol{x_0})}{\Delta S(\boldsymbol{x})}, \tag{2.92}$$

where ΔS_0 and ΔS are infinitesimally small surface areas connected by a tube of rays, and the mapping between $\boldsymbol{x}$ and $\boldsymbol{x_0}$ is defined by the same ray tube. Combining Eqs. (2.90), (2.91), and (2.92), we obtain

$$w_I(\boldsymbol{x}) = \lim_{\Delta S_0 \to 0, \, N \to \infty} I_0(\boldsymbol{x_0}) \frac{\Delta S_0(\boldsymbol{x_0}) N}{\Delta N_0(\boldsymbol{x_0}, N)} = \frac{I_0(\boldsymbol{x_0})}{p_0(\boldsymbol{x_0})}, \tag{2.93}$$

where we have used that $\Delta N_0(\boldsymbol{x_0}, N) = \Delta N(\boldsymbol{x}, N)$ if $\boldsymbol{x}$ and $\boldsymbol{x_0}$ are connected by a ray (tube). The expectation value of a Monte Carlo ray tracer will thus be the power incident on a plane, if each ray carries a power defined by Eq. (2.93). Since this does not involve the phase of the rays, we will denote this type of simulation by incoherent Monte Carlo ray tracing.

To calculate the *amplitude* of the field using a Monte Carlo ray tracer, one must adjust the weight of the rays. From Eq. (2.81) we see that by choosing

$$w(\boldsymbol{x}) = \lim_{\Delta S \to 0, \, N \to \infty} \sqrt{\frac{2 I_0 \Delta S_0 N}{\Delta N_0 \, \hat{\boldsymbol{n}} \cdot \hat{\boldsymbol{\rho}} n}} \left(\frac{\mu_0}{\epsilon_0}\right)^{1/4} \sqrt{\frac{\Delta S N}{\Delta N}}, \tag{2.94}$$

the expectation value of $w(\boldsymbol{X})$ is

$$\begin{aligned}
\langle w(\boldsymbol{X}) \rangle &= \iint_{S_p} \lim_{\Delta S_0 \to 0, \, N \to \infty} \sqrt{\frac{2 I_0 \Delta S_0 N}{\Delta N \, \hat{n} \cdot \hat{\rho} n}} \left(\frac{\mu_0}{\epsilon_0}\right)^{1/4} \sqrt{\frac{\Delta N}{\Delta S N}} \mathrm{d}^2 \boldsymbol{x}, \\
&= \iint_{S_p} \sqrt{\frac{2 I(\boldsymbol{x})}{\hat{n} \cdot \hat{\rho} n}} \left(\frac{\mu_0}{\epsilon_0}\right)^{1/4} \mathrm{d}^2 \boldsymbol{x}, \\
&= \iint_{S_p} |\boldsymbol{E}(\boldsymbol{x})| \mathrm{d}^2 \boldsymbol{x}. \tag{2.95}
\end{aligned}$$

In a simulation, different factors of the ray weight defined by Eq. (2.94) are implemented at different steps of the simulation. At the starting surface, every ray is assigned a weight proportional to the square root of the intensity divided by the local ray density

$$w_s = \sqrt{\frac{2 I_0 \Delta S_0 N}{\Delta N_0}} \left(\frac{\mu_0}{\epsilon_0}\right)^{1/4}. \tag{2.96}$$

If one has access to the electric field instead of the intensity one can rewrite Eq. (2.96), using Eq. (2.80), as[8]

$$w_s = \sqrt{\frac{\Delta S_0 N}{\Delta N_0}} \sqrt{n_0} |\boldsymbol{E_0}|, \tag{2.97}$$

---

[8]This assumes that the Poynting vector of the field is parallel to the local surface normal.

with n_0 the refractive index at S_0. At the final surface, upon incidence on a pixel, the weight is adjusted using the cosine of the angle between the ray direction and the surface normal, $\hat{\boldsymbol{n}} \cdot \hat{\boldsymbol{\rho}}$, and the local refractive index

$$w_i = \sqrt{\frac{1}{\hat{\boldsymbol{n}} \cdot \hat{\boldsymbol{\rho}} n}}. \tag{2.98}$$

Furthermore, the reciprocal of the local ray density is needed

$$w_p = \sqrt{\frac{\Delta SN}{\Delta N}}. \tag{2.99}$$

As stated in Section 2.4.2, this can be obtained by tracing neighbouring rays. So far we have assumed that the power carried by a ray remains constant, but the effect of transmission losses can easily be incorporated into an additional weight w_f, depending on the adjusted Fresnel equations, Eqs. (2.73, 2.74).

To obtain the full field we must also calculate the phase of the electric field. This is done by including the phase function $\exp[\mathrm{i}k_0 \mathrm{OPL}]$ into the ray weight w, where OPL is the optical path length of the ray (see Section 2.4.2). If the local wavefront and the pixel are not parallel, the phase will vary quickly over the pixel. We are not interested in the complex integral over the pixel, but in approximating the field at the centre of the pixel. Therefore we will approximate the field of the ray by a local plane wave and calculate its

Figure 2.4: Local plane wave approximation of a ray at a pixel.

contribution to the field in the centre of the pixel, $(\boldsymbol{x_p})$, using the adjusted phase function

$$\exp\left[\mathrm{i}\phi_p\right] = \exp\left[\mathrm{i}k_0\left(\mathrm{OPL} \pm d\, n \sin(\theta)\right)\right], \tag{2.100}$$

where θ is the angle between the incident ray and the local surface normal (see Figure 2.4) and

$$d = \sqrt{(x - x_p)^2 + (y - y_p)^2}, \tag{2.101}$$

is the distance between the centre of the pixel and the intersection point of the ray with the pixel, and the sign depends on the position of $\boldsymbol{x}$ with respect to $\boldsymbol{x_p}$. Using the final ray weight

$$w = w_s w_i w_p w_f \exp\left[\mathrm{i}\phi_p\right], \tag{2.102}$$

we can approximate the complex electric field at the centre of a pixel by averaging the electric fields of all properly weighted rays which arrive at the pixel, and dividing the complex result by the area of the pixel.

This Monte Carlo ray tracing procedure can be summarized in the following formula

$$\boldsymbol{E}(\boldsymbol{x}_p) \approx \frac{1}{NS_p} \sum_{j=1}^{N} \Pi\left[\boldsymbol{x}_p, \boldsymbol{x}(\boldsymbol{x}_j)\right] w(\boldsymbol{x}_j), \tag{2.103}$$

$$= \frac{1}{NS_p} \sum_{j=1}^{N} \Pi\left[\boldsymbol{x}_p, \boldsymbol{x}(\boldsymbol{x}_j)\right] w_s(\boldsymbol{x}_j) w_i(\boldsymbol{x}_j) w_p(\boldsymbol{x}_j) w_f(\boldsymbol{x}_j)$$
$$\exp\left[i\phi_p(\boldsymbol{x}_j)\right], \tag{2.104}$$

$$= \frac{1}{NS_p} \sum_{j=1}^{N} \Pi\left[\boldsymbol{x}_p, \boldsymbol{x}(\boldsymbol{x}_j)\right] \sqrt{\frac{\Delta S_0(\boldsymbol{x}_j)N}{\Delta N_0(\boldsymbol{x}_j)}} \sqrt{n_0} \boldsymbol{E_0}\left(\boldsymbol{x}_j\right)$$
$$\sqrt{\frac{1}{\hat{\boldsymbol{n}}\cdot\hat{\boldsymbol{\rho}}(\boldsymbol{x}_j)n}} \sqrt{\frac{\Delta S(\boldsymbol{x}_j)N}{\Delta N(\boldsymbol{x}_j)}} w_f(\boldsymbol{x}_j)$$
$$\exp\left[ik_0\left(\text{OPL}(\boldsymbol{x}_j) \pm d\left[\boldsymbol{x}(\boldsymbol{x}_j), \boldsymbol{x}_p\right] n \sin(\theta(\boldsymbol{x}_j))\right)\right], \tag{2.105}$$

where $\boldsymbol{x}_j$ is the position at the initial plane of a randomly sampled ray with index j, $\boldsymbol{x}(\boldsymbol{x}_j))$ is the position of this ray at the final plane, $\Pi(\boldsymbol{x}_p, \boldsymbol{x})$ is a rectangular function that evaluates to one if this ray position is within the pixel centred around $\boldsymbol{x}_p$, and S_p is the surface area of a pixel. Note that for geometrical optics to be valid every point on the first plane is mapped to a unique point on the second plane, which is why we can use $\boldsymbol{x}_j$ to indicate a unique ray.

Like the geometrical optical propagator, this adjusted Monte Carlo ray tracing propagator is thus able to propagate the complex electric field.

2.5. Exit pupil diffraction integrals

In section 2.4 it was explained how one can propagate the electric field through an optical system using geometrical optics, and that geometrical optics is not accurate when the amplitude of the field changes abruptly. Since imaging systems are designed to focus the light from an object point into a small spot, such a fast changing amplitude occurs within the focal region of an imaging system. As a result, geometrical optics is not valid in this domain. To calculate the field in the focal region on can use the exit pupil diffraction integral (EPDI), which combines geometrical optics with a diffraction integral.

In most imaging systems there exists an element that limits the bundle of rays that can pass through the system. Usually this is the aperture, a non-transparent screen with a hole in it. The exit pupil is the image of this aperture by all the optical elements after the aperture. The exit pupil plays a crucial role in the EPDI since it is the boundary between the geometrical optical simulation and the diffraction simulation.

The first step in an EPDI simulation is the propagation of the field of the object (e.g., a point source) to the exit pupil using ray tracing. Rays hitting the aperture are blocked, resulting in a sharp cut-off of the field. This violates one of the conditions of geometrical optics and is therefore a source of diffraction. In this thesis surfaces that produce such sharp cut-offs are called diffracting surfaces. In the EPDI all the causes of diffraction in an optical system are modelled in the exit pupil.

The second step in an EPDI simulation is the propagation of the field from the exit

pupil to the focal region. Since the exit pupil is by definition in the same uniform medium as the focal region, one can apply a diffraction integral for this propagation step. The

Figure 2.5: Schematic representation of the exit pupil diffraction integral. The rays are first traced to the exit pupil (depicted in blue), which is the (in this case virtual) image of the aperture (shown in black) by the lens behind it. The field is then propagated using a diffraction integral, which basically computes the superposition of spherical waves illustrated in the figure.

Rayleigh-Sommerfeld diffraction integral (see section 2.3.1) could be used for this purpose, but the Richards-Wolf integral [5] is also frequently used. The Richards-Wolf integral is an integral over plane waves. The direction, amplitude, and initial phase of a plane wave starting at a point in the exit pupil depend on the same quantities of the ray through this point.

By combining the geometrical propagation with a diffraction integral, the exit pupil diffraction integral accounts for the full geometry of the optical systems and for the diffraction at the aperture. It is a powerful method for calculating the intensity distribution in the focal plane of a system, the so-called point spread function (PSF). The EPDI is therefore used by many commercial ray tracers (e.g., CODE V, OSLO and ZEMAX).

2.6. RAY-BASED DIFFRACTION METHODS

A limitation of the exit pupil diffraction method is that it assumes that all diffraction in the optical system effectively takes place in the exit pupil. This assumption is not always sufficiently accurate when modelling systems with multiple consecutive apertures or when propagating a coherent diffracting field (e.g., a laser beam) through an optical systems with an aperture. This limitation is overcome by a group of methods which we will denote as ray-based diffraction methods. These methods rely on ray tracing from the first to the final surface of the simulation, while adjusting the standard ray-tracing procedure to account for the effects of (multiple) diffraction. For all ray-based diffraction methods this involves the generation of secondary rays at diffracting surfaces (e.g., the aperture plane), but the theory behind these secondary rays and the quantity carried by the rays differ between the methods.

In the Gaussian beam decomposition (GBD) method (see, e.g., [8]) the field is written as a superposition of Gaussian beams. For every Gaussian beam a small set of rays is propagated through the optical system. The field of the beam after this propagation

step can be calculated from the positions and directions of the corresponding rays. The total field is simply the sum of the fields of the individual Gaussian beams. The approximation in the propagation step is that the part of the system through which a Gaussian beam passes can be considered to be paraxial. This parabasal approximation is different from the standard paraxial approximation. The standard paraxial approximation uses the same paraxial system for the entire field, whereas in the parabasal approximation a different paraxial system is used for every beamlet. The smaller the extent of the beamlet, the more accurate this approximation is. A simulation with the GBD may have several steps, each consisting of decomposing, propagating, and a rebuilding the field. Apertures can be accounted for by cutting off part of the rebuilt field before the next decomposition step. The GBD is implemented in the commercial software package ASAP [23]. The beam synthesis propagation (BSP) in CODE V [9] (another commercial ray tracer) uses an unspecified beamlet decomposition that shares at least some of its properties with the GBD.

The vectorial ray-based diffraction integral (VRBDI) developed by Andreas *et al.* [10] decomposes an input field into plane waves or spherical waves. It then propagates these fields using ray tracing, connecting every point of the input grid to every point of the output grid. By adding the complex values (i.e., amplitude and phase) of all rays arriving at the same grid point in the output plane, the method can account for the diffraction introduced at the input plane. Andreas *et al.* mention that apertures could be accounted for by first calculating the field in the aperture using a VRBDI simulation, then applying the aperture function, and finally decomposing and propagating this field using a second VRBDI simulation, but the authors do not present such simulations [10].

The geometrical theory of diffraction proposed by Keller [24] models diffraction by distinguishing between the rays passing through an aperture and the rays that hit the edge of an aperture. The former 'ordinary rays' pass through the aperture undisturbed, whereas the latter generate cones of 'diffracted rays' emanating from the aperture edge. The initial values of these diffracted rays depend on the incident field and a diffraction coefficient determined using canonical solutions to diffraction problems. For these canonical solutions, the field after an aperture is split into an undisturbed field that passes through the aperture and a field originating from the edge of the aperture. The field from the edge is used to determine the diffraction coefficients. In regions where the ordinary or diffracted rays intersect (i.e., at caustics) the method described so far would predict an infinite amplitude of the field. Therefore an additional caustic correction factor is needed, which is again derived from canonical solutions. The geometrical theory of diffraction is mostly used for solving problems relating to waveguides and reflector antennas [25] for radio waves. An overview of the method is given by James [25].

Some of the other ray-based diffraction methods are stable aggregate of flexible elements [26–30] and Heisenberg uncertainty ray bending [31–33]. The former provides accurate results but has so far only been developed for two-dimensional space. The latter, which initiates secondary rays in the neighbourhood of the aperture edge, does treat diffraction in three-dimensional space. However, a distribution of the secondary rays over the radiation cones that consistently gives correct results could not be determined conclusively [31–33].

The methods which are the main topic of this thesis, Wigner-based ray tracing and

Huygens-Fresnel path integration, are also ray-based diffraction methods that initiate secondary rays at diffracting surfaces. A thorough discussion of these methods and related literature is provided in Chapters 3, 4, and 5.

2.7. WIGNER DISTRIBUTION FUNCTION

The Wigner distribution function (WDF) is a transformation of a complex function, introduced in 1932 by Eugene Wigner [11]. It is a function of both a variable and its conjugate (e.g., space and spatial frequency), and is therefore sometimes seen as the local frequency spectrum of a function. This concept may be uncommon in optics, but it should be familiar to anyone who has ever read music from a score. The WDF was first introduced into optics in a 1964 paper by Dolin [12] and later again by Walther [34] as a function that could describe the spatial coherence of a field, yet still remain invariant along straight lines (i.e., rays). In 1978, Bastiaans showed that upon propagating through a paraxial optical system, the WDF experiences the same coordinate transformation as geometrical optical rays [35]. Since the WDF is a representation of the complex scalar field, it was presented as a link between the wave optical description of light (discussed in Section 2.3.2) and geometrical optics (treated in Section 2.4) [35]. This link is at the core of Wigner-based ray tracing, the topic of Chapter 3. This section presents the basic properties of the WDF and discusses how the WDF changes upon propagation through a paraxial system.

The mathematical definition of the WDF is

$$W(\boldsymbol{x}, \boldsymbol{\xi}) = \iint E\left(\boldsymbol{x} + \frac{\boldsymbol{x'}}{2}\right) E^*\left(\boldsymbol{x} - \frac{\boldsymbol{x'}}{2}\right) \exp\left[-\mathrm{i}\boldsymbol{\xi} \cdot \boldsymbol{x'}\right] \mathrm{d}^2 \boldsymbol{x'}, \tag{2.106}$$

where W is the WDF, $\boldsymbol{x} = (x, y)$ represents a point in a plane in ordinary space, $\boldsymbol{\xi} = (\xi, \eta)$ is a point in its reciprocal plane, E is a scalar complex field, E^* its complex conjugate, and the integral is over all the elements of the vector $\mathrm{d}^2 \boldsymbol{x'} = \mathrm{d}x' \mathrm{d}y'$. The WDF is always real, but not necessarily positive (see e.g., [18]). When used in optics, the WDF can serve as a representation of the scalar complex electric field, which is how we will use it from now on.

As an example we discuss the WDF of a plane wave. The electric field of a plane wave in a plane (with a constant coordinate z) is given by

$$E(\boldsymbol{x}) = A \exp^{\mathrm{i}\boldsymbol{k} \cdot \boldsymbol{x}}, \tag{2.107}$$

with $\boldsymbol{k} = (k_x, k_y)$ the vector containing the two dimensional spatial frequencies and A the amplitude. Applying the transformation Eq. (2.106) gives

$$W(\boldsymbol{x}, \boldsymbol{\xi}) = \iint A e^{\mathrm{i}\boldsymbol{k} \cdot (\boldsymbol{x} + \boldsymbol{x'}/2)} A^* e^{-\mathrm{i}\boldsymbol{k} \cdot (\boldsymbol{x} - \boldsymbol{x'}/2)} e^{-\mathrm{i}\boldsymbol{\xi} \cdot \boldsymbol{x'}} \mathrm{d}^2 \boldsymbol{x'}, \tag{2.108}$$

$$= \iint |A|^2 e^{\mathrm{i}\boldsymbol{x'} \cdot (\boldsymbol{k} - \boldsymbol{\xi})} \mathrm{d}^2 \boldsymbol{x'}, \tag{2.109}$$

$$= (2\pi)^2 |A|^2 \delta(\boldsymbol{\xi} - \boldsymbol{k}), \tag{2.110}$$

where we have used the exponential representation of the dirac function

$$\delta(\boldsymbol{x}) = \iint \exp\left[-2\pi \mathrm{i}\boldsymbol{x'} \cdot \boldsymbol{x}\right] \mathrm{d}^2 \boldsymbol{x'}. \tag{2.111}$$

So in the phase-space representation of the WDF, the plane wave will be a delta function in frequency coordinates with a uniform amplitude along the spatial coordinates.

If E represents the scalar electric field, the integral over the frequency coordinates of the WDF gives the intensity, and the integral over the spatial coordinates gives the intensity of the plane wave expansion of the field [35]

$$I(\boldsymbol{x}) = \frac{1}{(2\pi)^2} \iint W(\boldsymbol{x}, \boldsymbol{\xi}) \, \mathrm{d}^2 \boldsymbol{\xi}, \tag{2.112}$$

$$|\tilde{F}(\boldsymbol{\xi})|^2 = \iint W(\boldsymbol{x}, \boldsymbol{\xi}) \, \mathrm{d}^2 \boldsymbol{x}, \tag{2.113}$$

with $\tilde{F}(\boldsymbol{\xi})$ the complex amplitude of the plane wave corresponding to spatial frequency $\boldsymbol{\xi}$. Applying these integrals to the WDF of a plane wave, Eq. (2.110), we find an intensity of $I = |A|^2$ and an intensity of the plane wave expansion of $|\tilde{F}(\boldsymbol{\xi})|^2 = (2\pi)^2 |A|^2 \delta(\boldsymbol{\xi} - \boldsymbol{k})$. Although the inverse Wigner transformation exists (see Appendix C.1), we will not use it in this thesis since Eq. (2.112) already provides us with a direct way to obtain the intensity from the WDF.

So far it was shown that the WDF is simply another representation of the electric field from which one can easily obtain the intensity. What makes this representation particularly useful, is that the propagation of a WDF through a paraxial optical system is governed by some very simple rules. We first derive the WDF of a field propagated using the Fresnel integral and then give the more general solution for a field propagated by the Collins integral (see Section 2.3.2 for a description of these integrals).

The Fresnel diffraction integral, given by Eq. (2.50), describes paraxial propagation through a uniform media. Let E_0 be the field before propagation and let W_0 be its WDF. The WDF of the Fresnel propagated field is then

$$W(\boldsymbol{x}, \boldsymbol{\xi}) = \iint \frac{-\mathrm{i}k \exp(\mathrm{i}kz)}{2\pi z} \iint E_0(\boldsymbol{x_0}) \exp\left[\frac{\mathrm{i}k}{2z}\left|\boldsymbol{x} + \frac{\boldsymbol{x}'}{2} - \boldsymbol{x_0}\right|^2\right] \mathrm{d}^2 \boldsymbol{x_0}$$

$$\frac{\mathrm{i}k \exp(-\mathrm{i}kz)}{2\pi z} \iint E_0^*(\bar{\boldsymbol{x}}_0) \exp\left[\frac{-\mathrm{i}k}{2z}\left|\boldsymbol{x} - \frac{\boldsymbol{x}'}{2} - \bar{\boldsymbol{x}}_0\right|^2\right] \mathrm{d}^2 \bar{\boldsymbol{x}}_0$$

$$\exp\left[-\mathrm{i}\boldsymbol{\xi} \cdot \boldsymbol{x}'\right] \mathrm{d}^2 \boldsymbol{x}', \tag{2.114}$$

$$= \frac{k^2}{4\pi^2 z^2} \iint \iint \iint E_0(\boldsymbol{x_0}) E_0^*(\bar{\boldsymbol{x}}_0)$$

$$\exp\left[\frac{\mathrm{i}k}{2z}\left(|\boldsymbol{x_0}|^2 - |\bar{\boldsymbol{x}}_0|^2 + 2\boldsymbol{x} \cdot (\bar{\boldsymbol{x}}_0 - \boldsymbol{x_0})\right)\right]$$

$$\exp\left[\frac{\mathrm{i}k}{2z}\boldsymbol{x}' \cdot \left(2\boldsymbol{x} - \boldsymbol{x_0} - \bar{\boldsymbol{x}}_0 - \frac{2z}{k}\boldsymbol{\xi}\right)\right] \mathrm{d}^2 \boldsymbol{x}' \mathrm{d}^2 \boldsymbol{x_0} \mathrm{d}^2 \bar{\boldsymbol{x}}_0. \tag{2.115}$$

By using the dirac delta function, Eq. (2.111), and its property

$$\delta\left(\frac{\boldsymbol{x}}{a}\right) = a^2 \delta(\boldsymbol{x}), \tag{2.116}$$

the integral over $\boldsymbol{x}'$ reduces to

$$\iint \exp\left[\frac{\mathrm{i}k}{z}\boldsymbol{x}' \cdot \left(\boldsymbol{x} - \frac{\boldsymbol{x_0} + \bar{\boldsymbol{x}}_0}{2} - \frac{z}{k}\boldsymbol{\xi}\right)\right] \mathrm{d}^2 \boldsymbol{x}' = \frac{4\pi^2 z^2}{k^2} \delta\left(-\boldsymbol{x} + \frac{\boldsymbol{x_0} + \bar{\boldsymbol{x}}_0}{2} + \frac{z}{k}\boldsymbol{\xi}\right). \tag{2.117}$$

Using this delta function and the coordinate transformations (with unit Jacobian)

$$x_0 = \tilde{x}_0 + \frac{\tilde{x}'_0}{2}, \tag{2.118}$$

$$\bar{x}_0 = \tilde{x}_0 - \frac{\tilde{x}'_0}{2}, \tag{2.119}$$

gives the final WDF

$$
\begin{aligned}
W(\boldsymbol{x}, \boldsymbol{\xi}) &= \frac{k^2}{4\pi^2 z^2} \iint \iint E_0\left(\tilde{x}_0 + \frac{\tilde{x}'_0}{2}\right) E_0^*\left(\tilde{x}_0 - \frac{\tilde{x}'_0}{2}\right) \\
&\quad \exp\left[\frac{\mathrm{i}k}{z}\left(\tilde{x}_0 \cdot \tilde{x}'_0 - \boldsymbol{x} \cdot \tilde{x}'_0\right)\right] \\
&\quad \frac{4\pi^2 z^2}{k^2} \delta\left(\frac{z}{k}\boldsymbol{\xi} - \boldsymbol{x} + \tilde{x}_0\right) \mathrm{d}^2\tilde{x}_0 \mathrm{d}^2\tilde{x}'_0,
\end{aligned} \tag{2.120}
$$

$$
\begin{aligned}
&= \iint E_0\left(\boldsymbol{x} - \frac{z}{k}\boldsymbol{\xi} + \frac{\tilde{x}'_0}{2}\right) E_0^*\left(\boldsymbol{x} - \frac{z}{k}\boldsymbol{\xi} - \frac{\tilde{x}'_0}{2}\right) \\
&\quad \exp\left[-\mathrm{i}\boldsymbol{\xi} \cdot \tilde{x}'_0\right] \mathrm{d}^2\tilde{x}'_0,
\end{aligned} \tag{2.121}
$$

$$= W_0\left(\boldsymbol{x} - z\frac{\boldsymbol{\xi}}{k}, \boldsymbol{\xi}\right). \tag{2.122}$$

So the WDF of the field after Fresnel propagation over a distance z (W), is the WDF of the field before propagation, evaluated at the coordinate $\boldsymbol{x}_0 = \boldsymbol{x} - z\frac{\boldsymbol{\xi}}{k}$. A propagation through a uniform media thus affects the WDF through a coordinate transformation

$$\boldsymbol{x} = \boldsymbol{x}_0 + z\frac{\boldsymbol{\xi}}{k}, \tag{2.123}$$

$$\boldsymbol{\xi} = \boldsymbol{\xi}_0, \tag{2.124}$$

where the index zero indicates the quantities before propagation. This is the same coordinate transformation as that for a ray with initial position $\boldsymbol{x}_0$ and initial direction $\boldsymbol{p}_0 = \frac{\boldsymbol{\xi}}{k}$, propagated over a distance z.

   This equivalence between the coordinate transformation of the WDF and of rays is not only valid for free space propagation, but holds for the propagation through any paraxial system [35]. So a system that, in matrix optics, can be described by

$$
\begin{pmatrix} \boldsymbol{x}_0 \\ \boldsymbol{p}_0 \end{pmatrix} = \begin{bmatrix} \mathbb{D}^T & -\mathbb{B}^T \\ (\mathbb{B}^{-1} - \mathbb{A}^T \mathbb{D} \mathbb{B}^{-1}) & \mathbb{A}^T \end{bmatrix} \begin{pmatrix} \boldsymbol{x} \\ \boldsymbol{p} \end{pmatrix}, \tag{2.125}
$$

which is the inverse of Eq. (2.52), will perform the following coordinate transformation on the WDF

$$W(\boldsymbol{x}, \boldsymbol{\xi}) = W_0\left[\mathbb{D}^T \boldsymbol{x} - \frac{1}{k}\mathbb{B}^T \boldsymbol{\xi}, k\left(\mathbb{B}^{-1} - \mathbb{A}^T \mathbb{D} \mathbb{B}^{-1}\right)\boldsymbol{x} + \mathbb{A}^T \boldsymbol{\xi}\right], \tag{2.126}$$

where W_0 is the WDF before and W is the WDF after propagation through the system.

This coordinate transformation is identical to that of Eq. (2.125) if we choose

$$p_0 = \frac{\xi_0}{k}, \tag{2.127}$$

$$p = \frac{\xi}{k}, \tag{2.128}$$

which relates the directional cosines of the rays to the spatial frequencies of the WDF. A proof of this property, based on the Collins integral, Eq. (2.51), is given in Appendix C.3.

An optical system will generally contain at least one aperture, but the matrix formalism is unable to account for such apertures. We therefore still need to find the effect of an aperture on the WDF.

In this thesis we treat apertures in the thin element approximation (see Section 2.3.4), in which the effect of an element on the electric field is modelled by multiplying the field by a mask function B. We are therefore interested in the WDF of the field

$$E(\boldsymbol{x}) = B(\boldsymbol{x})E_0(\boldsymbol{x}), \tag{2.129}$$

where E_0 is the field before, and E the field after the element. By calculating the WDF of Eq. (2.129), one can show (see Appendix C.2) that

$$W(\boldsymbol{x}, \boldsymbol{\xi}) = \frac{1}{(2\pi)^2} \iint W_0\left(\boldsymbol{x}, \boldsymbol{\xi}'\right) W_B\left(\boldsymbol{x}, \boldsymbol{\xi} - \boldsymbol{\xi}'\right) \mathrm{d}^2\boldsymbol{\xi}', \tag{2.130}$$

where W_B is the WDF of the mask function and W_0 is the WDF of the field before the thin element.

The definition for the WDF given in Eq. (2.106) is valid for coherent fields, but it can be adapted to partially coherent or incoherent fields by substituting $E\left(\boldsymbol{x} + \frac{\boldsymbol{x}'}{2}\right)E^*\left(\boldsymbol{x} - \frac{\boldsymbol{x}'}{2}\right)$ by the statistical average of the same quantity (which is a representation for fields of all degrees of spatial coherence). All relations derived in this section also hold for this WDF of partially coherent fields, which makes the WDF particularly useful for propagating such fields.

We have seen that the WDF after propagation through a paraxial optical system is given by a coordinate transformation of the WDF before the system, and that the WDF behind a thin element is given by a convolution between the WDF before the element and the WDF of the element itself.

3

WIGNER-BASED RAY TRACING

The topic of this chapter is a simulation method that uses the Wigner distribution function to incorporate wave-optical effects in ray tracing software. The method is used to calculate the intensity distribution in the focal region, and we show that it is accurate for paraxial systems but produces unphysical results in the presence of aberrations. The cause of these anomalies is explained using an analytical model.

3.1. INTRODUCTION

Because the coordinate transformations of the WDF and of paraxial rays are identical (see Section 2.7), the WDF remains constant along paths followed by paraxial rays. To our knowledge, the first usage of this property for ray-based diffraction simulations, was in the field of computer graphics [38]. Oh *et al.* rendered images that showed diffraction patterns, e.g., of a light source behind a square aperture casting a diffraction pattern on a cubic box. They also computed PSFs for two camera lenses with different numerical apertures (NAs, the product of the refractive index and the sine of the half cone of light in the image space) [38]. Although some of the diffraction patterns were validated using Fourier optics, the PSFs were only evaluated qualitatively. The method by Oh *et al.* was further developed in the computer graphics community (e.g., [14, 39, 40]), but PSF calculations were never validated.

The before mentioned research raises the question whether a ray-based transport of the WDF can be used for the simulation of the propagation of coherent scalar electric fields through optical systems. It is this question that we set out to address in the current chapter. First, the method under investigation is presented in Section 3.2. The numerical method is then applied to calculate diffraction in free space and the PSF of a thick and a thin lens in Section 3.3. To substantiate the conclusions drawn from those numerical simulations, Section 3.4 investigates an analytical model based on the same principles as the numerical model. In Section 3.5 the results of the numerical and analytical model are discussed. We end the chapter with the conclusions in Section 3.6.

3.2. METHOD

The principles described in references [14, 38–40] are implemented in a simulation method, which we call Wigner-based ray tracing (WBRT). The aim of WBRT is to simulate the propagation of an electric field from an input plane, through an optical system consisting of, e.g., free space, lenses and apertures, and calculate the intensity distribution at an output plane. A simulation consists of several steps, which are illustrated in the flow chart in Fig. 3.1 and described in the following paragraphs.

Figure 3.1: A flow chart illustrating the steps of WBRT.

The first step of WBRT is to derive the WDF of the electric field in the input plane. For some frequently used input fields, the WDF can be calculated analytically. The WDF of a plane wave is for instance given by Eq. (2.110). WBRT performs a Monte Carlo sampling of the region of the four dimensional phase space for which the WDF is non-zero ($W(\boldsymbol{x}, \boldsymbol{\xi}) \neq 0$). The algorithm chooses a (user defined) number of random points within this region of phase space, using a probability distribution function that is uniform in

spatial and frequency coordinates. Every such point in phase space is associated with a ray with the same spatial coordinates $\boldsymbol{x}$ and directions defined by Eq. (2.128), that caries the local WDF.

The second step is to use these rays to propagate the WDF through an optical system, e.g., a series of lenses in free space. The lenses will affect the ray paths according to Snell's law, but in WBRT they will not introduce diffraction unless they are combined with an aperture. Within the paraxial theory this procedure is exact, since the WDF remains constant along the paths of paraxial rays. The underlying assumption of WBRT, is that this property is also a good approximation beyond the paraxial domain.

At an aperture or (complex) mask, the ray tracing is interrupted. The WDF of the field after a mask is, as mentioned in Section 2.7, given by

$$W(\boldsymbol{x}, \boldsymbol{\xi}) = \frac{1}{(2\pi)^2} \iint W_0\left(\boldsymbol{x}, \boldsymbol{\xi}'\right) W_B\left(\boldsymbol{x}, \boldsymbol{\xi} - \boldsymbol{\xi}'\right) \mathrm{d}^2 \boldsymbol{\xi}'. \tag{3.1}$$

This means that the WDF directly after a mask can be calculated by a convolution of the input WDF (W_0) with the WDF of the mask (W_B) over the frequency variables [35]. Since the mask function can be complex it can be used to describe aberrations within the thin element approximation. Such a procedure has for instance been used to study the effect of aberrations in the exit pupil [41], but requires an additional (ray-based) method to calculate the aberrations. In WBRT aberrations are implicitly captured by the ray tracing procedure whereas diffraction is explicitly introduced by convolving the WDF of the field and the aperture function. Note that as a result of the convolution in Eq. (3.1), the WDF is smeared out over the frequency variables but remains localized in spatial coordinates. In our ray-based model this is implemented by letting every ray initiate a set of new rays. These rays have the same initial position but differ in their directions and carry a value of the WDF in accordance with Eq. (3.1). The set of secondary rays is created using a Monte Carlo procedure similar to the one used for the creation of the initial rays. This follows the example of the work by Oh *et al.* [38] and makes the method very suited for parallel computing. In this thesis we treat apertures using the Kirchhoff boundary conditions (see Section 2.3.4). As a result, the mask function, $B(\boldsymbol{x})$, is a binary function with the value 1 inside, and 0 outside of the aperture. The WDF of a square aperture can be calculated analytically and the result is given in Section 3.4. For the WDF of circular apertures of radius a, Bastiaans derived the integral [42],

$$W_B\left[\begin{pmatrix} x \\ 0 \end{pmatrix}, \begin{pmatrix} \xi \\ \eta \end{pmatrix}\right] = 16a^2 \Re\left[\exp[-\mathrm{i}2\xi x] \int_{x/a}^1 \frac{\sin\left[2a\eta\sqrt{1-s^2}\right]}{2a\eta} \exp[\mathrm{i}2a\xi s]\mathrm{d}s\right], \tag{3.2}$$

which we will use.

After the aperture the electric field will be propagated further through the optical system using the same procedure as for the propagation from the input plane to the aperture: every ray carries a value of the WDF and propagates through the optical system in accordance with the laws of geometrical optics. This procedure can be used for paraxial and non-paraxial propagation and its validity, particularly beyond the paraxial domain, is the topic of the following sections. If the system contains any further apertures they too can be treated in accordance with Eq. (3.1).

As a final step, the intensity distribution at the output plane (or any other plane of interest) is calculated. Since the intensity distribution is given by the frequency marginal of the WDF (see Eq. (2.112)), the intensity in a pixel can be calculated by summing the values of the WDF carried by each ray that hits the pixel. Using these principles, one can obtain an intensity distribution in a plane caused by the propagation of a coherent field through an optical system.

3.3. Results

The procedure described in the previous section is implemented in MATLAB, and used to simulate diffraction at an aperture in free space and the propagation of a plane wave, through a thick and through a thin lens, for wavefronts with and without aberrations.

3.3.1. Diffraction at an aperture in free space

The diffraction patterns of a circular aperture illuminated by a plane wave and a rectangular aperture illuminated by a point source are calculated using WBRT. The parameters of the two simulated systems are listed in Table 3.1 and the results shown in Figs. 3.2

Table 3.1: The parameters of the systems used for free-space diffraction simulations.

Parameter	system 1	system 2
source	point source	plane wave
distance to aperture	500 mm	-
aperture shape	square	circular
aperture size	1x1 mm	1 mm (radius)
propagation after aperture	1000 mm	100 mm
wavelength	600 nm	600 nm

and 3.3. The cross sections of the intensity patterns also show the results of a numerical Rayleigh-Sommerfeld diffraction integral (see Section 2.3.1), which match well with those of the WBRT simulations apart from some fluctuations which are due to statistical noise. Considering WBRT is to some extend based on the paraxial approximation of the Rayleigh-Sommerfeld integral (i.e., the Fresnel integral), this result was to be expected for these paraxial problems and it mainly validates the numerical implementation of the method.

3.3.2. Thick lens

The results of the propagation of a plane wave through a thick lens are shown in Fig. 3.4a and 3.4b (the parameters of the thick lenses are listed in Table 3.2). Apart from the WBRT results, the figures also show the results of a Rayleigh-Sommerfeld exit pupil diffraction integral (RSED) [43] (see Section 2.5) obtained with a Zeiss internal optical simulation software package. For both figures the WBRT results are averaged over the polar angle. As a result the convergence is better at larger distances from the origin, because the average there is taken over a larger number of Monte Carlo samples. Fig. 3.4a shows the

(a) Normalized intensity distribution   (b) Cross section at y=0

Figure 3.2: The diffraction pattern of a spherical wave passing through a rectangular aperture, simulated using WBRT. The cross section shows the result from a numerical Rayleigh-Sommerfeld diffraction integral (RSDI) as a reference. Intensities are normalized with respect to the maximum intensity. The system parameters are listed in Table 3.1.

(a) Normalized intensity distribution   (b) Cross section at y=0

Figure 3.3: The diffraction pattern of a plane wave passing through a circular aperture, simulated using WBRT. The cross section shows the result from a numerical Rayleigh-Sommerfeld diffraction integral (RSDI) as a reference. Intensities are normalized with respect to the maximum intensity. The system parameters are listed in Table 3.1.

results for axial imaging by a singlet with an NA of 0.2 that introduces no third or fifth order spherical aberrations to the wavefront. The results from WBRT and the RSED are in very good agreement. For the other singlet, which has an NA of 0.04, and wavefront aberrations with root mean square (rms) values[1] of 0.23 wavelengths for defocus and 0.12 wavelengths for primary spherical aberrations, WBRT results in negative intensities, which is clearly unphysical. It also indicates that WBRT results in a distribution other than the WDF, since the marginals of the WDF are always non-negative (see, e.g., Ref. 18).

---

[1] The rms wavefront error or aberration is defined as the square root of the average of the squared wavefront errors.

Table 3.2: The parameters of the simulated thick lenses. The distance from the object to the first lens surface is D_0, R_1 and R_2 are the radii of curvature of the first and the second lens surface, D_1 is the distance between these surfaces, and D_2 is the distance from the last lens surface to the image plane (all in mm). The stars indicate conical surfaces. Alle lenses have a refractive index of 1.51 and all simulations are performed at a wavelength of 600 nm.

System name	D_0	R_1	D_1	R_2	D_2
Singlet with def. and sph. aber.	∞	281.67	10	-1603.53	465
Aberration-free singlet	800	333.69	10	-297.01^*	501
Spherical	∞	298.21	10	-1904.02	500
Asphere A	∞	298.14^*	10	-1890.19^*	500
Asphere B	800	333.69	10	-297.01^*	501
Asphere C	∞	298.21	10	-1904.03^*	500

(a) Aberration-free singlet   (b) Singlet with defocus and spherical aberrations.
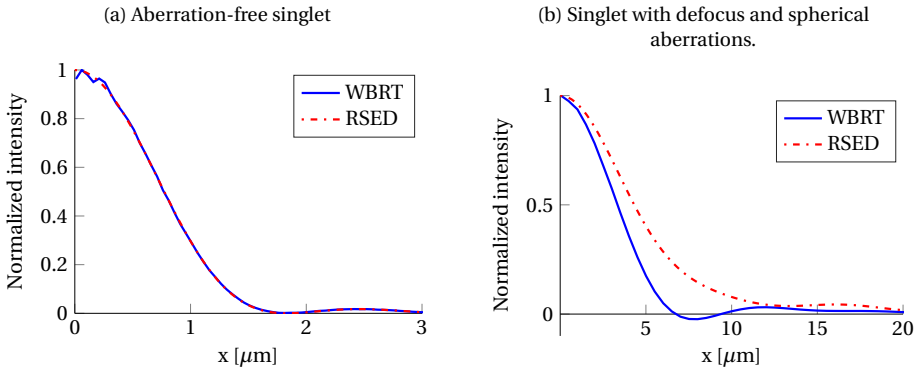
Figure 3.4: Cross sections of the PSF of a singlet with an aperture radius of 100 mm and an NA of 0.2 (left) and of a singlet with an aperture radius of 20 mm and an NA of 0. 04 (right). The wavefront of the singlet with an NA of 0.2 is practically free of primary aberrations, whereas the singlet with an NA of 0.04 has a defocus and spherical aberrations.

These results indicate that the validity of WBRT mainly depends on the wavefront aberrations. This is further supported by Figs. 3.5 and 3.6, where the agreement between the WBRT and RSED simulation is quantified in the L_{2A}-difference of the simulated PSFs, defined by

$$L_{2A} = \sqrt{\frac{\sum_p \left(|E'(p)| - |E(p)|\right)^2}{\sum_p |E'(p)|^2}},\tag{3.3}$$

where $|E'(p)|$ is the amplitude at pixel p predicted by the RSED and $|E(p)|$ the corresponding amplitude given by WBRT. The figures show the results for one spherical and three aspherical singlets with varying aperture radii. For these singlets, the radius at which the WBRT and RSED results start to diverge (see Fig. 3.5) varies between 10 mm and 120 mm, corresponding to NAs of 0.02 and 0.23. In Fig. 3.6 the same results are plotted, but with the rms wavefront error (excluding tip/tilt and defocus) instead of the aper-

ture radius on the horizontal axis. From rms wavefront errors of around 10^{-2} wavelength onwards, the L_{2A}-difference between the WBRT and RSED results starts to increase for all four singlets. Although the rms wavefront error at which this occurs is different for every system, simulations at different NAs and for different spherical and aspherical singlets show similar results: the results of WBRT and the RSED are in good agreement until spherical aberrations start influencing the PSF. Unless future research finds a way to solve this difficulty with an approach that differs significantly from the present one, it is safer not to use WBRT for thick lenses that suffer from aberrations. Other simulations

3
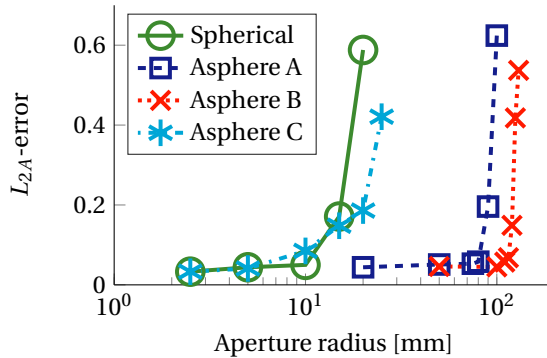
Figure 3.5: The L_{2A}-difference between WBRT and RSED simulations of four thick lens designs at varying aperture radii. All systems consist of a singlet followed by a circular aperture and a free-space propagation over roughly 500 mm (see Table 3.2 for further system parameters). The systems are illuminated at 600 nm by a plane wave or a point source 800 mm in front of the lens (asphere B). The error increases sharply at aperture radii of around 10 mm, 90 mm and 120 mm respectively.

Figure 3.6: The L_{2A}-difference between WBRT and RSED simulations of the four lens designs used in Fig. 3.5 simulated at the same aperture radii as for Fig. 3.5. The horizontal axis shows the rms wavefront error (excluding tip/tilt and defocus). For all four systems the L_{2A}-difference starts to increase significantly for rms wavefront errors onwards of 10^{-2} wavelength.

show that for systems with very small aberrations and NAs larger than 0.3 the results from WBRT and RSED also diverge, indicating the limits of WBRT due to nonparaxial effects. For the aberration-free wavefronts we did not encounter negative intensities.

3.3.3. THIN LENS

To study the effects of aberrations on WBRT, the simulation of a singlet is repeated using a thin-element model. Within the thin-element approximation all effects of the lens (i.e., phase changes or changes in ray direction) occur in a plane (see Section 2.3.4). The effect of a thin lens that introduces aberrations can be modelled by two separate stages. First, the lens will change the direction of each ray in accordance with perfect imaging. Second, the aberrations will cause small additional directional changes of the rays.

Results of the thin-lens model for lenses with NAs of 0.20, 0.45 and 0.9 are shown in Fig. 3.7 and Fig. 3.8. The WBRT results in the focal plane of aberration-free lenses are in very good agreement with those of reference methods at NAs of 0.20 and 0.45. At an

Figure 3.7: Cross section of the PSF of an ideal thin lens and the PSF of a thin lens that introduces spherical aberrations calculated using WBRT (left) and a cross section of the intensity in a defocused plane for the ideal thin lens (right). The lenses have an NA of 0.45, an aperture radius of 5 mm, and the wavelength is 600 nm. The result of the ideal thin lens is in good agreement with the corresponding Airy disk, the spherical aberration produces unphysical negative intensities. In the defocused plane the WBRT results do not match the EPDI results.

NA of 0.9 one notices only a small deviation between the WBRT results and those of the EPDI. Away from the focal area, the WBRT and EPDI results match when the NA is 0.20, but not when it increases to 0.45. When spherical aberration is introduced, WBRT gives unphysical results in the form of spatial regions with a negative intensity, even at an NA as low as 0.20. In summary, for an ideal thin lens WBRT produces the right PSF until very large NAs whereas defocused PSFs are only accurate at small NAs and spherical aberration can cause WBRT to produce negative intensities at all NAs.

3.3.4. CONVERGENCE

We simulated a thin and a thick lens with a varying number of rays to determine the convergence of WBRT. The difference between the WBRT results and those of an RSED
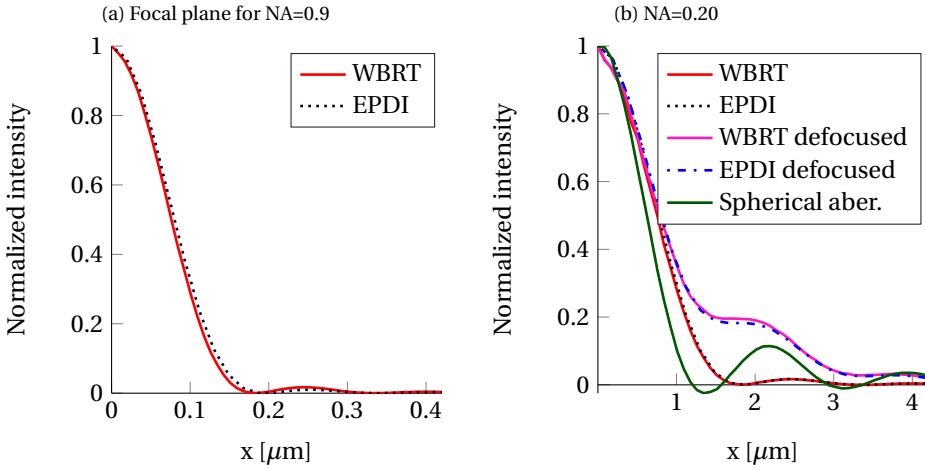
Figure 3.8: The PSF for a thin lens with a high NA (left) and a low NA (right). For the low-NA lens the intensity in the focal plane and in the defocused plane predicted by WBRT match those from the EPDI. The WBRT results for the low-NA lens with spherical aberrations show unphysical negative intensities. For the high-NA lens the WBRT results agree fairly well with those from the EPDI. Both lenses are illuminated by a plane wave with a wavelength of 600 nm and the aperture radii are 2 and 20 mm, respectively.

simulation is quantified using the L_{2I}-difference, defined by

$$L_{2I} = \sqrt{\frac{\sum_p \left(I'(p) - I(p)\right)^2}{\sum_p I'(p)^2}}, \tag{3.4}$$

where $I'(p)$ is the intensity at pixel p predicted by the RSED and $I(p)$ the corresponding intensity given by WBRT. The results, shown in Fig. 3.9, demonstrate that the L_{2I}-difference is inversely proportional to the square root of the number of rays (N). This convergence rate of order $\mathcal{O}(N^{-1/2})$ is a general property of Monte Carlo integration (see Section 2.4.3).

3.4. ANALYTICAL MODEL

The numerical results of Section 3.3 suggest that the negative intensities observed in WBRT simulations, are the effect of an inaccurate model of aberrations. In order to rule out other causes, e.g., inadequate sampling, we investigate an analytical model based on WBRT.

The analytical model is derived for the simplified case of a thin lens with only one spatial dimension perpendicular to the optical axis. The WDF of the one-dimensional plane wave with spatial frequency k_x incident on the lens, is given by

$$W_1(x_1, \xi_1) = 2\pi\delta(\xi_1 - k_x). \tag{3.5}$$

Figure 3.9: The L_{2I}-error, defined by Eq. (3.4), of a WBRT simulation is proportional to one over the square root of the number of rays. The thick lens (left) is Asphere A from Fig. 3.5 with an aperture radius of 20 mm, the thin lens is the same as used for Fig. 3.8 with an aperture radius of 2 mm. Both lenses have an NA of 0.2.

The one-dimensional aperture is a slit of width $2w$ with a WDF of

$$W_B(x,\xi) = \begin{cases} \frac{2}{\xi}\sin\left[\xi 2(w-|x|)\right] & \text{if } |x| \leq w \\ 0 & \text{if } |x| > w \end{cases}. \tag{3.6}$$

The result of the convolution (see Eq. (3.1)) between these two WDFs is

$$W_2(x_2,\xi_2) = \begin{cases} \frac{2\sin[(\xi_2-k_x)2(w-|x_2|)]}{\xi_2-k_x} & \text{if } |x_2| \leq w \\ 0 & \text{if } |x_2| > w \end{cases}. \tag{3.7}$$

The propagation through the thin (aberrated) lens to the image plane is treated in two steps. First, the effect of the aberrations is modelled by a coordinate transformation of the ray direction

$$W_3(x_3,\xi_3) = W_2(x_3,\xi_3 - f(x_3,k_x)), \tag{3.8}$$

where the function f, in accordance with geometrical optics, is the derivative of the wavefront aberration with respect to the spatial pupil coordinate (x_3). For spherical aberration this is $f(x_3,k_x) = ax_3^3$, with a a constant. The rest of the system can now be considered to be an ideal imaging system, which maps a spatial frequency in the pupil plane (ξ_3) to a position in the image plane (x_i). Characterizing this mapping by a magnification factor M, the intensity in the image plane is given by

$$I(x_i) = \int W_3(x_3,x_i/M)dx_3, \tag{3.9}$$

$$= \int_{-w}^{w} \sin\left[\left(\frac{x_i}{M} - f(x_3,k_x) - k_x\right)2(w-|x_3|)\right] \frac{2}{\frac{x_i}{M} - f(x_3,k_x) - k_x}dx_3. \tag{3.10}$$

For the aberration-free case this results in the sinc ($\mathrm{sinc}(t) \equiv \sin(t)/t$) pattern, familiar from the Fraunhofer diffraction pattern of a square aperture (see, e.g., [44, p. 76])

$$I(x_i) = 4w^2\mathrm{sinc}^2\left[(x_i/M - k_x)\,w\right]. \tag{3.11}$$

In the case of aberrations, an analytical evaluation of Eq. (3.10) is nontrivial. A numerical evaluation of the integral, using the integral function of MATLAB, for different types of aberrations is shown in Fig. 3.10, along with the result for the aberration free imaging. Negative intensities can be observed for spherical aberration and coma, whereas the results for the ideal and defocused systems remain strictly positive.

Figure 3.10: Numerical solutions of Eq. (3.10) for different aberrations. The half-width of the slit, w, and the magnification factor, M, are set to one. Both the result for coma and spherical aberrations (Spher.) show negative intensities, although the minimum intensity for spherical aberrations is only -0.02. When adding some defocus to the spherical aberration (Comb.), the negativity is more pronounced. The result for the ideal lens corresponds to Eq. (3.11). Intensity patterns for the ideal lens were calculated for a wide range of defoci. All these defocused intensity patterns were strictly positive. The figure shows one typical result (Def.).

3.5. DISCUSSION

The fact that the analytical model reproduces the effects observed in the other WBRT simulations, indicates that the unphysical results are inherent to the underlying model of WBRT and not related to the (numerical) implementation of the method. In simulations over a wide range of NAs, the negative intensities were only encountered in the presence of aberrations. It therefore appears that specifically the model of aberrations in WBRT is inaccurate.

Before proceeding to discuss the model of aberrations, it is helpful to take a step back and regard how diffraction is accounted for in WBRT. In every point in the focal region, rays originating from the entire aperture converge. The summation of the positive and negative values carried by these rays determines the local intensity. Since the WDF of an aperture is a fast oscillating function, the values and signs of the rays in WBRT vary quickly depending on their position and direction after the aperture. It is therefore not

surprising that small errors in the WBRT model can cause erroneous intensity distributions.

The most obvious erroneous result observed in WBRT, the negative intensities, can be explained by the simplified model of aberrations. In the thin element approximation, aberrations can be introduced using a phase mask. The effect on the WDF is a convolution between the WDF of the phase mask and the input field (Eq. (3.1)). The WDF of a function with a complex quadratic exponent, e.g., the quadratic phase function of paraxial defocus, reduces to a (double) delta function. As a result, the effect on the WDF of, e.g., defocus is a coordinate transformation. This coordinate transformation is implicitly implemented in WBRT by ray tracing and explicitly introduced in the analytical model by Eq. (3.8). However, coma and spherical aberration have third and fourth order phase functions. It was shown by Lohmann *et al.* [45], that the effect of such aberrations in the exit pupil on the WDF is twofold. First they lead to a coordinate transformation in accordance with geometrical optics (Eq. (3.8)), and second they lead to differential operators that change the shape of the WDF. For a cubic phase function, this results in a WDF concentrated around a quadratic function, but blurred by the Airy integral [46]. In WBRT the coordinate transformation is implemented but the blurring by the differential operators is neglected. Our results show that for the simulation of PSFs, this second effect is of crucial importance. But since in WBRT a ray only carries the local value of the WDF, one cannot apply the differential operators on the individual samples but must first rebuild the entire WDF. This would complicate the model and increase the computational load of WBRT to a great extent. Furthermore, an additional simulation method would be needed to calculate the aberrations. Some WDF simulations might benefit from modelling aberrations by coordinate transformations and differential operators or by a direct convolution with a WDF describing the aberrations, but we consider it infeasible for WBRT due to the above mentioned reasons. Contrary to hopes expressed elsewhere in the literature [14, 47], our results show that it is insufficient to model aberrations by mere coordinate transformations of the WDF.

Although WBRT is unsuited for systems with aberrations, it can be used for modelling aberration-free systems with moderate NAs. We have validated WBRT for the propagation of a coherent field through an aperture and a single lens, but expect the method to hold for any combination of paraxial lenses, mirrors and multiple apertures. As has been stated before, the effect of paraxial systems (i.e., systems for which all wavefronts, propagators, and optical surfaces are approximated by quadratic functions) on the WDF is a coordinate transformation [35], which is implemented in WBRT by ray tracing. Therefore, WBRT in the paraxial limit can be seen as the Monte Carlo integration method of paraxial wave optical propagators. For the Fresnel propagator for instance, this equivalence is shown numerically in Section 3.3.1 and analytically in Section 2.7.

Interestingly, our results show that WBRT produces accurate results for the focal field of thick lenses with NAs up to at least 0.2 (see Fig. 3.4a) and thin lenses with NAs as high as 0.45 (see Fig. 3.7). The quadratic approximation of paraxial optics is inaccurate at an NA of 0.2 and clearly violated at an NA of 0.45. It thus appears that WBRT can be used for the simulation of systems beyond the paraxial limit, if they are free of aberrations and there is no defocus. However, determining the exact limits of the validity of WBRT for aberration-free thin and thick lenses with and without defocus, requires further re-

search.

An advantage of WBRT is that the WDF can easily be adjusted for the representation of spatially partially coherent light. The WDF can be defined in terms of the mutual coherence function by replacing the product of the fields at the positions $x + x'/2$ and $x - x'/2$ in Eq. (2.106), by their correlation over time [18, 35]. A question that deserves a separate research (inspired by work of Alonso [48]) is whether the negative intensities observed in WBRT in the presence of aberrations are less pronounced for partially coherent fields. Compared to numerical implementations of paraxial wave propagation methods (i.e., the Collins integral treated in Section 2.3.2), WBRT has a relatively high computational demand but an extremely low memory usage, which makes it very suited for parallel computing. Simulation times can be significantly reduced by using importance sampling (i.e., placing more samples where the value of the WDF is higher) and performing calculations on the GPU [14].

3.6. CONCLUSIONS

We have implemented a Wigner-based ray tracing tool for diffraction calculations. The WDF of the input field is sampled using a Monte Carlo algorithm, the effect of lenses is assumed to be the coordinate transformations encountered in geometrical optics, and apertures are treated using the Kirchhoff boundary conditions. We evaluated whether this model can be applied for the simulation of optical systems. Our simulations lead to the conclusion that aberration-free systems with moderate NAs can be modelled correctly. In the presence of aberrations with rms wavefront errors as small as 10^{-2} wavelength or NAs of 0.3 and higher, the assumption that the propagation can be modelled by coordinate transformations breaks down. This conclusion is supported by the results from an analytical model. We thereby establish that the domain of validity of WBRT is limited to simulations in and near the paraxial domain.

4

SCALAR HUYGENS-FRESNEL PATH INTEGRATION

We present a method for simulating multiple diffraction in imaging systems based on the Huygens-Fresnel principle. The method accounts for the effects of both aberrations and diffraction and is entirely performed using Monte Carlo ray tracing. We compare the results of this method to those of reference simulations for field propagation through optical systems and for the calculation of point spread functions. The method can accurately model a wide variety of optical systems beyond the exit pupil approximation.

4.1. INTRODUCTION

In his book *Traité de la lumière* [13, p. xxvi-xxvii] from the seventeenth century, Christiaan Huygens proposed to model the propagation of light by considering secondary spherical light sources originating on the wavefront. More than a century later Augustin-Jean Fresnel added the idea of interference to this principle, changed the radiation distribution of the secondary sources, and showed that this could correctly model diffraction effects [13, p. xxvii-xxviii]. The Huygens-Fresnel principle was put on a solid mathematical basis with the introduction of the Kirchhoff diffraction integral [13, p. 421-422], which can be performed over an arbitrarily shaped surface. If diffraction occurs in a plane, the Kirchhoff diffraction integral can be modified to obtain the Rayleigh-Sommerfeld diffraction integral (RSDI) of the first kind, described in Section 2.3.1. Staying closer to the original concept of Huygens by performing the integral over a wavefront, the secondary sources can be shown to have a simple dipole-like form that only depends on the local field [50]. The previously mentioned methods are for scalar optics. For vectorial optics one can also perform a diffraction integral over the wavefront, which has been used to simulate diffraction in free space [51, 52].

Since all the above mentioned diffraction integrals involve highly oscillatory phase terms, they are difficult to evaluate numerically. Therefore numerical diffraction simulations often involve approximations of the phase term, like the Collins integrals described in Section 2.3.2 or Fourier optics [44]. Heisenberg uncertainty ray bending, first proposed by Heinisch and Chou in 1971 [53], completely neglects the phase term and limits the cone of radiation of the secondary sources based on their distance to the aperture edge and the Heisenberg uncertainty principle. Clearly this cannot predict destructive interference, but it can indicate the areas where diffracted light strikes using a ray tracer. Later, efforts were made to include destructive interference in Heisenberg uncertainty ray bending by coherently summing the rays [32, 33]. A correct distribution of the secondary rays over the radiation cones could however not be determined conclusively.

Another series of papers by Fischer, Prahl, and Duncan, implemented the Huygens-Fresnel principle in a Monte Carlo ray tracer. They mainly treated the propagation in two-dimensional space, where a point source emits a circular instead of a spherical wave. The simulations they performed were for free-space propagation, diffraction at an aperture, and imaging by a phase mask [54–56]. A later paper [57] treats the diffraction at a small (20 by 20 wavelengths) square aperture in the Fraunhofer regime and the propagation through a scattering phase plate [57]. The effect of scattering particles in the focal region was investigated by Ranasinghesagara *et al.* by ray tracing Huygens sources from the wavefront in the exit pupil to the focal region [58]. None of these papers [54–58] used the ray tracer to propagate through one or more lenses.

The effects of diffraction at an aperture in front of a lens were investigated by Ferré-Borrul and Bosch [59]. They included optical path lengths in an integral based on the Rayleigh-Sommerfeld integral to calculate the field near the focal region of a singlet. In this paper [59] and a subsequent paper [60] they only considered rotationally symmetric problems for a single lens and a single aperture. No details were given on the procedure to determine the optical path length, so it is unclear whether the method can be applied to non-rotationally symmetric problems. Another method, the vectorial ray-based diffraction integral (VRBDI) by Andreas *et al.* [10], simulates the propagation of

a diffracting beam (i.e., a laser) through a non-diffractive optical system. The method decomposes a field into point sources or plane waves and uses ray tracing to propagate these subfields through several lenses.

Like the work by Fischer *et al.*, the scalar Huygens-Fresnel path integration (HFPI) presented in this chapter combines Monte Carlo ray tracing with the Huygens-Fresnel principle. But like the work on Heisenberg uncertainty ray bending and the work by Ferré-Borrul and Bosch and Andreas *et al.*, it is used for simulating the effects of diffraction in optical systems. Unlike any of the previously mentioned methods, HFPI is developed to find the field in the focal plane of a multiple-diffracting optical system. The use of the Monte Carlo ray tracer greatly aids in this effort, since it eliminates the need of explicit (re)sampling.

We will introduce scalar HFPI in Section 4.2 and compare the results obtained with HFPI to those from reference methods in Section 4.3. A more detailed comparison between HFPI and beamlet propagation methods (i.e., the Gaussian beam decomposition and CODE V's BSP) is given in Section 4.4. Since HFPI is a computationally intensive method, we discuss several strategies to decrease the computational load in Section 4.5. The chapter is concluded by a discussion in Section 4.6 and the conclusions are presented in Section 4.7.

Scalar HFPI and vectorial HFPI, which is the topic of Chapter 5, are closely related and share many properties. Since light is a vectorial quantity, Chapter 5 provides a more fundamental derivation of HFPI, but the concept behind the method and its characteristics are probably easier understood by reading the present chapter.

4.2. Method

This section provides the physical reasoning behind scalar Huygens-Fresnel path integration and explains its implementation in a Monte Carlo ray tracer. The same reasoning also explains vectorial HFPI, the topic of Chapter 5, and the more formal derivation provided in that chapter is also valid for scalar HFPI. The aim of this section is to clearly explain HFPI, without getting lost in mathematical and physical details.

Figure 4.1: A graphical representation of Huygens-Fresnel Path Integration: At diffracting surfaces (e.g., the aperture) new sources are initialized and sampled using a ray tracer. The directional amplitude distribution of each source is determined using the Huygens-Fresnel principle (not explicitly shown in this figure).

That being said, we start this section with the Rayleigh-Sommerfeld diffraction inte-

gral (RSDI), Eq. (2.46), derived in Section 2.3.1 and repeated below for convenience

$$E(\boldsymbol{r_1}) = -\frac{\mathrm{i}k}{2\pi} \iint \frac{\exp(\mathrm{i}k\rho)}{\rho} \hat{\boldsymbol{\rho}} \cdot \hat{\boldsymbol{n_0}} \, E(\boldsymbol{r_0}) \, \mathrm{d}S_0. \tag{4.1}$$

The RSDI is over a flat surface S_0 with a forward pointing surface normal $\hat{\boldsymbol{n_0}}$, the vector $\boldsymbol{\rho}$ connects the point of integration $\boldsymbol{r_0}$ to the point of interest $\boldsymbol{r_1}$ and has a length ρ and normalized direction $\hat{\boldsymbol{\rho}}$. An interpretation of this integral using the Huygens-Fresnel principle is that the original field $E(\boldsymbol{r_0})$ initiates secondary sources of the form

$$G(\hat{\boldsymbol{n_0}}, \boldsymbol{r}_0, \boldsymbol{r}_1) = -\frac{\mathrm{i}k}{2\pi} \frac{\exp(\mathrm{i}k\rho)}{\rho} \hat{\boldsymbol{\rho}} \cdot \hat{\boldsymbol{n_0}}. \tag{4.2}$$

These fields are similar to those of point sources but they include an attenuation factor that depends on the direction between the observer and the surface normal.

Since we are interested in the propagation of a coherent field through an optical system, we could use diffraction integrals such as Eq. (2.46) to propagate the field from surface to surface, initiating new elementary fields at every surface. However, since these integrals have phase terms that vary on the scale of the wavelength, this can be numerically challenging for macroscopic systems. Furthermore, the RSDI is only valid for flat surfaces of integration, whereas optical surfaces generally are curved.

In the method presented here, this problem is addressed by a physical interpretation of the RSDI using Huygens's principle. The RSDI can be viewed as if every point on the surface S_0 initiates a new point source with a shape defined by Eq. (4.2) and an initial amplitude and phase given by the local field $E(\boldsymbol{r_0})$. Since the field defined by Eq. (4.2) has a slowly varying amplitude and a smooth surface of constant phase (i.e., the wavefront), it can be propagated accurately using geometrical optics (see Section 2.4). Such a propagation step can proceed through several optical elements (i.e., lenses and mirrors) until an element (e.g., an aperture) causes a sharp cut-off of the field. Since geometrical optics cannot accurately propagate the field past such a plane, we use the field determined by geometrical optics as the input field for another diffraction integral in the form of Eq. (2.46). The elementary fields associated with this integral are again propagated using geometrical optics until another diffracting surface or the output plane is reached. The electric field in the output plane is determined by the complex sum of all elementary fields.

We have implemented the principle, of initializing elementary fields at diffracting surfaces and propagating them between those surfaces using geometrical optics, in a Monte Carlo ray tracing program. The reader is referred to Section 2.4.3 for the details on Monte Carlo ray tracing, here we provide an overview of the entire HFPI simulation procedure.

The starting point of the simulation is a user defined coherent field. If this input field cannot be propagated using geometrical optics (e.g., a diffracting beam), the input field initializes elementary fields of the form of Eq. (4.2). These are sampled by starting rays with random directions and random positions in the input plane. Both the direction and position coordinates are uniformly sampled. The initial complex amplitude of a path ($w_s(m)$), where the integer m indexes a Monte Carlo sample point, is proportional to the

complex amplitude of the electric field (E) in the starting point ($\mathbf{r_0}(m)$), multiplied by an additional obliquity factor from the Huygens-Fresnel principle:

$$w_s(m) \propto \sqrt{n_0} E(\mathbf{r_0}(m)) \hat{\mathbf{n}}_0 \cdot \hat{\boldsymbol{\rho}}(m), \tag{4.3}$$

where $\hat{\boldsymbol{\rho}}(m)$ is the initial normalized direction of the ray, n_0 is the local refractive index, and $\hat{\mathbf{n}}_0$ is the forward facing surface normal of the input plane. In general, the initial path value also depends on the local ray density (see Eq. (2.97)), but for uniform sampling this is a constant factor omitted in Eq. (4.3). For input fields that can be propagated using geometrical optics (e.g., a plane wave), random geometrical rays are started in the input plane. The complex amplitudes of these rays correspond to the electric field in their starting point. For both the diffracting and the non-diffracting input fields, the samples are then traced through the optical system using regular ray tracing, until they reach a diffracting surface (e.g., the aperture).

During the ray tracing, the phasor of each path is determined using the optical path length and further amplitude factors can arise from the ray tracing procedure. The complex amplitude of a sample after a number of propagation steps, is given by

$$w(m) = w_s w_f \exp\left[ik_0 \text{OPL}\right], \tag{4.4}$$

with w_f the amplitude factor due to transmission losses (i.e., based on the Fresnel coefficients) and OPL the optical path length. To obtain the contribution to the local electric field, the complex amplitude of a path that terminates at a diffracting surface or at the final surface is multiplied with the additional factors $w_i(m)$ (Eq. (2.98)) and $w_p(m)$ (Eq. (2.99)) related to the angle of incidence and the intensity law of geometrical optics respectively. The factor w_p is not implemented in our numerical implementation of HFPI since it requires a method like tracing neighbouring rays (see Section 2.4.2) that is not standard present in Monte Carlo ray tracers. The final complex path amplitude is

$$w_e(m) = w_p w_i w_s w_f \exp\left[ik_0 \text{OPL}\right], \tag{4.5}$$

where the current implementation uses $w_p = 1$. All these factors are explained in detail in Section 2.4.3.

Every single path that reaches a diffracting surface will initiate a secondary Huygens-Fresnel source with a field of the form of Eq. (4.2). This secondary field is again sampled using Monte Carlo techniques, where the samples are distributed uniformly in angular coordinates, and propagated until the next diffracting surface (e.g., a pinhole) using regular ray tracing. Note that a path used for HFPI is split up into a bundle of new paths[1] with random directions at each diffracting surfaces whereas a geometrical ray would continue through such a surface in a straight line. To emphasize this difference, we prefer to use the term *paths* instead of *rays*. The initiation of secondary sources is independent for each individual path arriving at a diffracting surface. Therefore HFPI does not require the field at the diffracting surface to be explicitly calculated or resampled.

---

[1]Our current implementation only samples one of these secondary paths since the ray tracer cannot split a single ray into multiple rays. This is in principle only a difference in sampling that does not affect the outcome. It only increases the computational demand of the implementation.

The paths are traced further through the system, initializing new sources at every diffracting surface until the final plane is reached. In the final plane, e.g., the image plane, the electric field on a grid point is calculated by coherently summing all paths that arrive within a pixel. The phase of every path obtains a phase shift ϕ_d to approximate the phase at the centre of the pixel using a local plane wave approximation (see Section 2.4.3 for more details). The electric field at a pixel in the final plane, due to the coherent input field, is then given by

$$E(\boldsymbol{r_p}) = \frac{1}{A} \sum_{m=1}^{N} w_e(m) \exp[\mathrm{i}\phi_d(m)] \Pi_{S_p} [\boldsymbol{r}(m)] \,, \qquad (4.6)$$

where w_e is the complex amplitude of a path, $\boldsymbol{r_p}$ is the coordinate of the pixel center, $\boldsymbol{r}(m)$ the coordinate of the path at the final plane, A an overall normalization factor, and Π_{S_p} is a rectangle function that evaluates to one if a path passes through the pixel and to zero otherwise. The procedure is repeated with more paths until the noise levels in the results are deemed acceptable. The simulations presented in this chapter require 10^9 to 10^{10} paths.

The procedure outlined above is for the propagation of a coherent field through an optical system. The propagation of incoherent sources (see Section 2.3.3) can be modelled by choosing random points on the input plane and propagating the corresponding fields individually, using the methods described in this section. For every point on the input plane we can thereby calculate a field at the output plane. By taking the squared absolute amplitude of each of these fields and summing them, we obtain the total squared amplitude in the output plane due to the incoherent source.

4.3. RESULTS

We show the results of HFPI and compare them to those of reference methods for several cases. In Section 4.3.1 we show the results for a single diffraction occurring at an aperture in free space and within an optical system. Both simulations apply the Huygens-Fresnel principle at a single plane (the aperture). In Section 4.3.2 we use HFPI to propagate a Gaussian beam through an optical system, utilizing the Huygens-Fresnel principle at the input plane and the aperture. Section 4.3.3 shows the point spread function (PSF) for a system with cascaded diffraction, i.e., a system where the field diffracted by one surface, is again diffracted by one or more of the following surfaces. The system in Section 4.3.3 contains three diffracting surfaces.

For every simulation, the L_2-norm of the relative error of the calculated field is shown. This error is defined as

$$L_2 = \sqrt{\frac{\sum_p \left(|E'(p) - E(p)| \right)^2}{\sum_p |E'(p)|^2}} \,, \qquad (4.7)$$

where $E'(p)$ is the electric field at the pixel p as calculated by the reference simulation, $E(p)$ is the result from the HFPI simulation, and the sums run over all pixels. The dependence of the relative error on the number of paths will be discussed in Section 4.3.4.

4.3.1. DIFFRACTION AT A SINGLE APERTURE

The field after an aperture can be calculated using a diffraction integral like the Rayleigh-Sommerfeld diffraction integral (Eq. (2.46)). When diffraction occurs at an aperture in free space, one can directly apply this integral using the field at the aperture as the input field. In such cases, HFPI can be seen as a Monte Carlo integration of the RSDI. For the case of diffraction at an aperture in free space the factor w_p (Eq. (2.99)), which we normally approximate to be one, can be calculated analytically and accounted for in the HFPI simulation. Fig. 4.2 shows the results of such a simulation. In the figure the intensity of the field on the optical axis in the case of free space diffraction of a convergent spherical wave at a circular aperture is plotted. For this configuration, an analytical solution was derived by Aarts *et al.* [61]. The results from HFPI are in excellent agreement with the analytical formula (see Fig. 4.2). Note that the point of highest intensity does not

Figure 4.2: The intensity on the optical axis of a spherical wave (λ = 550 nm, the focal distance is 150 mm behind the aperture) after diffraction by a circular aperture with a radius of 0.313 mm. The analytical results were derived by Aarts *et al.* [61]. The L_2-norm of the relative error is 0.01.

coincide with the geometrical optical focal point. This so called focal shift is particularly pronounced at small Fresnel numbers [62], defined as

$$F_\# = \frac{a^2}{\lambda R},\tag{4.8}$$

where a is the aperture radius and R the distance from the aperture to the geometrical optical focal point. Since HFPI can be seen as a Monte Carlo integration of the RSDI, which is valid at small Fresnel numbers, it can accurately describe the propagation at small Fresnel numbers if the factor w_p (Eq. (2.99)) is accounted for.

Since the RSDI is derived using the Green's function for a homogeneous medium, it cannot be applied directly to apertures within optical systems. A common practice is to use an exit pupil diffraction integral (EPDI) (see Section 2.5). This group of methods uses geometrical optics to propagate the light to the exit pupil (the image of the aperture stop in image space), and uses a diffraction integral from there to propagate the field to the image plane. In contrast to exit pupil diffraction methods, HFPI introduces diffraction at the aperture where it is physically present inside the system. It was demonstrated

by Ferré-Borrul and Bosch [63, 64] that this can be important when modelling optical systems where the aperture is not in the image space, especially if the output plane significantly deviates from the Gaussian image plane.

In order to compare our results with those of an EPDI, we simulate the field in the image plane of a Cooke triplet. This imaging system, depicted in Fig. 4.3, has a numerical aperture of 0.2 in image space (see Table 4.1 for the system parameters).   Since the

Table 4.1: System parameters of the Cooke triplet designed for an object at infinity. The aperture, indicated by *, has a radius of 4 mm, the other elements do not clip any ray paths.

Surface number (j)	Radius [mm]	Distance [mm] (j to j+1)	Refractive index (after j)
0	∞	∞	1
1	21.5	2.00	1.62
2	-124.1	2.62	1
3*	∞	2.62	1
4	-19.1	1.25	1.62
5	22.0	4.69	1
6	328.9	2.25	1.62
7	-16.7	43.11	1
8	∞		

Figure 4.3: A schematic representation of the Cooke Triplet whose parameters are listed in Table 4.1. The exit pupil is shown in light blue.

system has a moderate NA and only a single diffracting surface, the exit pupil diffraction method provides a correct solution. The results of the methods, shown in Fig. 4.4, match very well.

4.3.2. FIELD PROPAGATION
By applying the Huygens-Fresnel principle in the input plane, HFPI is able to propagate general coherent fields[2] through an optical system. As an example, an aspherical singlet is used to image a Gaussian beam. The beam has a wavelength of 600 nm and a minimal

---

[2]Spatially incoherent sources (see Section 2.3.3) can be propagated using the procedure outlined at the end of Section 4.2.

Figure 4.4: The PSF of a Cooke triplet for a plane wave incident at 14° with a wavelength of 546.1 nm. The normalized intensity calculated using the RSDI in the exit pupil (left) and using HFPI (right) are in very good agreement. The L_{2A}-norm of the relative error is 0.06. Note that for this example we compute the difference between the absolute amplitudes (as defined by Eq. (3.3)). Since the object point is off-axis the phase of the PSF is varying very fast. This makes it numerically challenging to compare the phases computed with different methods (as is done by Eq. 4.7).

beam waist of 1.5 mm. The distance between the plane of minimal beam waist and the aperture of the system is 100 m. The circular aperture, with a radius of 20 mm, is immediately followed by an aspherical singlet with a focal distance of 100 mm. The image plane is in the focal plane of the singlet. The first surface of the singlet has a radius of 554.7 mm, the second surface has radius of -67.2 mm and a conical constant of -3.0. A schematic representation of the system is shown in Fig. 4.5. To generate a reference result, we con-


Figure 4.5: A schematic representation of the singlet and the Gaussian beam ($\lambda = 600$ nm) it images. Note that the distance between the starting plane and the aperture is 100 m and has not been drawn to scale.

volve the Gaussian input beam with the PSF of the system, which is calculated using the EPDI. Since the PSF is the response of the optical system to a delta function in the object plane, this provides us with the correct result. The results of the convolution with the PSF and the HFPI simulation are in good agreement, as can be seen in Fig. 4.6. Note that for HFPI, we have used the rotational symmetry to average the results in the output plane in order to reduce the effect of noise.

4.3.3. POINT SPREAD FUNCTION FOR DIFFRACTION AT MULTIPLE SURFACES
In contrast to the exit pupil diffraction method, HFPI can model diffraction at multiple surfaces within an optical system. This can be important in modelling systems with low

Figure 4.6: Cross section of the intensity in the image plane of a singlet imaging a Gaussian beam with a beam waist of 1.5 mm and a wavelength of 600 nm. The L_2-norm of the relative difference between results of the two methods is 0.08.

Fresnel numbers as encountered in microscopy. As an example we simulate the effect of placing a pinhole in an intermediate image plane. This technique is sometimes used in microscopy to limit the light from out of focus planes. The pinhole is usually about the size of an Airy disk to limit diffraction at the pinhole. Using HFPI, we are able to simulate the diffraction occurring at the pinhole and investigate the effect of decreasing the pinhole size. The system we simulate consists of two spherical lenses and a pinhole with a diameter of 0.2 mm in the intermediate focal plane between the lenses (see Fig. 4.7). The first lens has diameters of 0.4 mm, the second lens one of 0.8 mm and the lenses are 500 mm apart. We use an axial object at infinity with a wavelength of 500 nm. The second lens focuses the pinhole in a plane 50 mm after the lens. The systems parameters are listed in Table 4.2. Only the focal distances are listed for the lenses since the system can be accurately modelled using paraxial optics. For the same reason we use a numer-

Table 4.2: System parameters of the pinhole system for light with a wavelength of 500 nm.

Element	Diameter [mm]	Distance [mm] after the element
Lens (f = 250 mm)	0.4 mm	250 mm
Pinhole	0.2 mm	250 mm
Lens (f = 41.67 mm)	0.8 mm	50 mm
Image plane	-	-

ical implementation of the Collins integral (see Section 2.3.2) to propagate the field from the input plane to the image plane and thereby obtain a reference result. For the results shown in Fig. 4.8, the field at the edge of the pinhole has an intensity of 0.8 times the maximum intensity. In a cross section of the intensity in the output plane, the results from HFPI and the Collins integral overlap very well, whereas the exit pupil diffraction

Figure 4.7: A schematic representation of the pinhole system. The system parameters are listed in Table 4.2.

method gives a result that strongly deviates from these results.

Figure 4.8: Cross section of the intensity in the image plane after the pinhole system calculated using HFPI, the Collins integral and the exit pupil diffraction integral (EPDI). See Table 4.2 for a system description. The L_2-norm of the difference between results of HFPI and the Collins integral is 0.02.

4.3.4. CONVERGENCE

The convergence of scalar HFPI is investigated by repeating the simulations of the Cooke triplet (presented in Section 4.3.1) and the pinhole system (presented in Section 4.3.3) using fewer paths. As a measure for the convergence of the pinhole system we take the L_2-norm of the relative field difference between the HFPI and reference results, as defined in Eq. (4.7). For the Cooke triplet we use the L_2-norm of the amplitude differences, as defined in Eq. (3.3), since the phase in the image plane is varying very fast due to the off-axis illumination. The results of these simulations are shown in Fig. 4.9. The relative error (L_2) is inversely proportional to the square root of the number of paths (N_t) for both the Cooke triplet and the pinhole system. Note that the Cooke triplet has a single diffracting surface whereas the pinhole system has three diffracting surfaces. This suggests that the $1/\sqrt{N_t}$ convergence is a general property of HFPI simulations. Since the L_2-norm is proportional to the square root of the variance (defined by Eq. (2.88)), this corresponds to the statement in Section 2.4.3 that the variance of a Monte Carlo integral is inversely proportional to the number of samples (see Eq. (2.89)).
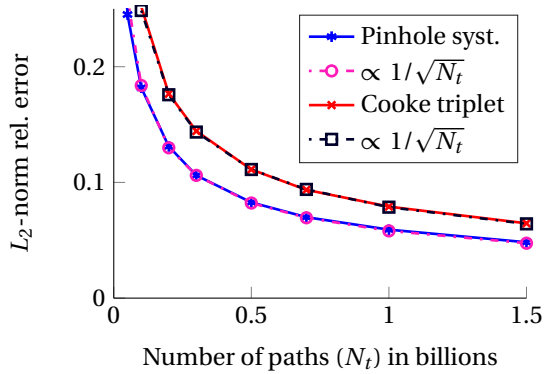
Figure 4.9: The L_2-norm of the relative error is proportional to one over the square root of the number of paths. This holds for both the amplitude error (Eq. (3.3)) used for the Cooke triplet and the field error (Eq. (4.7)) used for the pinhole system.

4.4. HFPI COMPARED TO BEAMLET PROPAGATION

In section 4.3 the results of scalar HFPI simulations were compared to those of the exit pupil diffraction integral and the Collins integral. These are established simulation methods, but both conceptually and in its domain of application, HFPI is closer related to the ray-based diffraction methods described in Section 2.6. In this section we will compare scalar HFPI to a MATLAB implementation of the Gaussian Beam Decomposition (GBD) [8] and CODE V's beam synthesis propagation (BSP) [9] for the simulation of the propagation of a Gaussian beam through a singlet.

We apply the three simulation methods for the propagation of a Gaussian beam (wavelength of 600 nm, beam waist of 1.5 mm) over 100 m to a circular aperture with a diameter of 20 mm, followed by a singlet with an NA of 0.1. All the methods perform a field decomposition at the aperture and do not make use of the exit pupil. The intensity at the aperture and the intensity in the image plane are shown in Fig. 4.10. Since all methods employ the Kirchhoff boundary conditions (see Section 2.3.4), we expect the field after the aperture to have a Gaussian profile that is sharply cut off at the aperture edge. However, both the GBD and BSP have smooth aperture edges. This is the result of using elementary beams with a finite width. A sharper aperture edge can be achieved by reducing the beam widths, but narrower beams will diverge stronger upon propagation, which can cause problems in the simulation [8]. The aperture intensity as predicted by the GBD also shows some rippling around the intensities predicted by the other methods. This is an artefact commonly encountered in this method, which can be reduced by increasing the sampling [8]. The small deviations between the methods in the aperture do not lead to substantial differences in the image plane (see Fig. 4.10). The error in the image plane, as a result of the deviations in the aperture, should correspond to the field produced in the image plane when propagating these deviations. A closer inspection of these differences suggests that their effect is limited due to their small amplitudes and because some of them are strongly diffracting.

The fastest method is BSP with a few core seconds, followed by the GBD with about

Figure 4.10: The normalized intensity of a Gaussian beam propagated through the singlet shown in Fig. 4.5 using three ray-based diffraction methods. The intensity in the aperture (left) changes smoothly for BSP and the GBD, whereas HFPI shows a sharp cut-off. In the image plane (right) there are only small differences in the intensity profiles of the three methods. The L_2-difference between BSP and the GBD is 0.04 in both the aperture and the image plane. The L_2-difference between the HFPI results and those of the other methods is 0.08 in the aperture and 0.06 in the image plane.

three core minutes. Regarding the GBD we must note that the method is currently implemented in MATLAB and one should be able to improve its computational efficiency by implementing the GBD in a dedicated ray tracing program and by improving the Gaussian fitting routine. HFPI requires about 8 core hours, which is significantly slower than the other methods. This is not surprising considering it uses a field decomposition in smaller, strongly diverging elementary fields. It is however very easy to parallelize HFPI, so practical simulation times can be significantly reduced.

In simulating the propagation of a Gaussian beam through an aperture and lens, all the investigated methods produced nearly identical results in the image plane. In the aperture, BSP and the GBD do not result in a hard cut off of the field and the GBD shows some rippling. HFPI requires significantly longer simulation times but does not suffer from these artefacts in the aperture.

4.5. INCREASING THE CONVERGENCE

The primary drawback of HFPI is its high computational demand. In this section several techniques to increase the convergence of HFPI simulations are discussed. First we treat the merit of variance reduction schemes, of which importance sampling is a prominent example. Then we present two variations on HFPI that use intermediate resampling, and a version of HFPI that associates infinitely extended plane waves to the rays. Finally we suggest some other methods to increase the convergence.

4.5.1. VARIANCE REDUCTION SCHEMES

Huygens-Fresnel path integration performs a Monte Carlo integration of a complex function. This section therefore investigates some of the commonly used techniques to increase the convergence of Monte Carlo integrals.

Monte Carlo integration approximates an integral by averaging random samples of the integrand:

$$w_N = \frac{1}{N} \sum_{i=1}^{N} \frac{f(\boldsymbol{x_i})}{p(\boldsymbol{x_i})} \approx \int f(\boldsymbol{x}) \mathrm{d}^m \boldsymbol{x}, \tag{4.9}$$

where $\boldsymbol{x_i}$ is an m-dimensional coordinate sampled from a random variable $\boldsymbol{X}$ with a probability density function (PDF) $p(\boldsymbol{x})$. The use of vectors serves to emphasize that this also holds for multi-dimensional integrals (see Section 2.4.3 for more details). The variance of such an estimator is given by Eq. (2.89):

$$V[w_N] = \frac{V[w(\boldsymbol{X})]}{N}, \tag{4.10}$$

with

$$w(\boldsymbol{X}) = \frac{f(\boldsymbol{X})}{p(\boldsymbol{X})}. \tag{4.11}$$

Many techniques to reduce the computational load of Monte Carlo integration (e.g., importance sampling) rely on reducing the variance of the estimator ($V[w(\boldsymbol{X})]$).

One way to obtain a small variance is by using a p that strongly correlates with the function f. Since p is a PDF, it should be strictly positive and integrate to one. Normalizing a PDF that correlates perfectly with f involves an integral like the one we are trying to evaluate (Eq. (4.9)) and therefore defeats the purpose. A PDF which reduces the variance to zero might not be practically achievable, but any PDF that reduces the variance of Eq. (4.11) and which is non-zero when $f(\boldsymbol{x})$ is non-zero, will lead to a better convergence. This can be as straightforward as using two uniform distributions in two domains of the integral, and placing more samples in the domain where $f(\boldsymbol{x})$ is larger. Intuitively this makes much sense: we simply place more samples in those parts of the integral that we know to contribute more to our result. This is why the method is known as importance sampling.

Some other variance reduction schemes change the estimator itself. Suppose we are integrating a monotonously increasing function $f(x)$ over the scalar variable $0 < x < x_{\max}$. Instead of integrating $f(x)$, one can integrate

$$f'(x) = \frac{1}{2} \left(f(x) + f(x_{\max} - x) \right), \tag{4.12}$$

which would strongly reduce the variance. This is a simple example of what is known as the antithetic variates method (see, e.g.,[65, p. 1158]). Another strategy modifies the estimator even stronger, by subtracting from $f(\boldsymbol{x})$ a function $h(\boldsymbol{x})$ that approximates $f(\boldsymbol{x})$ and for which one knows the integral:

$$\int f(\boldsymbol{x}) \mathrm{d}^2 \boldsymbol{x} = \langle \frac{f(\boldsymbol{X}) - h(\boldsymbol{X})}{p(\boldsymbol{X})} \rangle + \int h(\boldsymbol{x}) \mathrm{d}^2 \boldsymbol{x}, \tag{4.13}$$

where the brackets indicate the expectation value (Eq. (2.87)). If the variance of $f(X) - h(X)$ is lower than that of $f(X)$, this method of control variates (see, e.g., [65, p. 1157-1158]) will increase convergence.

The question is if these variance reduction schemes can be applied to HFPI. HFPI performs an integration of complex ray values over different paths between the source and the detector. The complex integral of HFPI can be written as two real integrals (one over the real and one over the imaginary part of the ray values). The variance in the samples is mainly due to the phase terms of these integrals, which depend on the optical path length and are highly oscillatory. These fast variations make it difficult to apply variance reduction schemes. The strategies described in this section require prior knowledge of the function to be integrated and can decrease the convergence if misapplied [65]. This makes them difficult to implement for HFPI, since the prior knowledge on the function is very limited, as will be shown using the example of importance sampling.

In the case of importance sampling, the PDF should follow the oscillations of the function in order to substantially reduce the variance of the estimator. But at the time of sampling (i.e., when a path is initiated or changes direction) the final phase of the ray is unknown, which means that one cannot construct a PDF that matches the oscillations of the function. It might be possible to generate a PDF that correlates with the amplitude term of the samples (e.g., a PDF that is lower for paths that pass through the edge of an aperture than for paths that remain close to the optical axis), but its effect on the variance will probably be fairly small. It therefore appears that a straightforward implementation of importance sampling to HFPI may not be very promising.

A method that can improve the convergence of HFPI is called stratified sampling. It divides the domain of integration into subdomains, and samples each subdomain individually. The total variance of this estimator is the sum of the variances of all the subdomains [65]. Uniform stratification divides the domain of integration into equally sized subdomains and samples the integrand with the same number of samples in every domain. Uniform stratification can only decrease the variance [66] and could be implemented in HFPI by dividing the cone of all possible ray directions into subcones and initiating a ray into each subcone[3]. To enhance the gains from uniform stratification one can combine stratified sampling with algorithms that adjust the stratification on the fly. These adaptive techniques can substantially increase the convergence but can also suffer from systematic errors [67]. Alternatively to stratified sampling, one could apply quasi-Monte Carlo sampling (e.g., digital nets) which could be described as a more sophisticated version of stratified sampling (see, e.g., [68]).

Although a detailed study of Monte Carlo integration techniques is beyond the scope of the current research, we can draw some preliminary conclusions. The use of variance reduction schemes seems to have little potential for HFPI. However, the implementation of adaptive sampling techniques may substantially increase HFPI's convergence.

4.5.2. Stepwise and bidirectional HFPI

This section presents two methods that can reduce the computational load of an HFPI simulation by splitting the simulation into multiple steps. The first method is a step-

---

[3]The implementation as described here does require ray splitting, where a single ray initiates multiple new rays. This is not possible in our current implementation.

wise HFPI simulation, where the system is split into subsystems and HFPI simulations of the subsystems are performed consecutively. The second method is bidirectional HFPI which merges a simulation from the object to the aperture with a simulation from a point in the image plane to the aperture.

Stepwise HFPI

In a stepwise HFPI simulation, the system is split into subsystems and HFPI simulations are performed for every subsystem. The output of the previous step (i.e., subsystem) is used as the input for the following HFPI step.

Before starting a stepwise HFPI simulation one must divide the system into subsystem. It makes sense to make the divisions at diffracting surfaces, but in principle one is free to place them almost anywhere. For the pinhole system of Fig. 4.7, which will be used as an example, one could for instance split the system at the pinhole and the final aperture to create three subsystems, or at one of those planes to create two subsystems.

After creating the subsystems, the first simulation step is a regular HFPI simulation from the input plane to the output plane of the first subsystem. After tracing sufficiently many paths, one obtains the field at this intermediate output plane (e.g., at the pinhole of our example system). This field then serves as the input field for the second simulation step. The simulation of this subsystem is identical to the simulation of a general input field (e.g., a Gaussian beam) in standard HFPI simulations (see Section 4.3.2). Any following subsystem is treated similarly, each time using the output of the previous subsystem to modify the starting values of the paths in the following system, until one obtains the output field of the final subsystem. The results of two stepwise HFPI simulations of the pinhole system are shown in Fig. 4.11. For one of the stepwise HFPI simulations the system was split into two parts at the pinhole. Its results are in very good agreement with those of a standard HFPI simulation. The other simulation consistent of three steps: one from the input to the pinhole, one from the pinhole to the final lens, and one from there to the output plane. The results show some small errors (not related to statistical noise), the origin of which is still undetermined.

Stepwise HFPI has some advantages over standard HFPI. A first advantage is that one can combine stepwise HFPI with other simulation methods. This allows the user to include models for specific optical components into an HFPI simulation. A second advantage of stepwise HFPI is that one can allocate the computational resources to the subsystems according to their need. If the first and third step converge quickly, one can simulate them using fewer paths and focus the computational time on simulating the second subsystem. This is not possible in standard HFPI, since every initiated path will be traced until the output plane of the system[4]. Such a redistribution of the computational load can be seen as a crude form of stratified sampling, which was described in Section 4.5.1. The drawback of stepwise HFPI is that it involves a discretization at the interfaces of the subsystems, which can lead to aliasing in the final result. This can to a large extent be avoided by determining the initial phase and amplitude of a path through interpolation between the fields in pixel centres. As long as one avoids splitting up a system in a plane with a strongly oscillatory field, stepwise HFPI is a more versatile version

---

[4]Using ray splitting in standard HFPI one can allocate more computational resources to the later subsystems than to the earlier ones (see Section 4.5.4), but this is significantly less flexible than stepwise HFPI.

(a) Normalized amplitude 2-step HFPI

(b) Phase 2-step HFPI

(c) Amplitude error 2-step HFPI

(d) Phase error 2-step HFPI

(e) Amplitude error 3-step HFPI

(f) Phase error 3-step HFPI
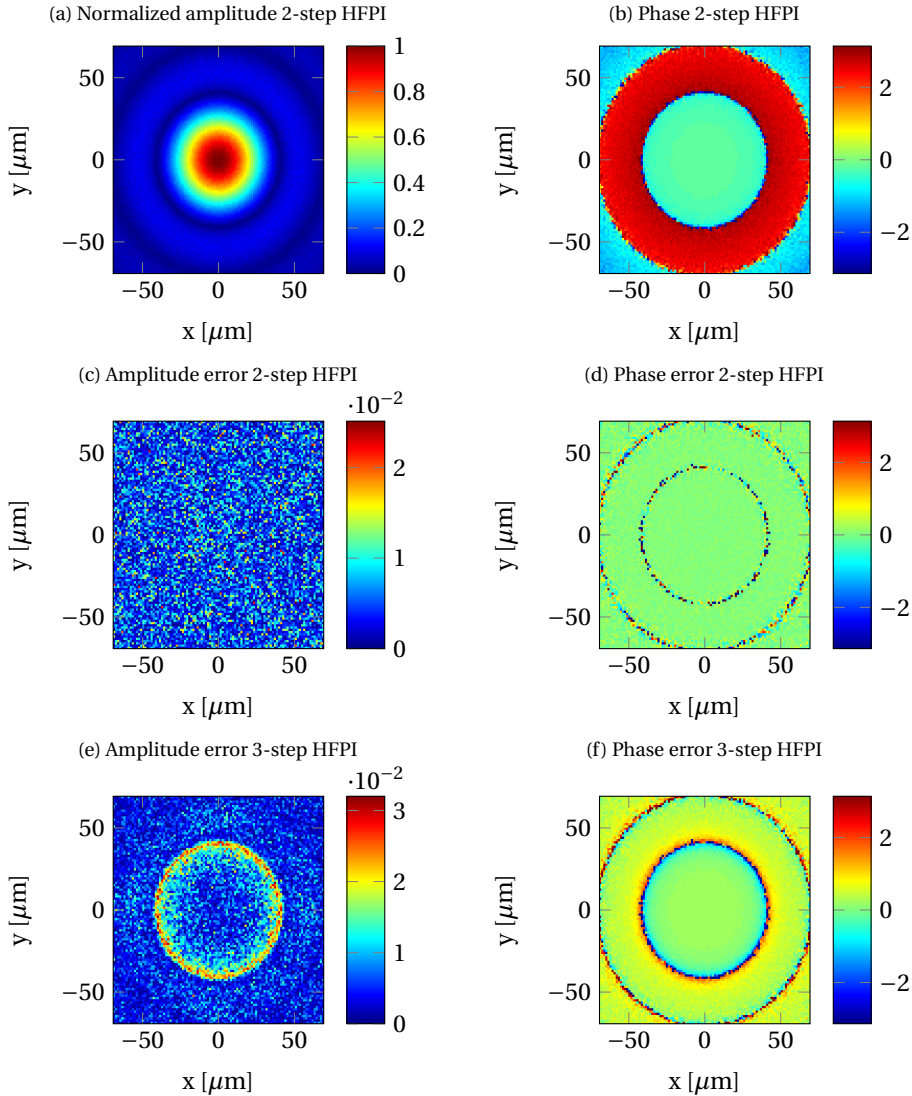
Figure 4.11: The amplitude and phase for stepwise simulations of the pinhole system with cascaded diffraction of Fig. 4.7. For the 2-step simulation the system is split at the pinhole, for the 3-step system it is split at the pinhole and the last aperture. The top row shows the normalized amplitude and phase of the 2-step HFPI. The second and third rows show the errors with respect to the standard HFPI simulation.

of HFPI that can provide computational benefits.

BIDIRECTIONAL HFPI

Bidirectional HFPI is based on the observation that a ray follows the same path when traced forward, as when traced backward (i.e., from image to object). Standard HFPI

simulations trace rays from the object to the aperture and then continue to the image plane. In bidirectional HFPI, this is split into a forward simulation from the object to a diffracting plane, and backwards simulations from each point of interest in the image plane to the same diffracting plane.

Let us consider how the field in a point of interest in the image plane is calculated using bidirectional HFPI. First a forward simulation is performed to calculate the field at the centres of the pixels in the aperture due to the field in the input plane. We denote the calculated complex field at pixel m by $A_{m,\text{fw}}$. Then the backwards simulation is performed, tracing paths from the point of interest in the image plane until the aperture. This also results in a complex value for every pixel in the aperture: $A_{m,\text{bw}}$. Multiplying the two complex values for one pixel

$$A_m = A_{m,\text{fw}} A_{m,\text{bw}}, \tag{4.14}$$

gives the contribution to the field in the point of interest, from all paths that pass from the object through that particular pixel. The total electric field in the point of interest $\boldsymbol{r_i}$ is calculated by summing the contribution from all pixels:

$$E(\boldsymbol{r_i}) = \sum_m A_{m,\text{fw}} A_{m,\text{bw}}. \tag{4.15}$$

This method of bidirectional HFPI is schematically depicted in Fig. 4.12.

Figure 4.12: A graphical representation of bidirectional HFPI. Both the forward and backwards simulation result in a complex value at every pixel in the aperture plane. The field in the point of interest is then calculated by multiplying the two values at every pixel and then summing over the pixels (see Eq. (4.15)). The figure shows how the field on the optical axis in the image plane due to an off-axis object point is calculated.

Bidirectional HFPI calculates the field at distinct (user defined) points in the output plane, e.g., at ten unequally spaced points on the line $y = 0$. In contrast, a regular HFPI simulation calculates the field on a uniform square grid in the output plane. If the user does not require knowledge of the field on the entire grid (e.g., due to rotational symmetry of the system), bidirectional HFPI can significantly decrease the computational time by only tracing the paths necessary to calculate the field in the points of interest. Like stepwise HFPI, bidirectional HFPI involves a discretization of the field in the intermediate plane.

The Cooke triplet, shown in Fig. 4.3, is simulated with bidirectional HFPI and the forward and backward simulation are performed from object to aperture and image plane to aperture respectively. The simulation is performed for different field positions and with different discretizations at the aperture. The results are shown in Fig. 4.13. For this system, the results agree with the HFPI results when using a grid at the aperture of 40 by

40 pixels or more. One can check the convergence of a bidirectional simulation by either redoing the simulation with a different number of grid points, or by eliminating part of the grid points.

(a) Cross section of PSF for an object at 0°.

(b) Cross section of PSF for object at 6.9°.

Figure 4.13: Comparison between the HFPI results and the bidirectional HFPI results for the Cooke triplet of Table 4.1. The bidirectional simulations are performed with a grid of 10x10, 40x40 and 100x100 pixels in the aperture plane, which is the plane where the forwards and backwards simulations end. The origin of the plot for the off-axis simulation, Fig. b, lies at $y = 6.050$ mm. The results of a simulation with the exit pupil diffraction integral are also shown.

For a system with multiple diffracting surfaces, one can either include a diffracting surfaces in the backward (or forward) simulation or combine bidirectional HFPI with stepwise HFPI described in Section 4.5.2. The results of both methods for the pinhole system (see Fig. 4.7), which has three diffracting surfaces, are shown in Fig. 4.14.

We can draw some interesting conclusions from the bidirectional HFPI simulation of the Cooke triplet (Fig. 4.13). The only diffracting plane in this simulation is at the interface between the forward and the backward simulation. Although we performed the

Figure 4.14: A cross section of the PSF for the pinhole system of Fig. 4.7. For the bidirectional method (bidir.) the interface between the forward and the backward simulation is at the pinhole. For the combined method (comb.) the stepwise simulation consists of a first simulation from the object to the pinhole followed by a simulation from the pinhole to the second lens (which is the final diffracting surface), and the backward simulation is from the image plane to the second lens. Both methods use a grid of 101x101 for the backward simulation. The results for plane wave HFPI (see Section 4.5.3) are also shown.

simulations using Monte Carlo ray tracing, one could also obtain the field in the aperture by standard ray tracing combined with ray aiming. The forward simulation can be performed by shooting a ray from the object point to every grid point in the aperture. Equivalently, the backwards propagation can be achieved by tracing rays from the point of interest to the grid points (or vice versa). The success of bidirectional HFPI thus hints at the potential of an implementation of HFPI that employs standard ray tracing combined with ray aiming instead of Monte Carlo ray tracing. Such a non-stochastical implementation of HFPI might require the tracing of fewer rays, but can lead to wrong results when the aperture is not sampled densely enough.

4.5.3. PLANE WAVE HFPI

The concept of plane wave HFPI (PW-HFPI) is to approximate the spherical wavefronts of the secondary HFPI sources by plane waves. This may appear to be a very bold approximation, but it turns out to be very effective within the focal region of most systems.

Suppose we are simulating a system with an aperture in image space and we are interested in finding the field in the focal region. In an HFPI simulation, rays would be propagated from the object to the aperture, where they would initiate secondary rays. In Fig. 4.15 one such ray and the corresponding secondary rays are drawn. In a PW-HFPI simulation the primary ray continues propagating until the output plane, where it is assumed to carry a plane wave that affects all pixels.

The error made by approximating the spherical wave of radius R by a plane wave can be calculated using the quantities shown in Fig. 4.15. The distance from a point at the edge of the aperture to a point a distance x away from the optical axis in the image plane

Figure 4.15: Plane wave HFPI approximates the spherical field of a secondary source as a plane wave. Usually the region of interest is only a few wavelengths large, much smaller than x pictured above.

is

$$s = \sqrt{R^2 + d^2} \approx R + \frac{d^2}{2R}, \tag{4.16}$$

where we have used a second order Taylor expansion around $d = 0$. When using the plane wave approximation, this length is approximated by R. The error in the length is then

$$\Delta_s = \frac{d^2}{2R} = \frac{x^2 R}{2z^2}, \tag{4.17}$$

where we have used that

$$\frac{x}{d} = \frac{z}{R}, \tag{4.18}$$

for $\sqrt{d^2 - x^2} \ll z$. In practise we are only interested in a very small region around the optical axis. The Airy disk is the circular area that includes the central peak and the first zero of a PSF. It has a radius approximated by [16, p. 469]

$$x_{\text{Airy}} = 0.61 \frac{z\lambda}{a}, \tag{4.19}$$

with λ the wavelength in the medium. Taking twice this radius as our area of interest, applying Eq. (4.17) and dividing by the wavelength, one obtains the following error for the optical path length (expressed in wavelengths)

$$\Delta_{s_\lambda} \approx 0.7 \frac{\lambda R}{a^2} = \frac{0.7}{F_\#}, \tag{4.20}$$

with $F_\#$ the Fresnel number [13, p. 417]. The plane wave approximation is thus valid for systems with a Fresnel number substantially larger than one.

The results of a PW-HFPI simulation, for the PSF calculation of the earlier presented Cooke triplet (see Table 4.1), are shown in Fig. 4.16. By using PW-HFPI instead of HFPI, the simulation times can be reduced from over 2 core-hours to about two core minutes while improving the accuracy (see Table 4.3). This performance gain is due to the fact that in a PW-HFPI simulation each ray contributes to all pixels, whereas for a regular HFPI simulation each ray only affects a single pixel.

(a) Amplitude

(b) Absolute error of complex field

Figure 4.16: Results for plane wave HFPI simulations of a Cooke triplet (see Table 4.1, $F_\# = 865$) for a plane wave incident at 0° with a wavelength of 546.1 mm. The error is with respect to the exit pupil diffraction integral. The PW-HFPI simulation uses 10^5 paths and has an L_2-error of 2.0%.

Table 4.3: Errors and convergence times (on a 2.5 GHz Intel Xeon CPU) of plane wave HFPI compared to HFPI. The methods are used to simulated the PSF of the Cooke triplet (see Table 4.1) for on-axis illumination at a wavelength of 546.1 mm. The listed L_2-errors are with respect to results from the EPDI.

Method	Number of paths	CPU time	L_2-error
HFPI	10^9	2h17m	3.5 %
PW-HFPI	10^5	2m11s	2.0 %
PW-HFPI	10^4	13.6s	6.6 %
PW-HFPI	10^3	1.5s	19.9 %

To investigate the effect of the Fresnel number on the accuracy of PW-HFPI, the aperture size of the Cooke triplet is reduced from 4 mm to values between 0.04 mm and 1.0 mm, thereby reducing the Fresnel number from 865 to values between 0.06 and 32. For PW-HFPI, the L_2-differences with respect to the EPDI results increase with decreasing Fresnel numbers, whereas they remain constant for the HFPI results (see Fig. 4.17). Assuming that the results for the Cooke triplet can be generalised to other systems, PW-HFPI is accurate when simulating systems with a Fresnel number higher than roughly 10. For the results shown in Fig. 4.17, the error of PW-HFPI is purely in the phase of the field. The difference between the amplitude of the PW-HFPI and the EPDI results, quantified in the L_{2A}-difference of Eq. (3.3), is in fact constant at 0.006 for all Fresnel numbers.

Plane wave HFPI can also be applied in systems with multiple diffracting surfaces. The result for such a simulation of the pinhole system (see Fig. 4.7) are plotted in Fig. 4.18. For this simulation the first two diffracting surfaces initiate secondary sources according to the standard HFPI procedure. At the final diffracting surface however, the algorithm does not initiate secondary sources and the paths are allowed to propagate further to the image plane (unless they are blocked by the aperture), where a plane wave that affects all pixels is associate to every ray. The error made in this approximation depends

Figure 4.17: The influence of the Fresnel number on the accuracy of PW-HFPI is demonstrated with simulation results for the PSF of the Cooke triplet (see Table 4.1) with aperture sizes varying from 0.04 mm to 1.0 mm (corresponding to $0.06 \leq F_\# \leq 32$). The amplitude predicted by PW-HFPI matches those simulated using HFPI and the EPDI for Fresnel numbers as small as 0.06 (top left), whereas the phase of the PW-HFPI simulation does not match that of the other methods for Fresnel numbers as high as 1 (top right). The bottom plot shows the L_2-difference with respect to the EPDI, calculated over a square area with sides of $2\lambda/\text{NA}$, for HFPI and PW-HFPI. The illumination is by a plane wave parallel to the optical axis with a wavelength of 546.1 mm.

on the Fresnel number of the subsystem from the final diffracting surface to the image plane, which in this case is 6.4. From the results obtained for the Cooke triplet (shown in Fig. 4.17) we can expect PW-HFPI to give accurate results for the amplitude of the field, but have some residual errors on the phase. Indeed the L_2-difference between a PW-HFPI simulation with 10^6 rays and those of the reference method is 0.09, whereas the L_{2A}-difference is only 0.02. As with the Cooke triplet with a low Fresnel number (see Fig. 4.7), the error of the PW-HFPI results are mainly in the phases at the edge of the PSF (see Fig. 4.18).

Plane wave HFPI can be interpreted as if every point in the exit pupil emits a plane wave in a direction perpendicular to the local wave front. This is exactly the physical interpretation Wolf gave for the exit pupil diffraction integral now known as the Richards-
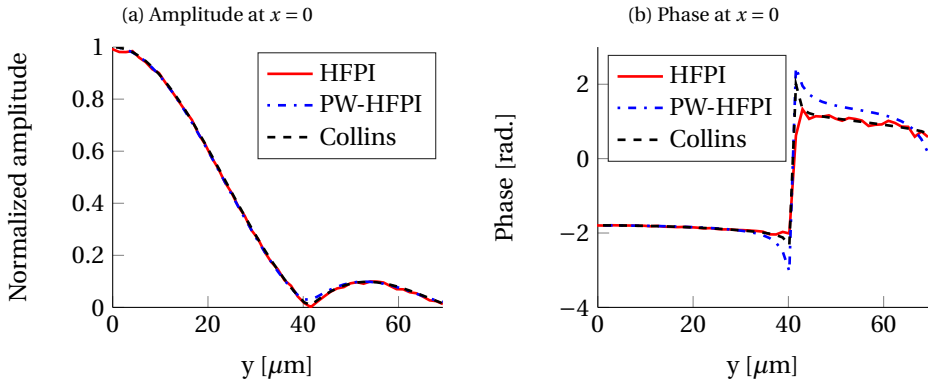
Figure 4.18: A cross section of the amplitude (left) and phase (right) of the PSF for the pinhole system of Fig. 4.7. The L_2-difference between the PW-HFPI results and the results of the Collins integral is 0.09 whereas it is only 0.02 for the HFPI results. Note that the error of the PW-HFPI results is mainly in the phase.

Wolf integral [5]. For systems with a single diffracting surface PW-HFPI is thus approximately a Monte Carlo integration of the Richards-Wolf integral[5], which is known to produce inaccurate results for systems with low Fresnel numbers [69].

The success of PW-HFPI appears to make the other methods described in this section irrelevant, but the plane wave simulation does have its limitations: it can only be applied to simulate the field near a focal plane. Still, in many cases one is interested in getting the field at a final focal plane, and in such cases, this approach is extremely powerful.

4.5.4. IDEAS FOR OTHER CONVERGENCE STRATEGIES

Intuitively one would expect that diffraction from an aperture is primarily induced by the electric field close to the aperture edge, whereas the field passing through the center of the aperture is left undisturbed. The geometrical theory of diffraction and Heisenberg uncertainty ray bending (both briefly introduced in Section 2.6) are both based on this principle. It is worth considering applying this concept to HFPI in order to achieve faster simulation times. However, the example of Heisenberg uncertainty ray bending demonstrates that this is far from straightforward.

One idea for HFPI is to divide the aperture into an undisturbed part, through which the rays can propagate according to geometrical optical laws, and a diffracting part where secondary sources are initiated. These two regions can have a smooth transition, along the lines of the division proposed by Baladron-Zorita [70], where part of the power of the ray propagates undisturbed and the rest of the power initiates a secondary source. If this were successful one can envision combining it with the approximation, as done by Lit [7] in combination with the geometrical theory of diffraction, that only the undisturbed field can be diffracted by apertures further along in the optical system. However, no undisturbed field should be allowed to propagate until a focal region since the

---

[5]The Richards-Wolf integral is actually formulated for calculating the vectorial PSF. In Chapter 5 it will be shown that PW-HFPI is also applicable to vectorial optics.

undisturbed geometrical optical rays would intersect there, rendering geometrical optics invalid. Such an extension of HFPI could have the potential to considerably reduce the computational load, particularly in the case of cascaded diffraction, but does require substantial further research.

Two more straightforward methods are presented below. They have not been studied in detail since they require a substantial modification of the current numerical implementation of HFPI. When implementing HFPI anew, they may however be useful.

We call the procedure of splitting one initial ray into multiple secondary rays ray splitting. For HFPI ray splitting would occur at the diffracting surfaces. As with stepwise HFPI, it allows dedicating computational resources to subsystems according to their need. Ray splitting is however less flexible than stepwise HFPI since, once initiated, a ray will continue through all the subsequent subsystems. Ray splitting would therefore allow the later part of the system to have a higher sampling than the earlier part, but not the other way around.

We denote another strategy that involves initiating multiple rays as grid tracing. Where for ray splitting every ray only carries a local field, grid tracing would involve approximating the field between rays, thereby reducing the number of rays one needs to trace. The disadvantage of this procedure is an increase in the complexity of the algorithm.

4.6. DISCUSSION

With Huygens-Fresnel Path Integration we perform diffraction calculations beyond the paraxial domain and beyond the exit pupil approximation. The method introduces diffraction in the planes where it physically occurs (e.g., in the aperture plane instead of in the exit pupil) and can account for multiple diffracting surfaces. By choosing the input field as one of the diffracting surfaces, HFPI can model the propagation of fields which cannot accurately be propagated using geometrical optics (e.g., a Gaussian beam in its focal plane). Since HFPI uses ray tracing, the effects of aberrations and thick optical components are automatically included in the model. Therefore HFPI can, for instance, be used to simulate the propagation of a Hermite-Gaussian beam through an aberrated system with multiple diffracting surfaces.

Another advantage of HFPI is its robustness. It does not require resampling in intermediate planes, a common cause of errors in simulations. Since the sampling in the input plane is performed using a Monte Carlo method, it is very easy to see if the simulation has converged to a solution. The fact that a single method is used for simulating the entire system greatly limits the input required from a user. In fact, the only critical choice when using the method is determining which surfaces should be considered as diffracting surfaces.

A drawback of the method is the computational cost. Calculating the PSF for the pinhole system with cascaded diffraction (see Section 4.3.3) required 10^9 paths and 4 core-hours (on a 2.5 GHz Intel Xeon CPU) in our current implementation. Since the error in the calculated amplitude is inversely proportional to the square root of the number of paths, one can obtain a result with about ten times the statistical noise by tracing 10^7 paths in a few core minutes. The general structure of the PSF can already be seen from such a simulation. Practical simulation times can be further reduced by distributing a simulation over multiple cores. It is very easy to parallelize the method since every path

is traced completely independently and the data transmission is small. This is particularly relevant in view of the increasing availability of cloud computing services, which provide a large number of cores at very low effort and cost, even for infrequent users. Although no cloud services were used for the work in this thesis, simulations in this chapter were distributed over multiple computers and every computer used multiple cores.

The simulation time can also be reduced by adjusting the method or its numerical implementation. Regarding the numerical implementation, importance sampling seems to be of little use due to the highly oscillatory nature of HFPI, but adaptive, stratified, or quasi-Monte Carlo sampling techniques may significantly increase the convergence of HFPI. A substantial adjustment of the method would involve dividing the field into an undisturbed and a diffracted part and only initiating secondary rays in the diffracted part. This has not yet been attempted and requires substantial further research. Some other alterations of the method have been investigated and shown to be effective. Bidirectional and stepwise HFPI split the simulation into multiple steps. They can reduce the simulation time for specific cases, but have the disadvantage of having to use an explicit intermediate resampling. The success of bidirectional HFPI indicates that a non-stochastical implementation of HFPI could be interesting for diffraction simulations. Plane wave HFPI reduces the simulation times by about two orders of magnitude, by associating a plane wave to every path in the image plane. It can be used for simulating the field in the focal plane of systems with single or multiple diffracting surfaces.

Although HFPI is a robust method for propagating the scalar field through a system of lenses, mirrors and (multiple) apertures, there are some underlying assumptions that limit the domain of applicability. The method as described in this chapter is limited to propagating the scalar field, but this limitation will be removed in Chapter 5. Another restriction is that all apertures should be much larger than the wavelength because of the use of the Kirchhoff boundary conditions. Furthermore, the light reflected at lenses is neglected and only the forward propagating field is taken into account. Finally, the diffraction is modelled to enter the system at distinct surfaces and it is assumed that all propagation distances are much larger than the wavelength. As a consequence we do not expect HFPI to be able to fully describe diffractive optical elements. However, we do anticipate the method to be able to handle phase and amplitude gratings within the thin element approximation.

4.7. CONCLUSION

We have introduced Huygens-Fresnel path integration, a method that incorporates diffraction effects in a Monte Carlo ray tracing program, and we validated it against reference methods. By using the Huygens-Fresnel principle consecutively we can simulate the effect of multiple diffracting surfaces in paraxial and non-paraxial systems. The theory for the method was developed by combining the Rayleigh-Sommerfeld integral with ray tracing. Since the computational load is the main drawback of HFPI, we presented several approaches to improve the convergence and demonstrated a decrease in simulation times by over two orders of magnitude. This makes HFPI not only a robust method for simulating (cascaded) diffraction and aberrations, but also a practically applicable simulation tool.

5

VECTORIAL HUYGENS-FRESNEL PATH INTEGRATION

The ray-based simulation method presented in Chapter 4 is extended to vectorial optics. Since light is a vectorial phenomenon, this extension offers the possibility of a more fundamental derivation of the method. The result is a Monte Carlo based simulation method that accounts for not only aberrations and (cascaded) diffraction, but also polarization. This is important when simulating a system with a high numerical aperture.

5.1. INTRODUCTION

One of the main limitations of Huygens-Fresnel path integration as presented in Chapter 4 is that it treats light as a scalar phenomenon. This chapter provides the extension of HFPI to vectorial optics.

The main motivation for the development of a vectorial method is the desire to simulate systems with high numerical apertures. These systems are characterised by the occurrence of large angles between the optical axis and the rays in the image space. Since the orientation of the electric field (i.e., the polarization) is perpendicular to that of the rays, the polarization is far from uniform in the image space and a vectorial theory is needed to accurately describe the field in the focal plane.

A second motivation for developing vectorial HFPI is that it allows for a more fundamental derivation of the underlying theory. Such a derivation is provided in Section 5.2, where HFPI is expressed as a cascaded diffraction integral.

The numerical implementation of vectorial HFPI, explained in Section 5.3, is validated in Section 5.4. This is done by comparing HFPI results to those of BSP and VRBDI, two other ray-based diffraction methods (see Section 2.6).

In Section 5.5 it is shown that combining HFPI with a plane wave approximation, as introduced for scalar HFPI in Section 4.5.3, can reduce the computational load by more than two orders of magnitude.

The discussion in Section 5.6 conceptually compares HFPI with a range of other methods. The properties of vectorial HFPI are not discussed in detail since they are very similar to those of scalar HFPI, which were already presented in Section 4.6. The main findings presented in this chapter are summarized in Section 5.7.

5.2. THEORY

The core idea of our method is to successively propagate the kernel of the m-theory diffraction integrals (see Section 2.3.1) using geometrical optics. In Section 5.2.1 it will be shown that these diffraction integrals can be interpreted as a decomposition of the electric field into the electric fields of magnetic dipoles and that these elementary fields are of such a form that they can be propagated using geometrical optics. In Section 5.2.2 we will apply this diffraction integral multiple times, to arrive at a cascaded diffraction integral for multiple diffraction in an optical system. In Section 5.2.3 we elaborate on the choice for the m-theory diffraction integral and the relation between the vectorial and the scalar theory for HFPI.

5.2.1. GEOMETRICAL PROPAGATION OF THE INTEGRAL KERNEL

Our derivation is based on the integrals of the m-theory, Eqs. (2.47) and (2.48), repeated below

$$E(\boldsymbol{r_1}) = -\frac{\mathrm{i}k}{2\pi} \iint \frac{\exp(\mathrm{i}k\rho)}{\rho} \hat{\boldsymbol{\rho}} \times (\hat{\boldsymbol{n_0}} \times \boldsymbol{E}) \, \mathrm{d}S_0,$$

$$H(\boldsymbol{r_1}) = \frac{-\mathrm{i}k^2}{2\pi\omega\mu} \iint \frac{\exp(\mathrm{i}k\rho)}{\rho} \hat{\boldsymbol{\rho}} \times \left[\hat{\boldsymbol{\rho}} \times (\hat{\boldsymbol{n_0}} \times \boldsymbol{E})\right] \, \mathrm{d}S_0,$$

where $\hat{\boldsymbol{\rho}}$ is the vector from the point of interest $\boldsymbol{r_1}$ and the integration variable $\boldsymbol{r_0}$ and $\hat{\boldsymbol{n}}_0$ is the surface normal pointing towards the half-plane containing $\boldsymbol{r_1}$. These integrals are only valid within a homogeneous medium. Since optical systems are only piecewise homogeneous, these diffraction integrals cannot describe the propagation through an entire system. Applying separate diffraction integrals for every homogeneous region (e.g., a lens or free space) is numerically challenging and requires the solution of boundary value problems at every interface (e.g., between glass and air) to determine transmission losses. Furthermore, interfaces are frequently curved, whereas Eqs. (2.47) and (2.48) require a flat plane to eliminate a term proportional to the curl of the electric field from the integral. To simulate optical systems comprised of multiple optical media we therefore choose a different approach based on a physical interpretation of these diffraction integrals.

Formula (2.47, 2.48) are integrals over $\boldsymbol{r_0}$ of the following two vector fields

$$\boldsymbol{E}^{(m)}(\boldsymbol{r_0}, \boldsymbol{r_1}, \boldsymbol{E}(\boldsymbol{r_0})) \equiv -2\mathrm{i}kG_H\hat{\boldsymbol{\rho}} \times (\hat{\boldsymbol{n}}_0 \times \boldsymbol{E}(\boldsymbol{r_0})), \tag{5.1}$$

$$\boldsymbol{H}^{(m)}(\boldsymbol{r_0}, \boldsymbol{r_1}, \boldsymbol{E}(\boldsymbol{r_0})) \equiv \frac{-2\mathrm{i}k^2}{\omega\mu}G_H\hat{\boldsymbol{\rho}} \times \left[\hat{\boldsymbol{\rho}} \times (\hat{\boldsymbol{n}}_0 \times \boldsymbol{E}(\boldsymbol{r_0}))\right], \tag{5.2}$$

with G_H the Green's function of the Helmholtz equation given by Eq. (2.39). These are equal to the magnetic and electric far field of a magnetic dipole $\boldsymbol{m}$ in $\boldsymbol{r_0}$ (which is why the name m-theory was introduced by Karczewski and Wolf [72])

$$\boldsymbol{E}^{(md)}(\boldsymbol{r_0}, \boldsymbol{r_1}, \boldsymbol{m}(\boldsymbol{r_0})) = -k^2\sqrt{\frac{\mu}{\epsilon}}G_H\hat{\boldsymbol{\rho}} \times \boldsymbol{m}, \tag{5.3}$$

$$\boldsymbol{H}^{(md)}(\boldsymbol{r_0}, \boldsymbol{r_1}, \boldsymbol{m}(\boldsymbol{r_0})) = -k^2 G_H\hat{\boldsymbol{\rho}} \times (\hat{\boldsymbol{\rho}} \times \boldsymbol{m}), \tag{5.4}$$

(see, e.g., Jackson [73]), with dipole moment

$$\boldsymbol{m}(\boldsymbol{r_0}) = \frac{2\mathrm{i}}{k}\sqrt{\frac{\epsilon}{\mu}}\hat{\boldsymbol{n}}_0 \times \boldsymbol{E}(\boldsymbol{r_0}). \tag{5.5}$$

The field predicted by the m-theory, Eq. (2.47), is therefore the same as the field of infinitely many magnetic dipoles, located in the integration plane S_0

$$\boldsymbol{E}(\boldsymbol{r_1}) = \iint \boldsymbol{E}^{(md)}(\boldsymbol{r_0}, \boldsymbol{r_1}, \boldsymbol{m}(\boldsymbol{r_0}))\, dS_0, \tag{5.6}$$

where the strength and direction of the dipoles are given by Eq. (5.5).

The diffraction integrals of the m-theory can be seen as a threefold process. First, the decomposition of the field in S_0 into magnetic dipoles according to Eq. (5.5), then the propagation of these dipole fields through a homogeneous medium to the point $\boldsymbol{r_1}$ using Eq. (5.3) or (5.4), and finally adding up all these fields coherently.

Since these subfields are equivalent to the fields of magnetic dipoles, they have the properties of physical electromagnetic fields (e.g., being divergence free). Although this can help in physically understanding the diffraction integral, the decomposition is primarily useful if the subfields are easier to propagate than the original field. It turns out that the elementary fields of the m-theory integrals can be propagated through a series

of homogeneous media (i.e., an optical system) using geometrical optics, due to several properties of the dipole subfields. A more detailed description of geometrical optics is provided in Section 2.4, but for the current discussion it suffices to know that in geometrical optics the field is transported using rays along paths that are perpendicular to surfaces of equal phase, and that the field associated with a ray locally has the properties of a plane wave.

The first useful property of the dipole subfield is that the phases of $E^{(md)}$ and $H^{(md)}$ are determined by the phase of the Greens function ($\exp(ik\rho)$). The surfaces of constant phase are therefore spheres around the point r_0. The phase changes fastest along the gradient of the phase ϕ,

$$\nabla\phi = k\hat{\boldsymbol{\rho}}, \tag{5.7}$$

which defines vectors perpendicular to the surfaces of constant phase and with a magnitude equal to the wave number. This is consistent with the concept of a ray in geometrical optics. The second useful property is that both $E^{(md)}$ and $H^{(md)}$ are perpendicular to $\hat{\boldsymbol{\rho}}$, the direction of phase propagation. Furthermore, the local electric and magnetic field can be obtained from each other by

$$E^{(md)} = -\sqrt{\frac{\mu}{\epsilon}}\hat{\boldsymbol{\rho}} \times H^{(md)}, \tag{5.8}$$

$$H^{(md)} = \sqrt{\frac{\epsilon}{\mu}}\hat{\boldsymbol{\rho}} \times E^{(md)}, \tag{5.9}$$

which are also the equations that relate the fields of a plane wave (see, e.g., Jackson [73]). We thus conclude that the fields of a magnetic dipole, Eqs. (5.3, 5.4), and thereby those of the kernel of the m-theory, Eqs. (5.1, 5.2), can be calculated using a geometrical optical propagator $\tilde{G}$. Using this propagator, the electric field of Eq. (5.3) can be written as

$$E^{(md)}(r_0, r_1, m(r_0)) = \tilde{G}\left[\frac{-k^2}{4\pi}\sqrt{\frac{\mu}{\epsilon}}m(r_0), r_0, r_1\right]. \tag{5.10}$$

In free space this propagator has the simple form

$$\tilde{G}[A, r_0, r_1] = (r_1 - r_0) \times A\frac{\exp(ik|r_1 - r_0|)}{|r_1 - r_0|}, \tag{5.11}$$

which reduces Eq. (5.10) to Eq. (5.3). The advantage of propagating the magnetic dipole field by a geometrical optics operator is that the propagation is no longer restricted to a single homogeneous medium. Whereas Eq. (5.3) was limited to free-space propagation, Eq. (5.10) can propagate the field through a series of optical components (e.g., lenses).

Since we are not so much interested in propagating dipole fields as we are in the m-theory diffraction integral, we use Eq. (5.10) to write the diffraction integrals of Eqs. (2.47, 2.48) as

$$E(r_1) = \iint_{S_0} \tilde{G}\left[\frac{-ik}{2\pi}\hat{\boldsymbol{n}}_0 \times E(r_0), r_0, r_1\right] dS_0, \tag{5.12}$$

$$H(r_1) = \iint_{S_0} \sqrt{\frac{\epsilon}{\mu}}\hat{\boldsymbol{\rho}} \times \tilde{G}\left[\frac{-ik}{2\pi}\hat{\boldsymbol{n}}_0 \times E(r_0), r_0, r_1\right] dS_0. \tag{5.13}$$

Again, the advantage of using Eq. (5.12) over the integral of Eq. (2.47) is the ability to propagate the field through an optical system with multiple lenses. Since Eq. (5.12) is based on the free-space integral of Eq. (2.47), the implicit assumption is that the field in the medium after the plane of integration can be calculated as if the medium were to continue to infinity. At optical interfaces (i.e., lenses) we neglect the part of the field that is reflected and only consider the forward propagating field (with an amplitude adjusted by the Fresnel equations). Although the Fresnel equations provide the reflected fields and these could in principle be propagated using our method, their influence is typically small whereas the associated computational cost would be substantial. We therefore only consider the direct path, neglecting (multiple) reflections. Under this minor assumption, the geometrical propagation of the integral kernel thus allows us to propagate a polarized field at the input surface S_0, through a (sub)system.

5.2.2. CASCADED DIFFRACTION INTEGRAL

By combining the diffraction integral with a geometrical optics propagator we can calculate the field after a set of lenses and account for the diffraction caused by the initial field (e.g., if the initial field is a top hat or a diffracting laser beam). One can accurately describe the propagation through a series of elements, as long as the geometrical optical operator is valid. In broad terms, geometrical optics cannot accurately describe the propagation of fields that change strongly in amplitude or phase over a small spatial dimension other than the phase change in the direction of propagation (for more details see Section 2.4). Apertures often cause these abrupt changes in the field. In such a plane one needs to perform another decomposition into secondary sources and propagate them using geometrical optics. Labelling the initial plane by S_0 and the second diffracting plane by S_1, we can summarize this mathematically by

$$
\boldsymbol{E}(\boldsymbol{r_2}) = \iint_{S_1} \iint_{S_0} \tilde{G} \left[\frac{-\mathrm{i}k}{2\pi} \hat{\boldsymbol{n}}_1 \times \right.
$$
$$
\left. \tilde{G} \left[\frac{-\mathrm{i}k}{2\pi} \hat{\boldsymbol{n}}_0 \times \boldsymbol{E}(\boldsymbol{r_0}), \boldsymbol{r_0}, \boldsymbol{r_1} \right], \boldsymbol{r_1}, \boldsymbol{r_2} \right] \mathrm{d}S_0 \mathrm{d}S_1. \tag{5.14}
$$

Any further planes where the geometrical optical propagator breaks down, can be treated in a similar manner. The magnetic field can be obtained by changing the final propagator to the one defined by Eq. (5.13)

$$
\boldsymbol{H}(\boldsymbol{r_2}) \approx \iint_{S_1} \iint_{S_0} \sqrt{\frac{\epsilon_2}{\mu}} \hat{\boldsymbol{\rho}} \times \tilde{G} \left[\frac{-\mathrm{i}k}{2\pi} \hat{\boldsymbol{n}}_1 \times \right.
$$
$$
\left. \tilde{G} \left[\frac{-\mathrm{i}k}{2\pi} \hat{\boldsymbol{n}}_0 \times \boldsymbol{E}(\boldsymbol{r_0}), \boldsymbol{r_0}, \boldsymbol{r_1} \right], \boldsymbol{r_1}, \boldsymbol{r_2} \right] \mathrm{d}S_0 \mathrm{d}S_1. \tag{5.15}
$$

The result is a cascaded diffraction integral that fully accounts for the diffraction introduced at multiple surfaces.

5.2.3. DISCUSSION

With some minor adjustments, the derivation presented here can also be used for scalar HFPI. In that case one would use the Rayleigh-Sommerfeld diffraction integral (see Sec-

tion 2.3.1) instead of the m-theory. This results in elementary fields with the form of Eq. (4.2), repeated below

$$E^{(RS)}(\boldsymbol{r_0}, \boldsymbol{r_1}, E(\boldsymbol{r_0})) = -2\mathrm{i}kG_H\hat{\boldsymbol{\rho}} \cdot \hat{\boldsymbol{n_0}}E(\boldsymbol{r_0}). \tag{5.16}$$

This is related to the elementary vectorial fields, Eq. (5.1), by the transformation

$$\hat{\boldsymbol{\rho}} \times (\hat{\boldsymbol{n_0}} \times E(\boldsymbol{r_0})) \rightarrow \hat{\boldsymbol{\rho}} \cdot \hat{\boldsymbol{n_0}}E(\boldsymbol{r_0}). \tag{5.17}$$

Both these sources radiate strongest in the direction perpendicular to the surface of integration, but their attenuation for other directions is slightly different. Since both Eq. (4.2) and (5.1) depend on G_H, the Greens function for the Helmholtz equation, their wavefronts and attenuation upon propagation are identical. The formalism of the cascaded diffraction integral developed in this section can therefore also be applied to scalar HFPI, resulting in

$$E(\boldsymbol{r_1}) = \iint_{S_0} \tilde{G}^s \left[\frac{-\mathrm{i}k}{2\pi} E(\boldsymbol{r_0})\hat{\boldsymbol{n_0}}, \boldsymbol{r_0}, \boldsymbol{r_1} \right] \mathrm{d}S_0. \tag{5.18}$$

where for free space the propagator is a spherical wave with an amplitude that reduces for directions away from the surface normal

$$\tilde{G}^s [A\hat{\boldsymbol{n_0}}, \boldsymbol{r_0}, \boldsymbol{r_1}] = A\hat{\boldsymbol{n_0}} \cdot \hat{\boldsymbol{\rho}} \frac{\exp(\mathrm{i}k\rho)}{\rho}. \tag{5.19}$$

Scalar HFPI is based on the Rayleigh-Sommerfeld and vectorial HFPI on the m-theory integral(s), but these are only two of the many diffraction integrals in optics. One advantage of the RSDI and the m-theory integrals over many other diffraction integrals is that their kernels only depend on the electric field. Many vectorial diffraction integrals also need the magnetic field and many scalar integrals depend on the derivative of the electric field. Additionally tracing the magnetic field or calculating the derivative of the electric field would needlessly complicate the algorithm. Many vectorial diffraction integrals (e.g., the Stratton-Chu [74] and the vectorial Kirchhoff integral [73, p. 482–485]) are incompatible with the concept of HFPI due to a more fundamental problem. Since in HFPI the kernel of the integral is interpreted as the field of a secondary source, this field should be divergence-free. Second, since geometrical optics is used to propagate this field, the local electric and magnetic field should be perpendicular to the direction of phase propagation. For many vectorial diffraction integrals this is not the case.

An obvious exception in the vectorial case are the integrals of the e-theory, which are analogue to those of the m-theory but for the magnetic instead of the electric field. Another diffraction integral that can be combined with the concept of HFPI is the Richards-Wolf integral [5]. This integral is closely linked to the plane wave HFPI introduced in Section 4.5.3. Section 5.5 will show the results for vectorial plane wave HFPI.

In summary, the formalism of a cascaded diffraction integral with a geometrical optical propagator is also applicable to scalar HFPI. The theory does require specific properties of the kernel of the diffraction integral: the field should only depend on the electric field and not on its derivative or on the magnetic field, the kernel of the integral should be divergence free, and the local electric and magnetic field of the kernel should be perpendicular to the direction of phase propagation. This makes the theory incompatible with many diffraction integrals.

5.3. Implementation in a Monte Carlo ray tracer

In Section 5.2 we derived a method to compute a cascaded diffraction integral that simulates the effects of diffraction and aberrations. In this section we will present the individual steps of a vectorial HFPI simulation using a Monte Carlo ray tracer.

Before any simulation one must define the optical system and the initial electric field. Furthermore, one must define at which planes diffraction has to be taken into account (e.g., at an aperture), and at which planes one wants to calculate the electric field.

The first step of the actual simulation consists of initiating random rays at the initial surface. If the input field can be propagated using geometrical optics, i.e., it is non-diffracting, the set of possible initial rays is defined by the field. Using the Monte Carlo algorithm one chooses random starting positions for the rays, which have directions perpendicular to the wavefront and a polarisation parallel to that of the input field. If the input field does not propagate according to the laws of geometrical optics (e.g., a Gaussian beam in its waist) the diffraction integral of Eq. (5.12) should be applied. The initiated rays have random starting positions and directions, and carry a field according to

$$E_{ray} = \sqrt{\Delta S_0 \Delta N n_0} \frac{\mathrm{i}k}{4\pi} \hat{\boldsymbol{\rho}} \times (\hat{\boldsymbol{n}_0} \times E(\boldsymbol{r_0})), \tag{5.20}$$

where $E(\boldsymbol{r_0})$ is the field in the starting point of the ray, ΔN the total number of rays distributed over the surface ΔS_0, $\hat{\boldsymbol{n}_0}$ the forward-pointing normal of the starting plane and $\hat{\boldsymbol{\rho}}$ the normalized ray direction.

After initiating the rays, they are propagated through the system using standard ray tracing procedures. At optical interfaces, the transmission coefficients are determined using the adapted Fresnel equations for ray tracing, Eqs. (2.84) and (2.85), and the weight of the rays are adapted accordingly.

At a diffracting surface, a ray is terminated and its complex contribution to the local electric field is determined by adjusting the ray weight with the factors w_i (Eq. (2.98)) and w_p (Eq. (2.99)). Since w_p is an amplitude factor related to the intensity law of geometrical optics, one needs a procedure like tracing neighbouring rays to determine w_p. This is not implemented in our Monte Carlo ray tracer, so we have to neglect this term, leading to a small error in the amplitude of the field. In Section 5.4 the effect of this error is quantified for a system with an NA of 0.7. After determining the local E-field, the diffraction integral of Eq. (5.12) is applied on the local field of a single ray. Note that the integral of Eq. (5.12) is over the total electric field in the plane, whereas we apply the integral to the fields of every single ray individually (and sum the resulting fields later). Due to the linearity of optics these two procedures (taking the integral over the sum of local fields or taking the sum of the integrals over local fields) are equivalent. In practice Eq. (5.12) is implemented by changing the direction of the incoming ray randomly, but one could achieve faster convergence by initiating a set with a user-defined number of new rays with random directions at the point of intersection of the ray and the diffracting surface. The weight of the new ray is adjusted according to Eq. (5.20), where $E(\boldsymbol{r_0})$ is the previously calculated contribution of the terminated ray to the local electric field.

These secondary rays are propagated further through the optical system. Any additional diffracting surface is treated identically to the first one, terminating the current rays and generating new rays instead. Finally, at the last (or any intermediate) surface,

the field is determined by having every ray contribute to the pixel it intersects. The weight of each ray is again adjusted using the factors w_i (Eq. (2.98)) and w_p (Eq. (2.99)). Furthermore, since we are now approximating the field at the centre of the pixel, the phase of each ray is shifted using the local plane wave approximation of Eq. (2.100).

Following this procedure, a Monte Carlo ray tracer can propagate an electric field through an optical system and include the aberrations introduced by the system and the diffraction introduced at specific planes of the system.

5.4. RESULTS

We compare results obtained using vectorial HFPI with those of two other simulation methods described in Section 2.6. The first method is the vectorial ray-based diffraction integral (VRBDI) of Andreas *et al.* [10] for which the authors have published their MATLAB code online [75]. The second reference method is Beam Synthesis Propagation (BSP) in the commercial software package CODE V [9]. The first simulated system is a low NA configuration, with diffraction of the input field, whereas the second and third system have higher NAs and demonstrate the effects of single and multiple diffraction respectively.

The first system is one of the example systems of the VRBDI MATLAB code [75] and consists of a singlet with in its focal plane a ring aperture illuminated by an x-polarised plane wave with a wavelength of 632.8 nm. A schematic representation of the system is shown in Fig. 5.1 and its parameters are listed in Table 5.1. This system can be modelled

Figure 5.1: A schematic representation of the ring aperture and singlet. The parameters are listed in Table 5.1, the results of an HFPI simulation of the system are shown in Fig. 5.2.

using HFPI by placing secondary sources in the ring aperture. The results of HFPI and the absolute difference between the HFPI results and those using the MATLAB code by Birk Andreas [75] are shown in Fig. 5.2. Note that these results and the other results presented in this section are normalized by the squared amplitude. The statistical nature of HFPI can be clearly seen in the difference plots. The difference between the methods is around 2% of the maximum amplitude for all three electrical components, even though the y-component of the field is around six orders of magnitude smaller than the

Table 5.1: System parameters of the singlet with a ring aperture [75] (see Fig. 5.1). The aperture is placed 1 mm before the front focal plane of the singlet.

Quantity	Value
Ring aperture radius	1.25 mm
Ring aperture width	0.01 mm
Distance aperture to lens	300 mm
Thickness lens	3 mm
Refractive index lens	1.5155
Radius of curvature first surface	308.5 mm
Radius of curvature second surface	-308.5 mm
Distance lens to output plane	100 mm
Wavelength	632.8 nm
Numerical aperture	0.004

x-component. The L_2-difference between the two methods is defined as

$$L_2\text{-dif} = \frac{\sqrt{\sum_{p=1}^{N_p} \sum_{i=x,y,z} \left(\left| E_{i,p}^{(2)} - E_{i,p}^{(1)} \right| \right)^2}}{\sqrt{\sum_{p=1}^{N_p} \sum_{i=x,y,z} \left| E_{i,p}^{(1)} \right|^2}}, \tag{5.21}$$

where the superscript (1) indicates the results of the reference method, the superscript (2) those of HFPI, the subscript p runs over all N_p pixels, and the subscript i runs over the three Cartesian components of the field. An overall phase difference (i.e., a piston phase term) between the two methods is corrected manually before calculating Eq. (5.21). For this example, the L_2-difference between the results of HFPI and those of the MATLAB code by Birk Andreas is 2.4%. The simulation times and the number of paths drawn for this simulation and the other simulations presented in this section are listed in Table 5.2. The computational loads for the simulations presented in this chapter are between 12 and 306 core hours. Since the method is almost perfectly parallizable, the duration of

Table 5.2: The number of paths and simulation times for the HFPI simulations. The simulations were performed on multiple (6-8) computers using a total of 36 to 48 cores simultaneously. Every path was traced from object to image plane.

System	Number of paths	Core hours
ring aperture	10^9	12
microscope	10^9	23
microscope corrected HFPI	$1.36 \cdot 10^9$	45
cascaded diffraction	10^{10}	306

the simulation can be a fraction of that time if parallel or distributed computing is used.
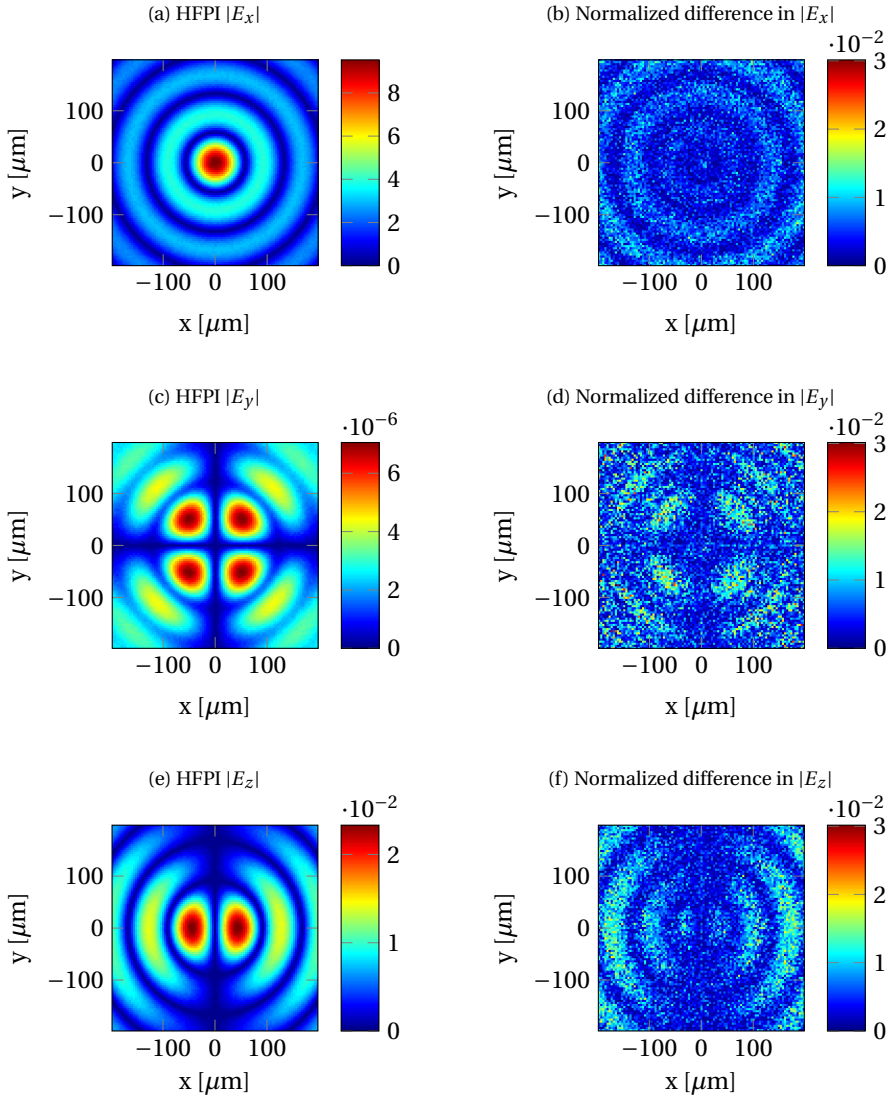
Figure 5.2: The results from HFPI for a plane wave illuminating a ring aperture, followed by a singlet (see Fig. 5.1 and Table 5.1 for the system description). The amplitudes of the electric field according to the HFPI simulation are shown in the left column, and the difference between these results and those of the model of Andreas *et al.* in the right column. The difference plots are normalized w.r.t. the maximum amplitude of the corresponding HFPI result. The L_2-difference (Eq. (5.21)) between the methods is 2.4%.

Furthermore, like for scalar HFPI, the L_2 error is inversely proportional to the square root of the number of paths (see Fig. 5.3). Decreasing the number of paths (thereby decreasing the computational load) by 2 or 3 orders of magnitude increases the L_2 error by

Figure 5.3: The L_2-norm of the relative error for a vectorial HFPI simulation is proportional to one over the square root of the number of paths. For small path numbers the error is smaller than the theoretical curve because the L_2-norm of the relative error can never exceed two.

roughly a factor of 10 or 32. As an example, we present such results for the ring aperture (Table 5.1) in Fig. 5.4. The plots are noisy but may suffice for many practical applications.

Figure 5.4: The amplitude of the z-component of the electric field in the output plane of the singlet with a ring aperture (see Table 5.1). The simulations are performed using three (left) or two (right) orders of magnitudes fewer rays than the simulation shown in Figure 5.2. Consequently the results are much noisier than the result shown in Figure 5.2, and the simulation times are reduced to roughly one and roughly six core minutes. This could offer a better compromise between accuracy and simulation time than the simulation for Figure 5.2.

In the previous example almost all the energy was concentrated in the main polarisation direction because the system had a very low effective numerical aperture of 0.004. We now consider the results for a microscope with an NA above 0.7 for which the effects of polarisation become much more pronounced. The system corresponds roughly to the example microscope system of CODE V which is based on a patent by Tojyo [76]. It is depicted in Fig. 5.5 and its parameters are listed in Table 5.3. The system is in reversed

Figure 5.5: A schematic representation of the reversed microscope. The system is in reversed order compared to the normal usage of microscopes. This results in a system with a high NA, and therefore clearly visible polarisation effects, in the image plane. The exit pupil and the rays at the exit pupil are depicted using dashed lines.

order compared to the original design and the normal usage of microscopes. Such a reversed configuration can be beneficial in the design process [77, p. 257–258] and, more importantly for the current discussion, makes the polarisation effects clearly visible. The results of an HFPI simulation of the system are shown in Fig. 5.6. As explained in Sections 5.2 and 5.3, propagating the amplitude of the field using a Monte Carlo ray tracer, which is designed to propagate intensities, introduces an error. To propagate the amplitude correctly, the Monte Carlo ray tracers would need a correction factor w_p, which is defined by Eq. (2.99) and is proportional to the ray density. For the system discussed above, w_p does not vary significantly for different rays, and this error has little effect. For the current system, the neglected term does vary noticeably and as a result one can see some systematic differences between the results of HFPI and BSP in Fig. 5.7. To verify that these differences are indeed mainly due to neglecting w_p, we adjust the system such that we can compensate for this error. This is done by letting the diffraction occur in the exit pupil, which is the image of the aperture in image space, instead of at the aperture itself. The two w_p factors of a ray (one for the propagation to the diffracting surface and one for the following propagation step to the image plane) are now both functions of positions of the rays in the exit pupil only. In such cases, our implementation enables us to account for w_p, by adjusting the amplitudes of the rays that pass through the exit pupil using a thin element with a spatially varying absorption. The right column of Fig. 5.7 shows that the differences between the BSP simulation and the HFPI simulation with correct w_p are significantly lower than for the HFPI simulation that assumes w_p to be 1. This may appear to suggest that assuming w_p to be constant is insufficient for simulating high NA systems but, as the cross plots in Fig. 5.8 demonstrate, the errors due to this approximation are mainly in the secondary polarisation directions and might be accept-

Table 5.3: System parameters of the reversed microscope. The aperture is indicated by ∗ and has a radius of 4.1 mm.

Surface number (j)	Radius [mm]	Distance [mm] (j to j+1)	Refractive index (after j)
0	∞	156.466582	1
1	-4.893054	1.235052	1.6164
2	331.759004	3.773815	1.7546
3	-7.999296	12.115041	1
4*	∞	0.000000	1
5	∞	4.439448	1
6	17.811462	2.602460	1.4984
7	-6.811396	1.000946	1.7918
8	65.868865	0.140215	1
9	11.474481	3.940087	1.4984
10	-6.105357	1.470398	1.5236
11	-14.404937	0.100095	1
12	7.120780	2.202081	1.4984
13	14.369366	0.100095	1
14	4.017020	4.660603	1.7919
15	2.431637	0.672514	1
16	∞	0.170000	1.5246
17	∞		

able for a wide range of applications. Furthermore, as mentioned before, the method can be adapted to account for w_p by tracing closely neighbouring rays and tracking the area spanned between these rays.

In the last example, diffraction is introduced in multiple planes. Moreover, the field in the second diffracting plane is affected by the diffraction in the first plane, i.e., cascaded diffraction occurs. The system consists of a pinhole placed in the focal plane of a singlet (see Table 5.4 and Fig. 5.9) followed by the previously described reversed microscope. The y-polarised plane wave ($\lambda = 546.1$nm) passing through the singlet generates a PSF in the plane of the pinhole. There, the semi-circular pinhole with a radius of 40 μm, completely blocks the lower half of the PSF. The pinhole is then imaged by the reversed microscope of the previous example (see Table 5.3). Both the object NA of the microscope and the image NA of the singlet are 0.018. The system has three diffracting planes: the aperture of the singlet, the semi-circular pinhole, and the physical aperture of the microscope. If we apply the cascaded diffraction integral presented in Section 5.2.2 to this problem, we need to perform three consecutive integrals. This means that the fields are first decomposed into secondary sources at the aperture of the singlet, then decomposed again at the pinhole, and for a final time at the aperture of the microscope. This is implemented in the Monte Carlo ray tracer by randomly changing the direction of every ray, and adjusting its field using Eq. (5.20), in each of the aforementioned diffracting

Figure 5.6: A magnetic dipole (see Eq. (5.3)) directed along the x-axis, imaged by the microscope with the parameters listed in Table 5.3. The three rows show the amplitudes of the three components of the electric field in the focal plane, as predicted by HFPI.

planes. The results of HFPI and the difference with a simulation with BSP are shown in Fig. 5.10.

Simulating the triple diffraction is computationally more demanding than the previous simulations which is why one can observe stronger statistical noise (due to the Monte Carlo nature of the method). This noise can be reduced at the cost of a longer simulation time by sampling more rays. The cascaded diffraction simulation shows similar differences between our implementation of HFPI and BSP as the simulation of the
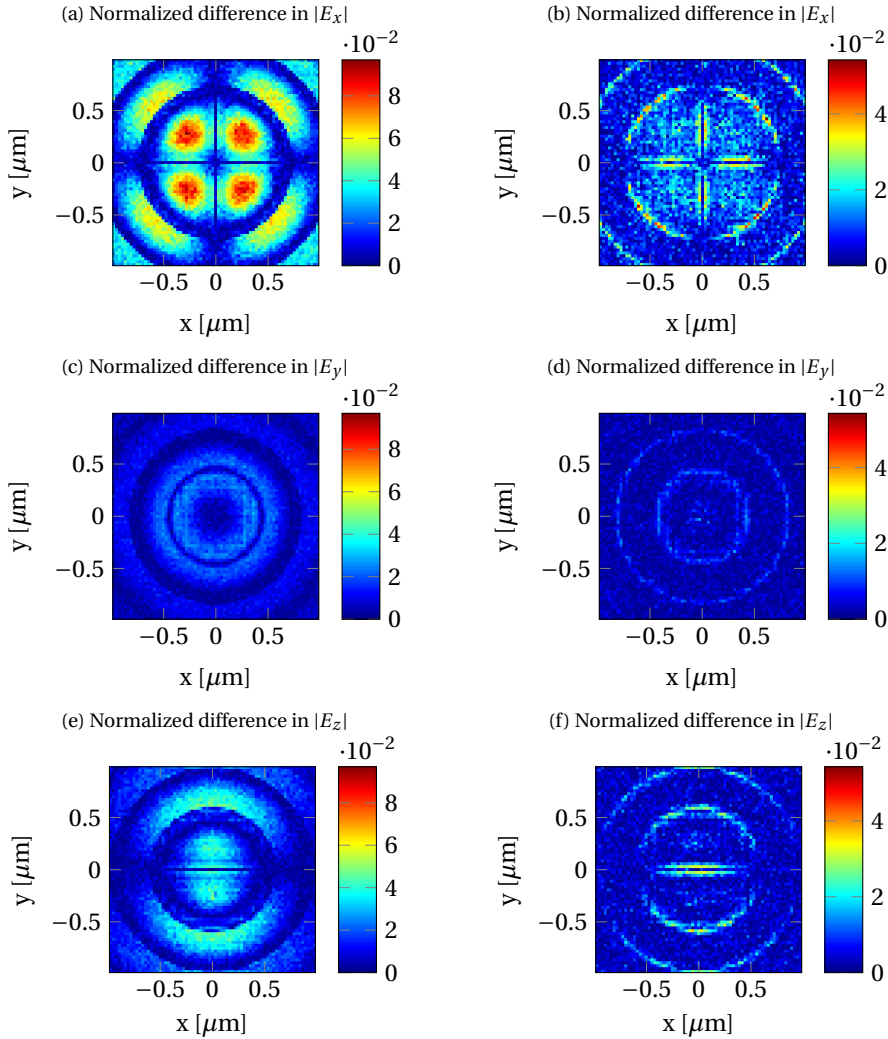
Figure 5.7: The normalized differences between two HFPI simulations and a BSP simulation of a microscope imaging a magnetic dipole. The left column shows the differences between the results of our HFPI implementation (shown in Fig. 5.6) and those of a BSP simulation. For the right column the amplitude factor w_p, which is usually neglected in our implementation of HFPI, is accounted for. The plots show the differences between the HFPI and BSP simulation. They are normalized w.r.t. the maximum amplitude of the corresponding HFPI result. Note that the left and right figures have different color scales. The L_2-differences between HFPI and BSP are 7.3% and 3.3% for respectively the uncorrected and the corrected HFPI.

microscope did (Figure 5.6), which suggests that the remaining differences are at least partly due to taking w_p constant. A simulation method that cannot describe cascaded diffraction would produce results very different from the ones shown in Figure 5.10. To
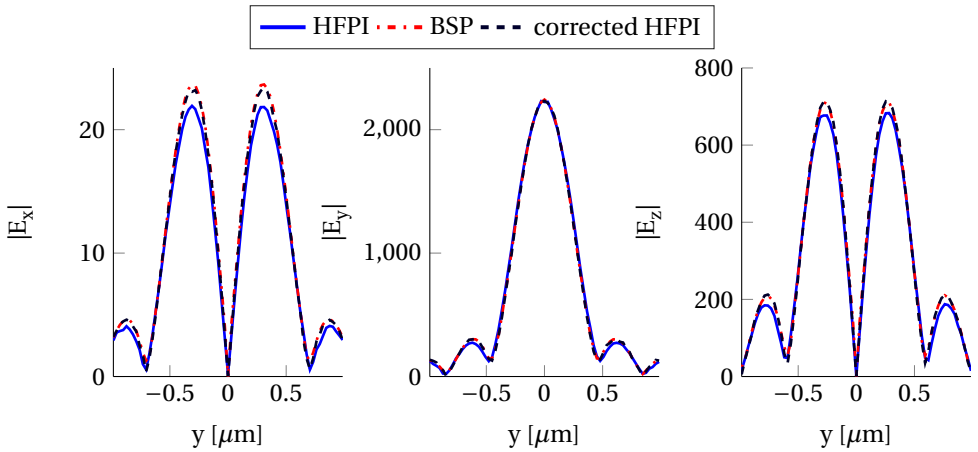
Figure 5.8: Cross sections of the amplitude of the field components at the focal plane of the microscope due to a magnetic dipole oriented along the x-axis positioned on the optical axis in the object plane. In order to include the peak amplitudes, the cross section of the x-component is at $x = 2.5\mu m$ and the other components are shown at $x = 0$. The plots show the results from the BSP simulation as well as the results from the standard HFPI simulation, which neglects an amplitude factor w_p (Eq. (2.99)), and the results from the HFPI simulation that accounted for w_p (corrected HFPI). It can be seen that, in comparison to the other simulations, the peak amplitudes of HFPI are a bit lower for the x- and z-polarisation and the PSF is slightly wider.

Figure 5.9: A schematic representation of the system with cascaded diffraction. The semi-circular aperture (shown at the top) is placed in the focal plane of the singlet and followed by the microscope (represented by the blue box) which is depicted in Fig. 5.5. The parameters of the singlet and the microscope are listed in Tables 5.4 and 5.3. The simulations are performed for a y-polarised plane wave with a wavelength of 546.1 nm.

illustrate this, we simulate the same system but only take the diffraction at the final aperture into account. This is done using an HFPI simulation with only one diffracting plane. In Figure 5.11 the result of this simulation is shown and compared to the result of the triple diffraction simulation of Figure 5.10. The L_2-differences between these simulations is around 50%. This shows the effect of (neglecting) cascaded diffraction.

5.5. VECTORIAL PLANE WAVE HFPI

It was shown in Section 4.5.3 that the convergence of HFPI can be greatly increased by associating a plane wave to every path in the image space. This concept of plane wave

Table 5.4: System parameters for the first part of the system showing cascaded diffraction. The second part of the the system is the microscope with the parameters listed in Table 5.3. The aperture is indicated by ∗ and has a radius of 4.1 mm. The simulation is performed for a y-polarised plane wave with a wavelength of 546.1 nm.

Surface number (j)	Radius [mm]	Distance [mm] (j to j+1)	Refractive index (after j)
0	∞	0.000000	1
1	244.210307	5.000000	1.5187
2	-244.210307	0.500000	1
3*	∞	0.000000	1
4	∞	234.000000	1
5	∞	0.000000	1

HFPI (PW-HFPI) is related to the Richards-Wolf integral and can be applied if the system has a large Fresnel number. By using PW-HFPI for the system demonstrating cascaded diffraction (depicted in Fig. 5.9) the simulation time can be reduced from 306 core hours for an HFPI simulation to 59 core minutes for an PW-HFPI simulation. As is evident from the plots in Fig. 5.12, the PW-HFPI simulation is also better converged.

Polarization effects are visible in systems with high numerical apertures, which, assuming the aperture is much larger than the wavelength, have high Fresnel numbers. The high NA thus requires the use of a vectorial method and guarantees that the error in using the plane wave approximation is small. This makes vectorial PW-HFPI a particularly useful method for the analysis of systems with large NAs.

5.6. DISCUSSION

The method presented here is an extension of the method presented in Chapter 4 to vectorial optics. As such it inherits the properties of scalar HFPI, which are discussed in Section 4.6. In the current section, we will therefore refrain from listing these properties again and focus on a comparison between HFPI and other simulation methods instead.

Huygens-Fresnel path integration has an interesting conceptual link with Feynman path integrals. As with the Feynman path integrals, the core idea behind HFPI is the complex addition of paths. Path integral methods are known in quantum mechanics, and the Schrödinger equation and its relativistic generalization, the Klein-Gordon equation, can be linked to scalar paraxial and non-paraxial optics [78]. Since the vast majority of paths interfere destructively, one challenge in numerical path integral methods is the exclusion of (most of) these non-contributing paths. In HFPI this is done by only sampling stationary paths (i.e., the paths of geometrical optical rays) between what we call diffracting surfaces.

On an abstract level, HFPI is also connected to the Gaussian Beam Decomposition (GBD) [8] and related methods like CODE V's BSP [9]. All these methods rely on decomposing a field into subfields, propagating these using geometrical optics, and rebuilding the field by coherent addition of the subfields. When the geometrical optical propagation fails, these methods decompose the field into new subfields, and can thereby ac-

(a) HFPI $|E_x|$

(b) Normalized difference in $|E_x|$

(c) HFPI $|E_y|$

(d) Normalized difference in $|E_y|$

(e) HFPI $|E_z|$

(f) Normalized difference in $|E_z|$
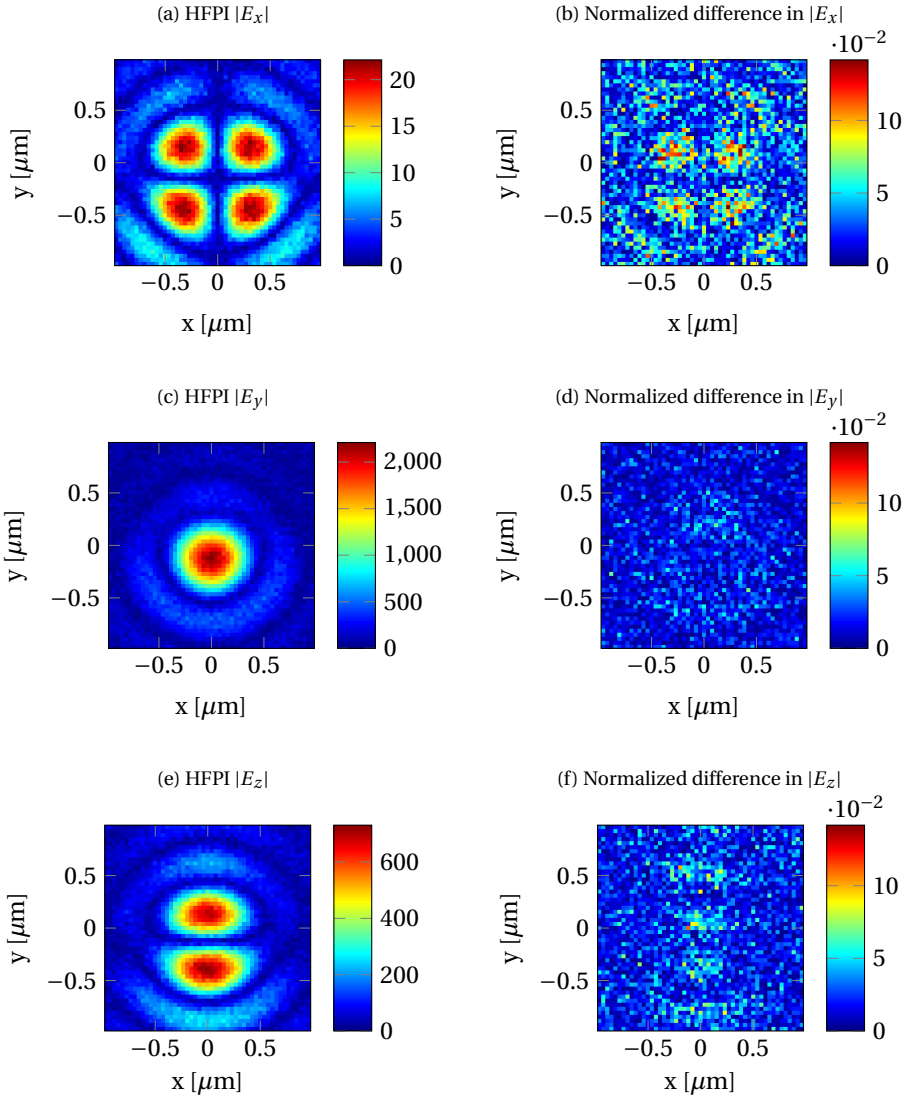
Figure 5.10: The amplitudes of the field components in the focal plane of the system demonstrating cascaded diffraction (left column) and the differences with the results from the BSP simulations (right column). The difference plots are normalized w.r.t. the maximum amplitude of the corresponding HFPI result. The L_2-difference between the methods is 13.2%.

count for diffraction. The methods differ in the type of subfield they use. Where the GBD, as the name implies, uses Gaussian subfields, HFPI uses the fields of magnetic dipoles. The fact that these dipoles are localized in a point makes the decomposition step in HFPI simple, but because the subfields are strongly diverging, propagating them

Figure 5.11: The figure demonstrates the effect of neglecting cascaded diffraction when simulating the system shown in Fig. 5.9. The left figure shows the amplitude of the y-component of the electric field in the focal plane of the system, calculated using HFPI with only the diffraction at the last aperture accounted for. The right figure shows the difference in the absolute amplitude between these results and the results shown in Fig. 5.10, where all three diffracting planes are accounted for.

5

Figure 5.12: Cross sections of the normalized amplitudes of the field in the focal plane of the system depicted in Fig. 5.9. The cross section of the x-component of the field is along the line $x = 0.4\mu m$ to include the peak amplitude. The other plots are along the line $x = 0$. The HFPI simulation has an L_2-difference of 13.2% and the PW-HFPI simulation of 5.4% with respect to the BSP simulation. A comparison with the plots for the microscope (Fig. 5.8) shows that the deviations between the HFPI and the reference simulations are very similar for both systems. As explained in Section 5.4, these deviations are due to approximating the factor w_p in Eq. (2.99) to be unity.

is computationally expensive compared to the GBD.

Another decomposition is the angular spectrum, used in Fourier optics, which decomposes a field into plane waves [44]. One could see the angular spectrum and HFPI as two extremes of the same concept, with HFPI sources perfectly localized in space but infinitely extended in frequency space, and the plane waves of the angular spectrum perfectly localized in frequency space but infinitely extending in space. The GBD adapts

an intermediate approach with sources that have a finite extent in both ordinary and reciprocal space.

The vectorial ray-based diffraction integral (VRBDI) by Andreas *et al.* [10] allows both the decomposition into point sources and into plane waves. As with HFPI, the field in VRBDI is propagated through an optical system using geometrical optics. Although conceptually similar, VRBDI and HFPI are implemented very differently with different applications in mind. HFPI uses a Monte Carlo ray tracer and was developed to find the field in the focal plane of a (multiple) diffracting optical system. VRBDI uses an algorithm that iteratively finds rays connecting the grid points on the input and the output plane, i.e., ray aiming. Furthermore VRBDI, in its current implementation, avoids focal regions and apertures. Andreas *et al.* [10] mention that apertures could be accounted for by first calculating the field in the aperture using a VRBDI simulation, then applying the aperture function, and finally decomposing and propagating this field using a second VRBDI simulation [10]. Based on our results, one can expect that such an intermediate resampling would allow VRBDI to calculate the electric field in the focal plane of an optical system. It is not clear to us how such a simulation would perform if the phase in the aperture is highly oscillatory, as is the case for many optical systems. Since HFPI does not require an explicit sampling grid, such intermediate oscillatory fields do not pose a problem for HFPI.

5.7. CONCLUSION

We present the theory for a ray-based vectorial diffraction simulation method based on a physical interpretation of the m-theory diffraction integrals, and compare its results to those of other methods. Huygens-Fresnel path integration is a robust method to simulate the effects of aberrations, (multiple) diffraction, and polarization. The method is implemented by making some small adjustments to an existing Monte Carlo ray tracing engine, thereby hugely extending its domain of application with relatively little programming effort. We have theoretically derived and numerically validated the error caused by using standard Monte Carlo ray tracing techniques, which were designed to propagate intensities, to propagate the amplitude of the electric field (i.e., the error caused by assuming $w_p = 1$). In our simulations this error was relatively small and did not justify the added complexity of changing the Monte Carlo ray tracer. The absence of explicit resampling in HFPI not only makes the method relatively easy to implement but also makes it robust and fairly easy to use.

6

CONCLUSION

The simulation of point spread functions is essential in the design of optical instruments. The exit pupil diffraction integral (EPDI) and Collins integral (or paraxial Fourier optics) are often used to this purpose, but sometimes we need to go beyond the approximations of these models. The Collins integral can only simulate systems consisting of ideal lenses. It does not give any information on the aberrations of the system, but can describe the effects of cascaded diffraction. The EPDI accounts for the aberrations of a system, but assumes all the diffraction to occur at a single surface in the image space. This thesis introduces two simulation methods for optical systems that aim to account for both aberrations and cascaded diffraction: Wigner-based ray tracing (WBRT) and Huygens-Fresnel path integration (HFPI).

The Wigner distribution function (WDF), when used as a representation of the scalar electric field, remains constant along paraxial rays. This property was used to develop WBRT, a Monte Carlo ray tracer that propagates the WDF using rays. Since rays are not limited to paraxial propagation, WBRT can be applied to nonparaxial problems (i.e., problems that include lenses with aberrations or rays with large angles with respect to the optical axis). The underlying assumption of WBRT is that taking the WDF to be constant along rays is a good approximation beyond the paraxial domain.

The results presented in this thesis show that this is only partly true. Although lenses with a numerical aperture (NA) of up to 0.2 were accurately modelled using WBRT, the method breaks down dramatically in the presence of aberrations. In these cases WBRT can predict unphysical negative intensities, even when simulating a singlet with an NA as low as 0.04 and a root mean square wavefront error below half a wavelength. Note that this does not mean that the WDF cannot be used beyond the paraxial domain. It means that WBRT, the simulation method using ray tracing to transport the WDF, produces inaccurate results in the presence of aberrations. Establishing the limit of its domain of application is the main result of this research into WBRT.

This result can be relevant to the work in computer graphics [14, 38–40] that inspired our investigation into WBRT. The goal of applying the WDF in computer graphics is to render photorealistic images by including diffraction effects. Being aware of the limits of

WBRT could be important to prevent rendering erroneous images. Considering the main goal in computer graphics is creating images that look similar to a photo, our results do not necessarily exclude the use of the WDF beyond the paraxial domain within computer graphics.

When physical accuracy is required the use of WBRT well beyond the paraxial domain, and in particular in the presence of aberrations, cannot be recommended. This is primarily a result regarding WBRT, but it can also serve as a warning regarding the use of the WDF beyond the paraxial domain. It was shown by Lohmann *et al.* [45] that the effect of aberrations on the WDF is twofold. First they lead to a coordinate transformation in accordance with geometrical optics, and second they lead to differential operators that change the shape of the WDF. We have shown that for PSF simulations, only modelling this primary effect does not approximate the correct answer but can result in unphysical negative intensities.

That is not to say that WBRT has no potential. The WDF can be a convenient representation of spatially partially coherent light, so it should be possible to use WBRT for the propagation of such fields through a paraxial optical system. It might also be worthwhile to investigate if there are problems that can benefit from the relatively low memory usage of WBRT. One could also try to develop more elaborate methods to propagate the WDF through a system with aberrations. Generalized versions of the WDF that remain constant upon nonparaxial propagation through homogeneous media exist [79]. Furthermore, it has been shown that the main effect of refraction at a plane interface on such a generalized WDF, corresponds to the transformation undergone by a ray [80]. However, this is only approximately true and, to our knowledge, it has not yet been investigated to which extent this result can be generalized to curved surfaces. This generalization of the WDF could be a starting point for the development of a ray-based WDF transportation through aberrated systems, but a successful outcome is far from certain.

Figure 6.1: The original Huygens principle and Young's double slit experiment illustrating the Huygens-Fresnel principle and HFPI. Huygens principle initiates secondary spherical waves on the wavefront to construct a new wavefront. The Huygens-Fresnel principle initiates spherical waves on the wavefront and lets them interfere to obtain a diffraction pattern. HFPI uses the same spherical waves and samples them using rays.

HFPI on the other hand already provides a way to simulate the effects of cascaded diffraction in an optical system with aberrations. The concept of HFPI is a straightforward extension of the Huygens-Fresnel principle to optical systems. This principle states

that one can model the propagation of light by placing secondary spherical sources of light on the wavefront, as illustrated in Fig. 6.1a, and propagate and coherently sum these fields. The resulting model describes the effects of diffraction (see Fig. 6.1b). The concept of HFPI, depicted in Fig. 6.1c, is simply to use ray tracing to propagate these secondary sources. Since ray tracing, and thereby HFPI, can be used within an optical system (see Fig. 6.2), this allows for the simulation of diffraction and aberrations.

Figure 6.2: Illustration of a PSF calculation for a single lens using HFPI.

The Huygens-Fresnel principle was put on a solid mathematical basis by the Kirchhoff diffraction integral. Likewise, HFPI obtains its theoretical foundation using the Rayleigh-Sommerfeld diffraction integral for the scalar method, and the m-theory integrals for the vectorial method. These diffraction integrals can be interpreted as a decomposition of the field into point, or more specifically dipole, sources which are of such a form that they can be propagated using geometrical optics.

The choice to implement HFPI in a Monte Carlo ray tracer was motivated by their ability to quickly trace many rays and their suitability for parallel and distributed computing. Furthermore, using Monte Carlo sampling keeps the algorithm relatively simple. The use of a standard Monte Carlo ray tracer to transport the amplitude (as opposed to the intensity which it usually propagates) does lead to an error. For the PSF calculations performed in this thesis these errors were small.

The numerical results presented in this thesis demonstrate that HFPI can be used for simulating optical systems beyond the approximations used in the EPDI or the Collins integral. HFPI was validated by comparing its results to those of a range of simulation methods, including CODE V's BSP. Like HFPI, this commercially available ray-based diffraction integral is capable of simulating cascaded diffraction in aberrated systems. Compared to BSP, HFPI has several advantages. First, HFPI is very robust since it does not require an explicit resampling in diffracting planes nor does the user need to set any parameters for a decomposition of the field. Another advantage of HFPI is that all knowledge regarding the method is, with the publication of this thesis, in the public domain. It is our experience that in performing complex simulations (like those involving cascaded diffraction effects in optical instruments), it can be of crucial importance to understand the working principles and limitations of the simulation method one uses. A final advantage of HFPI is that it is relatively easy to implement. Like most ray-based diffraction integrals, HFPI requires a ray tracer capable of keeping track of optical path lengths. But other than that, HFPI only needs a Monte Carlo algorithm whereas other methods need an algorithm to accurately decompose a field into secondary sources. If one has access to a Monte Carlo ray tracer that can keep track of optical path lengths, everything is in

place to perform HFPI simulations of systems with cascaded diffraction.

Further developments of HFPI could focus on improving its accuracy and speed. As mentioned before, the use of a standard Monte Carlo ray tracer with minimal adjustments causes a small error. This error can be eliminated by adapting the Monte Carlo ray tracer such that it explicitly calculates the changes in the field strength due to the intensity law of geometrical optics. That would also allow the method to calculate absolute energies. This can for instance be done by tracing very small bundles of rays instead of individual rays, but this is likely to come at the cost of a higher computational load.

In this thesis several strategies were suggested to reduce the computational load of HFPI. We showed that plane wave HFPI can reduce the computational time by about two orders of magnitude compared to regular HFPI for system with large enough Fresnel numbers. For one simple optical system (i.e., the Cooke triplet) we determined a Fresnel number of ten to be the limit for a correct simulation of the phase, whereas the amplitude was correct for all Fresnel numbers. Further research is needed to determine whether these results are specific for this optical system or general properties of plane wave HFPI. The other strategies for reducing the computational load still require substantial further research. One option is to divide the field into an undisturbed and a diffracted part and only initiate secondary rays in the diffracted part. This has not yet been attempted and is probably not straightforward for HFPI. Bidirectional and stepwise HFPI were already demonstrated to work in principle, but the computational benefit for specific systems remains to be determined. It is also unclear what the sampling criterion for these methods is. Other methods that have been identified as potentially beneficial to the simulation times are adaptive sampling, stratified sampling, quasi-Monte Carlo sampling, and ray-splitting. Adaptive sampling requires a careful implementation to prevent systematic errors, whereas the other three techniques should be relatively easy to implement if the ray tracing engine allows such procedures. Improvements in computational efficiency are particularly relevant for demanding simulations, like the calculation of a field far away from a focal region.

Within this thesis we have applied HFPI to propagate the electric field from a plane in object space to the image plane of an optical system. In practise however, optical systems are often used to image three-dimensional scenes. The imaging of such scenes can be simulated using rendering methods based on Monte Carlo ray tracing which do not capture diffraction effects (see, e.g., [81]). Since HFPI is based on the same technique of Monte Carlo ray tracing, the principles presented in this thesis could be used to adapt rendering methods such that they correctly simulate diffraction effects. In theory such an adaptation is straightforward, but to develop a method that can render a scene with realistic computational effort would require more research.

The duration of an HFPI simulation depends strongly on the implementation of the ray tracer and the type of hardware used. Recent years have seen an increasing availability of cloud computing services, big improvements in the performance of graphics processing units (GPUs), and a reduction in the effort required to use either. Our current implementation of HFPI uses neither cloud computing nor GPUs, but we did distribute simulations over multiple computers using several cores on each computer. This demonstrates the perfect scalability of the method.

Even without the use of the most modern hardware, diffraction effects were simu-

lated by summing complex paths. This success encourages the research into path integration methods for other problems in optics.

This thesis set out to develop a method for the simulation of cascaded diffraction in optical systems for which one cannot neglect the aberrations. This has been achieved with the introduction of Huygens-Fresnel path integration. Although only one of many simulation tools available to physicists, it is our hope that this additional tool can enable the development of new optical instruments.

6

REFERENCES

[1] A. Malet, *Kepler and the telescope,* Ann. Sci. **60,** 107 (2003).

[2] H. Volkmann, *Ernst Abbe and his work,* Appl. Opt. **5,** 1720 (1966).

[3] A. Sommerfeld and J. Runge, *Anwendung der Vektorrechnung auf die Grundlagen der geometrischen Optik,* Ann. Phys. **340,** 277 (1911).

[4] S. A. Collins Jr, *Lens-system diffraction integral written in terms of matrix optics,* J. Opt. Soc. Am. **60,** 1168 (1970).

[5] E. Wolf, *Electromagnetic diffraction in optical systems. I. An integral representation of the image field,* P. Roy. Soc. Lond. A Mat. **253,** 349 (1959).

[6] M. Rataj, *Convolution formulation of Fresnel diffraction for cascaded optical systems,* Opt. Acta **31,** 137 (1984).

[7] J. W. Y. Lit and R. Tremblay, *Boundary-diffraction-wave theory of cascaded-apertures diffraction,* J. Opt. Soc. Am. **59,** 559 (1969).

[8] A. W. Greynolds, *Propagation of generally astigmatic Gaussian beams along skew ray paths,* (1986) pp. 33–51.

[9] Synopsys, *CODE V beam synthesis propagation,* https://optics.synopsys.com/codev /pdfs/CV_BSP_datasheet.pdf (2017), accessed June 10, 2018.

[10] B. Andreas, G. Mana, and C. Palmisano, *Vectorial ray-based diffraction integral,* J. Opt. Soc. Am. A **32,** 1403 (2015).

[11] E. Wigner, *On the quantum correction for thermodynamic equilibrium,* Phys. Rev. **40,** 749 (1932).

[12] L. Dolin, *Beam description of weakly-inhomogeneous wave fields,* Izv. Vuz. Radiofiz. **7,** 559 (1964).

[13] M. Born and E. Wolf, *Principles of optics,* seventh ed. (Cambridge University Press, The Edinburgh Building, Cambridge, CB2 2RU, UK, 1999).

[14] T. Cuypers, T. Haber, P. Bekaert, S. B. Oh, and R. Raskar, *Reflectance model for diffraction,* ACM T. Graphic. **31,** 122 (2012).

[15] H. Severin, *Zur Theorie der Beugung elektromagnetischer Wellen,* Z. Phys. **129,** 426 (1951).

[16] E. Hecht, *Optics*, 4th ed. (Pearson Education Inc., 1301 Sansome St., San Francisco, CA 94111, USA, 2002).

[17] K. Sarabandi, *Dyadic Green's function,* Lecture notes (2009).

[18] M. A. Alonso, *Wigner functions in optics: describing beams as ray bundles and pulses as particle ensembles,* Adv. Opt. Photon. **3**, 272 (2011).

[19] P. C. Chang, J. Walker, and K. Hopcraft, *Ray tracing in absorbing media,* J. Quant. Spectrosc. Ra. **96**, 327 (2005).

[20] A. v. Pfeil, F. Wyrowski, A. Drauschke, and H. Aagedal, *Analysis of optical elements with the local plane-interface approximation,* Appl. Optics **39**, 3304 (2000).

[21] G. Spencer and M. Murty, *General ray-tracing procedure,* J. Opt. Soc. Am. **52**, 672 (1962).

[22] F. Dekking, C. Kraaikamp, H. Lopuhaä, and L. Meester, *A modern introduction to probability and statistics* (Springer, Mekelweg 4, 2628CD Delft, The Netherlands, 2005).

[23] *Wave optics in ASAP*, Breault Research Organization, Inc., 6400 East Grant Road, Suite 350, Tucson, AZ 85715, USA (2008).

[24] J. B. Keller, *Geometrical theory of diffraction,* J. Opt. Soc. Am. **52**, 116 (1962).

[25] G. L. James, *Geometrical theory of diffraction for electromagnetic waves*, 3rd ed. (Institution of Engineering and Technology, UK, 1986).

[26] G. W. Forbes and M. A. Alonso, *Using rays better. I. Theory for smoothly varying media.* J. Opt. Soc. Am. A **18**, 1132 (2001).

[27] M. A. Alonso and G. W. Forbes, *Using rays better. II. Ray families to match prescribed wave fields,* J. Opt. Soc. Am. A **18**, 1146 (2001).

[28] M. A. Alonso and G. W. Forbes, *Using rays better. III. Error estimates and illustrative applications in smooth media,* J. Opt. Soc. Am. A **18**, 1357 (2001).

[29] G. W. Forbes, *Using rays better. IV. Theory for refraction and reflection,* J. Opt. Soc. Am. A **18**, 2557 (2001).

[30] M. A. Alonso, *Ray-based diffraction calculations using stable aggregates of flexible elements,* J. Opt. Soc. Am. A **30**, 1223 (2013).

[31] E. R. Freniere, G. G. Gregory, and R. A. Hassler, *Edge diffraction in Monte Carlo ray tracing,* in *Proc. SPIE*, Vol. 3780 (1999) pp. 151–157.

[32] V. Serikov and S. Kawamoto, *Numerical experiments in Monte Carlo modeling of polarization, diffraction, and interference phenomena,* in *Proc. SPIE*, Vol. 4436 (2001) pp. 80–88.

6

[33] E. R. Freniere, *Numerical experiments in modeling diffraction phenomena with Monte Carlo ray tracing*, in *Proc. SPIE*, Vol. 6289 (2006) p. 62890N.

[34] A. Walther, *Radiometry and coherence*, J. Opt. Soc. Am. **58**, 1256 (1968).

[35] M. J. Bastiaans, *The Wigner distribution function applied to optical signals and systems*, Opt. Commun. **25**, 26 (1978).

[36] B. M. Mout, M. Wick, F. Bociort, and H. P. Urbach, *A Wigner-based ray-tracing method for imaging simulations*, in *Proc. SPIE*, Vol. 9630 (2015) p. 96300Z.

[37] M. Mout, M. Wick, F. Bociort, J. Petschulat, and P. Urbach, *Ray tracing the Wigner distribution function for optical simulations*, Opt. Eng. **57**, 57 (2018).

[38] S. B. Oh, S. Kashyap, R. Garg, S. Chandran, and R. Raskar, *Rendering wave effects with augmented light field*, Comput. Graph. Forum **29**, 507 (2010).

[39] T. Cuypers, T. Mertens, P. Bekaert, S. B. Oh, and R. Raskar, *Interactive point spread function simulation with diffraction and interference effects*, in *VISIGRAPP 2011*, edited by G. Csurka, M. Kraus, L. Mestetskiy, P. Richard, and J. Braz (IMAGAPP, 2011) pp. 19–24.

[40] T. Cuypers, R. Horstmeyer, S. B. Oh, P. Bekaert, and R. Raskar, *Validity of Wigner distribution function for ray-based imaging*, in *Proceedings of IEEE International Conference on Computational Photography* (IEEE, Pittsburgh, 2011) pp. 1–9.

[41] W. D. Furlan, G. Saavedra, E. Silvestre, J. A. Monsoriu, and J. D. Patrignani, *Assessment of a Wigner-distribution-function-based method to compute the polychromatic axial response given by an aberrated optical system*, Opt. Eng. **43**, 753 (2003).

[42] M. J. Bastiaans and P. G. J. v. d. Mortel, *Wigner distribution function of a circular aperture*, J. Opt. Soc. Am. A **13**, 1698 (1996).

[43] J. J. Stamnes, *Waves in focal regions: propagation, diffraction and focusing of light, sound and water waves* (Taylor & Francis, New York, 1986).

[44] J. G. Goodman, *Introduction to Fourier Optics* (Roberts & Company, Englewood, 2005).

[45] A. W. Lohmann, J. Ojeda-Castañeda, and N. Streibl, *The influence of wave aberrations on the Wigner distribution*, Opt. Appl. **13**, 465 (1983).

[46] K. Brenner and A. Lohmann, *Wigner distribution function display of complex 1D signals*, Opt. Commun. **42**, 310 (1982).

[47] D. Dragoman, *Wigner distribution function applied to third-order aberrations*, Appl. Optics **34**, 161 (1996).

[48] M. A. Alonso, *Diffraction of paraxial partially coherent fields by planar obstacles in the Wigner representation*, J. Opt. Soc. Am. A **26**, 1588 (2009), [doi:10.1364/JOSAA.26.001588].

6

[49] M. Mout, M. Wick, F. Bociort, J. Petschulat, and P. Urbach, *Simulating multiple diffraction in imaging systems using a path integration method*, Appl. Optics **55**, 3847 (2016).

[50] D. A. B. Miller, *Huygens's wave propagation principle corrected*, Opt. Lett. **16**, 1370 (1991).

[51] J. F. McCalmont, *A vector Huygens-Fresnel model of the diffraction of electromagnetic waves*, Ph.D. thesis, University of Arizona (1999).

[52] A. S. Marathay and J. F. McCalmont, *Vector diffraction theory for electromagnetic waves,* J. Opt. Soc. Am. A **18**, 2585 (2001).

[53] R. P. Heinisch and T. S. Chou, *Numerical experiments in modeling diffraction phenomena,* Appl. Opt. **10**, 2248 (1971).

[54] D. G. Fischer, S. A. Prahl, and D. D. Duncan, *Monte Carlo modeling of spatial coherence: free-space diffraction,* J. Opt. Soc. Am. A **25**, 2571 (2008).

[55] S. A. Prahl, D. D. Duncan, and D. G. Fischer, *Monte Carlo propagation of spatial coherence,* in *Proc. SPIE*, Vol. 7187 (2009) p. 71870G.

[56] S. A. Prahl, D. G. Fischer, and D. D. Duncan, *Monte Carlo Green's function formalism for the propagation of partially coherent light,* J. Opt. Soc. Am. A **26**, 1533 (2009).

[57] S. A. Prahl, D. D. Duncan, and D. G. Fischer, *Stochastic Huygens and partial coherence propagation through thin tissue,* in *Proc. SPIE*, Vol. 7573 (2010) p. 75730D.

[58] J. C. Ranasinghesagara, C. K. Hayakawa, M. A. Davis, A. K. Dunn, E. O. Potma, and V. Venugopalan, *Rapid computation of the amplitude and phase of tightly focused optical fields distorted by scattering particles,* J. Opt. Soc. Am. A **31**, 1520 (2014).

[59] J. Ferré-Borrull and S. Bosch, *Formal description of diffraction in optical systems: calculations and experimental evidence,* Appl. Phys. Lett. **85**, 2718 (2004).

[60] S. Bosch, *Ray tracing and scalar diffraction calculations of wavefronts, caustics and complex amplitudes in optical systems,* J. Mod. Optic. **58**, 146 (2011).

[61] R. Aarts, J. Braat, P. Dirksen, S. van Haver, C. van Heesch, and A. Janssen, *Analytical expressions and approximations for the on-axis, aberration-free Rayleigh and Debye integral in the case of focusing fields on a circular aperture,* J. Eur. Opt. Soc.-Rapid **3** (2008), 10.2971/jeos.2008.08039.

[62] Y. Li, *Dependence of the focal shift on Fresnel number and f number,* J. Opt. Soc. Am. **72**, 770 (1982).

[63] J. Ferré-Borrull and S. Bosch, *Exact calculation of the point spread function of an optical system: effect of the aperture stop in an intermediate space,* in *Proc. SPIE*, Vol. 5249 (2004).

[64] J. Ferré-Borrull and S. Bosch, *Formal description of diffraction in optical systems: Calculations and experimental evidence,* Appl. Phys. Lett. **85**, 2718 (2004).

[65] F. James, *Monte Carlo theory and practice,* Rept. Prog. Phys. **43**, 1145 (1980).

[66] A. Carpentier and R. Munos, *Adaptive stratified sampling for Monte-Carlo integration of differentiable functions,* in *Advances in Neural Information Processing Systems 25*, edited by F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger (Curran Associates, Inc., 2012) pp. 251–259.

[67] W. Jarosz, *Efficient Monte Carlo methods for light transport in scattering media*, Ph.D. thesis, University of California (2008).

[68] A. B. Owen, *Quasi-Monte Carlo sampling,* in *Monte Carlo Ray Tracing: Siggraph 2003 Course 44*, edited by H. W. Jensen (SIGGRAPH, 2003) pp. 69–88.

[69] C. J. R. Sheppard, *Validity of the Debye approximation,* Opt. Lett. **25**, 1660 (2000).

[70] O. Baladron-Zorita, H. Zhong, M. Kuhn, and F. Wyrowski, *A smooth field decomposition applied to modelling of scattering phenomena,* in *Proc. SPIE*, Vol. 9630 (2015).

[71] M. Mout, A. Flesch, M. Wick, F. Bociort, J. Petschulat, and P. Urbach, *Ray-based method for simulating cascaded diffraction in high-numerical-aperture systems,* J. Opt. Soc. Am. A **35**, 1356 (2018).

[72] B. Karczewski and E. Wolf, *Comparison of three theories of electromagnetic diffraction at an aperture. part I: Coherence matrices,* J. Opt. Soc. Am. **56**, 1207 (1966).

[73] J. D. Jackson, *Classical electrodynamics*, 3rd ed. (John Wiley & Sons, Inc., New York, NY, 1999).

[74] J. A. Stratton and L. Chu, *Diffraction theory of electromagnetic waves,* Phys. Rev. **56**, 99 (1939).

[75] B. Andreas, *Vectorial ray-based diffraction integral (VRBDI),* version 1.7, MathWorks® File Exchange, http://nl.mathworks.com/matlabcentral/fileexchange/52210-vectorial-ray-based-diffraction-integral–vrbdi- (2016), accessed 6 May, 2017.

[76] T. Tojyo, *Microscope objective lens system,* US4208099 (1980).

[77] W. J. Smith, *Modern lens design*, 1st ed. (McGraw-Hill, New York, 1992).

[78] A. V. Gitin, *The Feynman path integral and its optical applications,* Optik **125**, 3249 (2014).

[79] S. Cho, J. C. Petruccelli, and M. A. Alonso, *Wigner functions for paraxial and nonparaxial fields,* J. Mod. Optic. **56**, 1843 (2009).

[80] J. C. Petruccelli and M. A. Alonso, *Propagation of nonparaxial partially coherent fields across interfaces using generalized radiometry,* J. Opt. Soc. Am. A **26**, 2012 (2009).

6

[81] M. Pharr and G. Humphreys, *Physically based rendering*, 2nd ed. (Elsevier, 30 Corporate Drive, Suite 400, Burlington, MA 01803, USA, 2010).

A

LIST OF ASSUMPTIONS AND CONVENTIONS

Although some assumptions are specific to certain parts of this thesis, many assumptions are used for the entirety of this work (except in Section 2.2). By definition, the same holds for the conventions used in this thesis. As an aid to the reader, these assumptions and conventions, a list of physical and other quantities, as well as a list of abbreviations, are listed in this chapter.

A.1. LIST OF CONVENTIONS

Below is a list of conventions used in this thesis.

- The Sytème Internationale (SI) is used and accordingly the Maxwell equations are used in their SI format.

- All vectors are written in bold (e.g., $\boldsymbol{E}$).

- Normal vectors have a hat (e.g., $\hat{\boldsymbol{n}}$).

- Matrices are indicated by capital letters in blackboard bold (e.g., $\mathbb{A}$).

- The negative time-convention is used ($\exp(-\mathrm{i}\omega t)$). The physical field can be obtained by multiplying the complex quantity (e.g., $\boldsymbol{E}$) with this phase term and taking the real part.

- The integral over a surface is indicated by $\mathrm{d}S$, the integral over a volume by $\mathrm{d}V$.

- An integral over all components of a vector is indicated by the differential sign and the vector (e.g., $\mathrm{d}^2\boldsymbol{x} = \mathrm{d}x\mathrm{d}y$ if $\boldsymbol{x} = (x, y)$).

- The surface normal is given by $\hat{\boldsymbol{n}}$. For closed surfaces, it is the outwards pointing normal, for planes it points along the main direction of propagation (usually along the z-axis).

- The following is a list of matrix operators

    - Complex conjugate: $\mathbb{A}^*$
    - Transpose: $\mathbb{A}^T$
    - Inverse: $\mathbb{A}^{-1}$

- The expectation value is indicated by brackets (e.g., $< w(\boldsymbol{X}) >$).

A.2. LIST OF PHYSICAL QUANTITIES
The physical quantities used in this thesis are listed in this section.

Table A.1: Physical quantities

Notation	Physical quantity	Units
ϵ	Permittivity of medium	Farad per meter
ϵ_0	Permittivity of free space	Farad per meter
μ	Permeability of medium	Newtons per ampere squared
μ_0	Permeability of free space	Newtons per ampere squared
λ	Wavelength in the medium	Meters
ω	Frequency	Per second
$\boldsymbol{B}$	Magnetic flux density	Newtons per ampere per meter
$\boldsymbol{D}$	Electric displacement field	Seconds amperes per meter squared
$\boldsymbol{E}$ & E	Electric field	Volts per meter
$\tilde{F}$	Plane wave expansion	Volts meter
$\boldsymbol{H}$	Magnetic field	Amperes per meter
I	Intensity (Eq. (2.78))	Joules per second per metre squared
$\boldsymbol{J}$	Current density	Amperes per meter squared
$\boldsymbol{J_b}$	Bound current density	Amperes per meter squared
$\boldsymbol{J_f}$	Free current density	Amperes per meter squared
k	Wavenumber in medium	One over meter
k_0	Wavenumber in free space	One over meter
$\boldsymbol{M}$	Magnetization	Amperes per meter
n	Refractive index (Eq. (2.60))	Dimensionless
$\boldsymbol{P}$	Polarization density	Seconds amperes per meter squared
q	Charge density	Coulomb per meter cubed
q_b	Bound charge density	Coulomb per meter cubed
q_f	Free charge density	Coulomb per meter cubed
s	Eikonal	Meters
$\boldsymbol{S}$	Poynting's vector (Eq. (2.68))	Joules per second per metre squared
t	Time	Seconds
T_s & T_p	Fresnel coefficients (Eqs. (2.73, 2.74))	Dimensionless
T_s' & T_p'	Fresnel coefficients for ray tracing (Eqs. (2.84, 2.85))	Dimensionless
U	Energy	Joules

Table A.2: Other quantities

Notation	Quantity
$\delta(x)$	Delta function
θ	Angle, usually w.r.t. $\hat{\boldsymbol{n}}$
$\Pi(x)$	Rectangle function
$\boldsymbol{\rho}$	Vector between two points ($\boldsymbol{\rho_1} = \boldsymbol{r_1} - \boldsymbol{r_0}$)
$\hat{\boldsymbol{\rho}}$	Normalized $\boldsymbol{\rho}$ & ray directions (Eq. (2.61))
ρ	Vector length ($\rho = |\boldsymbol{\rho}|$)
$\sigma[X]$	Standard deviation of X
ϕ	Phase (in radians)
$\boldsymbol{\xi} = (\xi, \eta)$	Spatial frequencies
B	Mask function
$F_{\#}$	Fresnel number
$\boldsymbol{G}$ & G	Green's function
G_H	Green's function for the Helmholtz equation (Eq. 2.39)
i	Imaginary unit
$\boldsymbol{k} = (k_x, k_y)$ or $\boldsymbol{k} = (k_x, k_y, k_z)$	Spatial frequencies
N	Total number of rays
L_2	L_2 difference between the complex amplitudes of two scalar fields (Eq. (4.7))
L$_2$-dif	L_2 difference between the complex amplitudes of two vectorial fields (Eq. (5.21))
L_{2A}	L_2 difference between the absolute amplitudes of two scalar fields (Eq. (3.3))
L_{2I}	L_2 difference between the intensities of two scalar fields (Eq. (3.4))
$\hat{\boldsymbol{n}}$	Surface normal (outwards or forwards facing)
p	Probability density function of stochastic variable
$\boldsymbol{p} = (p_x, p_y)$ or $\boldsymbol{p} = (p_x, p_y, p_z)$	Direction cosines
$\Re$	Real part
$\boldsymbol{r} = (x, y, z)$	Cartesian coordinates
S	Surface
V	Volume
$V[X]$	Variance of X
$W(\boldsymbol{x}, \boldsymbol{\xi})$	Wigner distribution function
w	Ray weight (for Monte Carlo simulations)
X	Stochastic variable
$\boldsymbol{x} = (x, y)$	Cartesian coordinates without z

A.3. LIST OF ABBREVIATIONS

This section lists the abbreviations used in this thesis.

Table A.3: Abbreviations

Abbreviation	Meaning
BSP	Beam synthesis propagation
EPDI	Exit pupil diffraction integral
GBD	Gaussian beam decomposition
HFPI	Huygens-Fresnel path integration
NA	Numerical aperture
OPL	Optical path length
PDF	Probability density function
PSF	Point spread function
PW-HFPI	Plane wave Huygens-Fresnel path integration
rms	Root mean square
RSDI	Rayleigh-Sommerfeld diffraction integral
RSED	Rayleigh-Sommerfeld exit pupil diffraction integral
VRBDI	Vectorial ray-based diffraction integral
WBRT	Wigner-based ray tracing
WDF	Wigner distribution function

A.4. LIST OF ASSUMPTIONS

Below is a list of the assumptions used throughout this thesis (except in part of Section 2.2).

- All materials are

  - homogeneous

  - linear

  - source-free

  - non-magnetic ($\mu = \mu_0$)

  - non-absorbing

- Apertures are treated using the Kirchhoff boundary conditions (see Section 2.3.4).

- All propagation distances are much larger than the wavelength (the near field is neglected).

- All solutions are stationary.

- All sources are fully coherent (spatial and temporal).

B

OVERVIEW OF SIMULATION METHODS

The following table gives a very brief overview of the simulation methods treated in this thesis. The table also lists the main application of each method, along with a brief remark regarding the domain of validity. The aim of the table is to give a very compact overview and is therefore far from listing all the applications, approximations, and limitations of each method. For more information, the reader is referred to the listed sections.

Table B.1: Simulation methods

Method	Main application/ principle assumption	Inventor/Important contributor	See Section
Beam synthesis propagation (BSP)	Diffraction in optical systems	CODE V	2.6, 4.4
Bidirectional HFPI	Needs further research	This thesis	4.5.2
Collins integral	Paraxial propagation through lenses	S.A. Collins	2.3.2
Exit pupil diffraction integrals (EPDI)	Field in image space of optical system, diffraction assumed at exit pupil	R.K. Luneburg	2.5
Fresnel integral	Paraxial free-space propagation	A.J. Fresnel	2.3.2
Gaussian beam decomposition (GBD)	Diffraction in optical systems using a parabasal approximation	A.W. Greynolds	2.6, 4.4
Geometrical optics	Propagation through optical systems, valid if field varies slowly (except in propagation direction)		2.4
Huygens-Fresnel path integration (HFPI)	Scalar/Vectorial diffraction in optical systems	This thesis	4/5
m-theory diffraction integral	Vectorial free-space propagation from a plane	H. Severin	2.3.1
Monte Carlo ray tracing	Typically used for intensity calculations neglecting interference		2.4.3
Plane wave HFPI	Scalar/Vectorial diffraction in optical systems, valid for large Fresnel numbers and near the focal plane	This thesis	4.5.3/ 5.5
Ray tracing	Numerical implementation of geometrical optics		2.4.2
Rayleigh-Sommerfeld diffraction integral (RSDI)	Scalar free-space propagation from a plane (a vectorial formulation also exists)	J.W.S. Rayleigh, A.J.W. Sommerfeld	2.3.1
Rayleigh-Sommerfeld EPDI	EPDI using the RSDI, no additional assumptions		2.5
Richards-Wolf EPDI	EPDI using plane waves, valid for large Fresnel numbers and near the focal plane	E. Wolf	2.5, 4.5.3
Stepwise HFPI	Needs further research	This thesis	4.5.2
Vectorial ray-based diffraction integral	Diffraction in optical systems, all diffraction at input plane	B. Andreas	2.6
Wigner-based ray tracing (WBRT)	Scalar diffraction in aberration-free optical systems with moderate NAs	This thesis	3

C

ADDITIONAL PROOFS REGARDING THE WIGNER DISTRIBUTION FUNCTION

This appendix gives the derivations for some of the properties of the WDF used in this thesis. More specifically we will treat the inverse WDF (Section C.1), the WDF of the field after a thin element (Section C.2), and the WDF of a field propagated by the Collins integral (Section C.3).

C.1. INVERSE WIGNER DISTRIBUTION FUNCTION

Because the WDF contains all the information of the original electric field, it is possible to perform a back transformation using

$$E(\boldsymbol{x})\, E^{*}\,(\bar{\boldsymbol{x}}) = \frac{1}{(2\pi)^2} \iint W\left(\frac{\boldsymbol{x}+\bar{\boldsymbol{x}}}{2}, \boldsymbol{\xi}\right) \exp\left[i(\boldsymbol{x}-\bar{\boldsymbol{x}})\cdot\boldsymbol{\xi}\right] \mathrm{d}^2\boldsymbol{\xi}. \tag{C.1}$$

One can obtain the original field, except for an overall phase factor ϕ, by

$$E(\boldsymbol{x}) \exp\left[\mathrm{i}\phi\right] = \frac{E(\boldsymbol{x})E^{*}(\boldsymbol{0})}{\sqrt{E(\boldsymbol{0})E^{*}(\boldsymbol{0})}} \tag{C.2}$$

$$= \frac{\iint W\left(\frac{\boldsymbol{x}}{2}, \boldsymbol{\xi}\right) \exp\left[\mathrm{i}\boldsymbol{x}\cdot\boldsymbol{\xi}\right] \mathrm{d}^2\boldsymbol{\xi}}{\sqrt{\iint W(\boldsymbol{0}, \boldsymbol{\xi}) \mathrm{d}^2\boldsymbol{\xi}}}. \tag{C.3}$$

The proof of Eq. (C.1), and thereby of the fact that the frequency marginal gives the intensity (Eq. (2.112)), is obtained using the difference variable $\boldsymbol{x'} = \boldsymbol{x} - \bar{\boldsymbol{x}}$ and Eq. (2.111)

for the dirac function. Starting with the left hand side of Eq. (C.1) we obtain

$$E(\boldsymbol{x})E^*(\bar{\boldsymbol{x}}) = E\left(\frac{2\boldsymbol{x}+\bar{\boldsymbol{x}}-\bar{\boldsymbol{x}}}{2}\right)E^*\left(\frac{2\bar{\boldsymbol{x}}+\boldsymbol{x}-\boldsymbol{x}}{2}\right) \tag{C.4}$$

$$= \iint E\left(\frac{\boldsymbol{x}+\bar{\boldsymbol{x}}+\boldsymbol{x}'}{2}\right)E^*\left(\frac{\boldsymbol{x}+\bar{\boldsymbol{x}}-\boldsymbol{x}'}{2}\right)$$
$$\delta(\boldsymbol{x}'-\boldsymbol{x}+\bar{\boldsymbol{x}})\mathrm{d}^2\boldsymbol{x}' \tag{C.5}$$

$$= \iint\iint E\left(\frac{\boldsymbol{x}+\bar{\boldsymbol{x}}+\boldsymbol{x}'}{2}\right)E^*\left(\frac{\boldsymbol{x}+\bar{\boldsymbol{x}}-\boldsymbol{x}'}{2})\right)$$
$$\frac{1}{(2\pi)^2}\exp\left[-i\boldsymbol{\xi}\cdot(\boldsymbol{x}'-\boldsymbol{x}+\bar{\boldsymbol{x}})\right]\mathrm{d}^2\boldsymbol{\xi}\mathrm{d}^2\boldsymbol{x}' \tag{C.6}$$

$$= \frac{1}{(2\pi)^2}\iint W\left(\frac{\boldsymbol{x}+\bar{\boldsymbol{x}}}{2},\boldsymbol{\xi}\right)\exp\left[i\boldsymbol{\xi}\cdot(\boldsymbol{x}-\bar{\boldsymbol{x}})\right]\mathrm{d}^2\boldsymbol{\xi}, \tag{C.7}$$

which concludes the proof.

C.2. WIGNER DISTRIBUTION FUNCTION OF A THIN ELEMENT

In Section 2.7 we stated that the WDF of the field after a thin element is given by

$$W(\boldsymbol{x},\boldsymbol{\xi}) = \frac{1}{(2\pi)^2}\iint W_0\left(\boldsymbol{x},\boldsymbol{\xi}'\right)W_B\left(\boldsymbol{x},\boldsymbol{\xi}-\boldsymbol{\xi}'\right)\mathrm{d}^2\boldsymbol{\xi}', \tag{C.8}$$

if the field is described by

$$E(\boldsymbol{x}) = B(\boldsymbol{x})E_0(\boldsymbol{x}). \tag{C.9}$$

In this section we will provide a proof for this statement.

The proof starts by writing out the explicit WDF of Eq. (C.9) and then uses the inverse transform given by Eq. (C.1):

$$W_1(\boldsymbol{x},\boldsymbol{\xi}) = \iint B\left(\boldsymbol{x}+\frac{\boldsymbol{x}'}{2}\right)E_0\left(\boldsymbol{x}+\frac{\boldsymbol{x}'}{2}\right)B^*\left(\boldsymbol{x}-\frac{\boldsymbol{x}'}{2}\right)E^*\left(\boldsymbol{x}-\frac{\boldsymbol{x}'}{2}\right)\exp\left[-i\boldsymbol{\xi}\cdot\boldsymbol{x}'\right]\mathrm{d}^2\boldsymbol{x}'$$

$$= \frac{1}{(2\pi)^4}\iint\iint W_0\left(\boldsymbol{x},\boldsymbol{\xi}'\right)\exp[i\boldsymbol{x}'\boldsymbol{\xi}']\mathrm{d}^2\boldsymbol{\xi}'\iint W_B\left(\boldsymbol{x},\bar{\boldsymbol{\xi}}\right)\exp[i\boldsymbol{x}'\bar{\boldsymbol{\xi}}]\mathrm{d}^2\bar{\boldsymbol{\xi}}\exp[-i\boldsymbol{\xi}\boldsymbol{x}']d\boldsymbol{x}'$$

$$= \frac{1}{(2\pi)^4}\iint\iint\iint W_0\left(\boldsymbol{x},\boldsymbol{\xi}'\right)W_B\left(\boldsymbol{x},\bar{\boldsymbol{\xi}}\right)\exp\left[i\boldsymbol{x}'\cdot(\boldsymbol{\xi}'+\bar{\boldsymbol{\xi}}-\boldsymbol{\xi})\right]d\boldsymbol{\xi}'d\bar{\boldsymbol{\xi}}d\boldsymbol{x}'$$

$$= \frac{1}{(2\pi)^2}\iint\iint W_0(\boldsymbol{x},\boldsymbol{\xi}')W_B(\boldsymbol{x},\bar{\boldsymbol{\xi}})\delta(\boldsymbol{\xi}'+\bar{\boldsymbol{\xi}}-\boldsymbol{\xi})d\boldsymbol{\xi}'d\bar{\boldsymbol{\xi}}$$

$$= \frac{1}{(2\pi)^2}\iint W_0(\boldsymbol{x},\boldsymbol{\xi}')W_B(\boldsymbol{x},\boldsymbol{\xi}-\boldsymbol{\xi}')d\boldsymbol{\xi}'$$

where W_B is the WDF of B and W_0 the WDF of E_0.

C.3. WIGNER DISTRIBUTION FUNCTION OF THE COLLINS IN-TEGRAL

In Section 2.7 it was stated that the WDF of a field after propagation by the Collins integral

$$E_1(\boldsymbol{x}) = \frac{-\mathrm{i}k}{2\pi\sqrt{|\mathbb{B}|}} \iint E_0(\boldsymbol{x_0}) \exp\left[\mathrm{i}k\left(\frac{1}{2}\boldsymbol{x_0}^T\mathbb{B}^{-1}\mathbb{A}\boldsymbol{x_0} - \boldsymbol{x_0}^T\mathbb{B}^{-1}\boldsymbol{x} + \frac{1}{2}\boldsymbol{x}^T\mathbb{D}\mathbb{B}^{-1}\boldsymbol{x}\right)\right]\mathrm{d}^2\boldsymbol{x_0}, \quad \text{(C.10)}$$

is given by

$$W(\boldsymbol{x},\boldsymbol{\xi}) = W_0\left[\mathbb{D}^T\boldsymbol{x} - \frac{1}{k}\mathbb{B}^T\boldsymbol{\xi}, k\left(\mathbb{B}^{-1} - \mathbb{A}^T\mathbb{D}\mathbb{B}^{-1}\right)\boldsymbol{x} + \mathbb{A}^T\boldsymbol{\xi}\right], \quad \text{(C.11)}$$

with W_0 the WDF of E_0. In this section we will provide a proof for this property.

Before starting the derivation, we must introduce two properties of the matrices of the Collins integral that will be used frequently in this section

$$\left(\mathbb{B}^{-1}\mathbb{A}\right)^T = \mathbb{B}^{-1}\mathbb{A}, \quad \text{(C.12)}$$

$$\left(\mathbb{D}\mathbb{B}^{-1}\right)^T = \mathbb{D}\mathbb{B}^{-1}, \quad \text{(C.13)}$$

which are a result of the Collins integral being a second order approximation. The first step in the actual derivation is calculating the WDF of Eq. (C.10)

$$W(\boldsymbol{x},\boldsymbol{\xi}) = \iint E_1\left(\boldsymbol{x} + \frac{\bar{\boldsymbol{x}}'}{2}\right) E_1^*\left(\boldsymbol{x} - \frac{\bar{\boldsymbol{x}}'}{2}\right) e^{-\mathrm{i}\boldsymbol{\xi}^T\bar{\boldsymbol{x}}'}\mathrm{d}^2\bar{\boldsymbol{x}}'$$

$$= \frac{k^2}{(2\pi)^2|\mathbb{B}|} \iint \iint \iint E_0(\boldsymbol{x_0'}) \exp\left[\mathrm{i}k\left(\frac{1}{2}\boldsymbol{x_0'}^T\mathbb{B}^{-1}\mathbb{A}\boldsymbol{x_0'}\right.\right.$$

$$\left.\left. -\boldsymbol{x_0'}^T\mathbb{B}^{-1}\left(\boldsymbol{x} + \frac{\bar{\boldsymbol{x}}'}{2}\right) + \frac{1}{2}\left(\boldsymbol{x} + \frac{\bar{\boldsymbol{x}}'}{2}\right)^T\mathbb{D}\mathbb{B}^{-1}\left(\boldsymbol{x} + \frac{\bar{\boldsymbol{x}}'}{2}\right)\right)\right]$$

$$E_0^*(\bar{\boldsymbol{x}}_0) \exp\left[-\mathrm{i}k\left(\frac{1}{2}\bar{\boldsymbol{x}}_0^T\mathbb{B}^{-1}\mathbb{A}\bar{\boldsymbol{x}}_0 - \bar{\boldsymbol{x}}_0^T\mathbb{B}^{-1}\left(\boldsymbol{x} - \frac{\bar{\boldsymbol{x}}'}{2}\right)\right.\right.$$

$$\left.\left. +\frac{1}{2}\left(\boldsymbol{x} - \frac{\bar{\boldsymbol{x}}'}{2}\right)^T\mathbb{D}\mathbb{B}^{-1}\left(\boldsymbol{x} - \frac{\bar{\boldsymbol{x}}'}{2}\right)\right)\right]\exp\left[-\mathrm{i}\boldsymbol{\xi}^T\bar{\boldsymbol{x}}'\right]\mathrm{d}^2\boldsymbol{x_0'}\mathrm{d}^2\bar{\boldsymbol{x}}_0\mathrm{d}^2\bar{\boldsymbol{x}}'$$

$$= \frac{k^2}{(2\pi)^2|\mathbb{B}|} \iint \iint \iint E_0(\boldsymbol{x_0'})E_0^*(\bar{\boldsymbol{x}}_0)$$

$$\exp\left[\mathrm{i}k\left(\frac{1}{2}\boldsymbol{x_0'}^T\mathbb{B}^{-1}\mathbb{A}\boldsymbol{x_0'} - \frac{1}{2}\bar{\boldsymbol{x}}_0^T\mathbb{B}^{-1}\mathbb{A}\bar{\boldsymbol{x}}_0 + \left(\bar{\boldsymbol{x}}_0 - \boldsymbol{x_0'}\right)^T\mathbb{B}^{-1}\boldsymbol{x}\right)\right]$$

$$\exp\left[\mathrm{i}k\left(-\frac{1}{2}\left(\boldsymbol{x_0'} + \bar{\boldsymbol{x}}_0\right)^T\mathbb{B}^{-1} + \boldsymbol{x}^T\mathbb{D}\mathbb{B}^{-1} - \frac{1}{k}\boldsymbol{\xi}^T\right)\bar{\boldsymbol{x}}'\right]\mathrm{d}^2\boldsymbol{x_0'}\mathrm{d}^2\bar{\boldsymbol{x}}_0\mathrm{d}^2\bar{\boldsymbol{x}}'.$$

Using the following coordinate transformations

$$\boldsymbol{x_0'} = \boldsymbol{x_0} + \frac{\boldsymbol{x}'}{2},$$

$$\bar{\boldsymbol{x}}_0 = \boldsymbol{x_0} - \frac{\boldsymbol{x}'}{2},$$

and using Eq. (2.111) for the delta function, we obtain

$$
\begin{aligned}
W(\boldsymbol{x}, \boldsymbol{\xi}) = {} & \frac{k^2}{(2\pi)^2 |\mathbb{B}|} \iint \iint E_0\left(\boldsymbol{x_0} + \frac{\boldsymbol{x'}}{2}\right) E_0^*\left(\boldsymbol{x_0} - \frac{\boldsymbol{x'}}{2}\right) \\
& \exp\left[ik\left(\frac{1}{2}\left(\boldsymbol{x_0} + \frac{\boldsymbol{x'}}{2}\right)^T \mathbb{B}^{-1}\mathbb{A}\left(\boldsymbol{x_0} + \frac{\boldsymbol{x'}}{2}\right) - \frac{1}{2}\left(\boldsymbol{x_0} - \frac{\boldsymbol{x'}}{2}\right)^T \mathbb{B}^{-1}\mathbb{A}\left(\boldsymbol{x_0} - \frac{\boldsymbol{x'}}{2}\right) - \boldsymbol{x'}^T \mathbb{B}^{-1}\boldsymbol{x}\right)\right] \\
& \delta\left[\frac{k}{2\pi}\left(-\boldsymbol{x_0}^T \mathbb{B}^{-1} + \boldsymbol{x}^T \mathbb{D}\mathbb{B}^{-1} - \frac{1}{k}\boldsymbol{\xi}^T\right)\right] \mathrm{d}^2 \boldsymbol{x_0}' \mathrm{d}^2 \bar{\boldsymbol{x}}_0 \\
= {} & \frac{k^2}{(2\pi)^2 |\mathbb{B}|} \iint \iint E_0\left(\boldsymbol{x_0} + \frac{\boldsymbol{x'}}{2}\right) E_0^*\left(\boldsymbol{x_0} - \frac{\boldsymbol{x'}}{2}\right) \\
& \exp\left[ik\left(\frac{1}{2}\boldsymbol{x'}^T \mathbb{B}^{-1}\mathbb{A}\boldsymbol{x_0} + \frac{1}{2}\boldsymbol{x_0}^T \mathbb{B}^{-1}\mathbb{A}\boldsymbol{x'} - \boldsymbol{x'}^T \mathbb{B}^{-1}\boldsymbol{x}\right)\right] \\
& \frac{(2\pi)^2}{k^2}|\mathbb{B}|\delta\left[-\boldsymbol{x_0}^T + \boldsymbol{x}^T \mathbb{D} - \frac{1}{k}\boldsymbol{\xi}^T\mathbb{B}\right] \mathrm{d}^2 \boldsymbol{x_0}' \mathrm{d}^2 \bar{\boldsymbol{x}}_0 \\
= {} & \iint E_0\left(-\frac{1}{k}\mathbb{B}^T\boldsymbol{\xi} + \mathbb{D}^T\boldsymbol{x} + \frac{\boldsymbol{x'}}{2}\right) E_0^*\left(-\frac{1}{k}\mathbb{B}^T\boldsymbol{\xi} + \mathbb{D}^T\boldsymbol{x} - \frac{\boldsymbol{x'}}{2}\right) \\
& \exp\left[ik\left(\left(-\frac{1}{k}\mathbb{B}^T\boldsymbol{\xi} + \mathbb{D}^T\boldsymbol{x}\right)^T \mathbb{B}^{-1}\mathbb{A}\boldsymbol{x'} - \boldsymbol{x'}^T \mathbb{B}^{-1}\boldsymbol{x}\right)\right] \mathrm{d}^2 \boldsymbol{x'} \\
= {} & \iint E_0\left(-\frac{1}{k}\mathbb{B}^T\boldsymbol{\xi} + \mathbb{D}^T\boldsymbol{x} + \frac{\boldsymbol{x'}}{2}\right) E_0^*\left(-\frac{1}{k}\mathbb{B}^T\boldsymbol{\xi} + \mathbb{D}^T\boldsymbol{x} - \frac{\boldsymbol{x'}}{2}\right) \\
& \exp\left[ik\left(-\frac{1}{k}\boldsymbol{\xi}^T\mathbb{A}\boldsymbol{x'} + \boldsymbol{x}^T \mathbb{D}\mathbb{B}^{-1}\mathbb{A}\boldsymbol{x'} - \boldsymbol{x}^T \mathbb{B}^{-T}\boldsymbol{x'}\right)\right] \mathrm{d}^2 \boldsymbol{x'} \\
= {} & \iint E_0\left(-\frac{1}{k}\mathbb{B}^T\boldsymbol{\xi} + \mathbb{D}^T\boldsymbol{x} + \frac{\boldsymbol{x'}}{2}\right) E_0^*\left(-\frac{1}{k}\mathbb{B}^T\boldsymbol{\xi} + \mathbb{D}^T\boldsymbol{x} - \frac{\boldsymbol{x'}}{2}\right) \\
& \exp\left[i\left(k\left(\mathbb{A}^T\mathbb{D}\mathbb{B}^{-1} - \mathbb{B}^{-1}\right)\boldsymbol{x} - \mathbb{A}^T\boldsymbol{\xi}\right)^T \boldsymbol{x'}\right] \mathrm{d}^2 \boldsymbol{x'} .
\end{aligned}
$$

Comparing this result to the definition of the Wigner function

$$
W(\boldsymbol{x}, \boldsymbol{\xi}) = \iint E_o\left(\boldsymbol{x_0} + \frac{\boldsymbol{x'}}{2}\right) E_o^*\left(\boldsymbol{x_0} - \frac{\boldsymbol{x'}}{2}\right) e^{-i\boldsymbol{\xi}^T\boldsymbol{x'}} \mathrm{d}^2 \boldsymbol{x'}, \tag{C.14}
$$

one can rewrite it into

$$
W(\boldsymbol{x}, \boldsymbol{\xi}) = W_0\left[\mathbb{D}^T\boldsymbol{x} - \frac{1}{k}\mathbb{B}^T\boldsymbol{\xi}, k\left(\mathbb{B}^{-1} - \mathbb{A}^T\mathbb{D}\mathbb{B}^{-1}\right)\boldsymbol{x} + \mathbb{A}^T\boldsymbol{\xi}\right]. \tag{C.15}
$$

The Wigner distribution function in the output plane can thus be obtained by a coordinate transformation of the Wigner function in the input plane. In matrix form, this transformation is given by

$$
\begin{pmatrix} \boldsymbol{x_0} \\ \boldsymbol{\xi} \end{pmatrix} = \begin{bmatrix} \mathbb{D}^T & -\frac{1}{k}\mathbb{B}^T \\ k(\mathbb{B}^{-1} - \mathbb{A}^T\mathbb{D}\mathbb{B}^{-1}) & \mathbb{A}^T \end{bmatrix} \begin{pmatrix} \boldsymbol{x} \\ \boldsymbol{\xi} \end{pmatrix}, \tag{C.16}
$$

or in the other direction

$$
\begin{pmatrix} \boldsymbol{x} \\ \boldsymbol{\xi} \end{pmatrix} = \begin{bmatrix} \mathbb{A} & \frac{1}{k}\mathbb{B} \\ -k(\mathbb{B}^{-1} - \mathbb{A}^T\mathbb{D}\mathbb{B}^{-1})^T & \mathbb{D} \end{bmatrix} \begin{pmatrix} \boldsymbol{x_0} \\ \boldsymbol{\xi} \end{pmatrix}. \tag{C.17}
$$

The final equation is equivalent to Eq. (2.52) when

$$p_0 = \frac{\xi_0}{k} \tag{C.18a}$$

$$p = \frac{\xi}{k}. \tag{C.18b}$$

This relates the directional cosines to the spatial frequencies. Within the paraxial approximation these relation correspond to the directional cosines for plane waves with the given spatial frequencies.

Although the derivation given here breaks down if $\mathbb{B}$ is not invertible (e.g., if the matrix describes imaging), one can show that the results still hold by taking the limit of $\mathbb{B}$ going to zero [18].

ACKNOWLEDGEMENTS

The ancient Greek had the notion of walking backwards into the future. I would say that doing a PhD research is like walking backwards, first through terrain still known by others, then a little step out into the unknown. Although I had to do the walking by myself, I had many people at my side along the way. Some gave me the equipment and skill needed for the journey, some gave an overview of the 'optics landscape' that lay behind me but in front of them, others encouraged me, helped me back on my feet or suggested new directions to continue the journey into the unknown. This thesis would not have been possible without their contributions, for which I would like to thank them.

First and foremost, I would like to thank Michael Wick who proposed both methods investigated during this project. He taught me how to test quick and dirty, which ended our research into Wigner-Based ray tracing soon after I came to Oberkochen. Thanks to Michael, always full of ideas and energy, this did not cause a long crisis in my PhD, but was the start of our fascinating journey into Huygens-Fresnel path integration (HFPI). My next supervisor at Zeiss, Andreas Flesch played an important role in bringing the research to an end. He was a pleasant and effective project manager who helped me focus on the essential parts. He also helped implementing HFPI with his knowledge of the in-house ray-tracing software. I would like to thank my final supervisor, Florian Bociort, for helping me as a fresh PhD student to start my research on a completely new topic. Florian was a true advisor, driving the project with suggestions and questions but giving me the freedom to choose the direction of the path forwards. Paul Urbach was, what the Germans call my 'Doktorvater'. His love and demand for analytical solutions brought us to new insights and his critical reading (regarding both content and form) raised the level of the presentation of our results. One cannot mention Paul, without also thanking Yvonne van Aalst for her organizational help.

I would like to thank the Zeiss CRT colleagues in general for always being willing to help, explain, and discuss. In particular I would like to thank Toufic Jabour for the many discussions on optics, Steffen Siegler for all he taught me about ray tracing and diffraction integrals, and Johannes-Maria Kaltenbach for sharing his extensive knowledge in person and through the many documents he wrote. With David Abreu I pleasantly collaborated on three-dimensional image simulations and had many discussions on varied topics. A great thanks goes out to George Firanescu and in particular Hans-Joachim Frasch for implementing HFPI in the Zeiss in-house simulation software. We occasionally were a bit late with our requests, but they were always instantly implemented nonetheless. This would not have been possible without the support of Arnim Henze who also helped to iron out the final errors by proofreading the thesis. Wolfgang Singer gave me a good head start when I came to Oberkochen through his support and openness for new ideas. This project would never have started without the work of Bernd Kleemann who initiated it together with Paul Urbach. Once the project had started, Jörg Petschulat provided me with a work environment that allowed me to completely focus

on the PhD. On top of that, his critical reading improved many manuscripts.

I am thankful to the entire optics group of the TU Delft for creating such a pleasant atmosphere and fun place to work. The team of professors was approachable and supportive, and allowed me to truly focus on my thesis. The technical staff always helped out quickly and with my fellow (PhD) students we had a lot of fun and some good collaborations. A particular thanks goes out to Zhe Hou, my companion in optical simulations and ray tracing.

This project was paid for by Zeiss and the European Union FP7 people program through the ADOPSYS project. I am grateful for their support and would like to acknowledge that without the EU an international project like this, for which I moved across borders twice, would never have been this uncomplicated. Zeiss provided me with a stimulating work environment and all the support I needed to do this research. I believe this willingness to invest in an international project with an uncertain and public outcome speaks for the company culture. The ADOPSYS partners have provided a very interesting training program, were wonderful hosts at all the events, and provided for many interesting discussions. The ADOPSYS students formed a very supportive and entertaining group.

I sincerely thank Joran Kuijper for designing the cover of this thesis. As a friend he understood what I wanted, as a designer he exceeded what I had envisioned.

Finally I would like to thank the committee members for reviewing this thesis, providing me with feedback that improved the quality of the work. I also appreciate that they found the time to attend the defense ceremony, which for some involved lengthy journeys.

In the above I have restricted myself to the main people and organizations that contributed technically or professionally to this thesis and my professional education. There are also a great many people who supported me personally, gave me the joy outside of working hours needed to get through the tougher parts of the research, and about whom I care much more than about this thesis. I hope you know that and do not need to be listed here.

ABOUT THE AUTHOR

Beat Marco MOUT

Marco Mout was born in January of 1987 in Hengelo the Netherlands. He soon moved to Soest, where he finished high school (gymnasium) in 2005 at the top of his class. He then spend a gap year in West-Africa, teaching at a rural high school in Ghana for six months before returning to Europe overland. Between 2006 and 2009 he studied mechanical engineering at Delft University of Technology (TUD), during which he spend one semester in Paris at École normale supérieure de Cachan. After graduating cum laude at the top of his class, he did a one year pre-master program to enter the master in applied physics during which he took several philosophy classes at Leiden University. During the academic year of 2010-2011 he was a full-time board member of the Interstedelijk Studenten Overleg (ISO). This group represents the interests of the students in the Dutch higher education system through contacts with members of parliament, the minister and ministry of Education, Culture and Research, and the national press. In 2011 Marco started a master in applied physics at the TUD and a one year job as junior project leader at the Platform Bèta Techniek. In 2013 he finished his master thesis at TNO, where he investigated the potential of adaptive optics for compensating alignment errors in optical systems. This sparked his interest in optics, which from early 2014 he pursued as a doctoral student in the optics group of the TUD. Marco was employed by Carl Zeiss AG for his PhD studies. He did most of his research at the Corporate Research and Technology department of Zeiss in Oberkochen (Germany). The fruit of that labour is the document at hand. Marco currently works as a research scientist at the Corporate Research and Technology department of Zeiss.

LIST OF PUBLICATIONS

6. **Marco Mout, Andreas Flesch, Michael Wick, Florian Bociort, Joerg Petschulat, Paul Urbach**, *Ray-based method for simulating cascaded diffraction in high-numerical-aperture systems*, Journal of the Optical Society of America A **35**, 1356–1367 (2018).

5. **Marco Mout, Michael Wick, Florian Bociort, Joerg Petschulat, Paul Urbach**, *Ray tracing the Wigner distribution function for optical simulations*, Optical engineering **57**, 014106 (2018).

4. **B.M. Mout, A. Flesch, M. Wick, F. Bociort, J. Petschulat, H.P. Urbach**, *Comparison of three ray based diffraction simulation methods*, European Optical Society Annual Meeting, (2016).

3. **Marco Mout, Michael Wick, Florian Bociort, Joerg Petschulat, Paul Urbach**, *Simulating multiple diffraction in imaging systems using a path integration method*, Applied Optics **55**(14), 3847-3853 (2016).

2. **B.M. Mout, M. Wick, J. Petschulat, F. Bociort, H.P. Urbach**, *A ray based method for diffraction simulations*, Optics-photonics Design and Fabrication, (2016).

1. **B.M. Mout, M. Wick, F. Bociort, H.P. Urbach**, *A Wigner based ray tracing method for imaging simulations*, SPIE Optical Design and Engineering VI **9626**, (2015).