



Delft University of Technology

The Structure of Two-Dimensional Light Its Singular Behaviour

van Gogh, M.A.

DOI

[10.4233/uuid:eed7439a-3533-4fe3-aa00-5b1f39eb8e26](https://doi.org/10.4233/uuid:eed7439a-3533-4fe3-aa00-5b1f39eb8e26)

Publication date

2025

Document Version

Final published version

Citation (APA)

van Gogh, M. A. (2025). *The Structure of Two-Dimensional Light: Its Singular Behaviour*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:eed7439a-3533-4fe3-aa00-5b1f39eb8e26>

Important note

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

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

THE STRUCTURE OF TWO-DIMENSIONAL LIGHT

ITS SINGULAR BEHAVIOUR

THE STRUCTURE OF TWO-DIMENSIONAL LIGHT

ITS SINGULAR BEHAVIOUR

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus prof. dr. ir. T.H.J.J. van der Hagen,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op donderdag 9 januari 2025 om 15:00 uur

door

Matthijs Alexander VAN GOGH

Master of Science in de Natuurkunde,
Universiteit Utrecht, Nederland,
geboren te Utrecht, Nederland.

Dit proefschrift is goedgekeurd door de

promotor: prof. dr. L. Kuipers
promotor: prof. dr. Y.M. Blanter

Samenstelling promotiecommissie:

Rector Magnificus,	voorzitter
Prof. dr. L. Kuipers,	Technische Universiteit Delft
Prof. dr. Y.M. Blanter,	Technische Universiteit Delft

Onafhankelijke leden:

Prof. dr. S.M. Witte,	Technische Universiteit Delft
Prof. dr. ir. H.L. Offerhaus,	Universiteit Twente
Dr. W. Löffler,	Universiteit Leiden
Dr. T. Idema,	Technische Universiteit Delft
Dr. A.J.L. Adam,	Technische Universiteit Delft
Prof. dr. A.F. Otte	Technische Universiteit Delft (reservelid)



Keywords: Optics, light, phase, polarization, chaos, singularities, phase transitions, near-field microscopy, diffusion

Printed by: Gildeprint - Enschede

Front & Back: Streamlines of the in-plane Poynting vector of a 2D random light field.

Copyright © 2025 by M.A. van Gogh

ISBN 978-94-6384-706-3

An electronic version of this dissertation is available at
<http://repository.tudelft.nl/>.

*To my friends and family
To those who believed in me*

CONTENTS

Summary	xi
Samenvatting	xiii
1 Introduction	1
1.1 About light	2
1.2 Electromagnetic Fields	3
1.2.1 Maxwell's equations and the wave nature of light	3
1.2.2 Plane waves and transverse modes.	4
1.2.3 Guiding light.	5
1.3 Singularities.	8
1.3.1 Phase singularities	9
1.3.2 Polarization and polarization singularities	11
1.3.3 Singularities as particles	13
1.4 Layout of this thesis.	15
2 Experimental setup	17
2.1 Beyond the diffraction limit.	18
2.2 Near-field microscopy	19
2.2.1 Capabilities of the near-field microscope	21
2.3 Implementation of the near-field microscope.	22
2.3.1 Heterodyne detection & gain.	23
2.4 Generation of 2D light	24
2.4.1 Photonic crystals.	24
2.4.2 A playground for random waves	25
2.4.3 Time-resolved measurements	26
2.5 Determining the polarization state	29
2.5.1 Fourier filtering	31
2.6 Determining singularity positions	32

3	Diffusion and lifetime of phase singularities	35
3.1	Introduction	36
3.2	Tracking of singularities	38
3.3	Diffusion of phase singularities	39
3.4	Lifetime of phase singularities	42
3.5	Faithful and unfaithful singularities	43
3.5.1	Diffusion	44
3.5.2	Lifetime	46
3.6	Anomalous diffusion	47
3.6.1	Fractional Brownian motion	48
3.6.2	Auto correlation of trajectory	50
3.6.3	Turning angles	52
3.7	Conclusions	53
3.8	Supplementary information	55
4	Melting of Singularity lattices	57
4.1	Introduction	58
4.2	Phase transitions and order parameters	59
4.3	Melting indicators	60
4.3.1	Indicators based on $g(r)$	60
4.3.2	Indicators based on the angular distribution function	62
4.3.3	Melting indicators for increasing amount of plane wave sources	64
4.4	Broadening in momentum	66
4.5	Angular broadening	70
4.5.1	Convergence as function of sources	75
4.6	Conclusions	77
4.7	Supplementary information	78
4.7.1	Entropy and excess entropy	78
5	Flow field singularities	81
5.1	The Poynting vector and its singularities	82
5.2	Origins of the singularities	83
5.3	Restrictions on the types of critical points	85
5.4	Relation between singularity type and charge	87

5.5	Experimental measurements of the Poynting field	90
5.5.1	Varying the in-plane wavenumber	94
5.6	Conclusions.	95
6	Measuring the optical near field in dual colour	97
6.1	Introduction	98
6.2	Modifications to the near-field microscope	100
6.3	Simultaneous observation of two colours of light in the near field	103
6.4	Separation of the in-plane electric field components	107
6.4.1	Limitations of the separating the field components	110
6.5	Technical analysis.	111
6.5.1	Comparison of Lock-in amplifiers	111
6.5.2	Signal-to-noise ratio & crosstalk	112
6.6	Outlook.	115
6.7	Conclusion	117
7	Conclusions	119
A	Supplementary information to Chapter 4	123
A.1	Stochastic broadening	123
A.1.1	Convergence and dependence on N_d	124
A.1.2	Consistency with angular broadening	127
A.2	Convergence of the melting indicators as function of amount of sources for angular broadening	128
	Bibliography	131
	Quotes Bibliography	145
	Curriculum Vitæ	147
	Acknowledgements	149

SUMMARY

In this thesis we have investigated properties of light, with the aim of gaining a deeper understanding of its behaviour at a fundamental level when confined to two dimensions. Interestingly, the main focus of this thesis is looking at occasions where there actually is no light at all. These points of perfect darkness possess some fascinating properties, such as having a so-called topological charge associated with them, as well as being singularities of the field. Singularities are points where a certain way of describing a field, such as light, no longer works properly. For most of this thesis, the parameter that breaks down is phase: when there is no light, the determination of the phase at that point ceases to work, and hence we are left with a phase singularity.

In order to observe these singularities we use a special microscope. The singularities, in two dimensions, are points: they are infinitesimally small, and in order to observe them we need to be able to look at the light field at very small length scales. The problem with using a regular microscope is that it is fundamentally limited in its resolution: regardless of how strong a magnification lens is used, a regular microscope can only tell two objects apart if they are separated further than a certain limit, which is known as the diffraction limit. And the behaviour of these singularities occurs at a scale much smaller than this diffraction limit. As such, the points of darkness would be obscured by the light surrounding them. Furthermore, a regular microscope observes the light that propagates towards the objective, while we are interested in the light field and its features when it is confined to a two-dimensional plane in a sample instead. As such, we require a special microscope in order to observe these entities. This specific microscope is called a near-field scanning optical microscope (NSOM), and instead of using lenses to collect light, it relies on a tiny needle-like probe that we bring extremely close to the surface of the sample. This distance is typically on the order of 20 nanometers, less than a thousandth of the thickness of a human hair. The probe is able to collect light in the field that is contained in the tiny region near the surface, which is called the near field, and contains all the information on the light inside the sample. Since this technique does not rely on lenses and conventional optics, it is not limited by the diffraction limit, and allows us to observe the wondrous world of light at extremely small scales. Details on this technique are found in Chapter 2.

When looking at how the singularities move around, their behaviour appears similar to that of regular particles. So for the first part of this investigation (Chapters 3 and 4) we view these phase singularities as if they are regular particles. Through this lens we can make use of all the regular tools of statistical physics and see if we can analyse their positions and movements in that way. In Chapter 3 we study their diffusive behaviour, where we experimentally track their positions in time. Through this we can observe their movement, which turns out to be in a non linear fashion. Additionally, since we

can track them from their moment of creation to their moment of "death", we are able to gather statistics on how long a singularity lives for on average. Furthermore, we show that splitting the singularities up into two different families, based on how they die, has a big impact on both their lifetime and diffusion. Chapter 4, while following the same assumption as Chapter 3, is fully numerical instead. By letting three plane waves interfere with each other in two dimensions, both the amplitude of the resulting field as well as the phase singularities form an orderly lattice, just like atoms do when the material is in a solid state, like ice for instance. We investigate if it is possible to have this solid-like state "melt" into a liquid-like state, like when ice melts and turns into water. Here we show a method for slowly introducing more waves in a symmetric fashion that result in a transition from a solid-like state to a liquid-like state.

In Chapter 5 we depart from the direct investigation of phase singularities, and turn towards a different type of singularity instead: flow-field singularities. Here we study the behaviour of singularities in the 2D Poynting vector field, which describes the flow of optical energy inside the sample. Singularities of this field signify an absence of in-plane energy flow: stationary points of the flow-field. Here we go into mathematical detail under which circumstances the singularities arise, and how to classify them through vector field topology, which dictates the behaviour of the flow around these singularities. We show that singularities of different origins are linked with their topologies, and how singularities with different origins are spatially correlated.

And finally in Chapter 6 we push the boundaries of our near-field microscope to new limits. Normally, the experiments are performed using a continuous wave (CW) laser as the light source, which illuminates the sample, allowing us to observe the resulting field in the sample for a single colour. But when the power of the laser is sufficiently increased, the light can interact with the sample, generating new light that has half the wavelength of the original light, a process that is called second harmonic generation. This light could not be detected by the microscope simultaneously with the original light. In this Chapter we show modifications that were made to the setup to allow for the simultaneous detection of light and its second harmonic. We show that we are able to achieve this while retaining sensitivity to amplitude, phase, and polarization.

SAMENVATTING

In dit proefschrift is er onderzoek gedaan naar de eigenschappen van licht, met het doel om diepere kennis te vergaren over diens gedrag op een fundamenteel niveau wanneer het licht beperkt wordt tot twee dimensies. Interessant genoeg ligt de focus van het onderzoek op het kijken naar plekken waar er eigenlijk helemaal geen licht is. Deze punten van perfecte duisternis bezitten een aantal fascinerende eigenschappen, zoals het feit dat er een zogeheten topologische lading mee geassocieerd is, en dat ze singulariteiten van het veld zijn. Singulariteiten zijn punten waar een bepaalde manier om het veld te beschrijven niet langer goed werkt. In deze dissertatie beschouwen we hoofdzakelijk de fase van het veld: wanneer er geen licht meer is werkt de methode om de fase van het lichtveld vast te stellen op dat punt niet meer, en hebben we te maken met een fase singulariteit.

Om deze singulariteiten te observeren gebruiken we een speciale microscoop. De singulariteiten, in twee dimensies, zijn punten: ze zijn infinitesimaal klein, en om ze te kunnen observeren moeten we in staat zijn om het licht veld te kunnen zien op hele kleine lengteschalen. Het probleem met het gebruiken van een gewone microscoop is dat deze fundamenteel gelimiteerd is in diens resolutie: ongeacht de sterkte van de vergrotingsfactor van het gebruikte objectief, een gewone microscoop kan alleen twee objecten van elkaar onderscheiden wanneer deze verder uit elkaar liggen dan een bepaalde limiet. Dit limiet staat bekend als de diffractielimiet. En het gedrag van deze singulariteiten gebeurt op een lengteschaal die vele malen kleiner is dan deze diffractielimiet. Dus deze punten van duisternis zullen verhuld worden door het omringende licht. Daarnaast observeert men door een gewone microscoop alleen het licht dat richting het objectief propageert, terwijl we eigenlijk geïnteresseerd zijn in het licht veld en diens eigenschappen wanneer het veld beperkt wordt tot een tweedimensionaal vlak in een preparaat. Om deze redenen hebben we een speciale microscoop nodig om deze singulariteiten te observeren. Deze specifieke microscoop heet een scannende nabije-veld microscoop, en in plaats van lenzen gebruiken om licht op te vangen gebruikt het een kleine naald-achtige probe die we extreem dicht bij het oppervlakte van het preparaat brengen. Een typische afstand hiervoor is in de orde van 20 nanometer, minder dan een duizendste van een (menselijke) haardikte. De probe is in staat om licht op te vangen van het veld dat gelimiteerd is tot een minuscule regio in de buurt van het oppervlakte, wat het nabije veld is genoemd, en alle informatie bezit over het licht in het preparaat. Aangezien deze techniek niet berust op het gebruik van lenzen en conventionele optica is het tevens niet gelimiteerd door de diffractielimiet, en stelt het ons in staat om de wonderbaarlijke wereld van licht op zeer kleine lengteschalen te observeren. De details van deze meettechniek zijn te vinden in Hoofdstuk 2.

Wanneer we kijken naar hoe de singulariteiten bewegen, dan lijkt het wel alsof ze bewegen zoals gewone deeltjes. Dus voor het eerste deel van dit onderzoek (Hoofdstukken 3 en 4) bekijken we deze fase singulariteiten alsof het gewone deeltjes zijn. Op deze manier beschouwd kunnen we gebruik maken van alle standaard middelen van de statistische fysica en zien of we hun posities en bewegingen op die manier kunnen analyseren. In Hoofdstuk 3 bestuderen we hun diffusieve gedrag, waar we experimenteel hun gedrag in de tijd kunnen volgen. Daarmee kunnen we ook hun bewegingen volgen, die niet lineair blijken te zijn. En aangezien we ze kunnen volgen vanaf hun moment van creatie tot het moment van hun "dood" zijn we ook in staat om statistisch te bepalen wat hun gemiddelde levensduur is. Verder laten we zien dat het scheiden van de singulariteiten in twee families, gebaseerd op hoe ze sterven, een grote impact heeft op zowel hun levensduur, alsmede hun diffusie. Hoofdstuk 4 volgt dezelfde lijn der redenatie als Hoofdstuk 3, maar is volledig numeriek. Wanneer drie vlakke golven met elkaar interfereren in twee dimensies zal het resulterende veld een ordelijk rooster vormen, zowel voor de amplitude als de fase singulariteiten, net als atomen in een materiaal die zich in de vaste fase bevindt zoals bijvoorbeeld ijs. We onderzoeken of het mogelijk is om zo'n vast-achtige staat te laten "smelten" naar een vloeistof-achtige staat, net als wanneer ijs smelt en water wordt. Hier laten we een methode zien voor het introduceren van meer golven op een langzame, symmetrische manier die resulteert in een transitie van een vast-achtige staat naar een vloeistof-achtige staat.

In Hoofdstuk 5 slaan we een andere weg in dan het direct onderzoeken van fase singulariteiten en onderzoeken we een andere soort singulariteit: stromingsveld singulariteiten. Hier onderzoeken we het gedrag van singulariteiten in het 2D Poynting vectorveld die de stroming van optische energie in het preparaat beschrijft. Singulariteiten van dit veld betekenen de afwezigheid van stroming in het vlak: stationaire punten van het stromingsveld. Hier gaan we in wiskundig detail in op de omstandigheden waarin deze singulariteiten ontstaan, en hoe ze te classificeren volgens vectorveld topologie wat het gedrag van de stroming rondom deze singulariteiten bepaald. We laten zien dat singulariteiten met verschillende oorsprongen verbonden zijn met hun topologieën, en hoe singulariteiten met verschillende oorsprongen ruimtelijk gecorreleerd zijn.

Als laatste verzetten we de limieten van onze nabije-veld microscoop in Hoofdstuk 6. Normaal gesproken worden de experimenten uitgevoerd met een continue golf laser die het preparaat belicht en ons in staat stelt om het resulterende veld in het preparaat te observeren voor één bepaalde kleur. Maar wanneer het vermogen van de laser voldoende verhoogt wordt kan het licht interacteren met het preparaat. Als gevolg hiervan kan er nieuw licht gegenereerd worden met de helft van de golflengte van het originele licht. Dit proces wordt tweede harmonische generatie genoemd. Dit tweede harmonische licht kon niet tegelijkertijd met het originele licht gedetecteerd worden door onze microscoop in de oorspronkelijke opstelling. In dit hoofdstuk laten we de modificaties zien die we gedaan hebben aan de opstelling om het mogelijk te maken om het originele licht en de tweede harmonische tegelijkertijd te observeren. We laten zien dat we hier toe in staat zijn, terwijl we de sensitiviteit voor amplitude, fase, en polarisatie behouden.

1

INTRODUCTION

*A wizard is never late, nor is he early.
He arrives precisely when he means to.*

Gandalf the Grey - The Fellowship of the Ring¹

In this chapter we start with the basics of light fields, introducing Maxwell's equations of electromagnetic fields and other basic principles on which this thesis is based. We explain what singularities are, how they appear in optical fields, and why they are interesting. We finish the chapter by giving a brief overview of the structure of the entire thesis, including a brief summary of each chapter.

1.1. ABOUT LIGHT

Light is one of the fundamental ingredients of life as we know it. It is necessary for everything we need to stay alive on our planet travelling through space. It allows us to see the world around us, the (terrestrial) flora uses light for energy through the process of photosynthesis, and with it allows us humans to thrive on this planet at all.

If there is one thing to be said about humans, it's that we are a curious bunch. Humanity has strove to get a better understanding of light since the dawn of ages. From prehistoric times trying to understand the day-night cycle [1], to cults worshipping the sun [2]. And as the centuries advanced, so too did our understanding of what light is, and how to utilize it [3–5]. In modern times, there is barely any piece of technology that does not utilise light in some form. Your TV remote uses infrared light to send signals to your TV, we use light to heat up our food in a microwave, and light is essential for GPS. Even though none of these examples use waves that we can see directly, they are all part of the electromagnetic spectrum. Only a small part of it is what we can see with our eyes, the visible spectrum, and which is what we may be used to calling light¹. But there is a lot of potential for applications in all parts of this electromagnetic spectrum [6].

One such application, without which the modern world would quickly cease to function, is telecommunication. Most pieces of technology are connected with the internet these days, including many household appliances [7], but the exchange of data does not happen automatically. This data exchange between different devices is done using light at a frequency just outside of our range of vision: in the near infrared at a wavelength of 1550 nm, which is called telecom frequency for obvious reasons. This frequency has of course not been chosen randomly, but rather because the signals are guided using optical fibres. And for this specific frequency, the losses in the optical fibre are extremely low, while for other frequencies the fibre starts to absorb more and more of this light instead [8]. And for long range communication, and especially if you want to send data across continents, then having the lowest possible losses is essential.

But since human curiosity and innovativity knows no bounds, we are always wondering how to push our technology to its next limit. With the invention of laser devices came rapid advancement in light-based applications, also known as photonics, leading to for instance making chips using lithography and trapping atoms via laser cooling [9], the development of the latter leading to the 1997 Nobel prize [10]. And in turn the laser cooling has enabled the creation of an atomic clock with femtosecond (10^{-15} s) accuracy [11] and is used in order to define the second [12]. And these clocks in turn paved the way for applications such as the global positioning system (GPS) [13], and even in high-frequency trading in financial world [14]. But the invention of the laser also lead to more mundane applications such as scanning a barcode in the supermarket and being a visual aid during presentations. The current holy grail of photonics is all-optical switching, in which the chips are fully controlled using light instead of with electrons [15]. However, light does not readily interact with other light, and thus a commercially viable all-optical switch has remained elusive as of yet.

¹Even though we apparently also like to call ultraviolet and infrared light as well, despite being invisible to us.

1.2. ELECTROMAGNETIC FIELDS

1.2.1. MAXWELL'S EQUATIONS AND THE WAVE NATURE OF LIGHT

Depending on at which length scale we consider light, there are different ways to describe it. When looking at large scales, the ray picture of light is a good way to describe it. In this picture, also known as geometrical optics, light can be regarded as a beam travelling in a specific direction [16]. When this ray impinges on an interface, it can subsequently undergo refraction or reflection. This picture is an excellent way to describe its behaviour when the wavelength of the light is much smaller than the objects that the ray of light encounters. And since the wavelength of visible light is in the order of hundreds of nanometers, that means that the way visible light interacts with everyday objects such as lenses and mirrors is well described by geometrical optics.

Theoretically, light fields can be described through two vector fields, namely the electric field \mathbf{E} and the magnetic field \mathbf{B} [17]. While these fields might be quickly associated with charged particles like electrons, and magnets, (classical) optics is also an electromagnetic process. And the way to describe these processes is via the famous Maxwell's equations, which relate the electric and magnetic field to each other. In vacuum, these equations are

$$\begin{aligned}\nabla \cdot \mathbf{E} &= \frac{\rho}{\epsilon_0}, \\ \nabla \cdot \mathbf{B} &= 0, \\ \nabla \times \mathbf{E} &= -\partial_t \mathbf{B}, \quad \text{and} \\ \nabla \times \mathbf{B} &= \mu_0 (\mathbf{J} + \epsilon_0 \partial_t \mathbf{E}),\end{aligned}\tag{1.1}$$

where ρ and \mathbf{J} are respectively the charge density and charge current, and ϵ_0 and μ_0 are respectively the permittivity and vacuum permeability.

Assuming that the medium of propagation is not magnetized or polarized, is linear and has no free current or charges (meaning $\rho = \mathbf{J} = 0$, which tends to hold for optics), these equations can be reduced to:

$$\begin{aligned}\nabla \cdot \mathbf{E} &= 0, \\ \nabla \cdot \mathbf{B} &= 0, \\ \nabla \times \mathbf{E} &= -\partial_t \mathbf{B}, \quad \text{and} \\ \nabla \times \mathbf{B} &= \mu_0 \epsilon_0 \partial_t \mathbf{E}.\end{aligned}\tag{1.2}$$

Combining these equations with some identities from calculus, we arrive at what is arguably the basis of optics:

$$\begin{aligned}\nabla^2 \mathbf{E} &= \epsilon_0 \mu_0 \frac{\partial^2 \mathbf{E}}{\partial t^2}, \quad \text{and} \\ \nabla^2 \mathbf{B} &= \epsilon_0 \mu_0 \frac{\partial^2 \mathbf{B}}{\partial t^2},\end{aligned}\tag{1.3}$$

which are of the form of a wave equation:

$$\nabla^2 \psi = \frac{1}{v^2} \frac{\partial^2 \psi}{\partial t^2}. \quad (1.4)$$

By comparing the wave equation to the formulas for the electric and magnetic field, we quickly see that in this case $v = 1/\sqrt{\epsilon_0 \mu_0}$, which is of course the speed of light (in vacuum) c . So this shows that electric and magnetic fields actually propagate as waves at the speed of light and reveals the wave nature of light. Hence light actually has to be regarded as a wave phenomenon. And with it comes all the properties of waves, such as interference, and require us to describe light using parameters such as wavelength (or frequency). Additionally this means that these electromagnetic waves can interfere with other electromagnetic waves. This wave behaviour is the behaviour that is observed when looking at the famous double slit experiment [18] for instance, in which light can be observed at positions on a screen which would be impossible when regarding light as simple rays.

Finally, the particle-wave duality of quantum mechanics states that every particle has wave properties, and vice-versa [19]. And this particle nature of light becomes apparent only when looking at the smallest scales. Here we discover that the energy carried by light is quantised in the form of photons, which can be seen through processes such as photoluminescence where a photon is absorbed, exciting an electron to a higher energetic state, which then relaxes to a lower energetic state, emitting a photon once again [20]. In this thesis, we will only use the wave nature of light, but it is important to keep in mind that it does not describe everything properly.

1.2.2. PLANE WAVES AND TRANSVERSE MODES

One of the simplest solutions to the wave equation above is the electromagnetic plane wave. This is a wave that travels in a specific direction and oscillates with a certain frequency, which can be described mathematically as:

$$\psi(\mathbf{r}, t) = A e^{i(\mathbf{k} \cdot \mathbf{r} - \omega t + \phi_0)}, \quad (1.5)$$

where \mathbf{k} is the wave vector indicating the direction in which the wave travels, \mathbf{r} is the spatial position, ϕ_0 its phase offset, A its amplitude, ω the angular frequency of its oscillation, and i is the imaginary unit. Eqn. 1.5 also introduces the mathematical concept of a field: it has a value for all spatial positions and time. Since the values are a complex number everywhere, this is then a complex scalar field.

A simple plane wave just describes a single propagating wave, but generally light fields consist not of a singular plane wave, but rather a multitude of them. Each of these different plane wave can have different amplitudes, phases and frequencies. The superposition of all of these different plane waves then forms the full electromagnetic field. In this case the field can then be mathematically described through a summation over all these different waves:

$$\psi(\mathbf{r}, t) = \sum_{\mathbf{k}} A_{\mathbf{k}} e^{i(\mathbf{k} \cdot \mathbf{r} - \omega_{\mathbf{k}} t + \phi_{\mathbf{k}})}. \quad (1.6)$$

Here $A_{\mathbf{k}}$, $\omega_{\mathbf{k}}$ and $\phi_{\mathbf{k}}$ are respectively the amplitude, angular frequency and initial phase for each different plane wave, which can be different for all of them. But when we consider monochromatic waves, i.e. waves with only one frequency, where $\omega_{\mathbf{k}} = \omega$ for all \mathbf{k} , Eqn. 1.6 can be simplified to

$$\psi(\mathbf{r}, t) = \sum_{\mathbf{k}} A_{\mathbf{k}} e^{i(\mathbf{k} \cdot \mathbf{r} - \omega t + \phi_{\mathbf{k}})} = e^{-i\omega t} \sum_{\mathbf{k}} A_{\mathbf{k}} e^{i(\mathbf{k} \cdot \mathbf{r} + \phi_{\mathbf{k}})} \equiv e^{-i\omega t} \psi(\mathbf{r}). \quad (1.7)$$

So it is possible to split the time dependence from the spatial dependence. This then implies that the amplitude of the field remains constant in time, only varying in space.

Using this assumption of monochromaticity, Maxwell's equations can be written as:

$$\begin{aligned} \nabla \cdot \mathbf{E} &= 0, \\ \nabla \cdot \mathbf{B} &= 0, \\ \hat{\mathbf{k}} \times \mathbf{E} &= c\mathbf{B}, \quad \text{and} \\ \hat{\mathbf{k}} \times \mathbf{B} &= -\frac{1}{c}\mathbf{E}. \end{aligned} \quad (1.8)$$

Note that the third and fourth equations contain $\hat{\mathbf{k}}$ which is the unit vector for the direction of propagation and is defined as $\hat{\mathbf{k}} \equiv \mathbf{k}/|\mathbf{k}|$. This is the result of the equations having been Fourier transformed first. The first two equations simply demonstrate that there are no sources or sinks for either the electric or magnetic field and therefore imply energy conservation. The other two relate the direction of propagation to the direction of the electric and magnetic field. Specifically, by definition the vector resulting from a cross product will always be perpendicular to both vectors of the cross product. So this tells us that the electric field, magnetic field, and direction of propagation are all perpendicular to one another. This is what is called a transverse electromagnetic wave. For a plane wave travelling in free space, this will always hold, but this can change when boundaries and materials become involved.

When light is confined, it is useful to split the light fields into transverse electric (TE) and/or transverse magnetic (TM) modes. Consider a plane, which could be the plane of observation or a physical surface. The normal vector and the \mathbf{k} -vector of the incoming plane wave together span a surface. In the case of a TE mode, the electric field will only have components that are perpendicular to this surface. So there will be no electric field in the direction of propagation. This usually means that all electric field components will lie in the plane of observation, excluding the direction of propagation, which can also have a component in that plane. In the case of TM light, the above holds for the magnetic field instead. For the electric field this then means that there will be a component perpendicular to the plane of observation, and in the direction of the \mathbf{k} -vector.

1.2.3. GUIDING LIGHT

For technological applications, and especially for on-chip environments, we will need to move beyond freely propagating light. We want to be able to control the flow of light and steer it to where it needs to be. In this section we introduce the concept of a waveguide, which is used to accomplish this.




Figure 1.1: Schematic example of a simple slab waveguide, consisting of three layers with each layer having a different refractive index n_i , with $n_2 > n_1$ & n_3 .

PLANAR WAVEGUIDES

In the ray picture of light, when a light ray hits an interface between materials of different refractive indices, part of the light gets reflected, and part of it gets refracted into the new medium according to Snell's law. However, when going from a material with refractive index n_1 to a material with refractive index n_2 , where $n_1 \geq n_2$, a critical angle of incidence $\theta_c = \arcsin(n_2/n_1)$ exists above which the wave no longer gets refracted, but rather only reflected: total internal reflection. The result is that it can no longer leave the original medium as long as the angle of incidence remains above the critical value: a channel is created through which the light is guided, which is called a waveguide.

There are many different configurations that could accomplish this, but all rely on the same principle of sandwiching a material with a higher refractive index in between materials of lower refractive index. By sandwiching a slab of high refractive index material between two lower index materials, a 2D (planar) waveguide is created. The light is free to propagate in the plane of the slab, but is prohibited from leaving through the lower refractive index materials. A simple schematic of such a slab waveguide is shown in Fig. 1.1. Here a slab waveguide with three layers is shown, where the refractive index n_2 is higher than both n_1 and n_3 , allowing for total internal reflection to occur.

MODES

While so far we have treated the waveguide from the ray picture of light, Maxwell's equations have shown that the wave picture is the true nature of light. And the plane wave solutions to Eqn. 1.3 were under the assumption that the light could freely propagate. But here we have added constraints on the directions of propagation instead. So the solutions to the wave equations will need to be reconsidered under these new boundaries. Below a quick derivation is shown for solutions for this case. For a more in depth treatment and derivations, the reader is referred to any standard electrodynamics textbook such as Ref. [17] or Ref. [21].

Only the wave equation for the electric field,

$$\nabla^2 \mathbf{E} = \mu_0 \epsilon_0 \frac{\partial^2 \mathbf{E}}{\partial t^2}, \quad (1.9)$$




Figure 1.2: Two schematic examples of on-chip optical waveguides. (a) shows a ridge waveguide, while (b) shows a rib waveguide.

will be considered, since the solution for the magnetic field is analogous. Assuming a simple time-harmonic field, this wave equation can be reduced to

$$\nabla^2 \mathbf{E} + k^2 \mathbf{E} = 0, \quad (1.10)$$

where $k = nk_0$ and $k_0 = \omega\sqrt{\mu_0\epsilon_0}$, which is the wave number in vacuum. By convention, we take the $x-y$ plane to be the unconfined plane, allowing for propagation, and we take the plane to be infinite. Without loss of generality we can set the direction of propagation to be the x direction, with a propagation constant of k_{\parallel} . As a result, we find translational invariance in the y direction, leading to $\partial_y \mathbf{E} = 0$. This then yields

$$\frac{d^2 \mathbf{E}}{dx^2} + (k^2 - k_{\parallel}^2) \mathbf{E} = 0, \quad (1.11)$$

which is a simple second order differential equation for each electric field component separately.

Under the boundary conditions of a guided mode there is no net transport of energy through the interface. Using this condition, the solutions to the differential equation are found to be given by a set of discrete oscillatory solution, very akin to the solutions to the particle in a box from basic quantum mechanics. These discrete solutions are called the modes of the field.

Implicitly we have used E_z here, but the same computation can be done for H_z as well. Here once again a distinction can be made between TE and TM light, with the solutions under both cases leading to these discrete solutions. These TE and TM modes are labelled by their index with TE_0 and TM_0 being the fundamental modes of the field. Additionally, each mode has its own in-plane propagation constant.

1D OPTICAL WAVEGUIDES

The concept of a slab waveguide can easily be moved to one dimension lower by adding an additional axis of confinement. One example of this is the ridge waveguide, in which a rectangular strip of a width w and height h is placed on a substrate. This ridge waveguide is shown in Fig. 1.2 (a). Above the substrate and strip is simply air, which has a refractive index of $n_{air} \approx 1$ ², which will always³ be lower than the material used for both the ridge and the substrate. When the material of the ridge is of a refractive index higher than that of the substrate, a narrow channel for the light is created in which it is confined in all directions except the direction along the ridge, making this a 1D waveguide. Fig. 1.2 (b) depicts another type of 1D waveguide, which is the rib waveguide. In this case, the strip sits not directly on top of the substrate, but rather on a layer of thickness d of the same material. This type of waveguide will be used in Ch. 6.

Like with the slab waveguide, the 1D waveguide also only supports a discrete set of modes to propagate. And since an additional axis of confinement is added, there are discrete states in both the confinement directions, which are labeled as TE_{ab} and TM_{ab} , with $a, b \in \mathbb{N}$. Unlike the case of the slab waveguide, analytical solutions to Maxwell's equations in the case of 1D waveguides do not readily exist, and as such the modes that are supported by the waveguide can only be found numerically.

While in a planar waveguide we lose confinement in one direction when compared with a 1D waveguide, what is gained is a possible spread in the in-plane k -vectors. While in a 1D waveguide the light can only propagate in 1 direction, for a 2D waveguide waves can in principle come from all in-plane directions and interfere with each other. Each of these interfering waves can additionally have a different amplitude and phase. If we then have a sufficient amount of these waves coming from all directions, it creates what is known as a random wave field, or chaotic light [24]. It is here that a lot of interesting physics arises, since a random wave field is the most generic wave field that one can create [25]. Any physics we can derive from this random field will hold more generally for other fields, and as such makes for a fascinating playground for investigating the behaviour of light at a fundamental level [24, 26–28].

1.3. SINGULARITIES

One thing that has always fascinated physicists and mathematicians alike is the existence of singularities. Singularities are points where a parameter becomes undefined, where the mathematical description breaks down [29]. These special points might sound like they are rare, and that one would have to look hard to catch a glimpse of these elusive entities. But in reality they are actually extremely common in nature. If you've ever sat on a terrace in the sun with a glass, then you probably noticed the curious pattern that was formed on the table by the light that passed through the glass. And surely you have noticed the mesmerizing pattern of dancing lines at the bottom of a swimming pool on a

²The refractive index of vacuum is exactly $n_{vac} = 1$. The refractive index of air is ever so slightly higher, but so close to it that it can be taken to be $n_{air} = 1$ for all intents and purposes.

³Also not strictly true: for instance metamaterials [22] and plasma's such as the Earth's ionosphere [23] can have $n < 1$.




Figure 1.3: Amplitude (left) and phase (right) of a vortex beam of order 1. The amplitude vanishes exactly in the center, leading to a phase defect, which is observed as a point defect where all possible phase values meet.

sunny day. Both of these examples actually belong a class of singularities known as caustics [30]. They are lines where rays of light all become parallel to one another, and the intensity at that line diverges: it becomes infinite. At least, if not for the diffraction limit, which is probably for the best once you consider the implications of infinite brightness.

And along with these everyday life examples of singularities, there exist a multitude more. Black holes are another well known example, since they contain a singularity of space-time at their core⁴. In this thesis we concern ourself mainly with singularities of the optical field, where one of the describing parameters becomes undefined. Below we explain a few such types which we can observe.

1.3.1. PHASE SINGULARITIES

When draining a bathtub or sink, undoubtedly you will have noticed the swirling vortex that arises when the water level gets low enough⁵. Here the direction of the flow of the water takes on all possible directions as it circles the drain [35, 36]. In light, something similar can happen to the phase: points in the light field where the phase takes all possible values of 0 to 2π around the point. Here the phase at the point itself is undefined, making the point a phase singularity [37, 38], the study of which garnered a lot of interest [24, 26, 39–46].

Recall for a monochromatic field for instance that we are able to describe it as

$$E_i(\mathbf{r}, t) = A(\mathbf{r}, t)e^{i\phi(\mathbf{r}, t)}, \quad (1.12)$$

where E_i represents one of the electric field components. In order to reconcile the phase singularity with this equation, nature then demands that $A(\mathbf{r}, t)$ becomes exactly zero. In the case of an electric field component, which can be regarded as a complex scalar

⁴This is not the event horizon, which is a singularity of the Schwarzschild metric, but can be removed via a coordinate transformation, and is therefore not a true singularity [31–33].

⁵Contrary to popular belief, the direction of the swirl is not fully dictated by the Coriolis force and hence dependent on the hemisphere it is in; the scale is simply too small for it to have any significant effect, and the geometry of the vessel is much more important instead [34].




Figure 1.4: Two phase singularities of opposite sign. Depending on whether the phase increases or decreases when we trace it going clockwise we can assign a positive or negative charge to them.

field, these singularities can be found at positions where both the real and imaginary components vanish independently. In 3D space, there are three degrees of freedom (the three spatial directions) and two constraints (the real and imaginary part both vanishing). So there are two equations to be satisfied and three parameters, leaving one parameter free after solving the equations. From this, we see that the solution space is one dimensional, meaning that the singularities in this case are line defects. In the case of a 2D space, however, we then find that the solution space for the singularities is reduced to zero dimensions, making it a point defect instead.

A good way to illustrate this principle is by looking at the amplitude and phase of a Laguerre-Gaussian beam. The profile of the beams are characterised by two parameters, l and p , with l the azimuthal index and p the radial index. A Laguerre-Gaussian beam with $l = 1$ is commonly referred to as a vortex or donut beam. In this beam, the amplitude profile is radial and only depends on distance from the centre of the beam, and its amplitude is exactly zero in the centre of the beam cross-section. In Fig. 1.3 we show the amplitude and phase of such a beam. On the left the distinctly donut shape of the amplitude can be seen that gives the beam its commonly used name, and on the right the phase is shown. In the phase clearly a point can be observed where all possible phases meet, which is the phase singularity.

Interestingly, a charge can be assigned to these points. This charge is based on the winding number using Cauchy's integral theorem, which is defined as a closed loop integral over the phase:

$$\oint_C d\phi = 2\pi s, \quad (1.13)$$

where s is the total charge of the singularities that is enclosed. To put this mathematical statement in visual terms: take any closed loop that can be drawn in the plane. This loop can be shrunk like contracting a loop made of a rope. But the one caveat is that the

contour of the loop cannot cross a singularity. Now if it is possible to fully contract the loop, then that means there is no singularity enclosed, since the loop must have crossed every point enclosed, and hence the integral evaluates to zero. However, if there is a singularity enclosed, then the loop cannot be fully contracted, and the integral is non-zero. The sign of the charge depends on the way the phase wraps around the singular point: if the phase increases when moving clockwise around the point, then by convention we call them positive, and if it decreases, then we call them negative instead. This is also shown in Fig. 1.4, where 2 singularities are depicted with opposite sign. On the left we see that the phase increases when moving clockwise around the point, making it a negative singularity, while the opposite holds for the right singularity. Generically, this charge is $s = \pm 1$, but can of course be higher if multiple singularities are enclosed, or if a higher order singularity (charge ≥ 2) is present.

The ± 1 charge is not just a simple charge, it is in fact a topological charge, since it is based on the winding number [47]. And since it is topological, it has to abide by the rules of topological charges as well. One such rule is that one cannot simply remove one singularity: charge neutrality needs to be preserved. The only way to remove a singularity is to have it annihilate with another singularity of the opposite charge. And similarly, the only way to create a singularity is to create it along with another singularity of opposite charge. This is known as pair creation/annihilation [42].

Now while the case of a vortex beam might be a specific example in which a single phase singularity is generated, they are actually a generic feature of random waves [24]. Anywhere in the random wave field where the waves interfere destructively, leading to a zero in the amplitude, phase singularities arise. One thing to note here is that when considering the vortex beam example, we are actually looking in the direction of propagation, as it is a cross-section of a 3D propagating beam. Furthermore, this vortex also carries orbital angular momentum (OAM), with the direction pointing in the direction of propagation [48].

1.3.2. POLARIZATION AND POLARIZATION SINGULARITIES

The phase singularity is a singularity of a scalar field, and so far we have gotten away with treating each electric field component as a separate complex scalar field. But light fields are vectorial in nature, and vector fields can contain singularities too. One such vector property of light is that of polarization, which describes how the electric field vector behaves in time in the plane, for which we consider the $x - y$ plane. Generically this vector traces an ellipse in time. A useful way to describe the polarization state is through the Stokes parameters:

$$\begin{aligned} S_0 &= E_x^2 + E_y^2, \\ S_1 &= E_x^2 - E_y^2, \\ S_2 &= 2\Re(E_x E_y^*), \quad \text{and} \\ S_3 &= -2\Im(E_x E_y^*). \end{aligned} \tag{1.14}$$


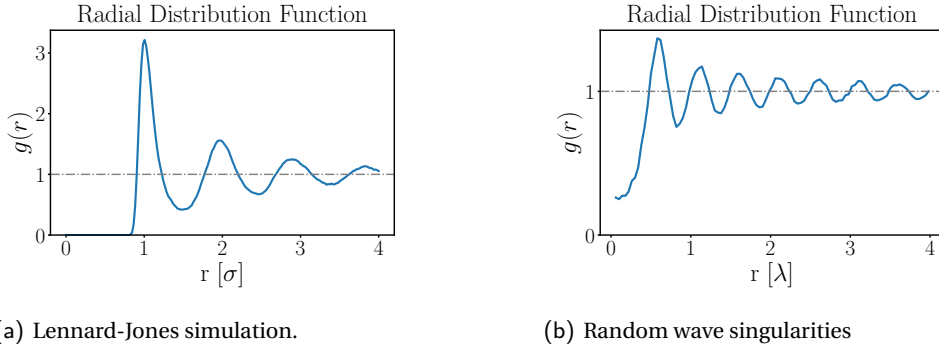


Figure 1.5: Schematic illustrations of (a) the polarization ellipse and (b) the Poincaré sphere. (a) shows the ellipse that the in-plane electric field vector \mathbf{E} traces in time. The major and minor axes of the ellipse are denoted by a and b respectively. (b) illustrates the Poincaré sphere spanned by the Stokes vectors S_{1-3} , with the radius of the sphere being given by S_0 . The two angles ψ and χ are illustrated with the blue and red areas respectively.

Here S_0 is the total intensity of the electric field, whereas the other three parameters describe the polarization state. S_1 describes how much of the electric field is linearly polarized along the x or y directions, with S_2 indicating the same, except for a rotation of 45 degrees with respect to the x -axis. Finally S_3 describes the amount of circular polarization. Together, the parameters S_{1-3} form a sphere with radius S_0 , known as the Poincaré sphere, and a point in this sphere uniquely characterizes the polarization state [49]. For visualization purposes, we tend to normalize the parameters S_{1-3} to S_0 , which we denote as $s_i \equiv S_i/S_0$. When the polarization state does not lie on the surface ($\sum_{i=1}^3 s_i^2 < 1$), the light is designated as not fully polarized. When instead it lies on the surface of the sphere ($\sum_{i=1}^3 s_i^2 = 1$), we call the light fully polarized. Additionally, the polarization state can be described using two parameters for fully polarized light, ψ and χ , which are analogous to the polar and azimuthal angles in the spherical coordinate system. Here we shall only consider fully polarized light.

When the light becomes fully circularly polarized ($s_3 = \pm 1$, $\chi = \pm\pi/4$), this polar angle ψ becomes undefined, leading to a singularity. This happens precisely when $s_1 = s_2 = 0$. Since we consider the system in 2D, and there are 2 equations to be satisfied for circular polarization to occur, the solutions are zero-dimensional, and hence discrete points in the plane. These singular points of circular polarization are also called C-points. Returning to viewing the polarization state through the lens of the ellipse that is being traced by the in-plane electric field vector reveals an intuitive view of the singularity: whereas an ellipse can be characterized by the angle of its major axis with respect to the x -axis, having a perfect circle leads to this angle becoming undefined.



(a) Lennard-Jones simulation.

(b) Random wave singularities

Figure 1.6: Comparison of the radial distribution function $g(r)$ for (a) a molecular dynamics simulation for a Lennard-Jones liquid with (b) the simulated $g(r)$ for phase singularities in a random wave field. For the Lennard-Jones liquid, the radius is plotted in units of the particle radius, while for the random wave field the radius is plotted in units of the wavelength. The Lennard-Jones liquid has been simulated through software provided by Ref. [51].

The other occurrence of a polarization singularity is when linear polarization occurs, and the minor axis of the ellipse becomes zero, and as a result the ratio of the minor and major axes, the ellipticity, becomes infinite. Since the requirement for this happening is the vanishing of s_3 , we have only 1 condition needing to be satisfied, while having 2 free parameters. As such linear polarization doesn't occur at a single point in space as C-points do, but rather appear as lines, and are also referred to as L-lines. This dimensionality of the singularities can also be seen when considering the Poincaré sphere once more: C-points occur only at two discrete points, the poles, while L-lines occur all along the equator and form a 1D space instead.

One thing to note is that, unlike phase singularities, polarization singularities don't carry a charge of ± 1 , but rather $\pm 1/2$ [50]. This difference in charge can be intuitively understood when considering that the shape of an ellipse is mirror symmetric over its major axis. As such, rotating it by π results in an ellipse that is indistinguishable from its unrotated form. As a result, the major axis rotates only by $\pi/2$ around a C-point, resulting in a charge of $\pm 1/2$ instead.

Polarization singularities will not be investigated further in this thesis, but serve both as an introduction to polarization as a concept, as well as an introduction to vector field singularities, which will be used in Chapter 5.

1.3.3. SINGULARITIES AS PARTICLES

One feature that makes singularities an interesting subject to study is that they tend to show remarkable similarities to real charged particles [52–56]. Of course they are not actually real particles, they are simply zeros in intensity of the light field, but it can be insightful to study in how far they are similar. And importantly, in what ways they are not similar as well. Understanding the differences and similarities between singularities and charged particles can lead to a deeper understanding of the behaviour of light.

For singularities in isotropic random waves, it is known that they behave liquid-like [24], which has been verified experimentally [25, 43, 57]. One way to appreciate this is by considering the pair correlation between the singularities. The pair-correlation function, also referred to as the radial distribution function $g(r)$, is a measure for how particles are distributed in space [58]. In particular it gives the probability distribution of finding another particle at a certain distance r from another particle. It is taken as an ensemble average over all particles, and can mathematically be defined as

$$g(r) = \frac{1}{N\rho} \left\langle \sum_{i \neq j} \delta(r - |\mathbf{r}_i - \mathbf{r}_j|) \right\rangle, \quad (1.15)$$

where N is the total amount of singularities, ρ is the average singularity density, δ is the Dirac delta function, $\mathbf{r}_{i,j}$ are the spatial positions of two singularities, and $\langle \dots \rangle$ denotes the ensemble average. The pair-correlation function $g(r)$ also gives an insight into the structure of the field. For instance, for a regular lattice, corresponding to what can be seen as a solid state of matter, one would expect to find sharp peaks in the pair-correlation function, since there are discrete distances at which another particle is found. For an ideal gas however, which is fully disordered, one would find that the distribution is unity everywhere, since all the positions are uncorrelated: there is no preferential direction at which one can expect to find another particle, they are simply everywhere. Intermediately one has liquids, which are neither fully disordered like a gas, nor as regular as a solid, which is also reflected in their correlation function: here we see a peak at around half a wavelength, after which it oscillates around unity with an amplitude that decays as the distance increases, eventually becoming uncorrelated at large distances.

In Fig. 1.6 we show a comparison between the pair-correlation function of a Lennard-Jones liquid (a), and that of phase singularities in a random wave field (b). The Lennard-Jones liquid describes a liquid of electrically neutral particles where the inter-particle interactions are governed by the Lennard-Jones potential [59, 60], and is perhaps the most well-studied potential in existence. This potential is strongly repulsive at an inter-particle distance smaller than their radius and softly attractive at distances beyond it. Molecular dynamics simulations using this potential can yield a liquid state, for which the $g(r)$ has been plotted in Fig. 1.6 (a). Here we observe that the $g(r) = 0$ for a distance smaller than the particle radius before exhibiting a peak just beyond the particle radius, indicating a high likelihood of finding another particle close to the reference particle. Beyond that the $g(r)$ shows oscillatory behaviour around unity, with the amplitude decaying as the distance from the reference particle is increased.

The pair-correlation function for the phase singularities displayed in Fig. 1.6 (b) shows very similar behaviour, but there is one immediate difference between the pair-correlation function for singularities here and those in a Lennard-Jones liquid: curiously, the $g(r)$ for singularities does not go completely to zero for $r \rightarrow 0$. This can be explained by considering the size of the particles under consideration: whereas the Lennard-Jones particles have a finite size, and cannot approach another particle closer than their diameter, phase singularities are actually zero dimensional objects. So while real particles have a finite size, singularities are actually infinitesimally small, and as such are able to approach each other to arbitrarily small distances [24].

1.4. LAYOUT OF THIS THESIS

In this thesis the emphasis lies on gaining a deeper understanding of the behaviour of singularities, using experiment, simulations and theory combined. The first part of this thesis is devoted to investigating in how far singularities and charged particles show similarities by taking a statistical physics approach. We take a look at the diffusive properties, and at how we can melt a lattice of singularities. After this we look toward the flow field, identify its singularities and origins thereof. Finally we push the experimental limits by showing that we are able to measure a field at telecom frequencies, as well as its second harmonic simultaneously.

- In Chapter 2 we introduce all the experimental tools that were used to perform the near field experiments for studying singularities and more.
- In Chapter 3 we investigate the diffusive properties of singularities by studying their dynamics in time, as well as their decay rate.
- In Chapter 4 we look at the melting of singularity lattices, and attempting to discern a mechanism that could act as an analogue to temperature to induce it.
- In Chapter 5 we look at the properties of the time-averaged Poynting vector field and its singularities.
- In Chapter 6 we present the first experiment in which we are able to measure the near field of both infrared and visible light simultaneously.
- In Chapter 7 we draw conclusions from the studies presented in this thesis and give an outlook on further lines of research.

2

EXPERIMENTAL SETUP

*Where the windows are breathing in the light
Where the rooms are a collection of our lives
This is a place where I don't feel alone
This is a place that I call my home*

That Home - The Cinematic Orchestra²

In this chapter we introduce the tools that are used in order to perform near-field experiments. We start by introducing the diffraction limit, as well as how it's possible to circumvent it. Next we show how we have implemented the near-field microscope as well as the type of samples that are used in this thesis. This includes how we are able to access the dynamical properties of 2D light fields, as well as their polarization state. Finally we show how we are able to detect optical singularities, which are deeply sub-wavelength entities.

2.1. BEYOND THE DIFFRACTION LIMIT

In a world where technology is ever shrinking, the structures that comprise it necessarily become smaller and smaller as well. And with it comes the need to be able to understand and control light at smaller scales as well. Nanophotonics is the field that concerns itself with the behaviour of light at the nanoscale, and it studies for instance the optical response of nanoantennas [61], or the structure of light inside of waveguides [62]. Understanding this behaviour leads to a plethora of applications, such as precise sensing [63, 64], better solar cells [65] and finding ways to manipulate and steer light at this scale [66]. What binds all these different applications and behaviours together is that they all involve optical effects on a scale much smaller than the optical wavelength used to induce them. In order to observe these optical effects, a way is needed to observe them.

Unfortunately, conventional microscopy is limited in the resolution that can be achieved due to the diffraction limit. Such a limit was first determined by Ernst Abbe already in 1873, who approximated the diffraction limit of a microscope to be

$$d = \frac{\lambda}{2\text{NA}}, \quad (2.1)$$

where λ is the wavelength of the light, and NA is the numerical aperture of the objective [67]. Functionally, this sets the limit at approximately half the wavelength [68]. So regular microscopy is no longer sufficient when studying the behaviour of light at length scales below half a wavelength. The diffraction limit describes the ability to resolve two nearby features and depends on the size of your aperture and the wavelength. The smaller the wavelength becomes, the smaller features become, and hence they can get closer before overlapping.

Since conventional microscopes pick up the light propagating through the objective, it is fundamentally bound the relation for freely propagating light, given by

$$k_x^2 + k_y^2 + k_z^2 = k_0^2. \quad (2.2)$$

This relation then also dictates the limit of the spatial frequencies that can be detected.

However, the in-plane momentum vector $k_{\parallel} = (k_x, k_y)$ can be larger than that of freely propagating light such that $k_{\parallel}^2 > k_0^2$. This can only happen when $k_z^2 < 0$, and thus k_z becomes a complex number instead. As a result, the out-of-plane contribution to the propagation of such a wave becomes proportional to $e^{-|k_z|z}$, leading to a field that decays exponentially away from the plane. This decaying field close to the surface is what is called the evanescent field, from the Latin word *evanescere*, which means *to vanish*. An illustration of this evanescent tail of light undergoing total internal reflection is shown in Fig. 2.1. The evanescent field does not propagate or transport energy away from the plane, but this evanescent does contain the full range of in-plane spatial frequencies, which can be higher than the frequencies of freely propagating light. Gaining access to this evanescent field thus allows for a spatial resolution beyond the diffraction limit through these high in-plane frequencies. This evanescent field can also be seen as a consequence of boundary conditions imposed by Maxwell's equations at an interface in combination with conservation of momentum. When light is confined to two-dimensions by total internal reflection, such as light propagating through a thin slab of material with a high




Figure 2.1: Illustration of the evanescent tail of a field. The red beam undergoes total internal reflection, leading to an evanescent field that decays exponentially away from the surface. The amplitude of the electric field away from the surface is depicted.

refractive index, laws of physics demands that the electromagnetic fields parallel to the interface must still be continuous at the interface. Thus the fields are not automatically zero, but rather are decaying exponentially across the boundary.

What is measured with a conventional microscope is what is also referred to as the far-field. The light that is collected is comprised solely of light not confined to a structure, such as light that gets scattered out of a sample, and is allowed to freely propagate. For this freely propagating light all k -vector components are real valued, and is bound by the diffraction limit as usual. But being able to detect the evanescent tail of the field, which is called the near field, allows access to light with imaginary k -vector components. This results in an increased k_{\parallel} , and hence allows for a resolution higher than the diffraction limit.

2.2. NEAR-FIELD MICROSCOPY

One early example of a proposed way to detect the near field is that of Edward Synge, who in 1928 already proposed the use of an opaque plate with a tiny hole in it [69]. By placing this plate close to the surface of the to-be-imaged sample, illuminating the sub-wavelength hole, and then measuring the light transmitted through the sample, one can measure with an illumination spot that is not limited by diffraction, but rather the size of the hole that is used. By subsequently scanning this hole over the surface, an image can be constructed with a resolution much higher than what the diffraction limit would allow. While at the time of proposal performing this experiment was not feasible due to the lack of technology to do it, it is still the core of what is currently used in near-field microscopy.

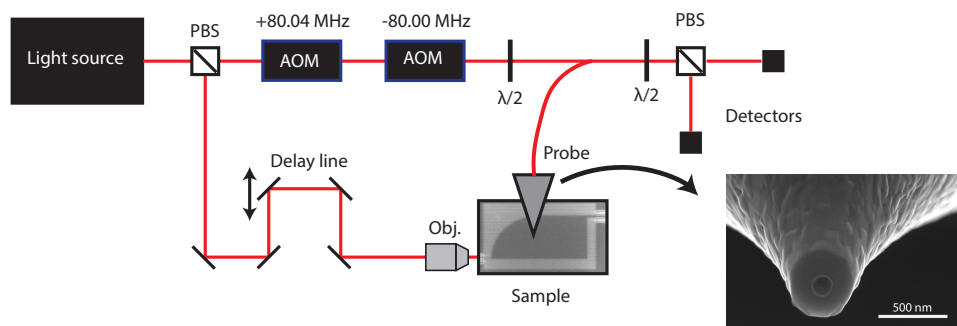


Figure 2.2: Schematic illustration of the near-field microscope used throughout this thesis. The two possible paths form a Mach-Zehnder interferometer, with the top path denoted as reference, and the bottom path as signal. The reference branch has its frequency shifted by 40 kHz through the use of the AOM's. Additionally, an SEM image of a near-field probe is shown.

In order to get the sub-wavelength resolution with the method proposed by Synge, the plate with the hole would need to reside extremely close to the surface. And while that might have been feasible, then additionally having to raster scan this plate over the sample surface without colliding with it suddenly becomes a monumental task. It would take over half a century before Synge's vision was realised and the first optical near-field microscope was made functional and a patent was filed by Dieter Pohl in 1982 [70]. However, they were not the first ones to show the viability of near-field microscopy, since Ash and Nicholls already a decade earlier managed to show superresolution using microwaves in 1972 [71]. In this case Pohl did not use an opaque sheet with a sub-wavelength hole in it, but rather a quartz crystal that was modified to have a sharp tip, and coated such that only the tip was transparent. In this way, the tip can be scanned over the surface, without having to deal with a cumbersome sheet. With this method he was able to achieve a resolution of $\lambda/20$ [72]. This revolutionary method is one of the implementations of scanning probe microscopy, which saw its inception merely a year earlier in 1981 with the invention of the scanning tunneling microscope (STM) by Binnig and Rohrer, earning them the 1986 Nobel prize [73]. This STM technique paved the way for a multitude of microscopes operating on the scanning probe principle, allowing for imaging at the nanoscale [74].

The implementation of Pohl is the first example of an aperture based NSOM, but other types of NSOM also exist, most prominently scattering type (s-NSOM) [75–77]. In this case a probe is used that is not hollow, but instead is made as sharp as possible. This is then put in the evanescent field to scatter light into the far-field where it can then be detected by a photodiode. In this thesis, we solely make use of an aperture based NSOM, which has been home-built and improved upon over the years [62, 78–81].

2.2.1. CAPABILITIES OF THE NEAR-FIELD MICROSCOPE

A near-field microscope is able to detect light fields at a sub-wavelength scale, where the probe can be used in two ways: collection mode and illumination mode [82]. In illumination mode, the aperture probe is used to locally illuminate the sample. This technique is generally used when local excitation is desired, for instance for the excitation of single molecules. Conversely, in collection mode, the aperture probe is used to locally pick up the light instead, and is used to map the spatial distribution of a confined field. Throughout this thesis the focus lies on the investigation of the spatial distribution of fields and their singularities, and hence only collection mode will be considered. A schematic representation of our near-field microscope is shown in Fig. 2.2, and in this section the function of all its different elements will be explained.

In its most basic form, an aperture-based near-field microscope in collection mode is able to record the optical field intensity. By raster scanning the probe over the surface, and recording the light transmitted through the probe on a photodiode, a map of the intensity can be recorded¹ [84]. While this is of course already a useful tool to have, a lot of information is discarded with this. In order to gain more information on the field, the incoming laser beam is split into two separate paths, where one goes to the sample where it then gets picked up by the tip (signal path), while the other bypasses the sample instead (reference path). The two paths are then recombined after the tip, forming a Mach-Zehnder interferometer. This then allows for sensitivity to the phase difference between the two branches. But in order to determine the phase and amplitude of the detected signal separately, the use of heterodyne detection is required in addition. By introducing a dual-phase lock-in amplifier (LIA), in combination with shifting the frequency of one of the two branches through the use of an acousto-optic modulator (AOM), we gain access to both the amplitude $|\mathbf{E}(x, y)|$ as well as the phase $\phi(x, y)$ of the measured field separately. This heterodyne detection scheme will be explained in more detail in Section 2.3.1.

But as has been mentioned before, light is vectorial in nature, and so far in the discussion about detecting light it has been treated as a scalar quantity. Of course access to the vector components separately is also desired. If the vectorial quantities of light are to be probed, polarization sensitivity is required as well in order to disentangle the different field components. By then utilising a pair of detectors, instead of a single detector, and a polarizing beamsplitter cube that splits the signal based on its polarization it becomes possible to also selectively detect E_x and E_y separately. Since the polarizing beamsplitter cube splits two orthogonal polarization states, two orthogonal polarization states are recorded on the two lock-in amplifiers. However, note that this does not immediately mean E_x and E_y , since the light that is picked up can undergo a rotation of polarization state before arriving at the beamsplitter and subsequent detectors due to birefringence. This birefringence can occur through the bends in optical fibres for instance. In order to compensate for this, adjustable waveplates are introduced in order to rotate the measured field back to a state where E_x and E_y are properly split on the detectors. Through this combination of techniques we are able to gain access to sub-wavelength features, while also gaining phase and polarization information.

¹This is a slight simplification, since the magnetic field is also picked up [83].

2.3. IMPLEMENTATION OF THE NEAR-FIELD MICROSCOPE

In the previous section we have discussed the principles of near-field microscopy and how it is possible to gain insight into the different components of the electric field, but not how it is implemented. The initial implementation by Pohl is based on scanning a tip with a sub-wavelength aperture over the surface, recording the throughput through that tip. This way is the basis of aperture-based near-field scanning optical microscopes (a-NSOM) in collection mode. The idea of using a probe was then iterated upon and improved by Betzig [85], where they introduced the use of pulled fibres. Here a single-mode optical fibre is pulled from two sides, while simultaneously being heated by a CO₂ laser, which under the correct circumstances (such as pulling force and laser power) leads to a fibre with a sharply tapered end. When we talk about near-field tips, this is what we are referring to. Additionally this tip can then be coated at the sides with aluminium. This revolutionary method is still the dominant method for creating tips today, but of course has since been further refined over the years [86]. The probes that are used in our lab are pulled fibres, with the pulling parameters chosen such that the end has a dumbbell shape in order to allow for a large opening angle to increase throughput. These fibres are then immersed in buffered acid before being coated with a 100 to 200 nm thick layer of aluminium to prevent light from leaking out of (or entering) the sides of the probe. Finally, using focused ion beam milling we create an aperture once more through which the light can enter the tip. For more exact details on the fabrication process, see Ref. [87].

While the idea of raster scanning an aperture probe over the sample sounds quite simple, one needs to keep in mind that probe needs to reside inside the evanescent field. In practice this means that the probe is suspended merely a few tens of nanometers above the sample, which needs to be kept constant during scanning. A constant height above the surface is maintained through the use of a shear-force feedback mechanism [88, 89]. The feedback mechanism has to be sensitive, as well as responsive, since tips are extremely fragile, and hitting the surface spells certain death for that tip. This feedback is achieved by gluing the probe to one prong of a quartz tuning fork, with the end of the tip sticking out only a tiny amount (about one fibre diameter). By only having the tip extend a tiny amount (around one fibre diameter) beyond the end of the prong, the tuning fork and tip combination forms a rigid system, and the tip does not start oscillating at a different frequency. This rigid system has a certain resonance frequency associated to it (usually around 33 kHz), and the system is subsequently also driven at this resonance frequency. When the tip approaches the surface, the resonance frequency shifts due to interactions with the surface, and this can be detected as a change in phase and amplitude. By giving it a desired phase setpoint, which translates to a certain distance to the sample, a feedback loop is created that aims to keep the distance fixed. This distance is controlled through a piezoelectric element, which allows for extremely precise movement. In this way, we are able to achieve a lateral step size of 15 ± 2 nm. Since a fixed distance is maintained between the probe and the sample, by recording the voltage on the piezo, which translates directly into a distance, we also get access to the topographical information. It acts similar to atomic force microscopy in that sense.

2.3.1. HETERODYNE DETECTION & GAIN

Additionally, we make use of a heterodyne detection scheme, where the reference path also goes through two acousto-optic modulators (AOM's), which shift the frequency of the light. This is done in such a way that the frequency is shifted by 80.04 MHz by the first one and 80 MHz by the second one. By making one a positive shift, and the other a negative shift, the reference path is shifted in frequency with respect to the signal path by 40 kHz. When the signal and reference paths are subsequently recombined, the resulting signal will then exhibit a beating of 40 kHz. Through the use of a lock-in amplifier (LIA), we can efficiently filter on this frequency to significantly improve the signal-to-noise ratio (SNR) by suppressing 1/f noise, and essentially limit the noise to be mainly shot-noise [90].

An additional benefit of using this heterodyne detection scheme is the so-called heterodyne gain. Since the voltage measured on the detector is proportional to the intensity of the combination of the electric fields of the signal (E_s) and reference (E_r) branches, we can write

$$V \propto I = [E_s \cos(\omega_s t + \phi) + E_r \cos(\omega_r t)]^2, \quad (2.3)$$

where ω_s (ω_r) is the frequency of the signal (reference) branch, and we take ϕ to be the phase difference between the two branches. Expanding this product then yields

$$V \propto |E_s|^2 + |E_r|^2 + 2|E_s||E_r| \cos(\omega_s t + \phi) \cos(\omega_r t). \quad (2.4)$$

The first two terms represent the DC voltage from the signal and reference branches separately. By using a LIA, these contributions are filtered out, since they do not vary in time. Instead we are left with the product of two cosines, which leads to a sum-difference term, where one varies with the difference frequency, and the other with the sum of the two frequencies. Since we filter on the difference frequency, the sum frequency is filtered out as well. The measured voltage after the LIA is then proportional only to the terms varying with the difference frequency:

$$V \propto |E_s||E_r| \cos(\phi) = \gamma |E_s|^2 \cos(\phi), \quad (2.5)$$

where $\gamma \equiv \frac{|E_r|}{|E_s|}$. Here we see that the measured voltage is a factor γ higher than when just measuring the intensity of the signal alone, and as such γ is called the heterodyne gain factor. By ensuring that $|E_r| \gg |E_s|$, we can obtain multiple order of magnitude in gain to boost our usually quite weak signal [91].




Figure 2.3: Examples of photonic crystals in one and two dimensions. (a) shows a one dimensional PhC, which is created by alternating layers with differing refractive index. (b) shows a 2D PhC, which is created through a hexagonal array of air holes ($n=1$) in a slab of material.

2.4. GENERATION OF 2D LIGHT

With the *how* of performing near-field microscopy out of the way, we can now turn our gaze to the *what*. As alluded to in the introduction of this thesis, light behaves differently when it is confined to two dimensions, when compared to a 2D cross-section of 3D propagating light. And with the potential applications aimed at on-chip control of light planted in the back of our minds, we turn towards investigating truly 2D light. In this section we continue to build upon Section 1.2.3, where we introduced the notion of waveguides, which gives us a way to confine propagation direction. We start by introducing photonic crystals as a way of selectively restricting propagation based on the frequency of the light. Afterwards we combine these photonic crystals with the planar waveguides introduced earlier in order to create a platform for 2D random waves.

2.4.1. PHOTONIC CRYSTALS

Photonic crystals (PhC's) are most intuitively understood by considering one-dimensional Bragg stacks. Here, layers of dielectric material with different refractive index are alternated, forming a periodic structure. This periodic modulation of the dielectric constant also leads to a region of frequencies that are no longer allowed to propagate through the material [92–94]. This range of forbidden frequencies is known as a stop gap, and is the optical equivalent of electronic bandgaps, where electron of certain energy ranges are forbidden from propagating through an atomic lattice. When an electromagnetic wave with a frequency in the stop gap range impinges on this Bragg stack, it is not allowed to propagate through the material, and is necessarily reflected instead. In this way it is possible to create an effective mirror with an extremely high reflectivity. This is an example of a 1D photonic crystal (see Fig. 2.3 (a)). The period a of the alternating layers and the refractive indices n_1 and n_2 determine the frequency range of this stop gap.

The higher-dimensional versions act essentially the same: through periodic variations in the effective refractive index one creates the conditions for the above to occur in high

dimensions [95, 96]. The band of frequencies that are allowed to pass through the crystal are called photonic bands, while the bands of frequencies that are forbidden from propagating through the crystal are called the photonic band gap. Photonic crystals also exist outside of fabricated materials, and appear in nature as well. The iridescent colours that can be observed in for instance the wings of morpho butterflies, or the chitin of certain beetle species, are actually due to naturally occurring photonic crystals [97].

In our case, we make use of a hexagonal array of holes patterned into our silicon slab. The holes have a refractive index of $n = 1$, since it contains air, whereas the silicon has a refractive index of around $n = 3.4$. This is schematically shown in Fig. 2.3 (b), where the white blue rectangle signifies the silicon slab, and the inlaid circles indicate the air holes patterned into the slab. The range of the photonic band gap is determined by the size of the holes, as well as the distance a between two adjacent holes.

One interesting and useful thing happens when part of the photonic crystal is removed, for instance one row of holes in the hexagonal lattice.

This configuration allows for the appearance of modes, which can lie inside the photonic band gap. These modes can then propagate through the channel that has been opened in the photonic crystal while remaining confined to this channel. This provides a way to guide the light through the channel in a controlled and low-loss way. And because of the small size of this channel (about 200nm in width) the mode profile of the light that propagates is additionally restricted. This channel is what is called a photonic crystal waveguide, and it's a promising way of guiding light in on-chip environments [98–100]. These photonic crystal waveguides are usually indicated through the size of holes missing: a waveguide comprised of one missing row is indicated as a W1 waveguide, and one where 3 rows are missing a W3 etc.

But besides the single mode PhC waveguide, we can also go bigger and open up a much wider gap in the photonic crystal and use it not to guide light, but rather trap it in a larger space instead. This amounts to combining the concept of a planar waveguide with 2D photonic crystals: the planar waveguide confines the light to 2D space, and then using a PhC boundaries are implemented at which light of certain frequencies is reflected. This is simply a way to create a cavity of a larger size than is common for PhC cavities, which quite often are just a few missing holes in a PhC and used for confining light in a tight space [101, 102].

2.4.2. A PLAYGROUND FOR RANDOM WAVES

In order to generate the wave chaos that is used throughout the majority of this thesis, we use what we call a chaotic cavity [103]. A SEM image of one such cavity is shown in Fig. 2.4. It is a silicon-on-insulator base, with 220 nm thick silicon on $2\text{ }\mu\text{m}$ of silica. The thickness is such that the slab of silicon only supports the fundamental TE and TM modes of the used monochromatic light source. The shape and size of the photonic crystal is chosen such that it forms a photonic band gap for the TE mode of telecom light ($\lambda = 1550\text{ nm}$ in air). In other words, the TE mode cannot pass through the boundaries of the cavity and is necessarily reflected, thus the sides act like a mirror. Note however that this is not the case for the TM mode, for which there is no band gap (for this con-




Figure 2.4: SEM image of a chaotic photonic crystal cavity. The quarter stadium shape is delimited by the photonic crystal formed through etching holes into the silicon slab. The inset on the right depicts the input facet of the cavity and provides a closer look at the photonic crystal structure. Additionally we introduce the axis label convention that will be used through this thesis here.

figuration of the PhC). Of course there still is some TM light reflected, since it is incident upon an interface. The shape of this cavity is that of a quarter-stadium, chosen such that the shape is asymmetric and leads to wave chaos [103]. It strongly resembles the stadium shape that is also used for microwave billiard experiments [25, 104], but taking only one quadrant eliminates the 2 symmetry axes that would otherwise be present. The removal of these symmetry axes also removes potential symmetries of the field, which could impede the generation of wave chaos.

For continuous wave (CW) laser light, the light enters the cavity and reflects off of the walls. Due to the asymmetric shape of the cavity, the light bounces all around the cavity and fills up the entire phase-space: eventually the resulting field is comprised of a superposition of plane waves of all possible propagation directions (k -vectors) and initial phases [103]. This is what we call wave chaos, or random light. The only thing that remains stable is the magnitude of the k -vector of the light, as that is a fundamental property of the light in the medium for a given frequency, and as such is independent of the photonic crystal design. When Fourier transforming this field, we thus find a ring in k -space, with a radius equal to the magnitude of the k -vector (see for instance the bottom right image of Fig. 2.5).

2.4.3. TIME-RESOLVED MEASUREMENTS

When using CW light to generate our fields, only a single frequency is present, and hence it does not evolve in time: we are observing a steady state system where the amplitude remains constant (See Ch. 1.2.2). However, we would also like to investigate the time dynamics of light inside our structure. So in order to be able to accomplish that, we need to broaden our frequency spectrum in order to introduce a field whose amplitude also varies in space and time. This is done using a pulsed laser system instead of a CW one.

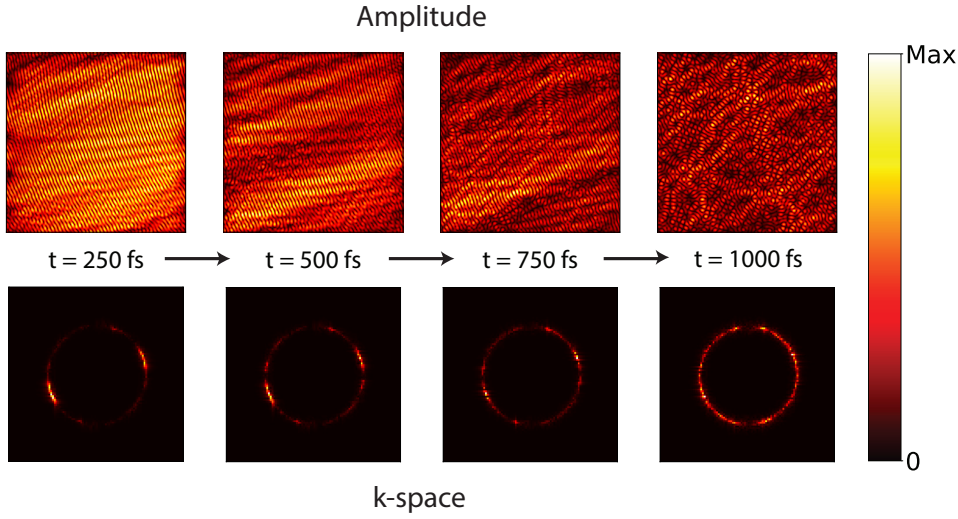


Figure 2.5: Evolution of a pulse as a function of time inside the chaotic cavity for the H_z field for 250 fs steps. The top row shows the amplitude images, while the bottom row depicts the k-space evolution. The progression from a structured field, and corresponding limited range of k-values to a fully chaotic isotropic random wave field can be clearly observed.

In order to obtain a beating of the two different branches of the interferometer, the branches now have to have the same length to within the coherence length of the broadband light field. This effective temporal overlap is achieved with an accurate delay stage [90, 105]. The delay stage is also key in providing the time resolution of our measurements. Once pulse overlap has been established, we can use that as a baseline measurement, a reference time. By then shortening the signal branch by a known amount, the pulse arrives earlier at the structure than it previously did. As such, the pulse has a longer time to propagate through the structure before the picked up light combines with the reference pulse. To put it mathematically, the light has $\Delta t = \frac{\Delta x}{v}$ longer to evolve in the structure before we measure it on the detector. Here Δx is the change in effective path length of the signal branch, equal to twice the distance of the movement of the delay stage. This results in a change in probing time of Δt . Please note that v is the speed of the light in the medium of propagation, which is different for free space and fibre. Since the light at the position of the delay stage is in air, the speed of light is simply equal to c for the computation of Δt . Using this motorised delay stage, we are able to reach a temporal step size of around 10 fs in this manner.

By taking measurements between every step of the delay line, it becomes possible to observe a pulse of light enter the cavity, and subsequently evolve into a random field as it has more and more time to fill out the phase space through sequential reflections at the photonic crystal boundaries. This process is shown in Fig. 2.5, where the evolution of the field as a function of time is depicted for a temporal step size of 250 fs. The top row shows the amplitude of the H_z field, while the bottom row depicts the reciprocal space. In the real space image the evolution from a somewhat ordered field propagating through the

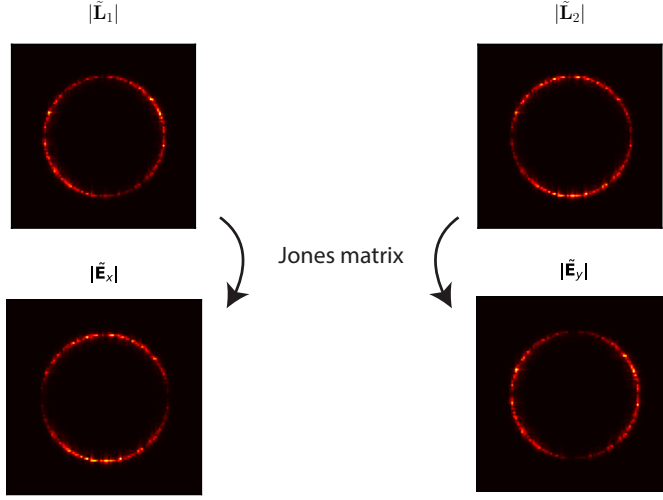


Figure 2.6: Application of a Jones matrix to experimentally measured data. The top two images show the reciprocal space distribution for the fields as initially detected with the LIAs. No clear minima of the k-space distribution is observed. The bottom row depicts the k-space distribution after applying a Jones matrix, leading to a k-space distribution as expected for the TE mode.

cavity to a fully random wave field can be observed. For $t = 250$ fs, the pulse has only just entered the cavity, and at most a few reflections have taken place. For each progressive time-step, more reflections have taken place, and as a result the field becomes more chaotic. After a picosecond the field has become (almost) fully random. This progression can be clearly observed in the reciprocal space as well. For $t = 250$ fs, only two areas mainly contribute to the field, signifying that the field mainly consists only of the incoming pulse, and an initial reflection in the opposite direction. As time progresses, the k-space distribution becomes more circle like, until an isotropic distribution is reached, indicating wave chaos.

One thing that needs to be noted here is that of course the measurement will always be an average over staggeringly many pulses. For the time-resolved measurements in this thesis a laser with a repetition rate of 80 MHz was used, and if a measurement took 15 minutes, which is a typical value, then a good 72 billion pulses will have passed through the cavity during it. Taking into account that with our typical scan speed of about 500 pixels a second, one pixel still ends up being the average of 16 thousand pulses. The assumption is that every pulse is identical and evolves in the cavity identically to the previous one. Since we are able to visualise and track pulses and their evolution, empirically we can confirm this assumption holds: if every pulse would evolve independent from every other pulse, the effective field formed would simply be uniform.

2.5. DETERMINING THE POLARIZATION STATE

In Section 2.2.1 we have described how we are able to achieve polarization sensitivity with a near-field microscope through the use of a polarizing beamsplitter cube and waveplates to compensate for polarization rotations. However, there is no a-priori way of knowing which waveplate settings will lead to proper separation of electric field components. Instead calibration is required, which will need to be redone every time something is changed in the measurement, such as a tip replaced, or the incoupling redone. This separation of electric field components is done instead by exploiting known mode profiles of a waveguide for instance. By coupling light into a single mode waveguide, for which one can numerically compute the mode profile, we can measure the fields on the two detectors, and rotate the waveplates until the desired field distributions are reached.

For a simple waveguide such as a W1 photonic crystal waveguide the available modes and their profiles are well understood [106–108] and hence forms a suitable candidate for this calibration. These known modes can then be used for determining the required waveplate settings in order to be able to split the in-plane electric field components. But for random fields in the chaotic cavity a different way to achieve a proper component separation can be used instead. For isotropic monochromatic random waves the expected fields and especially the fields in k-space are well understood. Since the fields are isotropic, that means that the waves are coming from all directions, meaning all orientations for the momentum vector are expected. Coupled with monochromaticity, where the k-space is limited to one radius only, the resulting momentum space has to look like a circle.

However, this isotropic assumption holds only for the out-of-plane oriented fields. So when considering the TE mode, this holds for H_z , and when looking at the TM mode this holds for E_z instead. The cavity is designed to have bandgap for TE light, but not TM light, and hence the TM light will not be a random wave field. So truly only for H_z do we expect to find this ring in k-space. But then what would this momentum space distribution light look like for the in-plane electric field components? Since a TE mode is considered, the restriction on the propagation direction applies: a transverse mode cannot have a component in the direction of propagation. In other words, when the field propagates solely into the x direction, then consequently we must find that $E_x(k_y = 0) = 0$. And similarly $E_y(k_x = 0) = 0$ has to hold. So when looking at the distribution in k-space, E_x will show a ring, but with its amplitude diminishing the closer to the $k_y = 0$ axis one goes, being exactly zero on the axis, whereas it shows a maximum for the $k_x = 0$ axis. This knowledge can then be used to separate the two electric field components by making use of virtual waveplates. In order to understand how these virtual waveplates work, the concept of Jones calculus and Jones matrices will be explained.

Polarized light can be described using Jones calculus, named after R.C. Jones, who formalised this in 1941 [109]. A polarization state can be represented using a Jones vector, and most optical elements can be represented with a Jones matrix, as long as the elements are linear [109]. The matrices describe the change to the polarization state as the light with a certain polarization vector propagates through the optical element. The two fields that are measured on the two lock-in amplifiers are a mix of E_x and E_y , which have




Figure 2.7: Example of a Fourier filtering applied to an experimentally determined H_z field, where we filter out the unwanted frequencies, such as the inner ring, which is the TM mode.

undergone a change in polarization. Hence the relation between the measured fields and the electric field components can be described with a Jones matrix as:

$$\begin{pmatrix} E_x \\ E_y \end{pmatrix} = J \begin{pmatrix} \text{LIA}_1 \\ \text{LIA}_2 \end{pmatrix}, \quad (2.6)$$

where J is a generic Jones matrix. Since we are only interested in a polarization transformation, and not in polarizing the light using a polarizer for instance, we have opted to employ the most general Jones matrix for a polarization transformation, which is that of the elliptical phase retarder. This Jones matrix is described as

$$J = e^{-\frac{i\eta}{2}} \begin{pmatrix} \cos^2 \theta + e^{i\eta} \sin^2 \theta & (1 - e^{i\eta}) e^{-i\phi} \cos \theta \sin \theta \\ (1 - e^{i\eta}) e^{i\phi} \cos \theta \sin \theta & \sin^2 \theta + e^{i\eta} \cos^2 \theta \end{pmatrix}, \quad (2.7)$$

where η represents the relative phase retardation between the fast and slow axis of a (virtual) waveplate, θ represents the angle of the fast axis with respect to the horizontal, and ϕ represents the circularity. Any polarization transforming element can be described using this matrix for certain values of those three parameters. For some combination of the variables, it should be possible to transform the two measured fields into E_x and E_y . Using the knowledge of what the fields should look like allows us to obtain the optimal values for these three variables. In Fig. 2.6 we show the Fourier space of the measured fields on both lock-in amplifiers before any virtual waveplate is applied, and after, where it is separated into the two in-plane electric field components. The difference here between before and after is clear: while the fields initially detected on the LIA's have no obvious minimum, the fields after applying the Jones matrix exhibits clear minima on the axes instead. With this method we are able to retrieve the E_x and E_y fields from the initially detected fields.




Figure 2.8: Illustration of the algorithm used to determine singularity positions for experimentally measured data. Image shows a zoomed in section of the measured phase, with the 4 numbered pixels showing the way the loop is constructed.

2.5.1. FOURIER FILTERING

When measurements are performed, there is still some TM light remaining in the cavity, even though it might not be (efficiently) trapped by the photonic crystal. But we are interested in random fields, and the TM contribution is decidedly not random. Consequently, it will need to be excluded from the analysis. This remaining TM light can be observed by considering Fig. 2.7: besides the expected distribution for the TE mode, the remnants of a second ring with a smaller radius can additionally be observed.

This difference in radius stems from the different propagation constants for the different modes: for the TE mode, a propagation constant of $k_{TE} = 11.5 \mu m^{-1}$ is found, while for the TM mode a value of $k_{TM} = 7.78 \mu m^{-1}$ is obtained. These propagation constants can also be given in terms of the effective refractive index for the different modes: the TE mode has an effective index of $n_{eff,TE} = 2.84$, while the TM sees an effective index of $n_{eff,TM} = 1.92$. The propagation constants of the TE and TM modes have not been analytically determined, but rather through the use of a 1D mode solver for slab waveguides, which looks for solutions to the Maxwell's equations [110].

It is therefore straightforward to get rid of the unwanted TM light through filtering in Fourier space, as is depicted in Fig. 2.7. By filtering out all of Fourier space besides a section around the radius $k_{TE} = 11.5 \mu m^{-1}$, we are able to get rid of all the unwanted frequencies that are unrelated to the TE mode. For this we use a Gaussian profile with a certain width σ , which serves as a way to soften the filter edges. A typical value for σ of 4 pixels was used in our filter. Subsequently transforming the filtered Fourier space data back to real space leaves us with only the TE field.

2.6. DETERMINING SINGULARITY POSITIONS

Now that we have also explained the *what* part of the near-field microscopy, and how we are able to measure selective fields, we conclude this chapter with a small part of analysis methods. Since singularities serve as a common object of investigation across most chapters, we will describe how we detect them experimentally. As described in Chapter 1, phase singularities are effectively 0-dimensional points in space, and are infinitesimally small. Then, one could wonder, how it is possible to see them with our setup anyway? Even with our sub-wavelength resolution microscope we "only" have a resolution of around $15nm$. For this we can return to Eqn. 1.13, which states that if one takes a closed loop somewhere in the field and integrates the phase over this loop, the result has to be equal to 2π time the total topological charge that the loop encloses. So if exactly one singularity with charge $q = +1$ is enclosed, then the result of the integral should be 2π . And if the loop contains no singularity, then the result of the integral is simply zero. So implementing this method for the measured field, we should be able to locate the singularities: continuously making loops and determining where the integral becomes non-zero, indicating an enclosed singularity. Of course care needs to be taken to not to make the loop too big, otherwise it might enclose two singularities with opposite charge. In that case the integral evaluates to zero, while still enclosing singularities.

Since our measurements do not consist of a continuous map, but rather a grid with pixels, all with their own phase value, the above procedure needs to be discretized. Discretizing this integral to a loop of N pixels yields

$$s = \frac{1}{2\pi} \sum_{i=1}^N \Delta\phi_{i+1,i}, \quad (2.8)$$

where position $N + 1$ is equal to position 1.

In Fig. 2.8 an example is shown of a 2×2 loop enclosing a singularity. It depicts a (zoomed in) image of the phase, which was determined experimentally, with a singularity visible in the middle. The 2×2 grid of pixels that contain this singularity are numbered one through four. Unfortunately, simply defining $\Delta\phi_{i+1,i} = (\phi_{i+1} - \phi_i)$ will not work, and will simply always result in zero, since all terms cancel each other out. Instead, we need to ensure that every individual difference in phase stays mapped on the interval $[-\pi, \pi]$ (or $[0, 2\pi]$ of course). This can be understood in the following way: if there is a singularity present, then clearly the phase field around that point will span the full phase range, and hence somewhere it has to cross over from $-\pi$ to π . Then necessarily, one of the $\Delta\phi_{i+1,i}$ has to cross that threshold, which needs to be taken into account when computing the phase difference between those two points. In Fig. 2.8, we can observe that between pixels 3 and 4 a phase difference larger than π is found. In order to remain in the $[-\pi, \pi]$ interval, the formula can be designed in the following way:

$$\Delta\phi_{i+1,i} = [(\phi_{i+1} - \phi_i + \pi) \bmod 2\pi] - \pi, \quad (2.9)$$

where first π is added to the phase difference and taking the term modulo 2π in order to shift it to the interval $[0, 2\pi]$, before subtracting π once more to get it into the desired interval.

The above is still defined for a pixel loop of arbitrary size, so how do we pick this size for our data? Ideally, we would like to take this loop as small as possible, where the smallest possible loop is 2x2 pixels. For larger loops there exists a risk that we enclose multiple singularities. Thankfully, singularities are infinitesimally small, and as such are much smaller than any pixel we can measure, and hence must be contained in a 2x2 loop. The above operation is then performed for every 2x2 pixel grid in the data, and we then know that there exists a singularity in that 2x2 pixel grid if the result is not equal to zero. Unfortunately, there is no way of knowing where exactly the singularity is positioned, so we define it to be in the upper left pixel of the grid. In this way we can map the positions and charges of the singularities that are present in the field.

3

DIFFUSION AND LIFETIME OF PHASE SINGULARITIES

Clouds are not spheres, mountains are not cones, coastlines are not circles, and bark is not smooth, nor does lightning travel in a straight line.

Benoît Mandelbrot³

In non-monochromatic random light fields, phase singularities move in time. Here we analyse how they move. By tracking the singularities in time and determining their diffusive properties we show that optical phase singularities in random light fields exhibit anomalous diffusive behaviour, signifying a deviation from regular Brownian motion. We also reveal a clear difference between two different families of singularities, affecting both the diffusion speed, as well as the lifetime of the singularities.

Finally, we investigate one specific model for anomalous diffusion that could give us insight in the physical mechanisms at play.



Figure 3.1: Simulated instances of 2D random walks. Every coloured trajectory is a different instance of a random walk.

3.1. INTRODUCTION

Non-monochromatic electromagnetic fields evolve in time, and their phase singularities will, in general, move as a result. This movement of the singularities can be visualised and studied by tracking them as they move, and studying their movement. As stated in Chapter 1, singularities show many parallels to real particles [111], so we can employ metrics from statistical physics that are normally used to analyse particle behaviour to analyse the movement of the singularities. One such avenue to explore is to consider their diffusive behaviour. Diffusion is often taken to mean the movement from a place of higher concentration to a place of lower concentration. This is what is known as Fick's first law [112]. While it is usually used in the context of ions or molecules, it can, for instance, be also applied to the spreading of ideas [113, 114] or finance [115, 116] as well.

When considering the movement and diffusion of particles, one of the simplest models that one can use to describe the observed behaviour it is that of random walks. In this model, a particle simply takes steps in a random direction every time step, and is also commonly referred to as a drunkard's walk [117]. The simplest case of a random walk is a particle on a 1D lattice, where the particle can simply take either a step to the left, or a step to the right with equal probability. This case is the most basic version of a random walk. In 2D you it can take a step in a random direction, with every direction being equally likely. In Fig. 3.1, we show an example of random walks, where each colour represents a separate instance of a simulated random walk. These random walks form the heart of Brownian motion, which describes the behaviour of randomly moving particles in a medium [118].

Since the steps are random every time, for every step the particle takes in one direction, it will on average also take a step in the exact opposite direction. This then immediately implies that on average, its mean displacement from the starting position is zero. But not every particle ends up back at the start, of course, since it is all up to chance. Instead, their positions after a certain amount of steps will form a Gaussian distribution, with the mean centered on zero. The width of this Gaussian distribution increases with the amount of steps taken, as the particles are able to move further away from the starting point.

But while the mean displacement, which is equal to the mean of the Gaussian distribution, might be zero, its square is not. The mean squared displacement (MSD) is not determined from the mean of the distribution, but rather its variance, which is non-zero for finite time or number of steps. Since this variance increases as a function of amount of steps taken, the MSD also increases. In dynamic systems, these steps are taken as time increases, and can be snapshots at certain times. So the variance, and thus the MSD are functions of time in these systems. In the case of a purely random walk, the MSD scales linearly, and can be given as

$$MSD(t) = 2dDt, \quad (3.1)$$

where d is the dimensionality of the random walk, D is the diffusion coefficient, and t is time [119, 120]. By fitting the gradient of the MSD curve, the diffusion constant D can be obtained, which depends on the dimensionality of the system. This is a solution to what is known as Fick's second law.

However, an interesting situation occurs when the scaling of the MSD in time turns out to not be linear, a situation called anomalous diffusion. When this happens, the MSD scaling typically changes to $MSD \propto t^\tau$, with $\tau \neq 1$. In that case, it can either be sublinear (subdiffusive, $\tau < 1$) or superlinear (superdiffusive, $\tau > 1$). Sublinear diffusion ($\tau < 1$) can occur when the particle encounters restrictions on its movement, which can for instance be a result of crowding effects, or geometric constraints [121]. Superlinear behaviour ($\tau > 1$) happens when the particle tends to move in the same direction for longer periods of time, as opposed to moving around randomly. This superlinear diffusion can occur for instance as a result of active transport [122], or when the step-size distribution has a heavy tail. The latter case is also called a Levy flight, and happens in nature for instance when bees are searching for flowers [123, 124]. Here the bee performs a random walk in a certain area with flowers, and once the bee decides to look elsewhere, they will fly in a (fairly) straight line towards a next flower patch. Along with the above examples, anomalous diffusion has been observed in a plethora of systems, including ultra-cold atoms [125], interstellar medium [126], and optical systems [127]. But the examples are not limited to physics only, and appear for instance in biology [128–131] and finance [132–134] as well, making the concept of anomalous diffusion broadly applicable.

In this chapter we investigate the dynamics of phase singularities by tracking them in both time and space, and show that they do not behave linearly, and hence exhibit anomalous diffusive behaviour. Additionally we investigate the lifetime of phase singularities to investigate how the singularity population declines as a function of time.




Figure 3.2: Trajectories of measured phase singularities. Each colour represents the trajectory of one singularity. The starting point of each trajectory has been shifted to the centre of the image.

3.2. TRACKING OF SINGULARITIES

In order to investigate how the singularities move in time, we need access to their trajectories. And to gain access to their trajectories, we are required to be able to track the singularities in time. This is done by mapping a series of field distributions for increasing time by varying the delay line in the setup (see Section 2.4.3). For each time step the positions for the singularities are determined (see Section 2.6), and with that we observe how the topological skeleton¹ of the field changes in time as well.

In order to compute dynamic quantities, a tracking algorithm is employed that compares the singularity positions of two consecutive measurements and determines where each singularity moved to in the second measurement. It takes into account charge (as charge should not change in time), which pairs have annihilated or have been created. Additionally, it attempts to resolve conflicts where it might be ambiguous which singularity moved to a certain position when there are multiple candidates. When it really is unclear which singularity is which, and no consistent resolution is found, the singularities under consideration are removed from the tracking instead. Finally, since we only have a finite scan range, it is possible for a singularity to move outside the measured area, or for a singularity to walk into it. These cases are also both removed. In this way, the data we obtain from the algorithm is kept as clean as possible, where we are confident that a measured path is indeed the path of an individual singularity.

¹Knowing the positions and charges of the singularities can be used to determine the general structure of the field [135, 136].

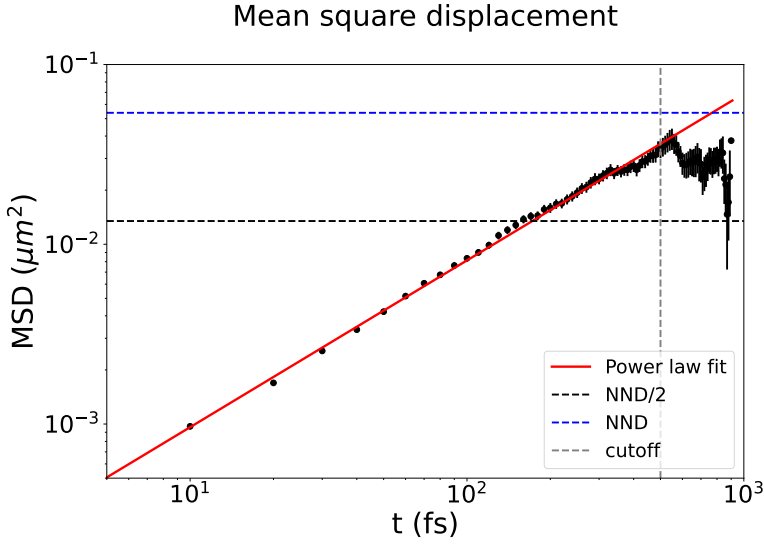


Figure 3.3: Mean square displacement for phase singularities as a function of time since their creation, based on approximately 2400 individual trajectories. The horizontal dashed black line signifies half the nearest neighbour distance (NND), and the horizontal dashed blue line is the nearest neighbour distance. The vertical dashed gray line indicates the point where so few singularities are left ($<2\%$) that the statistics are deemed unreliable. The red slope is the best fit based on an exponential ($f(t) = at^b$, with a and b variables). The exponent is found to be $b = 0.928 \pm 0.007$, which is sub-linear, and hence the singularities diffuse anomalously.

3.3. DIFFUSION OF PHASE SINGULARITIES

In order to investigate the diffusion of phase singularities in time, we have performed a measurement on a $15 \times 15 \mu\text{m}$ area in a chaotic cavity (see Section 2.4.2), using a spatial step size of 15 nm , and a temporal step size of 10 fs , over a temporal range of 1 ps . This measurement area typically contains on the order of 2400 phase singularities. To get an impression of the behaviour of the phase singularities, we have plotted several singularity trajectories in Fig. 3.2. Curiously, the trajectories exhibit a similar behaviour as the random walks shown in Fig. 3.1. It is therefore natural to investigate their movement in the same way as would be done for random walks: by computing their mean squared displacement.

One curious thing with our time-resolved data is that we are seemingly unable to simultaneously resolve both polarizations. We are able to select both the E_x and the E_y field, which can be inferred from the k-space, but not both at the same time. This is possibly a result of an elliptical tip, leading to polarization mixing in the fibre. So for the purpose of this analysis we have computed the E_x and E_y fields separate from each other, and from this computed the H_z field. While this may not be ideal, there is no reason to assume that the resulting fields are not the real fields.

Fig 3.3 shows the computed mean squared displacement as a function of time for all phase singularities combined on a double-log scale. The black dots indicate the data

points. We observe the MSD increasing as a function of time, with the error bars increasing in time as well, resulting from both the occasional annihilation of singularities as well as singularities leaving the finite scan range area: since we track singularities from birth to death in order to compute this curve, only long-lived singularities are left at high values of time, and hence the population starts decreasing in time, resulting in less statistics, and hence larger error bars, at higher times. Furthermore, we have introduced two horizontal dashed lines, indicating a diffused distance of half the nearest-neighbour distance (black) and the nearest-neighbour distance (blue). These values can be particularly interesting, since half the nearest-neighbour distance (NND) is the average distance a singularity has to travel before encountering another singularity, possibly leading to annihilation. The vertical dashed gray line is a cutoff, beyond which the statistics are deemed to be unreliable, because there are too few singularities left alive ($< 2\%$). While we do show the data points after the cutoff for completion, we will not discuss the behaviour for these long times, nor will they be taken into consideration for fitting purposes.

At first glance, the data points appear to form a straight line, which makes it natural to fit a power law to the data. The sloped red line is the best fit based on a simple power law function given by at^τ , where a is the slope of the curve, and τ the power, which fulfils the role of diffusion exponent. From the fit we obtain a diffusion exponent of $\tau = 0.928 \pm 0.007$, which shows slightly sub-linear behaviour, indicating a deviation from regular Brownian motion.

Interestingly, the diffusion shows sublinear behaviour, instead of linear behaviour. Sub-linear diffusion is indicative of some type of restriction on their movement [137]. Please note that the nonlinear mean squared displacement curve cannot be explained simply due to uncertainty in position when determining the position of the infinitesimally small singularities. While our position determination is indeed limited by the size of our pixels, we assume that their true position is distributed with a mean centered around the determined point, since they are distributed isotropically [24]. And even if there were to be a bias direction, it would not influence the power of the MSD, but only the slope. Hence we conclude that the sub-diffusive behaviour is inherent to the system of singularities in a 2D random field.

STEP SIZE ANISOTROPY

Another anomaly that we have found is that there is a discrepancy in the step size of the singularities in the two spatial directions x and y . We find that the average step in the x -direction is about 1.5 times larger than those in the y -direction: whereas the average step size in the x -direction is about 15 nm, the average for the y -direction is only 10 nm, based on a total of 31791 recorded individual steps. The cause of this is unclear as of yet. Using the code to analyse simple simulated random walks results in perfectly equal step sizes in all directions, as it should. Rotating the measured fields by 90 degrees additionally switches the step sizes accordingly. Furthermore, we have taken care that the pixel sizes are equal in both directions. As such, we have taken care to exclude that our analysis code is the cause of this discrepancy.

It is also highly unexpected given that the distribution of the singularities is expected to be isotropic. We have confirmed that the k-space distribution for H_z is a perfect circle with radius $k = |\mathbf{k}|$, and its pair-correlation function shows exactly what is expected of an isotropic random wave field [24]. So this field is the actual H_z field.

We have also checked if this is possibly introduced due to the drift compensation that is performed (see supplementary at the end of this chapter), since the drift is anisotropic. Having compared the anisotropy both with and without drift compensation, we find no difference between them, concluding that drift compensation is not the origin of this anisotropy.

Another possible cause for this anisotropy could be stemming from the time evolution of the field. If we consider the field at a too early time, then the laser pulse in the signal branch (see Section 2.4.3) is still partially travelling through the cavity from the entrance, and hence the field is not fully randomized yet at this time. This could be observed from the k-space distribution, where the k-space distribution would exhibit a non-uniform amplitude distribution. While we are able to observe this effect for low times (< 200 fs), computing the step sizes for $t \geq 500$ fs reveals no change in the observed step sizes: the anisotropy remains unchanged, while the amplitude in k-space has become uniform. As such we conclude that the anisotropy is not caused by a not yet fully randomized field.

So with everything indicating this being in the data, and not being introduced by our analysis, we have to conclude that there is an unknown physical mechanism at play here causing this anisotropy. The sole thing that could offer an explanation for this is the sample itself. The chaotic cavity that was used (see Chapter 2) is bigger in one direction than the other. However, that still raises the question how that could possibly influence the step size.

A clue to its origin might be found when considering the distribution of the steps, instead of simply the mean. While the steps in the x-direction are larger than those in the y-direction, there are approximately equal amount of steps in the positive and negative direction, and the mean step sizes are equal for both positive and negative directions: a total of 16348 steps were taken in the positive direction, and 15443 in the negative direction, both with a mean size of approximately 15 nm per step. On the other hand, when considering the y-direction, the step size is on average smaller, but there is an additional anisotropy present here: there are more steps in the positive directions, but the steps themselves are smaller: 18379 steps were taken in the positive direction with a mean step size of 8.6 nm. Conversely, the negative direction has fewer steps, but larger on average: only 13412 steps were recorded in the negative direction, but with a mean step size of 11.8 nm. This results in the average displacement still being equal between the positive and negative directions. When computing the above statistics without drift correction, we find that the mean step size in the y-direction is once again equal in magnitude (about 10 nm), but the amount of steps are unequal for the two directions, leading to a net displacement in that spatial direction instead. So while this does not explain the discrepancy in step sizes, it does hint at there being an underlying mechanism leading to this anisotropy.

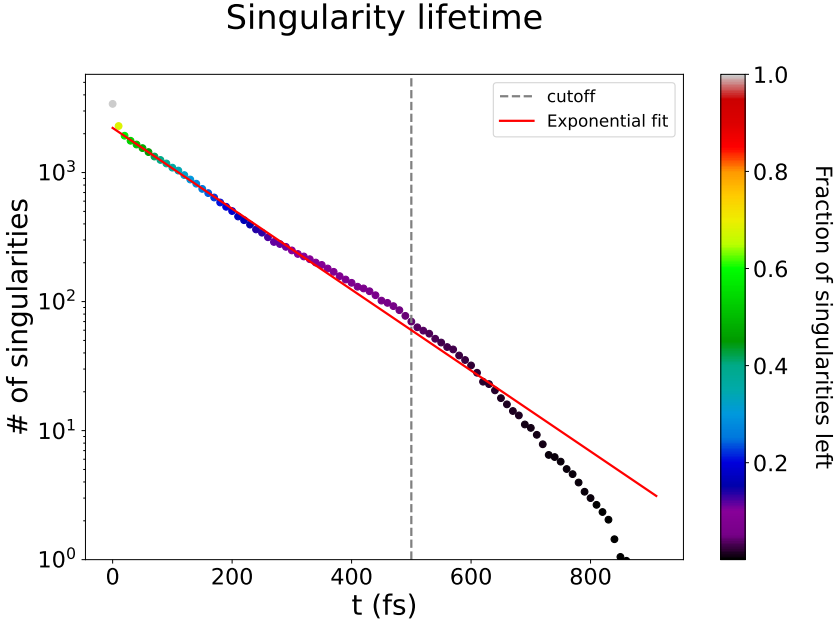


Figure 3.4: Number of singularities as a function of time since their creation, with a fit of the form $n_0 e^{-t/\tau}$, where τ is the average lifetime, and n_0 the initial population. The colourbar indicates the fraction of singularities remaining. After the initial decline it seems to overall match an exponential decay model. From a simple exponential fit to this exponential regime we determine an average lifetime of 142 ± 2 fs.

3.4. LIFETIME OF PHASE SINGULARITIES

Phase singularities are created in pairs, and can subsequently be annihilated in pairs as well (see Chapter 1). As a result, phase singularities do not live forever, but rather exist for a finite amount of time. As such, a natural question that arises is "how long do singularities exist for?". Analogous to Ref. [43], where their persistence as a function of wavelength shift was measured, we can measure how long they exist by tracking the decay in population as a function of time.

The way the time dependence of the population is measured is by gathering all the trajectories that started and ended within the measured time range that additionally stayed within the spatial boundaries of the scan range. From those that existed at earlier or later times, or left or entered the boundary during their lifetime we cannot determine how long they ultimately lived, and hence are discarded. In this way, only the trajectories that we have been able to follow for their entire duration are selected. Finally, we have also taken into account the finite size of the measurement area. The longer a singularity lives for, the higher the likelihood of it moving outside of the measurement area, and hence being discarded. This could then result in the lifetime being biased towards lower lifetimes. We have employed a box size correction for this as defined on page 48 of Ref.[138].




Figure 3.5: Schematic illustration of the two types of families of singularities. (a) shows faithful singularities, which are created and annihilated with the same partner. (b) shows unfaithful singularities, which are created with one partner, but annihilate with another singularity. The green circles indicate a creation event, while the red crosses indicate an annihilation event.

In Fig. 3.4 we show the phase singularity population as a function of time on a log scale. The data points are displayed as coloured dots, where the colour of the data points indicates the fraction of singularities left. Here the vertical dashed gray bar indicates a cutoff value where there is a statistical insignificant amount of singularities left ($< 2\%$ of initial population). A few things stand out immediately: there is a steep drop in singularity population right at the start (< 30 fs). After this initial decline, the curve turns less steep and appears to form a straight line, which implies an exponential decay.

Because the region $t \leq 30$ fs and the region $t > 30$ fs have such different slopes, we investigate them separately. For both an exponential of the form $n(t) = n_0 e^{-t/\tau}$ has been fitted, with $n(t)$ the population as function of time, n_0 the initial population and τ the lifetime. When fitting the first region, an average lifetime of 32 ± 6 fs is found. Fitting the exponential model to the second part of the curve yields an average lifetime of 140 ± 4 fs instead. This fit to the second region is indicated by the red slope. So it seems that if the singularities manage to survive the first few steps, their life expectancy increases drastically afterwards. If we simply fit the single exponential model to the entire dataset, we find an average lifetime of about 94 ± 5 fs for both regions combined.

3.5. FAITHFUL AND UNFAITHFUL SINGULARITIES

It has already been shown that there is an interesting distinction to be made between types of singularities, besides their charge [44, 139]. This division is based on the way the singularities are annihilated: specifically, whether they annihilate with the singularity they are created with (faithful), or another singularity (unfaithful) [43]. A schematic illustration of the two different families is shown in Fig. 3.5. It has already been shown that these two different families show different behaviour when considering their persistence under wavelength changes [43], so it is natural to assume that this division might

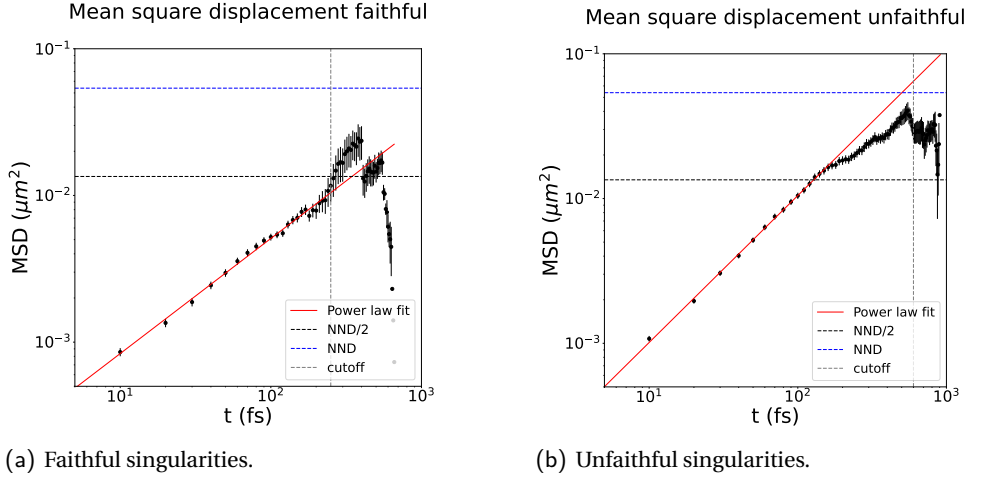


Figure 3.6: Comparison of the mean square displacement (MSD) between faithful (a) and unfaithful (b) singularities, with a fitted exponent of the form $f(t) = at^b$. For the faithful singularities we find a fitted exponent of $b = 0.785 \pm 0.013$, and they seem to not diffuse much further than half the nearest neighbour distance. The unfaithful singularities clearly diffuse further than half the nearest-neighbour distance a fitted exponent of $b = 1.015 \pm 0.010$ is found, which is slightly super-linear.

have an effect on their lifetime as well. Since faithful singularities are annihilated with their birth partner, it is intuitive to assume that they are never separated too far from their birth partner. The further the two are separated, the higher the chance of them colliding and annihilating with another singularity, making them unfaithful instead. For this reason we might also expect faithful singularities to live shorter than their unfaithful counterparts, using a similar reasoning as above. Unfortunately, this reasoning does not give an intuitive feeling for the way they diffuse in time. Therefore, we determine the diffusive properties for both singularity families separately.

3.5.1. DIFFUSION

Access to the individual trajectories of every singularity enables the determination of whether they are created and annihilated with the same partner or not. In this way we can divide them into their respective category based on the trajectory data. By dividing the phase singularities up into two families, we are able to investigate their diffusive behaviour separately. Fig. 3.6 shows the mean squared displacement curve of faithful and unfaithful singularities. It is immediately clear that there are qualitative differences between the two types of singularities. The most striking difference is the maximal mean squared displacement reached by faithful singularities versus unfaithful ones: the faithful ones only barely reach half the NND before reaching the phase singularity population cutoff where only few singularities are left. On the other hand, unfaithful singularities are actually able to reach the NND. Both families of singularities exhibit an MSD described by a power law in time.

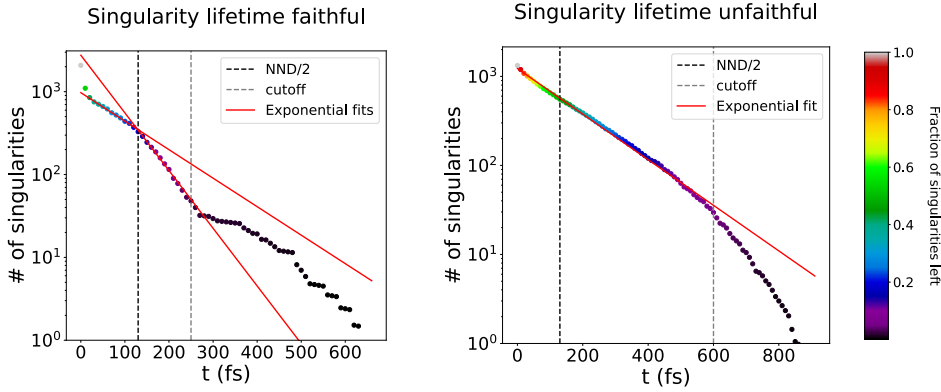


Figure 3.7: Comparison of lifetimes between faithful (left) and unfaithful (right) singularities, with a fit of the form $n_0 e^{-t/\tau}$, where τ is the average lifetime, and n_0 the initial population. The faithful singularities show a rapid drop in the first 20 fs, which then stabilizes briefly before dropping more rapidly again after having diffused about half the next nearest neighbour distance. We have fitted two separate exponentials (depicted as red slopes) for both $30 \text{ fs} \leq t \leq 140 \text{ fs}$ and $140 \text{ fs} \leq t \leq 250 \text{ fs}$. The first of these fits yields an average lifetime of $126 \pm 3 \text{ fs}$, while the second fit yields $62 \pm 2 \text{ fs}$ instead. The unfaithful singularities show a more steady decline on the other hand, and is well described by a single exponential, which is also depicted by the red slope. From the fit we estimate an average lifetime of $170 \pm 2 \text{ fs}$.

We can therefore fit a general power law of the form $f(x) = ax^\tau$, with the proportionality constant a and power τ as the two free fit parameters and compare the results. Faithful singularities follow a power law with an exponent of $\tau = 0.785 \pm 0.014$. But the unfaithful singularities follow a power law that is close to linear behaviour $\tau = 1.015 \pm 0.010$. Furthermore, there seem to be two regimes for the diffusion here: for about the first 100 fs, the diffusion is almost linear, even slightly super-linear, with the diffusion clearly slowing down after. Curiously, this slowing down happens right after the average diffused distance is half the nearest-neighbour distance, at which they should statistically start coming across other unfaithful singularities more often. So while it is unclear why this slowing down happens, it seems to be a physical process. We thus find a difference not only in the diffused distance, but also in the power law for the two different families, with the unfaithful singularities showing linear, or even slightly super-linear diffusion.

Another clear difference is that the threshold of 2% of the initial singularity population remaining is clearly reached at different times for the two types of singularities: the cutoff for the faithful singularities is at a much earlier time than that of unfaithful singularities. For faithful singularities is around 250 fs, while the cutoff for unfaithful ones lies around 600 fs instead. This already gives an indication that there is a strong difference in the lifetimes of the two families of singularities as well.

3.5.2. LIFETIME

Like for the diffusion, we can now also determine the lifetime of the two singularity families separately. Fig. 3.7 displays the singularity population as function of time for (a) faithful and (b) unfaithful singularities. Here time is the time that has passed since the creation of a singularity, and not the absolute time. The dots are the data points where the colours of the dots represent the fraction of singularities that are left after a certain time since their creation. The black vertical dashed line indicates half the nearest-neighbour distance. Finally, the red slope is a simple exponential fit of the form $n_0 e^{-t/\tau}$, where τ is the average lifetime, and n_0 the initial population.

For the unfaithful singularities we observe behaviour that is similar to that of all singularities combined. Only for the first 30 fs do we see a slightly steeper slope than for $t > 30$ fs, while afterwards the population looks like a straight line. Fitting an exponential decay model to the data for unfaithful singularities reveals an average lifetime of 170 ± 2 fs.

For faithful singularities, however, the pattern appears to be less straightforward. First we see a rapid drop in the first 30 fs, where about two-thirds of the initial population annihilates. Between 30 fs and 120 fs we observe a slower decay, as indicated by the decreased slope. And from 120 fs onwards until the cutoff value at approximately 250 fs the decay rate increases once more. Interestingly, $t = 120$ also corresponds to half the nearest-neighbour distance, as indicated by the vertical dashed black line. Both the 30 fs to 120 fs regime, as well as the 120 fs and 250 fs regime have been fitted with an exponential function, which are depicted as red slopes. Fitting the simple single exponential for the faithful singularities leads to an average lifetime of 70 ± 8 fs, but is strongly influenced by the rapid drop-off at low times. As such, fitting a single exponential is clearly not representative of the behaviour, and instead an exponential decay model is fitted to each of the sections individually. Fitting an exponential to the the initial steep decline from $t = 0$ fs to $t = 30$ fs reveals an average lifetime of 20 ± 4 fs. In this time span about 60% of the singularity population has annihilated. Of course it is a bit bold to make an exponential fit to only three data points, but it still exhibits different behaviour from the rest of the data points. Furthermore, these data points represent the behaviour for the majority of the faithful singularities. As such we felt compelled to extract a lifetime for this group as well. For the region $t = 30$ fs to $t = 120$ fs, we find that the average lifetime is 126 ± 3 fs. Finally, for $t > 120$ fs we observe that the population decline quickens once more, for which we determine an average lifetime of 62 ± 2 fs.

The initial steep drop off in lifetime is easily explained, apparently a lot of singularity pairs that are created then immediately annihilate again the next measurement frame with their birth partner, leading to most faithful singularities perishing in one or two frames. The singularities that have not annihilated within this time frame likely separated a bit further from their birth partner, leading to a slower decay. The second region of steep decline, which is from $t = 120$ fs onwards, seems to start after the MSD of the singularities is approximately half the nearest-neighbour distance. At this distance diffused you are statistically likely to start encountering singularities from other pairs. A possible explanation for the population decline to speed up from this point onward is that the longer the singularities exist here, the higher the chance of recombining with a singularity of another pair, which would make them unfaithful. As such, the popu-

lation of faithful singularities could be self-selecting for shorter lived singularities. The stark difference in lifetime between the two families also immediately explains the difference in population cutoff times. This is a direct reflection of the difference in lifetime between the two types, since these cutoff values are determined relative to the starting population of the family. A lower lifetime means the amount of singularities also drops below the threshold faster.

We can compare the obtained results of the phase singularity lifetime with the results found for the persistence as a function of wavelengths as seen in Ref. [44]. Both changes in time and wavelength lead to an evolution of the electromagnetic field, resulting in movement of the singularities under change of the parameter. As such, it is natural to compare the results of both these parameters to each other. When considering all singularities combined, we observe a similar pattern as for the amount of singularities as a function of wavelength shift: a quick decline for the first couple steps, followed by a steady decay strongly resembling simple exponential decay. When considering the two families of singularities separately, we find a high amount of faithful singularities for low times, which matches the results for their persistence as well. So the results found for the lifetime of singularities matches the behaviour found for their persistence under wavelength shifts.

3.6. ANOMALOUS DIFFUSION

From the analysis above, we have found that phase singularities in a random field diffuse anomalously. Assuming that they would indeed be physical particles in a liquid, as might be expected from their liquid-like spatial distribution (see Section 1.3.3), when leaving charge out of the equation, they should diffuse exactly according to Brownian motion, i.e. linear diffusion. This raises the question what causes the anomalous diffusion.

When splitting the full ensemble of singularities into faithful and unfaithful, the following argument could be made for the found results. When being created, the singularities will, necessarily, move away from each other. For faithful singularities, they will need to find their original partner again. This post selection of only faithful singularities therefore suggests an effective attraction, which will slow down their MSD curve, leading to subdiffusive behaviour. In the case of unfaithful singularities, they are not bound to their original partner, and no such effective attraction exists. In fact, their post selection prohibits a return to their original partner, so that an effective repulsion might result, leading to super-diffusive behaviour instead. These effective attractions/repulsions are of course not real, but rather resulting from the selection criteria.

While Brownian motion is no longer able to adequately describe the found results, we would still like to be able to model the behaviour. There are three main models that are generally used to describe anomalous diffusion, specifically sub-diffusive behaviour. All of them release a certain assumption that is made when considering regular Brownian motion. Continuous time random walks (CTRW) lets go of the assumption that each step takes the same amount of time and generally works for trapping scenarios. In this case particles can get stuck for a varying amount of time in between steps, such as in energy potentials [140]. The second class of models deals with geometric constraints, that can

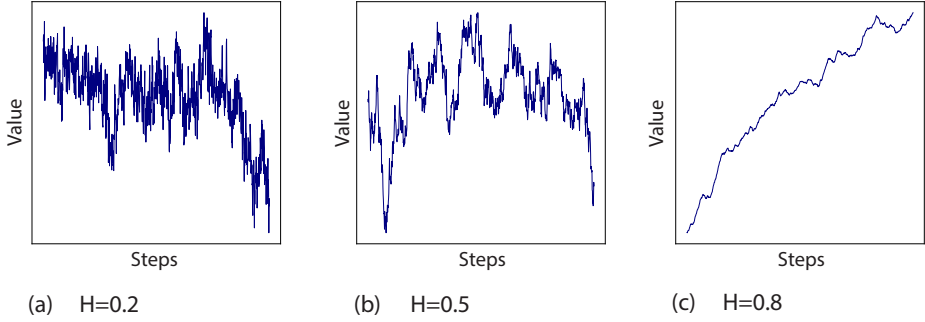


Figure 3.8: Examples of 1D trajectories of fractional Brownian motion for three different Hurst parameters. (a) shows the trajectory for a Hurst parameter of $H = 0.2$, where subsequent steps are anti-correlated. (b) shows the case of $H = 0.5$, which corresponds to regular Brownian motion. Finally (c) shows the trajectory for $H = 0.8$ where the steps are positively correlated.

limit the motions that are exhibited by the particles, such as in percolation clusters [141] and random walks on fractal structures [142, 143]. This tends to lead to sub-diffusive behaviour because dead-ends and other obstacles can frequently force the particle to go back the way it just came, leading to an anti-correlation in its movements. Finally we come to fractional Brownian motion, which completely lets go of the assumption that steps are independent from the previous steps. This type of behaviour can be observed in more complex, interacting systems exhibiting viscoelastic behaviour [144, 145]. Of these situations, the latter seems to be the most promising candidate for our data, since neither trapping events, nor geometric constraints are immediately obvious situations that would appear for phase singularities.

3.6.1. FRACTIONAL BROWNIAN MOTION

The one core assumption of regular Brownian motion is that every step is completely independent from every other step. Fractional Brownian motion (fBm), however, is a model to describe diffusion without that assumption, and as a result steps are allowed to be correlated. The fractal Brownian motion $B_H(t)$ is a continuous time Gaussian process and is defined via a Weyl integral, a fractional integral of white noise [146]:

$$B_H(t) = B_H(0) + \frac{1}{\Gamma(H + 1/2)} \left(\int_{-\infty}^0 [(t-s)^{H-1/2} - (-s)^{H-1/2}] dB(s) + \int_0^t (t-s)^{H-1/2} dB(s) \right). \quad (3.2)$$

Here, Γ is the Euler gamma function. H is the so-called Hurst parameter, which is a real number $H \in (0, 1)$ and the integrand $dB(s)$ is a white noise measure.

The parameter of interest in this is the Hurst parameter H , which determines the correlation between moves. From the integral defined above, one can easily determine that $H = \frac{1}{2}$ has to be a special case. For this specific value the first integral becomes identical

zero, and the second integral is just over the measure. This case corresponds to regular Brownian motion, where the moves are again uncorrelated.

When the Hurst parameter deviates from exactly a half, the behaviour starts to differ from Brownian motion. When $H > 1/2$ we find that moves become positively correlated. So when the previous step has been in a certain direction, there is a higher likelihood that the next step will be in the same direction as well. Conversely, when $H < 1/2$ we find a negative auto correlation between steps. So when the previous step was in one direction, there is a higher chance of the next one being in the opposite direction. The behaviour of fractional Brownian motion for three sample trajectories with different Hurst parameters is depicted in Fig. 3.8. Fig. 3.8 (a) shows the case of $H = 0.2$, where the steps are anti-correlated, (b) depicts $H = 0.5$, which is regular Brownian motion and the steps are uncorrelated, and (c) shows $H = 0.8$, where the steps are positively correlated. When comparing (a) and (c) to regular Brownian motion, it is clear that (a) has a much more rugged appearance, whereas (c) shows a much smoother curve instead. So the Hurst parameter is also used to describe the ruggedness of a curve, with higher values leading to a smoother curve.

Fractional Brownian motions have a lot of interesting properties, such as self-similarity ($B_H(at) \approx |a|^H B_H(t)$), and long-range dependence when $H > 1/2$, where all steps remain correlated to the initial step. Another thing to note is its fractal dimension: it has been shown that the Hausdorff dimension (a measure of roughness [147]) of the graph of $B_H(t)$ is equal to $2 - H$ [148], which then immediately implies that a lower Hurst parameter corresponds to a higher fractal dimension. In fact, when H approaches zero, its Hausdorff dimension approaches 2, corresponding to a surface. However, this only holds for the one dimensional case, as it is also known that in 2D, normal Brownian motion has a Hausdorff dimension almost surely equal to 2, where it has a dimension of 1.5 in the 1D case. This is expected, since 2D random walks, fractional or not, can reach every point on a 2D surface, and as such, its graph is simply a 2D surface. So it is expected that the fractional dimension of a 2D fractional Brownian motion matches its topological dimension, namely 2.

Fractional Brownian motion trajectories are defined through an integral as defined in Eqn. 3.2, which cannot readily be evaluated. Since fractional Brownian motion trajectories are a stochastic process, every instance of it is different, in the same vein that every instance of normal Brownian motion is different every time. Fortunately, it is possible to generate discrete instances of fBm through numerical approximations of the fractional integral. Three of such methods are Hosking's method [149], the Cholesky method [150], and the Davies-Harte method [151].

Fractional Brownian motion exhibits anomalous diffusion for values other than $H = 1/2$. In fact, it turns out that the diffusion exponent τ scales with the Hurst parameters as $\tau = 2H$ [152]. So for $H < 1/2$ we find sub-diffusive behaviour, for $H = 1/2$ we retrieve linear behaviour, and for $H > 1/2$ we have super-diffusive behaviour. This also makes sense when considering what the Hurst parameter actually means. For $H < 1/2$ the anti-correlation in consecutive steps leads to a slowing down of movement, and hence sublinear diffusion, similar to the case of geometric constraints. And for $H > 1/2$, the positive correlation in consecutive steps leads to the particle moving much quicker into a certain

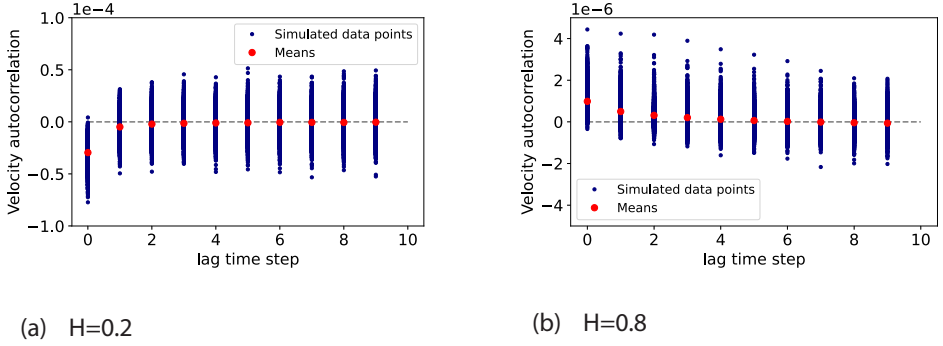


Figure 3.9: Auto-correlation function as a function of amount of steps taken for simulated fractional Brownian motion, for two different Hurst parameters. (a) shows the ACF for $H = 0.2$, which dips below zero for the first step, indicating a negative correlation, after which it quickly becomes uncorrelated. (b) shows the ACF for $H = 0.8$, which slowly drops from one to zero, indicating a much longer correlation length than the $H < 0.5$ case.

direction than it would if diffusing linearly. This then leads to super-diffusive behaviour. These two different possible correlations between steps can then also be interpreted in terms of the singularities under consideration. Unfortunately, for unfaithful singularities, there does not seem to be a clear interpretation for their behaviour in terms of Hurst parameters. For faithful singularities, however, a case can be made for them being experiencing an effective attraction towards their birth partner after initially moving away from each other. If we assume diffusive behaviour baseline, then it can be interpreted as having a negative correlation with respect to their initial trend, which would correspond to a Hurst parameter $H < 1/2$.

3.6.2. AUTO CORRELATION OF TRAJECTORY

A hallmark feature of fractional Brownian motion is the correlation between consecutive steps. These correlations can be quantified by looking at the auto correlation of the trajectories, where we look at the correlation of steps within one trajectory. This step auto correlation function (ACF) can be defined as

$$C_i = \langle s_i s_0 \rangle / \langle s_0^2 \rangle, \quad (3.3)$$

where s_i are the i th steps taken with respect to s_0 , and the brackets indicate a sliding ensemble average. Please note that this sliding ensemble average also averages over all possible initial steps s_0 . Of course per definition we find $C_0 = 1$. For a fully random walk, one expects that all steps are independent of each other, i.e. $C_0 = 1$ and $C_i = 0$ for all $i \neq 0$. For anomalous diffusion, this no longer has to be the case.

In this section we use the `yupi` python package in order to compute the auto correlation function [153]. The package computes the velocity auto correlation function, using the velocity vector instead of the displacement vector as used in Eqn. 3.3. However, when

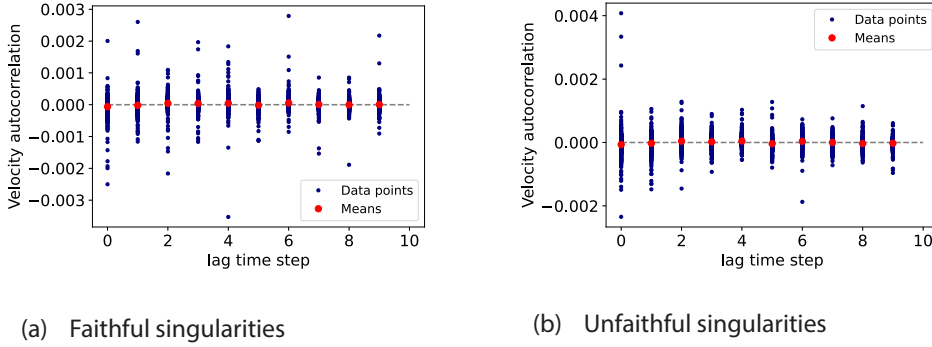


Figure 3.10: Auto-correlation function as a function of amount of steps taken for experimentally determined trajectories. (a) shows the ACF for faithful singularities, while (b) shows the ACF for unfaithful singularities. Neither shows a clear correlation function which matches the fractional Brownian motion case.

the time steps is taken to be unity, this is exactly the same as the step auto correlation as defined above. In Fig. 3.9 we present the computed auto correlation functions for two different Hurst parameters $H = 0.2$ (a) and $H = 0.8$ (b). They are representative of the behaviour of the ACF for the sub-diffusive ($H = 0.2$) and super-diffusive ($H = 0.8$) cases. In the case of a Hurst parameter $H < \frac{1}{2}$ (Fig. 3.9 (a)), the ACF dips into the negative values, signifying a negative correlation between steps, before quickly approaching zero as the steps increase. So in this case, the auto correlation is short-range negatively correlated, and increasingly uncorrelated for increasing number of steps. But when the Hurst parameter is $H > 1/2$ (Fig. 3.9 (b)), the ACF instead decays to zero over a much longer range. Thus, the auto correlation is positively correlated over a much longer scale, signifying its long-range correlated behaviour. In these cases the Hurst parameters are relatively far removed from $H = 1/2$, which makes their features in the ACF more pronounced. When considering values closer to $H = 1/2$, the ACF will still qualitatively exhibit the same behaviour, but less pronounced: for $H < 1/2$ the initial dip will be less prominent, and for $H > 1/2$ the correlation decays more rapidly. Since we know that the diffusion exponent and Hurst parameter are related as $\tau = 2H$ in fractional Brownian motion, and we have experimentally determined the diffusion exponent for the full ensemble of singularities to be $\tau = 0.911$, if fBM indeed properly described the motion of the phase singularities, a Hurst parameter $H \approx 0.45$ is expected. This value is much closer to $H = 1/2$, and as such we expect the ACF features to not be very pronounced for our data. The behaviour of the ACF, and not the exact values, being the hallmark for the two different cases is also why setting our time step size to unity in the determination of the velocity ACF is allowed. Using this unitary time step, we can determine the ACF for the experimentally determined singularity trajectories in order to explore the possibility that our data is described by fractional Brownian motion. Fig. 3.10 shows the auto correlation function for the experimentally determined trajectories, with (a) showing the ACF for faithful singularities, and (b) showing the ACF for unfaithful singularities. Neither families show a clear correlation function, with both fluctuating around zero, meaning no clear correlation seems to be

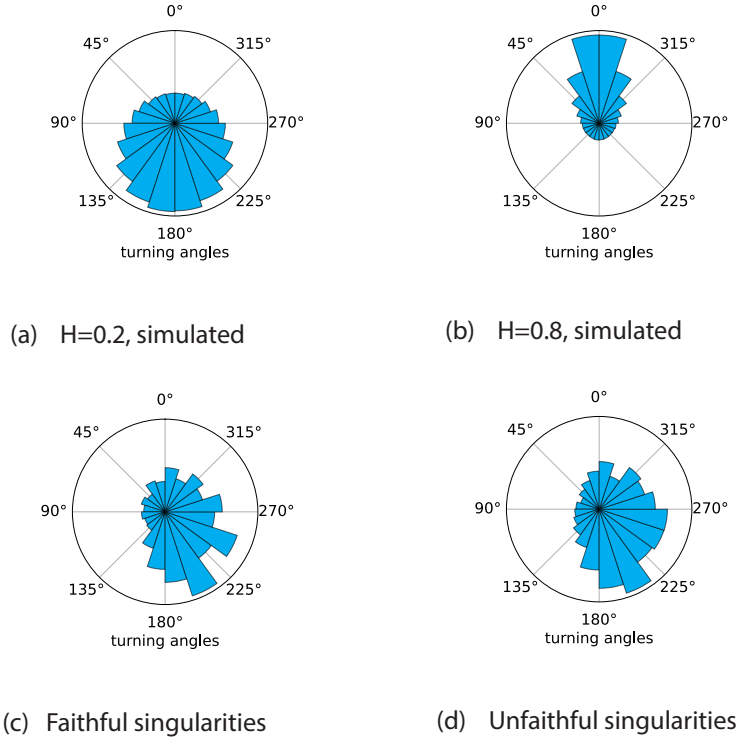


Figure 3.11: Turning angle distribution for four different cases. (a) shows the distribution for simulated fBM data with a Hurst parameter of $H = 0.2$. (b) shows the distribution for simulated fBM data with a Hurst parameter of $H = 0.8$. (c) and (d) show the distribution for the experimentally measured data. (c) depicts the turning angle distribution for faithful singularities, and (d) the distribution for unfaithful singularities.

present. Unfortunately, the spread around zero is also quite large, orders of magnitude more than for the simulated data. Because of this large spread, we cannot conclusively determine whether or not fractional Brownian motion is able to model the system properly. Thankfully, there are more ways to test the model for compatibility with the experimentally measured data. Once such way is investigating correlation between steps by considering the relative angle of consecutive steps, referred to as turning angles.

3.6.3. TURNING ANGLES

A more detailed way to investigate the deviation from Brownian motion is by considering the turning angles of the path. The turning angle is the angle a step takes with respect to the previous step. So taking a step in the exact direction of the previous step would give a turning angle of zero degrees, while the exact opposite direction would yield 180 degrees. It is similar to the ACF that was used before, but here only the angle with respect with the previous step is considered. When considering a purely random walk, one would expect

every angle to be equally likely for the next step, and hence the distribution is uniform over the entire angle range.

In Fig. 3.11 we have plotted the turning angle distribution for the measured trajectories of our singularities (c & d), as well as a comparison to the different fractional Brownian motion cases (a & b). Fig. 3.11 (a) shows the turning angles for fractional Brownian motion with a Hurst parameter of $H = 0.2$. The anti-correlation in steps is clearly visible as the majority of the steps are taken at an angle signifying the opposite direction from the previous step. Fig. 3.11 (b) shows the turning angles for fractional Brownian motion with a Hurst parameter of $H = 0.8$ instead, which clearly showcases the positive correlations in step direction.

In Fig. 3.11 (c) the turning angles for the experimentally measured data for faithful singularities is depicted. What clearly stands out here is that this angle distribution is neither uniform nor symmetric. There is a clear bias towards angles away from the previous direction, since most angles are between 180° and 270° . This does not seem to fit with fractional Brownian motion with a Hurst parameter $H < 0.5$, as the distribution is too narrowly focused when compared with Fig. 3.11 (a). However, it is consistent with subdiffusive behaviour.

Even more curious perhaps is the turning angle distribution for unfaithful singularities, which has been depicted in Fig. 3.11 (d). While the angle distribution is clearly broader than the distribution of the faithful singularities, the angles are also indicative of anti-correlated steps, which would yield subdiffusive behaviour. Interestingly, it looks like the bias is not evenly spread around 180 degrees, as one would expect it to be in the case of fractional Brownian motion. We find that there is also a bias towards the bottom right quadrant.

From the turning angle distributions we conclude that we find a clear deviation from Brownian motion. However, it also does not match the known distributions from fractional Brownian motion either. Hence from the turning angles we still cannot draw a clear conclusion with respect to the underlying physics, and more research is warranted.

3.7. CONCLUSIONS

In this chapter we have experimentally measured the non-monochromatic fields inside a chaotic cavity as a function of time by using a pulsed laser. From these time-dependent fields we have extracted the singularity positions and tracked them in time in order to study their dynamics. We have shown that the diffusion is non linear, but is sub-linear instead. Hence their motion cannot be described with regular Brownian motion.

Furthermore, there is a clear difference in behaviour between faithful and unfaithful singularities with unfaithful singularities diffusing faster and further than their faithful counterparts. Additionally, we have measured their lifetime, and found for unfaithful singularities the distribution of their lifetimes can be described quite accurately with an exponential decay. On the other hand, faithful singularities show a more complex decay behaviour.

As a potential model for the displayed singularity motion we have investigated fractional Brownian motion, which allows for correlations between consecutive steps. While promising, the computed velocity auto-correlation function does not give us conclusive evidence for or against the model. And by considering the turning angles, we see that it does not seem to match fractional Brownian motion either, but does offer more evidence against regular Brownian motion instead.

3.8. SUPPLEMENTARY INFORMATION

DRIFT COMPENSATION

One additional factor that could potentially impact the fidelity of the tracking and subsequent determination of the diffusive properties is drift. Since we measure consecutive frames in order to probe the time dynamics, it can take hours or even days to complete the entire measurement run. In that time, it is possible that for instance the sample shifts due to mechanical relaxation of screws, or that the tip does not return exactly to the starting point due to hysteresis. This drift can then influence the mean squared displacement as well. Consider the case of classical Brownian motion for instance. As noted, the MSD should be perfectly linear in this case. However, if drift is added on top of the random walk, then suddenly the walk is biased into a certain direction. This then leads to the particle moving quicker away from the origin than one would expect from Brownian motion, resulting in seemingly superdiffusive behaviour.

In order to prevent drift from skewing the determination of the diffusion exponent, we need to have a method to minimize its impact. In the data this is done by making the assumption that the mean displacement should remain zero. By then tracking the singularities in between consecutive frames, we can compute the mean displacement in between two frames for both spatial directions separately. If this is non-zero, then we assume that the mean displacement is caused by drift, and we then subtract it from the positions of the particles in order to determine what their positions should have been without drift. In this way we attempt to minimize the impact of potential drift in our measurement on the determination of the diffusion parameters.


4

MELTING OF SINGULARITY LATTICES

*Like reflections on the page
The world's what you create*

Wither - Dream Theater⁴

In two dimensions, three interfering plane waves with random initial phases always lead to a hexagonal amplitude pattern containing phase singularities in a regular pattern with the same symmetry. This ordered lattice of phase singularities subsequently "melts" as the amount of interfering plane waves is increased. In this chapter we numerically investigate the melting of these phase singularity lattices. We study this melting without changing the symmetry inherent to 3 plane waves in two ways. We show that broadening in longitudinal momentum, analogously to increasing temperature in molecular beams, does not lead to melting behaviour. However, introducing more plane waves in a way that does not break spatial symmetry abruptly, and simultaneously adding phase noise between these waves does lead to melting behaviour. We show that phase noise is a critical component for inducing phase transitions.



(a) 3 interfering plane waves

(b) 20 interfering plane waves

Figure 4.1: Comparison of the singularity structure that is formed by different numbers of plane waves with random initial phases interfering in two dimensions. The plane waves have equal amplitude and are equidistributed in k -space. The phase is plotted in false colour, and the light (dark) gray dots indicate the positions of positive (negative) singularities. Figure (a) shows three interfering plane waves, leading to a regular trigonal lattice structure. (b) shows 20 interfering plane waves, showing a disordered, seemingly liquid-like pattern instead.

4.1. INTRODUCTION

Of particular interest throughout all of physics is the phenomenon of phase transitions, where one state of matter transitions into another state. Especially discovering under which circumstances this occurs can give a lot of valuable insight into the behaviour of a system [154–157]. In previous chapters we have indicated that singularities can sometimes be treated as particles. In this chapter we will continue to entertain this notion and investigate if a phase transition can occur for phase singularity structures. It has already been shown that an effective interaction potential can be defined for phase singularities for two-dimensional random light [111]. When an effective interaction potential can be given, then it might also be possible to form a singularity crystal. As it turns out, three interfering plane waves with random initial phases in two dimensions will always lead to a hexagonal field pattern. When taking the singularity charges into account, this results in a trigonal phase singularity lattice, in fact forming a two-dimensional crystal. An example of such a trigonal lattice is shown in Fig. 4.1 (a).

Since we can form a lattice of singularities, representing a solid, crystalline state, can we then also "melt" this crystal to its "liquid" state? If we now consider the singularity structure that is formed for 20 interfering plane waves with random initial phases, as shown in Fig. 4.1 (b), the result is clearly not an ordered crystal anymore. So at some point, when increasing the number of interfering plane waves that interfere, a phase transition seems to have taken place.

In this chapter we numerically investigate the solid-to-liquid phase transition for singularity structures. We look into which variables are relevant to the occurrence of a phase transition. The work in this chapter was carried out in collaboration with bachelor student Luuk Crooijmans, and part of it was the subject of his bachelor thesis [158].

4.2. PHASE TRANSITIONS AND ORDER PARAMETERS

A usual way to describe and classify phase transitions is through the use of order parameters [159]. An order parameter is a quantity that distinguishes one phase from another: it is chosen such that it exhibits a sharp change in value at the transition. Additionally, depending on whether this change in order parameter value is continuous, or discontinuous at the transition point, the transition can be classified as first order (discontinuous) or second order (continuous) [160]. An example of such an order parameter is compressibility for water. When water is in its solid (ice) or liquid (water) state, it can be regarded as incompressible¹. Its volume won't change when the pressure is increased. However, in its gaseous state, it becomes compressible again. Hence the compressibility is an order parameter for the transition to its gaseous state, since it is normally zero, but becomes non-zero once the phase transition into the gaseous phase takes place. Similarly an order parameter can be defined for the transition from a liquid state to a solid state. However, for this case it is less straightforward and relies on for example correlations between atomic positions: whereas crystals possess spatial symmetries and the atoms do not tend to move around, liquids possess no such symmetries and the atoms are free to move around [163–166].

Generally, order parameters change as a function of a variable. This variable can be anything that affects the state of the system, with some common variables being, e.g., temperature and pressure [159]. A phase transition then occurs when a specific value for the variable is crossed and the system crosses over into its new state. This value is referred to as the critical value². It is at this value that the order parameter shows an abrupt change, either in its value (first order), or its derivative (second order). Ideally we would like to define an order parameter for the phase singularity network in order to quantify its phase transition from a solid lattice state (Fig. 4.1 a) to a liquid state (Fig. 4.1 b).

However, the usual way a phase transition from solid to liquid is quantified is based on the shear modulus. But since we are considering phase singularities of an electric field, they do not have a shear modulus to use as order parameter. Furthermore, since we consider monochromatic fields, we are not able to access dynamic quantities, and are restricted to snapshots of the field instead (see Chapter 1). This prevents us from using order parameters based on kinetics, such as those based on the mean-squared displacement [167]. Thankfully there are more tools to study phase transitions with besides order parameters, such as melting indicators, which are an established tool in condensed-matter physics [168, 169]. A number of indicators are available, all of which use a slightly different criterion derived from different characteristics of the different phases to determine whether you are in one phase, or the other.

¹It's not completely true, its bulk modulus is just so high that it's functionally incompressible under normal circumstances [161, 162].

²This definition is a bit simplistic, as there can exist multiple transitions in the system. In this case the critical value refers to the point where the system first undergoes a phase transition.

4.3. MELTING INDICATORS

Melting indicators are, in a sense, similar to order parameters, since they fulfil a similar role. But unlike order parameters, melting indicators cannot be used to determine the type of phase transition that takes place. However, melting indicators do show a rapid change around the point of the phase transition, and as such are helpful indicators for pinpointing a phase transition. We define the melting indicators, in analogy to order parameters, based on observables of the singularity structure: the radial distribution function $g(r)$ and the angular distribution function $c(\theta)$. The radial distribution function has been used extensively throughout this thesis (see Section 1.3.3). The angular distribution function can be seen as a direct companion to the radial distribution function $g(r)$. Instead of determining at what radial distance from a reference particle it is most likely to find another particle, this instead looks at which angles with respect to a reference angle (often the horizontal) it is most likely to find another particle [170]. Like $g(r)$, it can give valuable insight into the structure of the field. The angular distribution function can be defined as

$$c(\theta) = \frac{1}{N\rho} \left\langle \sum_{i \neq j} \delta(\theta - |\theta_i - \theta_j|) \right\rangle. \quad (4.1)$$

In a crystalline state the angles at which a nearby particle is likely to be found are sharp, since a lattice has a well defined structure. For a liquid on the other hand, there is no regular structure, and as a result the features are much less defined. From the radial distribution function and angular distribution function we can derive quantities that act as the melting indicators, since the the radial distribution function and angular distribution function give insight into the structure of the singularity network.

Another melting indicator is based on the singularity density. It is known from theory that in a liquid-like state of singularities, the density of vortices should converge to $\pi [\lambda^{-2}]$ [24]. For atoms it is known that the density of a solid is, in general (water being a well known exception), higher than the density of a liquid, so for our singularity structures, we expect the density of the crystalline state to be higher than that of the liquid-like state as well.

4.3.1. INDICATORS BASED ON $g(r)$

DIRECT METRICS

To start we consider two metrics that can be directly inferred from the radial distribution function. The first of these indicators is the position of the first (non-zero) minimum $r(g_{min})$, which is the first non-zero minimum following the first maximum. Note, however, that for a solid-like state any minimum between maxima should automatically be zero on account of the crystalline nature of the singularity distribution. Furthermore, the $g(r)$ depends also on the box size of the performed simulation. Were this size to be infinite, the $g(r)$ would theoretically consist of a series of delta peaks. Hence the size of the "valley" of zeroes between peaks also depends on the simulation size. Since we want to compare simulations with a finite size we pick the minimum here to be the first point

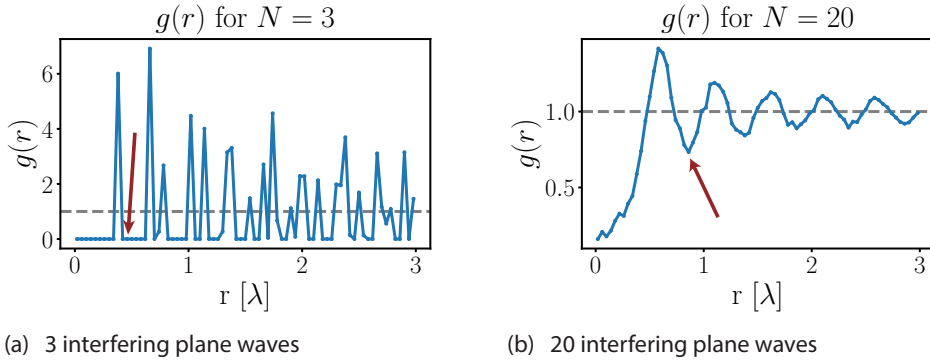


Figure 4.2: Comparison of the radial distribution function $g(r)$ for 3 interfering plane waves (a), and 20 interfering plane waves (b), both with random initial phases for all plane waves. (a) Shows a radial distribution function that resembles that of a crystal, with peaks at well defined distances, indicative of an ordered system. (b) Shows a $g(r)$ that resembles that of a liquid. In both plots the horizontal dashed gray line indicates $g(r) = 1$, corresponding to an uncorrelated spatial distribution of singularities. The arrows indicate the position of $r(g_{min})$.

4

where $g(r) = 0$ after the first peak. This can additionally be justified when considering the meaning of $r(g_{min})$ as the radius of the first coordination shell. This first coordination shell contains the nearest neighbours of a lattice site, and in the $g(r)$ these are contained within the first peak. As such defining $r(g_{min})$ such that it exactly contains this first peak is a natural choice.

The second is the ratio between the first minimum and first maximum $R \equiv \frac{g_{min}}{g_{max}}$ [171]. Fig. 4.2 depicts the radial distribution function $g(r)$ of the singularity network for both 3 interfering plane waves (Fig. 4.2 a) and 20 interfering plane waves (Fig. 4.2 b), respectively. For 3 plane waves we have a crystalline distribution, while for 20 plane waves the $g(r)$ is clearly liquid-like, and hence a phase transition has seemingly occurred as the amount of plane waves were increased. By considering Fig. 4.2, it is clear that the position of the first minimum (indicated with the arrows) shifts to larger radii when melting has occurred. For the ratio R , in the crystalline phase, we have high maxima and low minima. In fact, the g_{min} is zero in this phase, leading to a value for R of 0. For the liquid-like regime, we observe a ratio of around 0.5 instead. So both of these indicators increase when melting occurs.

COORDINATION NUMBER

Another indicator that can be defined based on the pair-correlation function is the coordination number N_C . The coordination number is the average number of vortices within the first coordination shell, which in turn is the area of the circle with radius $r(g_{min})$:

$$N_C = 2\pi \int_0^{r(g_{min})} r g(r) \rho dr, \quad (4.2)$$

where ρ is the average density of vortices in the field. When a transition from a solid-like to a liquid-like state is made, we know that the singularity density decreases [24].

However, since $r(g_{min})$ also increases, the radius of the first coordination shell also increases. Furthermore the $g(r)$ is consistently non-zero, whereas the $g(r)$ for a solid is only non-zero for a specific radius, and hence increases the result of the integral. When computing the coordination number for the cases of 3 and 20 interfering plane waves, we obtain a value of around $N_c = 3$ in the case of 3 waves, and $N_c = 7$ for the case of 20 waves. So we conclude that the coordination number increases when comparing the two singularity structures, despite the decreased singularity density.

PAIR EXCESS ENTROPY

Perhaps a less obvious indicator is the pair excess entropy s_2 . Entropy is a measure for all the possible microstates (configurations) of a system: the higher the amount of possible microstates, the higher the entropy of the system, and hence the more information about the system is required in order to pinpoint the microstate of the system [172, 173]. Excess entropy is the difference in entropy with respect to an ideal gas, for the same density and temperature. It can be written as $S = S_{ID} + S_{exc}$ and thus S_{exc} is always negative, since the entropy of an ideal gas is maximal since it has the maximal amount of possible microstates. The total entropy S thus consists of two contributions: the excess entropy S_{exc} , which stems from correlated particle positions, and the entropy of an ideal gas S_{ID} , which only contains uncorrelated particle positions. Since correlations between particle positions give information about the state of the system, the amount of possible microstates decrease, resulting in a decrease of the total entropy.

The general form for the entropy S can be written in the form of a multi-body correlation expansion, with each of the n -body correlations contributing to the total entropy. The one-body correlation term then of course being the uncorrelated term leading to the ideal gas contribution³. Pair excess entropy is the two-body correlation contribution to the excess entropy, and can be computed from the pair-correlation function as

$$s_2 \propto \int \{g(r) \ln[g(r)] - g(r) + 1\} d\vec{r} \propto \int \{g(r) \ln[g(r)] - g(r) + 1\} r^2 dr. \quad (4.3)$$

Close to phase transitions the pair entropy s_2 is the most significant contributor of the multi-body terms to the excess entropy [174, 175], and thus of the possible multi-body terms the most reliable indicator of a phase transition. In the limit of infinite temperature the excess entropy goes to zero, since in the thermal limit everything behaves as an ideal gas.

4.3.2. INDICATORS BASED ON THE ANGULAR DISTRIBUTION FUNCTION

Next we discuss the indicators based on the angular distribution. Fig. 4.3 depicts the angular distribution functions for the singularity structure formed through (a) 3 interfering plane waves and (b) 20 interfering plane waves. Fig. 4.3 (a) exhibits 6 well defined peaks, which is indicative of a 6-fold symmetry. Fig. 4.3 (b) on the other hand shows no clear

³For a more in-depth treatment of (pair) excess entropy, the meaning behind it, and a rough derivation of how to get from the general formula of entropy to a multi-body correlation expansion, please see the supplementary at the end of the chapter (4.7.1).

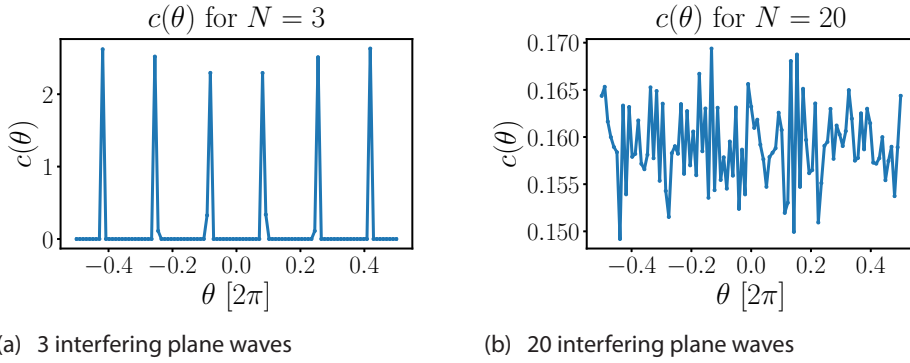


Figure 4.3: Comparison of the angular distribution function $c(\theta)$ for 3 interfering plane wave sources (a), and 20 interfering plane wave sources (b). (a) Shows an angular distribution function that exhibits peaks at 6 well-defined angles, which is indicative of a lattice with a 6-fold symmetry. (b) Shows a $c(\theta)$ with no clear preferential direction, which fits a disordered liquid-like state.

structure, meaning no clear angles at which you are more likely to find another singularity. This is befitting of a liquid-like state, since the distribution of singularities in a liquid is isotropic. Below we discuss which melting indicator can be defined from the angular distribution function.

ROTATIONAL INVARIANTS

This final melting indicator is not a single indicator, but rather a set of related variables, called rotational invariants:

$$q_n = \left\langle \frac{1}{N_C} \left| \sum_{j=1}^{N_C} e^{in\theta_j} \right| \right\rangle, \quad (4.4)$$

where N_C is the number of nearest neighbours (coordination number) and θ_j is the angle of nearest neighbour j with respect to the reference particle. The summation occurs over all particles within the first coordination shell of the reference particle. The brackets indicate the ensemble average, which is the average of the evaluation of the above function over every particle in the structure under consideration. Due to the symmetries of the function, q_n will only be non-zero if the system, on average, is a lattice with n -fold symmetry. For instance, a hexagonal lattice will have a non-zero q_3 and q_6 , since it is both three-fold and six-fold symmetric. Of course this is based on an ideal case, so in general the rotational invariants will not be zero outside of a perfect lattice, but their value will be significantly lower instead.

The way we use the rotational invariants is by considering the maximum value of q_n , in order to see if there is any symmetry present that could indicate a crystalline state. Since crystals can not possess a rotational symmetry higher than 6-fold, no value beyond q_6 has to be considered [176]. By looking at the maximum of the set, we do not need to consider

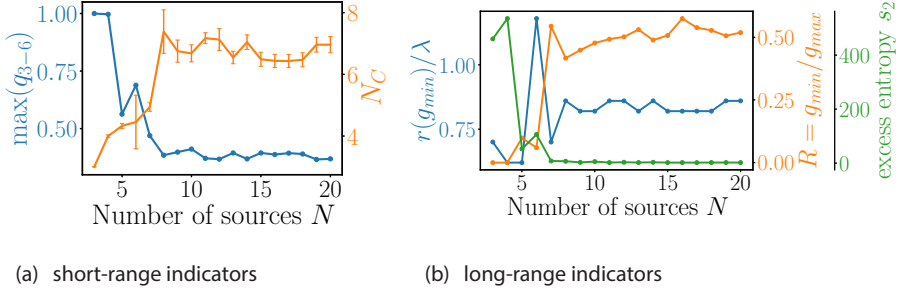


Figure 4.4: Melting indicators as a function of the number of plane wave sources. Simulated at a wavelength of $\lambda = 1.55 \mu\text{m}$. (a) shows the short-range indicators with the maximum of the rotational invariants in blue, and the coordination number N_c in orange. (b) shows the long-range indicators with the pair excess entropy s_2 in green, the location of the first minimum of $g(r)$ in blue, and the ratio between the first maximum and first minimum of $g(r)$ in orange. All of these indicators are consistent in showing that the transition from solid to liquid happens around 7 sources, with all of them stabilising around 8 sources.

all rotational invariants separately. Instead the maximum value will be sufficient to indicate the presence of a rotational symmetry: in the case of an ordered lattice structure being present, one of the indicators, and as a result the maximum of the set, will be high. In the case of a liquid-like state, no rotational symmetry will present, and hence the maximum of the set of q_n will be a low value instead. A figure of the behaviour of q_{3-6} for increasing amount of plane waves for each q_n separately can be found in Appendix A.

Please note that the melting indicators can be separated into two sets: those for short-range order of the angular distribution function and first coordination shell (rotational invariants and coordination number N_c), and long-range order which is based on the radial distribution function (s_2 , R and $r(g_{\min})$) [168].

4.3.3. MELTING INDICATORS FOR INCREASING AMOUNT OF PLANE WAVE SOURCES

To attain a more intuitive understanding of the melting indicators that were introduced above, we revisit Fig. 4.1 and compute the melting indicators for increasing number of plane wave sources. These plane waves are all equidistributed along a circle in reciprocal space. Considering Fig. 4.1 once more, we can already observe that the singularity density for (a) is higher than the density for (b). Thus when a transition occurs, we expect to see the density decrease from a value higher than π/λ^{-2} to around this value of roughly π/λ^{-2} .

Fig. 4.4 depicts the behaviour of the indicators as a function of increasing amount of plane wave sources. Fig. 4.4 (a) depicts the short-range indicators with the maximum of the rotational invariants depicted in blue, and the coordination number in orange. Fig. 4.4 (b) presents the long-range indicators, with the ratio between g_{\min} and g_{\max} depicted in orange, the position of the first non-zero minimum of the radial distribution function represented in blue, and the pair excess entropy depicted in green. We observe that all indicators vary greatly as a function of the number of plane waves that interfere.






Figure 4.5: Averaged vortex density for different amounts of interfering plane waves. The gray dashed line is the theoretical value of π that we expect in a true liquid-like state. Of interest is the $N = 6$ case, which clearly depicts a higher variance than the rest. This is explained through the fact that it seems to generate a bimodal distribution between liquid and crystalline states.

For 3 to 6 plane waves we have a solid-like state, i.e. an ordered lattice of singularities: a high value for the maximum rotational invariant and a low value for the coordination number for the short-range indicators is observed. Additionally we see that the excess entropy starts high, and already shows a rapid decline for 5 plane waves, while the other indicators exhibit a more gradual change as the amount of interfering plane waves are increased. The ratio R and position of first minimum $r(g_{min})$ both start low, and increase steadily, with $r(g_{min})$ exhibiting a peak at $N = 6$ waves. Having 6 interfering plane waves appears to be a special case: there seems to be a bimodal distribution in singularity structures, which explains the large margin of error in the coordination number at that point [138]. Investigation of this phenomenon is outside the scope of this thesis however, but a curious thing to investigate in the future. For 7+ waves we transition towards the liquid-like regime, and above 10 waves all indicators are starting to stabilize around the values that match with those for a liquid-like state: high coordination number, low excess entropy and low contrast in $g(r)$ peaks.

Finally, the average vortex density can be determined for the different number of interfering plane waves. In Fig. 4.5 we have plotted the global vortex density as a function of number of sources used, with the gray dashed line indicating the theoretical density of π/λ^2 that we expect for a liquid-like state [24]. We observe that for increasing number of sources, the theoretical value for the density is asymptotically approached. Interestingly, the convergence seems to be slow, and even though the other melting indicators show that we are in the liquid-like regime for 20 waves, the overall density is not yet at the theoretical value.



(a) k-space distribution

(b) phase distribution

Figure 4.6: Schematic illustration of the momentum broadening case for $N = 3$ symmetry points. The k -values (a) form a disk rather than a perfect circle. Note that here the 3 arrows do not depict a singular plane wave, but rather embody an ensemble of waves with different magnitudes in k -space, but launched from the same angle on the disk towards the center. The gradient shows the amplitudes when waves are launched from that k -value. Additionally, the red dashed curve shows the amplitude distribution as function of momentum. The phase difference (b) of each of these waves with respect to the closest source point is also drawn from a Gaussian distribution with a width of $\Delta\phi$.

From the behaviour of the melting indicators we can conclude that increasing the number of plane waves seems to lead to a phase transition from crystalline-like state to a liquid-like state. But inherently, changing the number of plane waves causes the underlying symmetry to change accordingly. And from crystallography, it is known that only certain symmetries can lead to a regular crystal, so perhaps it is not surprising that this melting occurs when the symmetries no longer support crystalline structures. The question then arises if it is possible to induce melting through a different method that preserves this initial spatial symmetry.

4.4. BROADENING IN MOMENTUM

With the melting indicators introduced, we can start exploring ways to induce a phase transition without breaking the 3-fold spatial symmetry present. Ideally we find a variable that we can tune that acts like an analogue to an effective temperature. When looking for a way to define an effective temperature in our system, we can look towards similar systems that also have a temperature, such as molecular beams. Molecular beams are, as the name implies, beams of molecules travelling in a well defined direction. They are usually generated by having a source of gas at high pressure, allowing it to expand through a small opening into another chamber that is kept at a lower, sometimes near-vacuum, pressure. This leads to a beam of particles that travel at similar velocities, and with a low collision rate between the particles. These beams can then be used for a variety of goods like fabricating thin films through molecular-beam epitaxy [177], or being turned into a Bose-Einstein condensate state through slowing the particles via a Stark

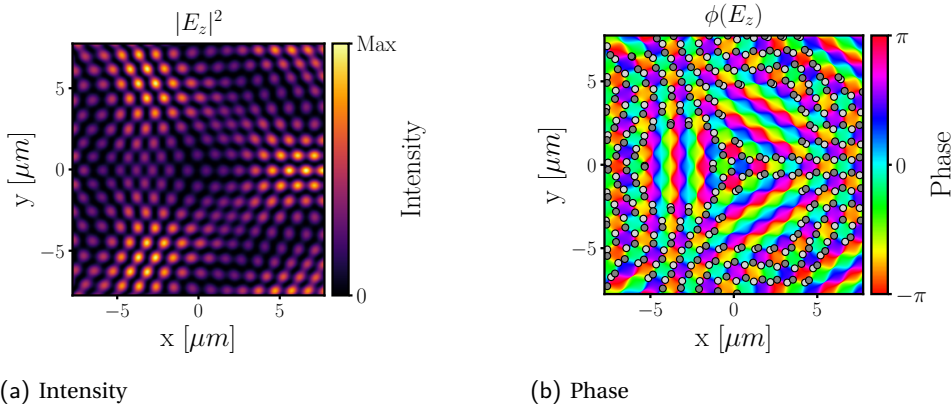


Figure 4.7: Instance of the intensity (a) and the phase (b) of the E_z field for simulated values with a momentum broadening of $\Delta k = 0.3$ and $\Delta\phi = 0.4$. The phase is plotted in false colour, and the light (dark) gray dots indicate the positions of positive (negative) singularities. From both the intensity and phase figures we can observe that the lattice is clearly deformed from its trigonal base, but at its core still resembles the trigonal lattice. Hence there is no indication that the distribution is liquid-like for these variable values.

or Zeeman slower[178]. We now consider the way temperature is commonly defined for these beams, which is through the spread of longitudinal velocities [179]. So while the beam can be highly energetic in the reference frame of the lab, it is not necessarily the case for the beam in its own frame of reference. In this way, ultracold molecular beams can be fabricated. The most obvious way to translate this into optics is to introduce a spread in momentum for light, which means a spread in k -values Δk . Thus we no longer stick with a ring in k -space that is infinitesimally thin, but rather introduce a disk (more precisely, an annulus) in k -space, with its centre at k_0 . However, we do keep ω constant here in order to preserve monochromaticity (this can of course only be done in simulations). We then introduce a Gaussian envelope for the amplitude of the waves, with the mean centered on k_0 , and a full width at half maximum (FWHM) of Δk . Like the amplitudes, the phases of the waves are also variable, and randomly selected from another Gaussian distribution with a FWHM of $\Delta\phi$. See Fig. 4.6 for a schematic illustration.

RESULTS

As an initial test, we have plotted the intensity (a) and phase (b) of one (representative) instance of the simulated field for $\Delta k = 0.3$ and $\Delta\phi = 0.4$, which is shown in Fig. 4.7. From both Fig. 4.7 (a) and (b) we can infer that the behaviour does not differ strongly from a lattice, and hence a solid-like state. While the trigonal pattern has been distorted, it can still clearly be observed, and thus clearly does not exhibit liquid-like behaviour.

To further investigate the possibility of melting through momentum broadening we have simulated the system for a range of different values for the phase noise $\Delta\phi$ and the spread in momentum values Δk . The phase noise was varied over a range of $\Delta\phi \in [0, 1.2]$ in units of 2π and the momentum spread over a range of $\Delta k \in [0, 0.3]k_0$. It is of note here that this spread in momentum values is wholly unphysical at this point for the largest values.

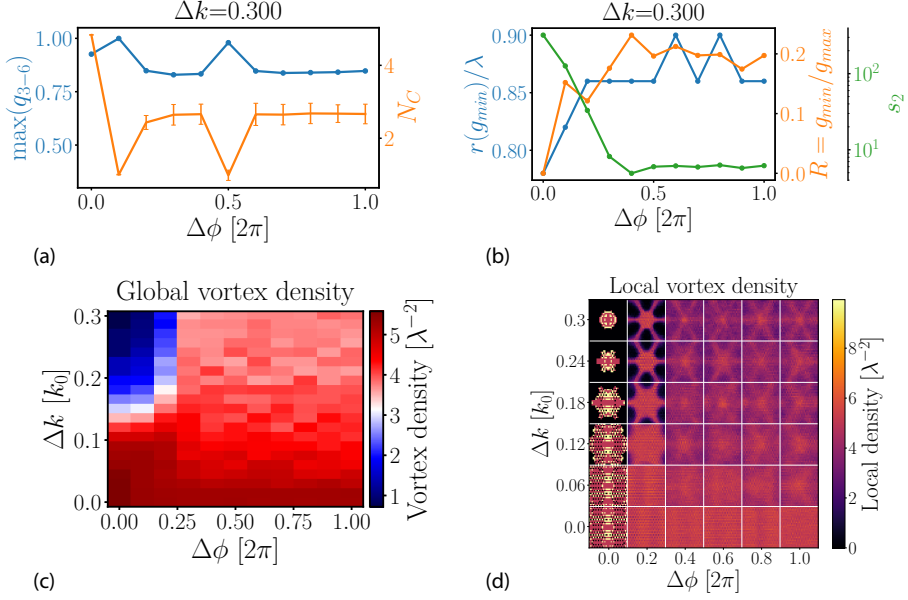



Figure 4.8: Melting indicators as a function of phase noise $\Delta\phi$ for a momentum spread value of $\Delta k = 0.3k_0$. (a) depicts the short-range indicators, while (b) shows the long-range melting indicators. Both of them show that, while they do vary over the range of values, neither shows a clear phase transition when looking at the values over which they vary. Vortex densities with (c) showing the global densities and (d) showing the heat maps of the local vortex densities. (c) clearly shows that there are almost no spots with global densities that match the expected densities for a liquid-like state, which would be coloured white. (d) supports this by showing that the spoke-like pattern does not seem to go away, even for higher values of momentum and phase spread. The area in (c) that is coloured white is shown in (d) to be a spoke like state, clearly not liquid-like.

There is no way in which this can be experimentally tested and have bound modes for all values involved. Despite this, it is still enlightening to know if melting can be induced through this method.


Fig. 4.8 (a) and (b) shows the computed melting indicators with (a) showing the short-range indicators and (b) showing the long-range indicators for the largest simulated momentum spread value, as a function of the phase noise. Considering Fig. 4.8 (a), we observe that both the maximum of the rotational invariants and the coordination number both show some variations, but neither get close to what we expect for a liquid-like state. For the coordination number we observe that it remains around a value of $N_C = 3$, while for a liquid-like state we would expect values around $N_C = 7$ instead. Likewise for the maximum of the rotational invariants, the value remains above 0.8, while for a liquid-like state we would expect values around 0.4 instead.

Similarly we observe in Fig. 4.8 (b) that the values do change, and even seem to transition around $\Delta\phi = 0.2$, but looking at the scale over which they change reveals that they do not get close to the expected values for a liquid-like state: the ratio R changes from about 0.775 to 0.875, which actually is consistent with the value for a liquid-like state as found in Fig. 4.4. The excess entropy also sharply decreases, signifying a deviation from an ordered lattice. Finally, the ration R only increases to 0.2, where we expect it to increase to around 0.5 for a liquid-like state. So it seems like the long-range melting indicators are not fully in agreement with each other, while the short-range indicators clearly show a crystalline state.

To further our understanding we additionally look at the local and global singularity densities. The local singularity density is obtained by computing the singularity density over a limited radius around each pixel of the image, averaged over typically at least 100 iterations for each variable combination. In this a heat map can be constructed, which indicates where the singularities tend to cluster or not. Fig. 4.8 (c) and (d) shows the vortex densities, with (c) showing the global densities, and (d) depicting the heat map with local vortex densities for the simulated values. In Fig. 4.8 (c) we clearly observe that most variable combinations significantly differ from the expected theoretical value of π for a liquid-state, which is indicated as white on the colour scale. However, looking at (d) we observe that the variable combinations close to the theoretical value occur in a region where the spread is clearly not uniform. Instead, it just happens to be that the areas with vortex concentrations and areas without vortices even out to the global density of a liquid-like state. So it seems that the lattice gets distorted enough to affect the long-range order, leading to some long-range melting indicators showing a transition, while the short-range order remains intact. As such, we have to conclude that this method of broadening does not lead to the melting of the lattice. So it seems that some spatial symmetry breaking is required in order for melting to occur.



(a) k-space distribution



(b) phase distribution

Figure 4.9: Schematic illustration of the angular broadening case for $N = 3$ symmetry points and $N = 36$ sources. The N sources (a) are given an amplitude A_n based on their position on the ring in k -space, and is determined by the angular distance to the nearest symmetry point via a Gaussian distribution with a tunable width of $\Delta\theta$, which is given by the blue curve in the circle. The phase difference (b) of each of these waves with respect to the closest symmetry point is also drawn from a Gaussian distribution with a width of $\Delta\phi$.

4.5. ANGULAR BROADENING

While broadening the ring in k -space radially to a disk did not lead to a phase transition, we now investigate broadening tangentially, i.e., along the ring in k -space.

As a baseline, we consider 3 monochromatic plane waves, which have been distributed equidistantly along the ring in reciprocal space. These three points will be referred to as the symmetry points. In order to broaden these symmetry points, we introduce more plane waves, which have also been equally distributed along the ring, but with additional restrictions on their amplitudes and initial phases. Their amplitudes depend on the distance to the nearest symmetry point. This amplitude distribution is chosen to be a Gaussian, with its centre at the symmetry point, leading to a decaying amplitude the further away from the symmetry point you get. The Gaussian amplitude distribution has certain width $\Delta\theta$, which we can tune. By tuning the width of this Gaussian, effectively more sources are slowly introduced into the system, but as a function of a continuous variable. Like the amplitudes, the initial phases of the waves are also variable, and randomly selected from another Gaussian distribution of width $\Delta\phi$, just like in the momentum broadening case. This Gaussian is centered around the phase of the symmetry point. Hence the added waves have a phase that is close to the phase of the symmetry point, but have a variation dependent on the variable width $\Delta\phi$. We refer to this method of broadening as angular broadening. In Fig. 4.9 we show a schematic illustration of angular broadening. The left side of the image portrays the effect of the Gaussian envelope on the amplitude of each source point and how we define $\Delta\theta$, illustrated for $N = 36$ sources. The right side of the figure illuminates how we define the phase spread $\Delta\phi$, which is through the full width at half maximum (FWHM) of the Gaussian distribution.

Since there are two variables that we can tune, namely $\Delta\phi$ and $\Delta\theta$, we can either consider their effects separate, or one can be made dependent on the other. For the initial




Figure 4.10: Comparison of two instances for variable combinations, with (a) and (c) respectively depicting the intensity and phase for $\Delta\theta = 0, \Delta\phi = 0$. (b) and (d) depict respectively the intensity and phase for $\Delta\theta = 0.2, \Delta\phi = 1.8$. (a) and (c) clearly exhibit an ordered lattice structure for both the intensity and phase singularities, and is clearly in a solid-like state. (b) and (d) do not exhibit a regular pattern, and are no longer clearly in a crystalline state.




Figure 4.11: Melting indicators as function of the angular spread $\Delta\theta$ for 720 sources. (a) shows the short-range indicators with the maximum rotational invariant in blue and the coordination number depicted in orange. (b) depicts the long-range indicators with the position of the first minimum of the $g(r)$ in blue, the ratio $R = g_{min}/g_{max}$ in orange and the pair excess entropy in green.

investigation of angular broadening, we let $\Delta\phi$ depend on $\Delta\theta$, as well as to the amount of symmetry points used:

$$\Delta\phi = \Delta\theta N_s^2, \quad (4.5)$$

where N_s is the amount of symmetry points, which we have fixed at $N_s = 3$ for this investigation.

As a first indication as to whether or not a state besides a crystalline-like state can be obtained, Fig. 4.10 displays a comparison of two instances for different variable combinations. Fig. 4.10 (a) and (c) show the intensity and phase respectively for the variable values $\Delta\theta = 0$, $\Delta\phi = 0$ and (b) and (d) depicting the intensity and phase respectively for the variable values $\Delta\theta = 0.2$, $\Delta\phi = 1.8$. For Fig. 4.10 (a) and (c) we clearly observe a perfectly ordered trigonal lattice structure, indicating a solid-like state. For Fig. 4.10 (b) and (d) on the other hand, there is no regular structure to be discerned, possibly indicating a liquid-like state. In order to see if the singularity network shown in Fig. 4.10 (d) is in a liquid-like state, we compute the melting indicators.

We present the obtained melting indicators in Fig. 4.11, which have been determined for a value range of $\Delta\theta \in [0.0, 0.2]$ in units of 2π , with the phase coupled to it as defined in Eqn. 4.5. Fig. 4.11 (a) depicts the behaviour of the short-range melting indicators, while Fig. 4.11 (b) displays the long-range indicators. Based on the short-range indicators, a change clearly happens between $\Delta\theta = 0.06$ and $\Delta\theta = 0.07$, with both indicators exhibiting an abrupt jump. But based on the long-range indicators, we observe a more gradual transition, with especially R slowly increasing before finally plateauing at around $\Delta\theta = 0.12$. This discrepancy between the changes in the short-range indicators and the long-range indicators could indicate that short-range order is lost more quickly in the transition than long-range order. Despite the differences, both exhibit a transition from solid-like state to liquid-like state. As such we can conclude that does method of broadening does lead to a transition from a solid-like state to a liquid-like state. Next we consider the two variables separately.

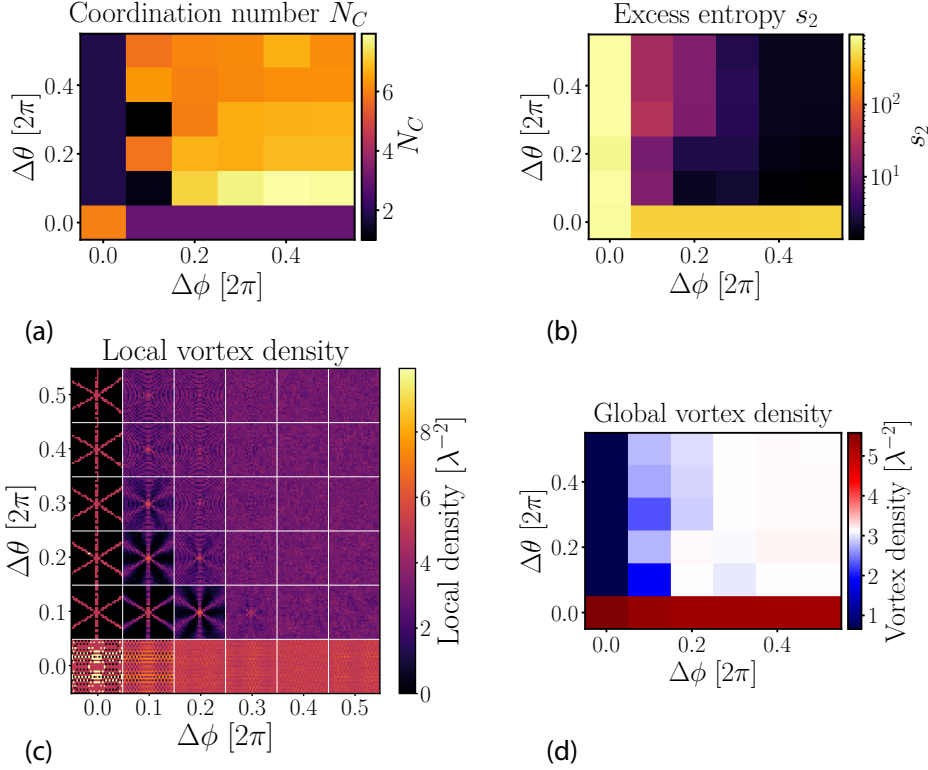


Figure 4.12: (a) coordination number and (b) excess entropy indicators as function of the phase spread $\Delta\phi$ and amplitude spread $\Delta\theta$. Note that the colour scale of the excess entropy is logarithmic due to high values for the non-liquid-like state that drops off quickly. (c) shows heat maps of local vortex densities, separated by white straight lines and (d) the global vortex densities, with white indicating the theoretical value for the liquid-like state. Each block is a separate heat map generated for the variables and shows the distribution of vortices over the space, averaged over the 100 iterations that were performed. We can identify three regions in this image. First we see that for $\Delta\theta = 0$, we retain the lattice structure, and for low values we stay close to it. At low values for $\Delta\phi$ we see that we move from a lattice to the spoke-like structure. Finally, at higher values of both we see the liquid-like state appear. This is also reflected in the global density, as crystalline states have a higher density, and the spoke states a lower density than a liquid. Finally we see for higher values of both variables that the density goes towards the expected density for a liquid.

In order to investigate the effect of the two variables separately, we now no longer assume that one variable is a function of the other. We have taken a range of $[0, 0.5]$ in steps of 0.1 for the two variables, for 720 sources and 100 iterations for each variable combination. Because it is cumbersome to display all of the melting indicators separately, we have selected one indicator from the short-range and long-range indicators each, alongside the local and global singularity densities. The selected indicators are representative of the complete collection of melting indicators. In Fig. 4.12 we show the results for the coordination number N_c (a), the pair excess entropy s_2 (b), and the local (c) and global (d) vortex densities. Note that the colour scaling of the excess entropy is logarithmic due to high values for the solid-like state, which subsequently diminishes rapidly. Fig. 4.12 (a) reveals that for both $\Delta\theta = 0$ and $\Delta\phi = 0$ a low coordination number is observed regardless of the other variable. Such a low coordination number is consistent with a solid-like state. When both variables are increased, we observe that the coordination number increases as well to values that are consistent with a liquid-like state. Considering the pair excess entropy s_2 , which is depicted in Fig. 4.12 (b), we observe that this melting indicator displays the same behaviour as the coordination number: for both $\Delta\theta = 0$ and $\Delta\phi = 0$ a high excess entropy is found, consistent with a solid-like state, and for both variables increased we find a sharp decline in the value of s_2 , consistent with a liquid-like state.

The local vortex density, which is depicted in Fig. 4.12 (c), reveals additional information: for $\Delta\theta = 0$ we observe that the local density exhibits a pattern that is consistent with a lattice, revealing that the structured formed through the phase singularities thus far remains in a solid-like state for all considered values of the variable $\Delta\phi$. However, for $\Delta\phi = 0$ we notice a spoke-like pattern instead. Only for higher values of $\Delta\phi$ do we observe that a transition is made to a state where the local density appears to be uniform, which is indicative of a liquid-like state. This is reflected in Fig. 4.12 (d): $\Delta\theta = 0$ shows a consistently high global density, which fits a solid-like state, and $\Delta\phi = 0$ instead has a low density since the resulting patterns appear to have few singularities overall. Once both variables are increased, the global density is near the theoretical density of an isotropic liquid-like state.

Of all these observations, the statements about $\Delta\theta = 0$ should not come as a surprise, as it simply corresponds to no broadening at all. But what does stand out is that for low values of $\Delta\phi$, which plays the role of phase noise, no melting behaviour is observed, suggesting that this is a crucial ingredient for melting to occur. What additionally stands out is that for higher phase spread, the transition into the liquid regime as a function of angular spread is much sharper as well. However, for low values (0.1) of both the angular and phase spread, no melting is found either. From this we can conclude that, while phase noise is a critical ingredient for melting, it does not lead to melting purely on its own. In other words, having a non-zero phase noise is a necessary, but not a sufficient requirement for melting to occur. So we conclude that both phase noise and angular broadening are a requirement. The necessity of phase noise can be understood in the following way. Assume that the sources would have the exact same phase as the symmetry point they are based on. For each symmetry point, the resulting field is mirror symmetric over the symmetry point, and the result is simply a beam that is focused onto the middle of the ring. And thus the result of all sources combined is a field comprised of

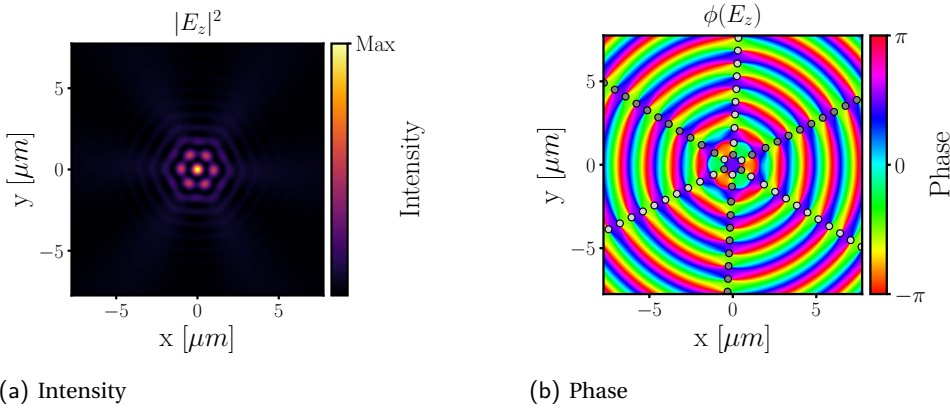


Figure 4.13: Instance of the E_z field intensity (a) and the phase (b) for simulated values of $\Delta\theta = 0.08$ and $\Delta\phi = 0$. Figure (a) shows clearly that intensity is concentrated in the centre and a 6-fold symmetric pattern around it, while (b) shows that the resulting interference leads to singularities only on the axes of the beams.

three beams that are focused onto the center. The final result is still a 3-fold symmetric field, but instead of a hexagonal structure, we find that the result becomes more circular with 3 lines of singularities, which cross exactly in the center. We have dubbed this singularity pattern a spoke-like pattern, as they look akin to spokes on a wheel. One instance of such a field is shown in Fig. 4.13. Fig. 4.13 (a) shows the intensity of the field, which has most of its intensity focused in the middle. Furthermore a 6-fold symmetric pattern can be seen around the centre of the image. Fig. 4.13 (b) shows the spoke-like pattern for the phase singularities that arises. The singularities are confined to the beam axes, leading to 6 distinct lines radiating outward from the center. Now clearly this no longer classes as a lattice structure, and as such is no longer a solid as we know it. However, it is not in a liquid-like state either. It rather becomes a state without clear analogy to solid-state physics. Furthermore, it seems like in order to get to the liquid-like state from the solid-like state, this spoke-like state regime has to be crossed first. Unless the phase noise is strongly increased when increasing the angular spread, this spoke state will have to be crossed before arriving at the liquid-like state. This can be observed from Fig. 4.12 (c). Thus it seems that this spoke-like state acts as an intermediate state.

4.5.1. CONVERGENCE AS FUNCTION OF SOURCES

So far a fixed value for the number of source points N has been used. One thing that should be investigated is the behaviour of the system as N is varied. We have computed the melting indicators as a function of $\Delta\phi$ for a fixed value of angular broadening $\Delta\theta = 0.08$ in order to investigate if the system converges to a common value of $\Delta\phi$ for which the phase transition can be observed. Fig. 4.14 depicts the behaviour of the coordination number N_c for increasing amount of sources up to $N = 1440$. It is clear that even for a large number of points, the behaviour does not seem to converge. Simulating even more points than that starts to become computationally prohibitively expensive. While the behaviour for only one melting indicator is shown here, the behaviour is consistent




Figure 4.14: Behaviour of the coordination number N_C as a function of the phase spread $\Delta\phi$ for the continuous broadening case, with a fixed $\Delta\theta = 0.08$. The different coloured lines indicate different amount of sources used to simulate the fields as indicated in the legend. It is clear that for increasing number of sources the phase spread needs to be increased in order to achieve the same effect on the coordination number, indicating that the system has not converged, even for a large amount of sources.

across all of the other indicators as well, and figures for the remaining indicators can be found in Appendix A. One possible explanation for this is that for increasing number of points, the points start to lie close together, making them effectively one point with a summation of their phases. It is possible that then we need to increase the phase spread to make up for that behaviour once more. Please note that while this may highlight a potential flaw in the model, it does not alter the conclusions that can be drawn from our investigation: while the values of the used variables might change depending on the amount of sources used, the behaviour remains consistent. For further investigation we recommend looking into different models. One such model of highlighted in Appendix A.

4.6. CONCLUSIONS

It is clear that melting singularity lattices is a complicated process. We have identified two variables that are able to change the system in a continuous way, yet neither of them is able to explain melting simply by them self. Simply introducing more waves into the system via the continuous broadening case, but not changing the phase of them with respect to the nearest symmetry point will simply focus the three beams instead due to symmetry. Changing the phase alone does not do anything to alter the three-fold symmetry that creates the hexagonal lattice, since their relative phases do not matter here. Unfortunately, this reasonably simple model strongly depends on the number of sources that are added, and we have found no indication of it converging at a significant rate. Attempting to broaden in momentum space as a more direct analogue to the temperature of molecular beams also does nothing to break the symmetry, even with the added phase distributions and hence cannot explain the observed melting behaviour.

The interesting part here is that a clear deviation from molecular beam behaviour is observed. Where normally temperature can be defined through a Boltzmann distribution, which can be related with momentum, here we have observed that increasing momentum spread does not lead to melting. And in fact, broadening the beams in either momentum or angle alone also does not suffice. It seems that phase noise is a required component of the puzzle here, because without it, no melting is observed. It alone can also not induce melting, however, so the phase noise by itself cannot break the symmetry of the system. So while a clear effective temperature analogue remains elusive, we can conclude that phase noise is required.

4.7. SUPPLEMENTARY INFORMATION

4.7.1. ENTROPY AND EXCESS ENTROPY

This section is based on ref [180], and follows its line of reasoning to qualitatively show how the multi-body correlation expansion of the entropy is derived. For a more in-depth treatment and derivation the reader is referred to ref. [180].

In the statistical definition it is related to the amount of microstates the system can take. Given a distribution of discrete states i , which system can be in with probability p_i , the entropy of the system can be calculated as

$$S = -k_B \sum_i p_i \ln(p_i). \quad (4.6)$$

As such, the entropy increases with the amount of possible microstates. In essence this represents a measure for the amount of information that is required in order to pinpoint the state the system is in: the more microstates are available, the more information is required in order to pinpoint the exact microstate being occupied.

Now discrete states are mainly relevant when considering quantum states, but we can take this into the classical realm by taking a Boltzmann distribution for the state probabilities, given as

$$p_i = e^{-E_i/(k_B T)/Z}, \quad (4.7)$$

where E_i is the energy of that state, and Z is the partition function, which normalizes the probability distribution. By taking the high-temperature limit, and the canonical ensemble, where we have a fixed volume V and amount of particles N , we can compute the entropy as

$$S_N/k_B = -\frac{1}{N!} \int \prod_i f_N \ln(h^{3N} f_N) d\vec{r}_i d\vec{p}_i, \quad (4.8)$$

where h is the Planck's constant, $f_N(\vec{r}_1, \vec{p}_1, \dots, \vec{r}_N, \vec{p}_N)$ is the probability density as function of positions \vec{r}_i and momenta \vec{p}_i .

The above equation can then be refactored in terms of n-body distribution functions $g_N^{(n)}$ as

$$S_N/k_B = s_1 + s_2 + s_3 + \dots, \quad (4.9)$$

where each successive term takes into account higher-order position correlations, so that s_1 only takes into accounts itself, s_2 the pair correlations, s_3 the three-body correlations etc. Working out these terms, we find that

$$s_1 = \frac{3}{2} - \ln(\rho \Lambda^3), \quad (4.10)$$

where Λ is the de Broglie wavelength, and ρ the particle density. And the pair correlation term equates to

$$s_2 = -\frac{\rho}{2} \int g_N^{(2)} \ln(g_N^{(2)}) d\vec{r}. \quad (4.11)$$

The entropy for an ideal gas can be computed to be

$$S_{id}/k_B = \frac{5}{2} - \ln(\rho\Lambda^3) = s_1 + 1 \quad (4.12)$$

The discrepancy of the 1-body term with the ideal gas stems from the $\frac{1}{N!}$ series that is added to the higher order terms. In order to reconcile this, we can simply add 1 to s_1 , and subtract 1 from the rest of the terms combined:

$$S_N/k_B = S_{id}/k_B + (s_2 - 1/2) + (s_3 - 1/6) + \dots \quad (4.13)$$

At this point it is natural to introduce the term excess entropy, which is simply the difference between the total entropy and the contribution of the entropy from an ideal gas

$$S_{ex} \equiv S_{tot} - S_{id} = (s_2 - 1/2) + (s_3 - 1/6) + \dots \quad (4.14)$$

By recasting the expressions above in the grand canonical ensemble, the canonical correlation functions $g_N^{(2)}$ can be exchanged for radial distribution functions $g(r)$. Finally, the excess entropy contribution from 2-body correlations can be found to be

$$s_2 \propto \int \{g(r) \ln[g(r)] - g(r) + 1\} d\vec{r} \propto \int \{g(r) \ln[g(r)] - g(r) + 1\} r^2 dr. \quad (4.15)$$

Now going back to what entropy means, we see that it gives a measure of disorder and (lack of) information about the system. The higher the entropy, the less we know about the exact microstate the system is in. From here it is easy to see that the excess entropy has to be a negative quantity. Indeed, an ideal gas is the most disordered a system is going to get, meaning the least amount of information about the system is available, and hence the entropy is maximized in this case. What the excess entropy tells us is how much information is *gained* due to correlations between particle positions. And information gained means entropy lost, since it restricts the amount of possible microstates.

Now investigating the integrand a little further, two different terms can be distinguished: $g(r) \ln(g(r))$ and $g(r) - 1$. The first of these terms is reminiscent of the original expression for entropy, and indeed represents the information about the state that is stored in the pair-correlations. The second of these arises when we moved to the grand canonical ensemble and stems from the fluctuation in number of atoms. This latter term is actually always positive, since fluctuations in particle number leads to more uncertainty, and hence less information about the system.

5

FLOW FIELD SINGULARITIES

Even the tiniest, blandest, simplest light holds a wonderful secret.

Emily - Stardew Valley⁵

In this chapter we investigate, both theoretically and experimentally, the singularities that appear in the in-plane Poynting vector field of random waves. We show that there is a restriction on the types of critical points allowed in the flow-field, impacting the allowed topological structures of the field. We find that there are three types of singularities that can arise in the transverse flow-field, namely magnetic, electric, and polarization types, of which electric type is shown to be a special case of the polarization type. Next, we give a mathematical proof that magnetic type singularities only exist as positively charged singularities in the flow-field, and reveal the mechanism by which they can still annihilate with another magnetic type singularity, while still abiding charge conservation of singularities in the magnetic field itself. Finally we consider the imaginary part of the Poynting vector, and show that its behaviour is fundamentally different from its real counterpart.

Parts of this chapter have been published in Optics Letters **45**, 9: 2600-2603 (2020) [181].

5.1. THE POYNTING VECTOR AND ITS SINGULARITIES

Besides considering the electric and magnetic fields by themselves as separate entities, one can also look at entities that are comprised of a combination of both fields at the same time. One such quantity is the Poynting vector, which describes the power flow of the electromagnetic field, and can be defined as

$$\mathbf{S} = \mathbf{E} \times \mathbf{H}, \quad (5.1)$$

with \mathbf{E} the electric field and \mathbf{H} the auxiliary magnetic field [182, 183]. This specific form is the one introduced in the original paper by Poynting published in 1884 [184]. Specifically, it is an expression of the energy flux density transported by the fields. To increase our understanding of the behaviour of light in structures, we want to also understand the behaviour of its flow via this Poynting vector. It is known that light in 2D shows a fundamental difference in the behaviour of its generic singularities from those present in a paraxial beam [43, 56]. It can be presumed, that such a confinement not only has consequences for singularities in the complex vectorial electromagnetic fields, but also for the singular behaviour exhibited by other physical observables, such as the energy flow, which is known to contain vortices, for example in the vicinity of the focus of a light beam [185] or near sub-wavelength apertures [45]. Since the Poynting vector is directly related to the optical momentum, potential applications could be utilizing the in-plane optical momentum present in many photonic systems for enhanced particle manipulation schemes [186].

Since optical cycles are extremely short¹, for most practical applications the instantaneous power is irrelevant, since it's the average over many cycles that is observed. Hence another common form of the Poynting vector is the time-averaged version. Considering telecom frequency light, and an average measuring time per pixel in our near-field microscope of 2 milliseconds, we collect roughly $4 \cdot 10^{11}$ cycles at every measured pixel. So clearly for our experimental work, the time averaged form is the one that is natural to consider. Assuming that the fields are time-harmonic and the medium in which it resides is linear, the time-averaged Poynting vector can be computed as

$$\mathbf{S} = \frac{1}{2} \Re(\mathbf{E} \times \mathbf{H}^*), \quad (5.2)$$

where the factor of 1/2 stems from the integration of the time-harmonic field over one period [88].

In this chapter we once again consider two-dimensional random waves as the basis of our investigation, and thus we will consider the in-plane Poynting vector only. This 2D projection of the Poynting vector onto the plane is called the transverse flow-field [38]. As with any 2D vector field, the flow field can exhibit singularities as well, which occur when the Poynting vector at a point in the light field vanishes, meaning there is no flow of energy at that point.

¹The frequency of visible light is generally hundreds of terahertz, resulting in a single cycle on the order of femtoseconds.

Once more we investigate TE slab mode, confined in a chaotic cavity. For a TE slab mode the out-of-plane component of the electric field is zero, and the cross product from Eqn. 5.2 can be expanded:

$$\mathbf{S} \propto \Re\{(E_y H_z^*, -E_x H_z^*, E_x H_y^* - E_y H_x^*)\}. \quad (5.3)$$

Additionally, for a TE mode the real part of the magnetic field is oriented only out-of-plane, and as such the in plane components will be fully imaginary. This can be seen considering that for a 2D TE mode, $E_z = 0$, and through the Maxwell's equations we find

$$\begin{aligned} H_x &\propto -k_z E_y, \\ H_y &\propto k_z E_x, \end{aligned}$$

with k_z purely imaginary for a bound mode. As a result, the out-of-plane part of the Poynting vector is purely imaginary. The absence of a real part of the Poynting vector out-of-plane is just a manifestation of the fact that the evanescent field does not transport energy. The in-plane components are then the energy flow inside the chaotic cavity itself:

$$\mathbf{S}_{\parallel, \text{TE}} = \Re\{H_z^*(E_y, -E_x)\}. \quad (5.4)$$

In order to experimentally determine the positions of singularities in the S_{\parallel} field a complex scalar field can be artificially constructed out of the two real vector components of the transverse flow-field:

$$\xi = S_x + \iota S_y, \quad (5.5)$$

where ι is the imaginary unit. This is similar to the construction of a complex Stokes field to locate polarization singularities [50, 187]. The mapping $(S_x, S_y) \mapsto S_x + \iota S_y$ is a homeomorphism between \mathbb{R}^2 and \mathbb{C} , meaning that the topological properties of the constructed field are identical to that of the original vector field [188]. After constructing this complex scalar field, the positions of its singularities can be determined in the exact same way that has been done in the previous chapters for phase singularities in the magnetic field.

5.2. ORIGINS OF THE SINGULARITIES

Now that we have determined how the positions of singularities in the transverse flow-field can be located, we can investigate under which circumstances the singularities arise. By looking at the formula for the Poynting vector, there appear to be three ways in which the Poynting vector can vanish identically. Either the out-of-plane magnetic field vanishes ($H_z = 0$), the in-plane electric field vanishes ($E_x = E_y = 0$), or the real part of the cross product vanishes, with neither the electric field nor the magnetic field being zero.

The first type can then be easily identified as being phase singularities of the magnetic field, and we call those magnetic type singularities. Similarly, the singularities that arise from zeros of the electric field are called electric type singularities. Both of these types can occur in a generic TE field [43, 189]. However, since the electric types require two field components to be zero, as opposed to only one for the magnetic types, they are expected to occur only rarely (comprising about 1% of singularities) [189].

Finally we can investigate the conditions in which $\Re(\mathbf{E} \times \mathbf{H}^*) = 0$, where we demand that $\mathbf{E} \neq 0$ and $\mathbf{H} \neq 0$. Under these conditions the inner product of the Poynting vector with the electric field can be taken:

$$\Re(\mathbf{E} \times \mathbf{H}^*) \cdot \mathbf{E} = \frac{1}{2} (\mathbf{E} \times \mathbf{H}^* + \mathbf{E}^* \times \mathbf{H}) \cdot \mathbf{E} = 0.$$

Discarding the factor of a half and using the commutation rules for the cross product, this can be rewritten as

$$(\mathbf{E} \times \mathbf{E}) \cdot \mathbf{H}^* + (\mathbf{E} \times \mathbf{E}^*) \cdot \mathbf{H} = (\mathbf{E} \times \mathbf{E}^*) \cdot \mathbf{H} = 0. \quad (5.6)$$

Since $(\mathbf{E} \times \mathbf{E}^*)$ and \mathbf{H} are always parallel and out-of-plane, and $H_z \neq 0$ by our set condition, we require that $(\mathbf{E} \times \mathbf{E}^*) = 0$. Computing the cross product, we find that the condition is set by

$$E_x E_y^* - E_x^* E_y = 2\Im(E_x E_y^*) = -S_3 = 0, \quad (5.7)$$

where we introduce the Stokes parameter S_3 , which gives the degree of circular polarization [190].

Physically speaking, $S_3 = 0$ means that the electric field is linearly polarized. However, since the parameter space is two-dimensional, two conditions are required to be met in order to find singularities, which are zero-dimensional entities. The second condition can be found by considering the expression for the components of the Poynting vector. When the condition for linear polarization is valid, it can be rewritten as

$$E_x = \frac{E_x^* E_y}{|E_y|^2} E_y = \alpha E_y, \quad \alpha \in \mathbb{R}, \quad (5.8)$$

given that $E_y \neq 0$. From here it can be deduced that

$$S_y = -E_x H_z^* = -\alpha E_y H_z^* = -\alpha S_x. \quad (5.9)$$

So the two components are linearly related, and the second condition can simply be found by looking for occurrences where one of the two components vanishes, which happens when

$$S_x = \Re(E_y H_z^*) = 0, \quad (5.10)$$




Figure 5.1: Generic critical points with charge $q = \pm 1$ of a 2D field. (a): Centre. (b): Saddle point. (c): Focus. (d): Node. (See [191])

which occurs only when the two fields are $\pi/2$ out of phase with each other.

So these types of singularities arise when the electric field is linearly polarized, and the magnetic field is then $\pi/2$ out of phase with the electric field, but under the condition that neither components of the electric field are equal to zero.

5

5.3. RESTRICTIONS ON THE TYPES OF CRITICAL POINTS

We now know under which circumstances the singularities arise, so a next logical step is to look at what these singularities look like in the vector field. In order to do this, we first investigate the allowed topological structures of the field based on the divergence and rotation of the field.

To start the investigation, we first show that the Helmholtz relation is satisfied to allow simplifications of the equations. In case of isotropic monochromatic random waves, it is easy to show that H_z satisfies the Helmholtz relation:

$$\nabla^2 \psi + k^2 \psi = 0, \quad (5.11)$$

where ψ is a generic (complex) scalar function. For a plane wave this scalar field can be written as $\psi = ae^{i(\mathbf{k} \cdot \mathbf{r} + \phi)}$, where \mathbf{k} is the momentum vector, i the imaginary unit, \mathbf{r} the position vector and ϕ a generic phase offset.

For this plane wave field it holds that the gradient is

$$\nabla \psi = i\mathbf{k}\psi. \quad (5.12)$$

We can then generalize the field ψ to an isotropic random wave field, where we have a sum of plane waves. Each of the components can have different \mathbf{k} -vector directions \mathbf{k}_n , a different amplitude a_n , and phase offset ϕ_n , with all vectors having same magnitude:

$$\psi = \sum_n a_n e^{i(\mathbf{k}_n \cdot \mathbf{r} + \phi_n)}, \quad |\mathbf{k}_n| = k. \quad (5.13)$$

Computing the gradient is straightforward, and yields

$$\nabla\psi = \imath \sum_n a_n \mathbf{k}_n e^{\imath(\mathbf{k}_n \cdot \mathbf{r} + \phi_n)}. \quad (5.14)$$

Finally, the Laplacian is given by

$$\begin{aligned} \nabla^2\psi &= \nabla \cdot \nabla\psi \\ &= \imath \sum_n a_n \mathbf{k}_n \cdot \nabla e^{\imath(\mathbf{k}_n \cdot \mathbf{r} + \phi_n)} \\ &= \imath \sum_n a_n \mathbf{k}_n \cdot (\imath \mathbf{k}_n) e^{\imath(\mathbf{k}_n \cdot \mathbf{r} + \phi_n)} \\ &= - \sum_n a_n |\mathbf{k}_n|^2 e^{\imath(\mathbf{k}_n \cdot \mathbf{r} + \phi_n)} \\ &= -k^2 \sum_n a_n e^{\imath(\mathbf{k}_n \cdot \mathbf{r} + \phi_n)} \\ &= -k^2 \psi, \end{aligned}$$

5

proving that the Helmholtz equation is satisfied for a scalar field consisting of isotropic random waves, which H_z is known to be [43].

Since a perfect 2D TE-mode is considered, we can write the in-plane electric field components in terms of the out-of-plane magnetic field H_z . By using the Maxwell's equations and the fact that the electric field components are time-harmonic, we find

$$\begin{aligned} \partial_x H_z &= \imath \omega \epsilon E_y, \\ \partial_y H_z &= -\imath \omega \epsilon E_x. \end{aligned}$$

Using the found relation between the electric field components and spatial derivatives of H_z , the equation for the Poynting vector can be rewritten as

$$\mathbf{S} \propto \Im(H_z^* \nabla H_z), \quad (5.15)$$

where we have additionally made use of the identity $\Re(-\imath f) = \Im(f)$, for a complex function f . Since the Helmholtz relation holds for H_z , it follows that

$$\begin{aligned} \nabla \cdot \mathbf{S}_{\parallel} &= \nabla \cdot \Im(H_z^* \nabla H_z) \\ &= \Im(\nabla H_z^* \cdot \nabla H_z + H_z^* \nabla^2 H_z) \\ &= \Im(|\nabla H_z|^2 - k^2 |H_z|^2) \\ &= 0 \end{aligned}$$

So the in-plane divergence has to be zero everywhere.

Similarly, the rotation of the field can be computed:

$$\nabla \times \mathbf{S}_{\parallel} = 2\Im(\partial_x H_z^* \partial_y H_z) \quad (5.16)$$

The above deduction has some important consequences for the singularities of the field. Since the time-averaged Poynting vector is computed via the real part, looking at the real parts of the above yields two conclusions: the rotation is not necessarily zero, but the divergence is always zero.

Fig. 5.1 shows the four generic topological configurations that can exist in a general two-dimensional vector field. By looking at the types of critical points of 2D vector fields, we can immediately conclude that only saddle points and centre points are allowed to exist, since nodes and foci have a non-zero divergence [24, 135, 192]. Considering this from a physical point of view, instead of a mathematical one, the reason for this conclusion is immediately obvious. If the divergence at some point would not be zero, it would mean that there are regions where there is a net flow of energy into or out of that region. So for a region with a positive divergence, there would be a net flow of energy out of that region, meaning it has to contain a source of energy. And the opposite holds for regions with a negative divergence. So requiring zero divergence is, in essence, simply a statement of conservation of energy, and since there are no sources or sinks of energy in the system under consideration, only saddles and centre points exist.

5

5.4. RELATION BETWEEN SINGULARITY TYPE AND CHARGE

Since the transverse flow-field is a real 2D vector field, its singularities are mathematically speaking critical points of the field [193]. The study of critical points and their classification in vector fields, where the field vanishes, such as the ones shown in Fig. 5.1 is the domain of vector field topology.

The Jacobian matrix, which is the matrix of partial derivatives and essentially a generalization of the gradient, is a way to approximate the local velocity field in a linear fashion. For the flow field \mathbf{S} the matrix is defined as

$$J = \nabla \mathbf{S} = \begin{bmatrix} \partial_x S_x & \partial_y S_x \\ \partial_x S_y & \partial_y S_y \end{bmatrix}. \quad (5.17)$$

This Jacobian matrix is then evaluated at a critical point in order to gather information about the local structure of the field. The Jacobian matrix is a square matrix, and as such has two (in two dimensions) eigenvectors and eigenvalues associated with it. These eigenvectors constitute a somewhat abstract way of stating that locally, when you move in the field along one of those eigenvectors, you will find that the velocity of the field will be oriented along this eigenvector. And the way you move along these vectors is determined by the eigenvalues of the Jacobian matrix. For instance, a positive and real valued eigenvalue λ_1 would indicate that moving away from the critical point along eigenvector 1, the velocity is found to be oriented in the direction of the movement, meaning away

from the critical point. And a negative real value means that the flow is oriented towards the critical point instead. But eigenvalues are not necessarily real; they can also be a complex number. In the case where the eigenvalue is a complex number, the real part retains its properties as defined above, and the imaginary part describes the movement perpendicular to the eigenvector instead, indicating an element of rotation.

Since the eigenvalues are determined by finding the roots of a second order polynomial for a 2D field, they can either be both real values, or a pair of complex conjugated values. In the case of two real valued eigenvalues a distinction can be made between both eigenvalues having the same sign, or opposite sign. When the eigenvalues have the same sign, that means the vector field is either oriented towards (negative), or away from (positive) the critical point along both eigenvectors, leading to a node of the field, since locally all flow is either directed away from, or towards the point. If the signs are opposite instead, then the vectors are oriented towards the critical point along one axis, but away from it along the other, leading to a saddle point of the field.

In the case of a complex pair, a distinction has to be made between having a real part, or not having a real part. In the case of a real part being present, the real parts will always have the same sign, but their imaginary parts will be opposite in sign. The real parts then imply a flow either away from, or towards the critical point, just like in the case of a node, but the complex part leads to an element of rotation in addition. This type of critical point is a focus. But if the real part is zero, then only the rotational part remains. Then close to the critical point, there is no flow towards, or away from the point, and only pure rotation remains, leading to a centre point. It should also be noted here that we make a distinction between 4 different cases, while in literature often 6 cases are treated instead. In that case they make the additional distinction between attractive and repulsive nodes and foci. For our purpose this extra distinction is unnecessary and is thus omitted for simplicity.

In order to investigate how charge and singularity type are related, we look at the eigenvalues of Eqn. 5.4 at points where magnetic type singularities occur. For this we have to look at the determinant of the Jacobian matrix of the flow field at the singular points. Using the definition for the time-averaged Poynting vector, the first term is found to be given by

$$\begin{aligned}\partial_x S_x &= \partial_x (\Re(E_y H_z^*)) \\ &= \frac{1}{2} \partial_x (E_y H_z^* + E_y^* H_z) \\ &= \frac{1}{2} (E_y \partial_x H_z^* + H_z^* \partial_x E_y + E_y^* \partial_x H_z + H_z \partial_x E_y^*).\end{aligned}$$

Because magnetic type singularities are considered, we can set $H_z = H_z^* = 0$. However, derivatives of these terms are not necessarily zero. Using once more that the fields are monochromatic and time harmonic, this can then be rewritten as

$$\begin{aligned}\partial_x S_x &= \frac{1}{2} (-i\omega |E_y|^2 + i\omega |E_y|^2) \\ &= 0.\end{aligned}$$

In a similar way the other three components for the Jacobian matrix can be computed. By additionally making use of the definition of the Stokes parameter S_3 once more, the Jacobian matrix is computed to be:

$$J = \begin{bmatrix} 0 & -\frac{\omega\epsilon}{2} S_3 \\ \frac{\omega\epsilon}{2} S_3 & 0 \end{bmatrix}. \quad (5.18)$$

Hence we find the determinant of the Jacobian matrix to be

$$\text{Det}(J) = \frac{\omega^2 \epsilon^2}{4} S_3^2 \geq 0, \quad (5.19)$$

which indicates that the determinant of the Jacobian matrix at a magnetic type singularity is always greater than, or equal to zero. A determinant of zero is a special case and will only occur during a saddle-node bifurcation event [193], where singularities are created or annihilated. So for all other cases, we find that it will always have a positive determinant. Since the determinant of a matrix is equal to the product of its eigenvalues, saddle points are the only type where the product of its eigenvalues is a negative number, because their eigenvalues are real valued and with opposite sign. All other possible eigenvalue pairs lead to a positive number when multiplied. Thus we can conclude that all magnetic type singularities are centre points of the transverse flow field, and hence the saddle points have to be comprised of polarization type singularities.

As a corollary of this proof, we can conclude that all magnetic singularities are positively charged in the flow field. However, recall that magnetic type singularities are simultaneously phase singularities of the magnetic field H_z . So they necessarily have to annihilate in pairs of magnetic type singularities in order to conserve topological charge in the magnetic field itself. But topological charge in the flow field is also required to be conserved, and hence two magnetic singularities cannot simply annihilate with each other, as they carry the same charge in the flow field. This seemingly impossible dilemma is solved by considering the criteria of having a zero determinant, which occurs when a creation/annihilation event takes place. From Eq. 5.4 we see that the determinant only vanishes when $S_3 = 0$, so when the electric field is exactly linearly polarized. This linear polarization is also a requirement for polarization type singularities, so by having two magnetic singularities meet at a point where $S_3 = 0$, we also satisfy the relation for having a polarization type singularity. By having two negatively charged polarization type singularities meet at that point as well, all four can be annihilated without changing the net topological charge of the flow-field [194].




Figure 5.2: Measured topological structure, with the phase of the constructed complex scalar field ξ shown in false colour. The phase corresponds to the direction of the Poynting vector, with zero being the positive x direction. Red (Blue) dots are magnetic (polarization) type singularities. A white (black) outline of the dots corresponds to a positive (negative) charge q .

5

5.5. EXPERIMENTAL MEASUREMENTS OF THE POYNTING FIELD

In order to experimentally investigate the singularities of the Poynting field, we have performed near-field measurements on the chaotic cavity sample (see Chapter 2) in order to generate isotropic and monochromatic random waves. We once again use light at free-space wavelength of 1550 nm, and perform a raster scan over a central $17\mu\text{m} \times 17\mu\text{m}$ area of the cavity with a lateral step-size of (17 ± 2) nm. Fig. 5.2 shows a section of the experimentally determined Poynting field, with the phase of the complex scalar field shown in false colour. It is of note here that the phase of the complex scalar field is directly related to the flow direction: the phase represents the angle of the transverse Poynting vector with respect to the positive x direction (horizontal), and hence represents the direction of the flow of energy. Furthermore we have plotted the singularities present in this field with the dots. Here red dots are magnetic type singularities and blue dots are the polarization type singularities, with the white (black) outline indicating a positive (negative) charge. The black curves (streamlines) serve as a visualization of the flow pattern of the field. The streamlines reveal the topological structure around the critical points by showing how the Poynting vector is oriented. For the positive singularities (magnetic type) we observe that the Poynting vector is oriented in a circle, which correspond to centre points of the field. For the negative singularities (polarization type) we observe that from two directions the Poynting vector is oriented towards the singularities, while in the other two directions it is oriented away from the singularity, which is indicative of a saddle point.

To gain further insight into the spatial distribution of the singularities that comprise the general structure of the flow-field, we then determine their pair correlation function $g(r)$. Fig. 5.3 (a) shows the pair correlations for all singularities found in our experi-

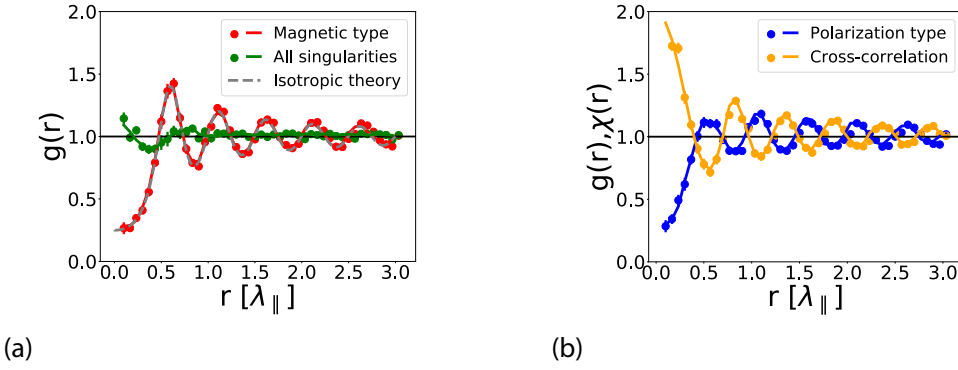


Figure 5.3: Pair correlations for a field confined to the plane. (a) The spatial distribution of magnetic-type singularities (red) and the overall correlations between all singularities (green). Points are experimental data, and the solid lines are from numerical calculations. The gray dashed curve depicts the theoretical distribution for scalar isotropic random waves [24]. Distance is given in units of effective in-plane wavelength λ_{\parallel} . (b) The pair correlation $g(r)$ of the polarization types (blue) and the cross-correlation $\chi(r)$ between the magnetic and polarization types (orange).

5

mentally determined flow-field (depicted as green dots) and for only the magnetic-type singularities with themselves (depicted as red dots). Additionally, the theoretical $g(r)$ for isotropic scalar random waves is plotted as a dashed gray curve. The magnetic type singularities exhibit a clear liquid-like behaviour, consistent with isotropic scalar random waves. On the other hand, when taking all singularities into account, the overall correlation appears to be much weaker. The correlation only deviates slightly from unity at short distances, after which it rapidly approaches unity, where unity signifies uncorrelated behaviour.

Fig. 5.3 (b) shows the pair correlation of only polarization types with themselves (depicted as blue dots) and the cross-correlation $\chi(r)$ of the magnetic and polarization types (depicted as orange dots), where we consider the pair correlations between two singularities of opposite type. The polarization-type singularities, like the magnetic type, exhibit a liquid-like correlation, but with a smaller first-peak amplitude. This damped oscillatory behaviour was already observed from the approximate model for the pair correlations of random waves [57]. Such striking difference with the isotropic wave model hints at a deeper physical distinction between the singularity types. The cross-correlation, in contrast, shows a strong anti-correlation. Distances where there is a higher probability to find a singularity of the same type have a lower probability to find one of the opposite type and vice-versa. This explains why each type exhibits a clear correlation amongst themselves, while the overall ensemble of singularities does not. We were not directly able to discern a deeper physical mechanism for this weakly correlated behaviour. But, since the two types of singularities are oppositely charged, it could be some kind of indication of charge screening for instance [24]. Furthermore, looking at the flow field patterns surrounding the singularities, a case could be made to explain the charge screening as well. It has been shown in Ref. [135] that all centre points are separated by an arm of a saddle point. This suggests that it is more likely for a saddle point to be in between two




Figure 5.4: Generation of an electric type singularity. The Stokes parameter S_3 , which is the degree of circular polarization, is plotted for three different wavelengths. Red dots indicate magnetic type singularities of the Poynting field, blue dots indicate polarization types and green indicates electric types. The black curves give the contours where $S_3 = 0$, which is linear polarization. We see that two $S_3 = 0$ curves cross at the point of a polarization type singularity, leading to an electric type, before splitting again.

5

centre points. This translates to a polarization type singularity separating two magnetic type singularities and hence a higher chance of being closer to an oppositely charged singularity, than a same charged one.

In addition to experimental data we have performed numerical simulations of the flow field by performing random wave simulations. The simulations generate a random H_z field, from which we can then compute the associated Poynting vector field for comparison with the experimental data. In Fig. 5.3 the numerical results are indicated with lines, and we observe extremely good agreement between the numerical and experimental results.

When considering the circumstances under which singularities exist in the flow field, electric types have been treated as a distinct class. It turns out that they will always be created and annihilated as polarization types instead. Precisely, the electric types are found where at the point of a polarization type two $S_3 = 0$ curves cross each other. This can be seen in the following way. As stated in the introduction, locations of linear polarization occur as lines, since they have a one dimensional solution space. When two such lines cross, the only way to satisfy the requirement for linear polarization to occur along both curves is by having the electric field vanish identically, leading to the electric type singularity discussed here. As such they are actually a subset of polarization types with additional restrictions. In Fig. 5.4 we have plotted the Stokes parameter S_3 , along with the singularities of the Poynting field for three separate wavelengths. The black curves indicate contours where $S_3 = 0$, meaning linear polarization. We can clearly see that over the course of the three measured wavelengths that two $S_3 = 0$ contours merge, changing the polarization type (blue dots) into an electric type (green dots), before splitting again and returning to the polarization type once more.

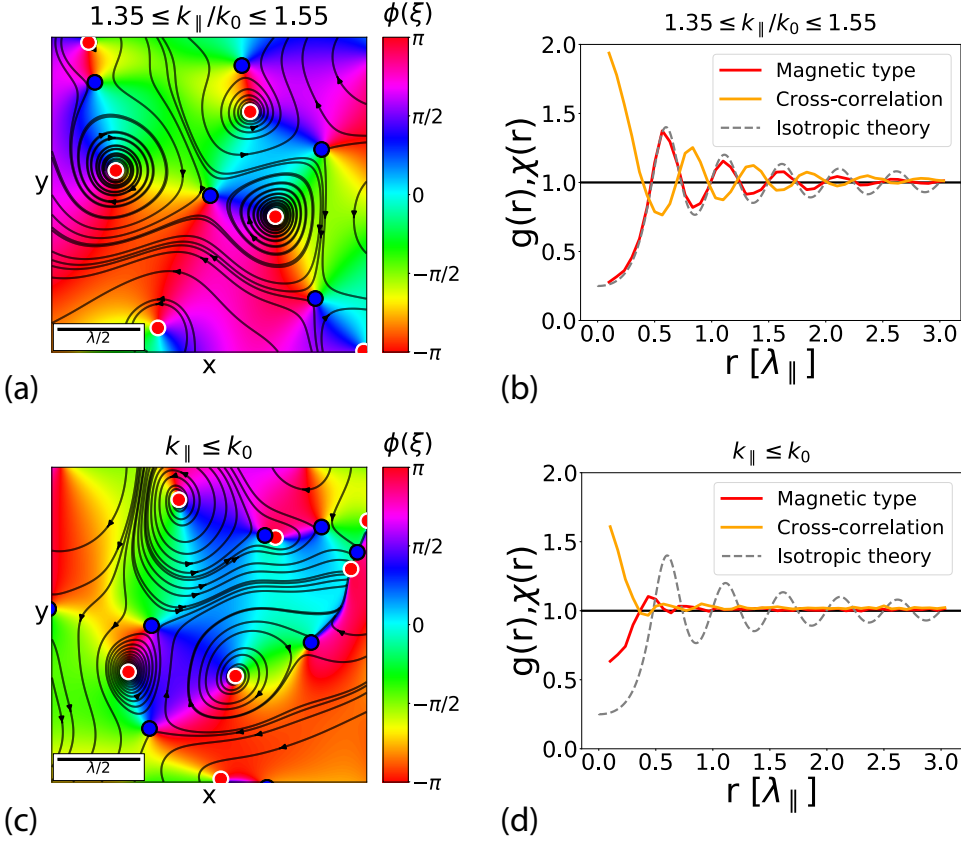


Figure 5.5: Simulated instance for fields with varying in-plane wavenumber, which are either inside ($k_{\parallel} \leq k_0$), or outside ($k_{\parallel} > k_0$) the light cone. (a) and (c) show the topological behaviour, with the phase of the scalar field shown in false colour. Red (Blue) dots are magnetic (polarization) type singularities. A white (black) outline of the dots are positive (negative) charged singularities. The white arrows in (c) highlight examples of topological behaviour that was not allowed in (a). (b) and (d) show the spatial distribution of magnetic-type singularities (red curve) and the cross-correlation (orange curve). The gray curve shows the theoretical distribution for isotropic random waves. Distance is given in units of effective in-plane wavelength.

5.5.1. VARYING THE IN-PLANE WAVENUMBER

To further explore the pair-correlation strength, we study the effect of easing the restriction of a discrete in-plane wavenumber numerically. By allowing for a range of in-plane wavenumbers to build up the field, k_z also has to vary. Therefore, the field is in general no longer invariant along the z -axis, meaning the out-of-plane divergence is not zero everywhere, which necessitates a non-zero in-plane divergence. Thus, the transverse flow-field may exhibit singularities with a topology differing from centers and saddles. If we were to ease the restriction of no divergence, we could expect to find foci and nodes as well for positive singularities. All negative singularities are still going to be saddle points.

Fig. 5.5 (a) and (c) depict simulations of the transverse flow field for different in-plane wavenumber ranges, utilizing the same false colour map and streamlines as in Fig. 5.2. Simulations are done by performing random wave simulations where we additionally allow the ensemble of momentum vectors which are used to initialize the plane waves to have a spread in magnitude in addition to their direction. The presented regions of the flow field are representative for the global field, the difference in apparent singularity density is coincidental. Fig. 5.5 (a) shows the field for a range of in-plane wavenumbers of $1.35 \leq k_{\parallel}/k_0 \leq 1.55$, which still lie outside the light cone and hence still represent bound modes. By observing the streamlines (depicted as the black curves) we see that they still trace a circle around positive singularities, indicating a centre point, and saddles around negative singularities. Hence we observe no significant deviations from centre points and saddle points for the simulated range of wavenumbers, which is easily explained by looking at the allowed k_z values. Even though the magnitudes of k_z are not restricted to a single value, because $k_{\parallel} > k_0$, k_z has to remain imaginary everywhere, leading only to evanescently decaying fields, which we know not to transport energy away from the surface. Hence the real part of S_z remains zero everywhere, and as a consequence the restriction of $\nabla_{\parallel} \cdot \mathbf{S}_{\parallel} = 0$ is still fulfilled. When we consider the case of $k_{\parallel} \leq k_0$, the situation changes drastically. Fig. 5.5 (c) clearly shows that the field structure around magnetic-type singularities can now also be foci and nodes (recall Fig. 5.1 for their topological structure), with two such points highlighted in the field with white arrows.

Fig. 5.5 (b) and (d) show the pair correlation function for the magnetic-type singularities (depicted in red), the cross-correlation (depicted in orange), and the theory for isotropic scalar random waves (gray dashed curve). The fulfilment of $\nabla_{\parallel} \cdot \mathbf{S}_{\parallel} = 0$ for $1.35 \leq k_{\parallel}/k_0 \leq 1.55$ is reflected clearly when considering Fig. 5.5 (b), where we see that the correlation for magnetic-type singularities still exhibits a liquid-like behaviour. However, its oscillations around unity decay more rapidly as compared to isotropic random wave theory, making for an effectively shorter correlation length given by the size of the chosen wavenumber spread.

Fig. 5.5 (d) depicts the correlations for $k_{\parallel} \leq k_0$. The correlations rapidly approaches unity, and do not exhibit any clear oscillations after approximately one wavelength in distance. Thus, we deduce that the singularities only exhibit extremely short nearest-neighbour correlations in the $k_{\parallel} \leq k_0$ regime. Since $k_{\parallel} \leq k_0$, k_z is real-valued, and hence $\partial_z S_z \neq 0$. And since the full flow-field is required to be divergence-free due to energy conservation, we now necessarily have $\nabla_{\parallel} \cdot \mathbf{S}_{\parallel} \neq 0$, leading to the topological restrictions

on the in-plane Poynting vector being lifted. It is also worth mentioning here that since k_z is now real, the field is no longer bound to the plane, and hence we are now looking at a 2D slice of a 3D field again instead of a 2D field, again highlighting the difference between a 2D slice of a 3D field and a 2D field.

One thing of note here is that the full range of topological behaviours is only found when the z -invariance is broken, regardless of confinement. If we assume that the field is z -invariant, such as in the case of a perfect Bessel beam, then despite having a real valued k_z , we still find $\partial_z S_z = 0$. And as a result, the in-plane Poynting field remains divergence-free, with the topological restrictions that follow from this.

5.6. CONCLUSIONS

In conclusion, we have investigated the topological properties and spatial distributions of singularities in the transverse flow-field of random waves, using theory, simulations and experimental data. Specifically, we demonstrated experimentally that by confining the wave field to two dimensions, the topological structure of the flow-field is restricted to saddle and centre type singularities. The topological restriction holds for any z -invariant light field and manifests itself in a liquid-like pair correlation of the magnetic-type singularities of the flow-field, as well as a distinct cross-correlation between all singularity sub-species. When loosening the restriction of a single transverse wavenumber, a significant difference occurs between fields originating from wavevectors fully outside or inside the light cone. For propagating beams such as a paraxial or non-perfect Bessel beam, the full range of possible 2D topological behaviour is found. The found results were corroborated by the change in the pair correlation function of the singularities for the two different cases. For a propagating field, no pronounced pair correlation is present. Conversely, the pair correlation function of any confined field is reminiscent of a liquid-like state, but with the amplitude of the oscillations around unity decaying faster than in the case of isotropic random waves.

6

MEASURING THE OPTICAL NEAR FIELD IN DUAL COLOUR

Small things were important. Seconds were small things, and if you heaped enough of those on top of one another, they became a man's life.

Towers of Midnight - Robert Jordan⁶

In this chapter we describe the experimental progress that was made to measure the near field simultaneously at near infrared frequency and its second harmonic. We show the adjustments that have been made to the near-field microscope in order to accommodate two highly different frequencies, and our choice of optical fibre to use for fabricating near-field probes. Furthermore, we present the first results of measurements where we have simultaneously mapped both IR and visible light in the near field. We show that amplitude, phase and polarization resolution is maintained for both frequencies. Finally, we end with a proposed future experiment and recommendations for measuring a fundamental and its second harmonic in the near field.

6.1. INTRODUCTION

In our normal day to day life, optics and light is fairly intuitive, especially as long as we restrict ourselves to the ray optics picture of light. Light propagates and may be reflected or refracted at an interface (as described by Snell's law), which is easy to visualize (but not per se easy to compute [195]). However, as tends to be common in physics, when we push something to an extreme, this simple view rapidly breaks down [196, 197]. For instance when we consider small length scales, the ray optics picture of light no longer holds [18]. Instead, we need to start regarding light as a wave phenomenon. In a similar vein, this intuitive understanding starts breaking down when the intensity of the light is significantly increased [198, 199]. At sufficiently high intensities, the light starts modifying the optical properties of the medium in which it propagates, which in turn affects the propagation of the light [200]. When this occurs we have entered the realm of nonlinear optics [201].

In general the optical response of a material to an applied electric field can be described through the polarization P as a power series in the electric field strength E :

$$P = \chi^{(1)}E + \chi^{(2)}E^2 + \chi^{(3)}E^3 + \dots, \quad (6.1)$$

where $\chi^{(i)}$ are constants of proportionality and $\chi^{(i)} \gg \chi^{(i+1)}$. These χ are known as the nonlinear susceptibilities and are a material property. Under normal circumstances all terms beyond $\chi^{(1)}E$ are negligible in comparison, leaving us with a linear system. Only when the field strength becomes sufficiently large do nonlinear terms become relevant, and each progressive term comes with its own possibilities for nonlinear effects. With the realization of the laser in 1960 by Maiman [202], high electric fields became readily accessible, and probing nonlinear effects became viable along with it. This in turn has lead to the Nobel Prize in physics for Bloembergen, Schawlow and Siegbahn for the development of laser spectroscopy in 1981 [203]. And many more Nobel prizes have been awarded for research made possible by lasers such as the realization of Bose-Einstein condensates (2001) and atomic clocks (1989). The $\chi^{(2)}$ term describes effects which involves two photons, and can lead to second harmonic generation (SHG), optical rectification (OR), sum-frequency generation (SFG), and its brother difference-frequency generation (DFG). These are all the combinations that can occur with two photon processes. It is easy to see that from here the amount of possible combinations starts to increase rapidly when considering more than 2 photons. For 3 photons, effects arise such as third harmonic generation (THG), four-wave mixing (FWM) and the optical Kerr effect [201], in which the refractive index of the medium becomes intensity dependent, leading to self-focusing of the light [200]. Thankfully, more than 3 terms rarely need to be considered.

Now while theoretically the $\chi^{(2)}$ is the first nonlinear term to become relevant, symmetry may force it to be zero instead. In systems with inversion symmetry, all even-order terms of nonlinear susceptibilities have to vanish identically [201]. So for instance a simple W1 waveguide, which is just a 2D photonic crystal with one row of holes missing (see Section 2.4.1), has a mirror symmetry axis through the centre of the waveguide, and hence second order nonlinear effects can be discounted.




Figure 6.1: Comparison of the throughput of infrared light through two different pulled fibres. (a) shows the light transmitted through an SMF28E fibre which is single mode for telecom frequencies. (b) shows the light transmitted through a SM980 fibre, which is single mode for visible frequencies. Both images were taken at 1 ms exposure time and at the same laser power in order to serve as a direct comparison for the throughput amounts in both fibre types.

But the statement above holds for the structure as a whole, but not necessarily at the local scale, or at interfaces where this inversion symmetry is broken [204, 205]. Local differences in structure can break this inversion symmetry on the local scale, allowing for even-ordered nonlinear effects to occur locally, as long as they add up to zero overall. For instance in a W1 waveguide, just considering the right hand side of the waveguide we find no such inversion symmetry, and hence $\chi^{(2)}$ effects are not forbidden when only considering this part. The same then holds for the other half of the waveguide. But when the two halves are combined, they necessarily add up to zero, as the generated SHG of each half of space is exactly in counter-phase with each other.

In this chapter we lay the foundation for an experiment that allows us to probe these local nonlinear effects by measuring a field and its second harmonic at the same time in the near field in order to investigate where and how these second-harmonic fields are generated in the sample. We will be showing the results of a proof-of-concept measurement, where we succeeded in measuring both IR and visible light simultaneously. With that, we prove that it is indeed possible in our system to probe the near field of both the fundamental and its second harmonic in the same measurement, without having to make adjustments to accommodate either frequency. Furthermore, we show that we are able to collect both colours simultaneously, while retaining amplitude, phase and polarization resolution. We include an extensive analysis of the measurement technique, as well as its limitations. Finally we conclude the chapter with an outline of the experiment for measuring a pulse and its second harmonic, and how to perform it in the future.

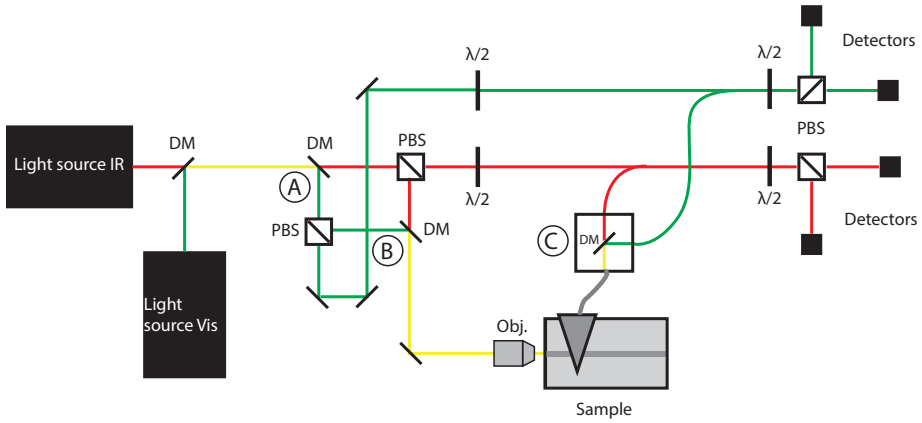


Figure 6.2: Schematic of the setup used for dual-colour NSOM. Here we make use of two different lasers, that are coupled into the sample at the same time, and are picked up by the same tip. Afterwards they are split again via a dichroic mirror and are recombined with their reference paths, and finally sent to their respective detectors. Green lines represent visible light only, red lines represent infrared light only and yellow lines represent combined visible and infrared light. The labels A, B and C highlight a few important elements for measuring dual colour, with C being the crucial addition to facilitate it.

6

6.2. MODIFICATIONS TO THE NEAR-FIELD MICROSCOPE

As a first step towards measuring dual-colour, we need to set the parameters of the experiment, and adjust the near-field microscope to accommodate the collection of two colours simultaneously. As stated, the ultimate goal is to measure both the fundamental and second harmonic of a pulse of infrared light at the same time. To this end we want to demonstrate that we are capable of measuring the fundamental frequency at 1550 nm, as well as its frequency doubled part at 775 nm simultaneously in the near field. For our proof-of-concept experiment, we have opted to use two CW lasers. For the IR laser a Santec tunable semiconductor laser (1480 nm - 1640 nm in air) has been used and as a visible laser a HeNe laser (632.8 nm in air). Since the wavelength of the HeNe is smaller than what we ultimately need to measure in the visible, if we can show sensitivity to both of these wavelengths, then it is clear that 775 nm in air is well within our range as well.

The choice of optical fibre from which the near-field probes are fabricated is crucial. No commercial fibre is designed to be single mode, and optimised for both IR and visible at the same time. While there are techniques to extend the single-mode regime to include both colours of interest, without sacrificing significant throughput, it is in practice hard to do [206]. Therefore we either utilise a fibre that is optimized for IR, or a fibre optimized for visible. The drawback of choosing the IR fibre is that the core diameter is large enough to facilitate multiple modes in the visible. As a result of this, the measured electromagnetic field gets projected onto the multiple modes that are allowed to propagate through the fibre, instead of only the fundamental. Only under very specific (and unrealistic) circumstances are we then able to retrieve phase and polarization information from this signal. On the other hand, when choosing a fibre designed for visible light, the core diameter is small enough to significantly attenuate the IR signal instead.

So ideally both colours can be transmitted through the fibre single-mode. To start, we have investigated the transmission of IR light through a probe made with a fibre designed for visible light. We have based our choice of optical fibre on the graphs presented in Figure 2 of ref. [206], where the transmission as function of wavelength and bending radius is shown for several types of optical fibre. Both 780HP (Fig. 2 (c)) and SM980 (Fig. 2 (d)) are shown to provide throughput across the desired frequency range, while remaining single mode making them promising candidates. We have opted to use SM980 fibres, since it shows comparable behaviour to the HP780 in the visible range, but is less strongly attenuated in the infrared regime.

To check that the transmission of IR light through this fibre is sufficient, we couple light in from the back of the tip (illumination mode), and measure the transmission through the apex of the tip with a camera setup. In Fig. 6.1 (a) and (b) we show the transmission of IR light through respectively an SMF28E fibre, which is optimized for 1550 nm, and an SM980 fibre. Both images were taken at 1 ms exposure time and the same laser power in order to allow for a fair comparison. From the figure it is clear that the both types of fibre are able to transmit IR light.

After this verification that our SM980 probes are able to transmit both visible light as well as IR light with the same probe, the next step is to utilise them to measure both colours at the same time in the near field. For this experiment modifications to the setup need to be made to accommodate this feature. Fig. 6.2 shows a schematic of the setup used to measure in dual-colour, with some areas labelled to refer to in this section. This can be compared with Fig. 2.2, which shows the near-field microscope setup at its core for a single colour. With the presence of a dichroic mirror (A), the light is automatically filtered into their respective reference paths at the point where the split between signal and reference paths is made. So this already made it possible to have two reference branches active at the same time. The signal path for the two different colours is briefly separated to allow for adjustments via waveplates, before being subsequently recombined and being sent to the sample (B). Both (A) and (B) were already present in the setup, and hence coupling two different beams into a sample was also already possible. However, we encounter a challenge when the light gets picked up with the probe: after the tip two fully disjointed paths are necessary that lead to either the IR or visible detectors. The pigtailed¹ fibre end of the probe can be connected to either the fibre coupler that goes to the IR detectors, or the fibre coupler that goes to the visible detectors, but not both simultaneously. So in order to be able to measure both at the same time, the optical signals of the different frequencies need to be de-multiplexed.

We de-multiplex the signal with a box, consisting of 3 fibre couplers on 3 of the 4 sides, and a dichroic mirror in the centre (C). With the dichroic mirror after the tip fibre, we can split the collected light into a visible and an IR channel once again with their respective detectors. While this is an effective way to re-split the different colours, the resulting loss in efficiency in the signal path is a disadvantage of this method. Measurements of incoming and outgoing power indicate both branches having a coupling efficiency of around 50%.

¹Pigtails are a type of connector that we use to couple optical fibres to either another optical fibre or to a fibre coupler to couple free space light into, or out of the fibre.

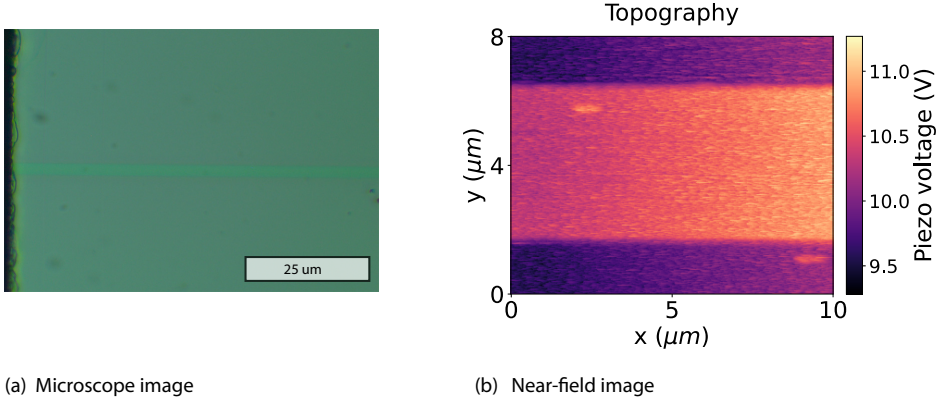


Figure 6.3: Images of a SiN rib waveguide that is 5 microns across. Figure (a) shows an optical microscope image of a waveguide taken with an objective at 50 times magnification. (b) shows a waveguide as measured on the topography channel in our near-field microscope. The lighter region around the centre of the image signifies the waveguide. Increased voltage means increased height.

6

Having made the necessary adjustments to the near-field microscope itself, we can now investigate compatibility with the electronics used to read out the detected signals. Since both signals are split into two detectors to allow for the detection of the individual electric field components, four detectors are simultaneously read out. We employ three lock-in amplifiers from Stanford Research systems (type SR830 DSP), and in addition we employ one lock-in amplifier from Zurich Instruments (type HF2LI). Unfortunately, that does mean that either IR or visible has to be read out using two different lock-in amplifiers. In Section 6.5.1 below, we have performed an analysis where we compare the two different types for both IR and visible light. We have opted to use a separate digital Zurich Instruments lock-in (type MFLI) for the visible instead of the third Stanford lock-in. So the IR signal is measured using 2 analogue lock-in amplifiers, while the visible signal is measured using two digital lock-in amplifiers.

Finally, we require enough data acquisition (DAQ) card slots for all the lock-in output channels. Ideally, we want to be able to record all possible information on the computer, but there were not enough slots available to accomplish this. By opting to not store the DC signals of the detectors we have reduced the amount of required channels to a number compatible with our DAQ card setup. In order to prevent this necessity of choice in the future, we make two separate recommendations. Currently we read out both the X and Y channels of the lock-in, as well as the amplitude channel. By computing the amplitude channel in software instead, it no longer needs to be read out directly, freeing up channels. Alternatively, adding an additional DAQ card circumvents this entirely.




Figure 6.4: Measured amplitude and phase of both IR and visible light simultaneously on one of the detectors for a SiN waveguide that is 5 microns in width. a (b) shows the measured amplitude of the IR (visible) field over a length of 15 microns along the waveguide, and 8 microns in the transverse directions. c (d) shows the measured phase of the IR (visible) field for the same spatial map. The black circles highlight the position of a phase singularity in each of the phase maps.

6.3. SIMULTANEOUS OBSERVATION OF TWO COLOURS OF LIGHT IN THE NEAR FIELD

To test the capabilities of our near-field microscope to simultaneously detect light of two different colours, we make use of a model system. A silicon nitride (Si_3N_4) rib waveguide is an ideal candidate for this, because it is transparent for both the desired colours, and can readily be simulated in order to compare the measured fields with simulations. The silicon nitride rib waveguide, consisting of a 300 nm thick layer of Si_3N_4 , with the rib protruding 20 nm above the surface. The SiN has a refractive index of $n = 2.04$ (2.00) at $\lambda = 632.8$ (1550) nm. This layer sits on a glass substrate with a refractive index of $n = 1.46$ (1.44) at $\lambda = 632.8$ (1550) nm. An image of the sample is shown in Fig. 6.3, where (a) shows an image taken with a regular microscope with an objective at 50 times magnification, and (b) shows what a waveguide looks like in the near field instead. The colour bar on the near-field image indicates the measured voltage, which corresponds to the tip position in the direction perpendicular to the surface of the sample, with higher voltage meaning tip further retracted. Since the tip maintains a constant tip-sample distance, this map translates to a height topography map: the waveguide can be observed as a horizontal band of higher voltage in the centre of the image. The width of the rib waveguides on our sample varies between 5 microns for the larger ones (pictured in Fig. 6.3), to less than 2 microns for the smaller ones. These rib waveguides are able to support guided modes in both the IR and visible regimes, which we have verified with mode simulations. These simulations are performed using a 2D mode solver, using a variational

effective index approximation [207]. For the initial experiment we are interested mainly in showing the viability of measuring the near field in dual-colour. As such we have opted to select the larger waveguide size of about 5 microns, which allows for more light to be coupled into the waveguide and hence provides more signal to be picked up. Since we are trying to couple beams of two different frequencies into the same waveguide using the same objective, we are going to encounter chromatic aberrations: the distances from the objective at which the beam is focused is going to be different for the two colours as a result of the refractive index of the objective being frequency dependent. Thus we will never be able to focus both beams onto the waveguide entrance perfectly. The question is how badly this affects our ability to couple to the sample, which needs to be answered experimentally.

In Fig. 6.4 we show the results of a measurement, where (a) shows the measured amplitude for one of the lock-in amplifiers for IR light, and (b) shows the measured amplitude for one of the lock-in amplifiers for the visible light. Fig. 6.4 (c) and (d) show the measured phase for the same field for respectively IR and visible. One striking feature for the amplitude images is that for both IR and visible, the highest amplitude is measured towards the sides of the waveguide. This can be understood when considering the modes that are supported by the waveguide. From the mode simulations, we expect to find the TE_{00} and TE_{01} for both colours, which then form a spatial beating pattern together due to their slightly different k -values. Since they lie close together in k -space, the wavelength of the beating pattern will be long range, since the wavelength of the beating is given by $\lambda = \frac{2\pi}{\Delta k}$, with Δk the difference in k -values of the mode. And indeed, if we measure over a much longer range along the waveguide, we observe the amplitude shifting from one side of the waveguide to the other. In the phase images it is also possible to discover phase singularities once again, showing that we have retained the phase resolution necessary to observe them as well.

This long-range shift of the amplitude between the sides of the waveguide is shown in Fig. 6.5, where we have measured the fields along the wavelength over a range of $180\mu\text{m}$. While the interference pattern between the TE_{00} and TE_{01} modes is the most clear, more spatial beating patterns can additionally be observed here, most notably a fairly high spatial frequency beating in the visible measurement. This higher spatial frequency is indicative of a larger difference in k -values of the mode, which likely stems from interference between a TE_{0x} and TE_{1x} mode.

The question of how badly chromatic aberrations affect our ability to couple to the sample has now been answered experimentally: since we are able to observe both frequencies of light in the same waveguide, chromatic aberrations are clearly not strong enough to meaningfully affect the quality of the signal.

In Figures 6.6 and 6.7 we have plotted the power spectral density of the reciprocal space at $k_y = 0$ for respectively IR and visible light for a scan range of $60\mu\text{m}$ along the direction of propagation. We have taken the $k_y = 0$ cross section, since the x direction is the direction of propagation in our measurements. As a result k_x is the main contributor to the k -value of the modes. In these plots the light cone is indicated by a vertical dashed line, and peaks in the spectrum with a black dot. In these graphs, we have filtered out everything below the light cone, as these cannot be attributed to guided modes.




Figure 6.5: Measured amplitude both IR and visible light over a length of $180\text{ }\mu\text{m}$ along the waveguide, which is $5\text{ }\mu\text{m}$ in width. For both we can clearly observe the beating pattern between the TE_{00} and TE_{01} modes, leading to a slow shift of amplitude between the top and bottom of the waveguide. Please note that the colour-bar indicates a normalized amplitude and that both images are normalized with respect to themselves only, and no relative amplitude information can be inferred here.

Each peak corresponds to a mode in the measured field. The attribution of the modes (see labels) has been achieved through mode simulations. Due to the high density of modes in the reciprocal space for this geometry, especially for the visible light, it is impossible to properly resolve all the individual modes in this measurement. For IR we expect to find 4 modes, whereas we are only able to discern two peaks in the power spectrum of Fig. 6.6. In the visible regime, we expect to find 16 modes, while in Fig. 6.7 we are only able to discern 4 peaks. Finally, in Fig. 6.6 we see signal at positions not indicated by modes. Since we do not expect to find guided modes there as supported by simulations, we attribute this to noise.

Caution needs to be taken when drawing conclusions from this graph: if the measurement has not been performed perfectly along the direction of propagation, the modes will not be exactly at $k_y = 0$ in the Fourier transformed data. Furthermore, modes that have a transverse profile that is not at a maximum in the centre of the waveguide, such as the TE_{01} mode, will not have a peak at $k_y = 0$, and hence might not be detected in this way. In fact, in a perfect measurement, the TE_{01} should be zero on the $k_y = 0$ axis and should not be detected at all in this way. To properly judge a Fourier spectrum and determine its modes, one will always need to consider a 2D selection of reciprocal space, rather than a 1D slice. But this slice is still useful to gauge where in the k -space the modes lie, and hence where we can zoom in to properly investigate. And more importantly, these plots show us that all the modes that we can find lie outside of the light cone. In other words, this is proof that what we have measured are actually bound modes, and not free space light. With this measurement we demonstrate our ability to measure both amplitude and phase for both colours at the same time in the near field on a waveguide of $5\text{ }\mu\text{m}$ in width.




Figure 6.6: Power density of the k-space at cross-section $k_y = 0$ for the measured infrared electromagnetic field on one lock-in amplifier. Measurement has been performed on a waveguide 5 microns across, and for a scan range of 60 microns along the waveguide. Modes can be clearly seen that lie outside of the light cone, indicating guided modes. Peaks are too close to resolve, but from the mode simulations we know which modes are expected to occur.

6




Figure 6.7: Power spectral density of the k-space at cross-section $k_y = 0$ for the measured visible electromagnetic field on the lock-in amplifier. Measurement has been performed on a waveguide 5 microns across, and for a scan range of 60 microns along the waveguide. The $TE_{0,x}$ modes are barely separable, while the $TE_{1,x}$ are simply too close together to be able to resolve them individually. From mode simulations we know which modes are possibly present, and the peaks are labelled to indicate them. Note that we have denoted the indices of the mode for combined peaks using a comma solely for visual clarity when denoting the range for the second index.

6.4. SEPARATION OF THE IN-PLANE ELECTRIC FIELD COMPONENTS

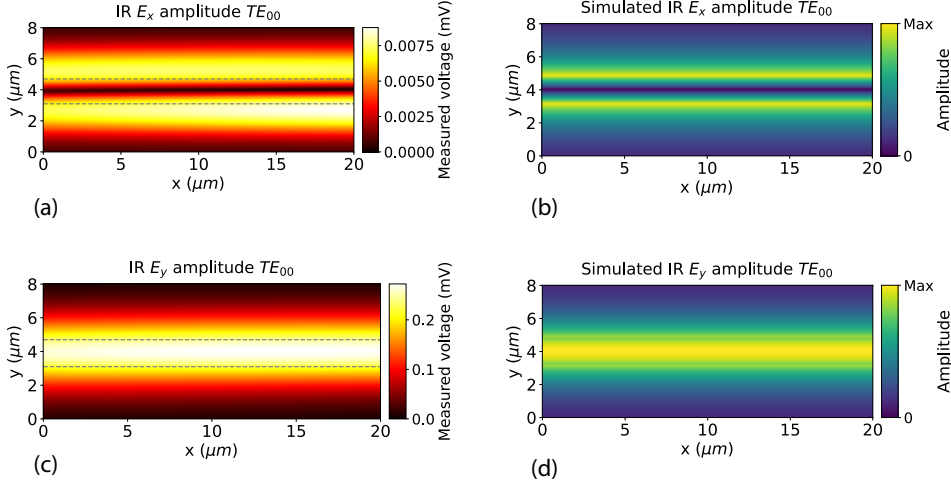


Figure 6.8: Amplitudes of the TE_{00} modes for IR fields measured and rotated using a Jones matrix for an elliptical phase retarded for optimal split into E_x and E_y compared with the fields of the simulated mode structure. a (b) shows the measured (simulated) E_x field, while c (d) shows the measured (simulated) E_y field. Note that the measured fields are in units of the measured voltage, and the scale differs greatly between the two fields. For ease of view we have opted to scale the colours along with them. The simulated fields are in arbitrary units instead.

Having coated probes ideally also allows us to discern the two in-plane electric field components [62]. By rotating a $\lambda/2$ waveplate along with polarizing beamsplitter cube, it is possible to redistribute the signal over the two detectors according to their polarization. This allows us to then also separate the measured fields into E_x and E_y (see Section 2.5). In order to check that we are indeed able to separate E_x and E_y , we want to compare the measured fields with the performed mode simulations. For this, we need access to the individual modes of the measured field. In order to select an individual mode, we Fourier transform the data and look for peaks in the k -spectrum, which correspond to the modes of the field. The k -values of the detected peaks are compared to mode simulations of the system in order to discern which peak corresponds to which mode. Then by selecting only the part of Fourier space corresponding to that particular mode and performing an inverse Fourier transform on that selection, we retrieve only that mode in real space.

To aid in this endeavour, zero-padding is used on the data by padding our data array with additional zeroes around the outside, and serves to more easily distinguish features from each other in Fourier space. It is a common data analysis tool used in image visualisation, and is used in many fields such as biomedical imaging [208, 209]. By introducing




Figure 6.9: Amplitudes of the TE_{00} modes for visible fields measured and rotated using the Jones matrix for optimal split into E_x and E_y compared with the fields for the simulated mode structure. Measurement has been performed on a waveguide approximately 1.6 microns across. a (b) shows the measured (simulated) E_x field, while c (d) shows the measured (simulated) E_y field. Note that the measured fields are in units of the measured voltage, and the scale differs greatly between the two fields. For ease of view we have opted to scale the colours along with them. The simulated fields are in arbitrary units instead.

more pixels on which to perform a Fourier transform, the pixel size in Fourier space is lowered, meaning smaller frequency bins. Please note that it is not a way to arbitrarily improve the k-space resolution of your data. If two frequencies are too close together to be resolved in the original image that is Fourier transformed, they won't be resolved using this technique, as it does not add any data. But it does allow for two peaks that are already resolvable to be more easily separated.

In addition, a windowing function is needed to avoid a step function when going from data to zeroes, which would introduce unwanted frequencies in Fourier space. Windowing functions are generally bell-curve shaped, with highest value in the center, and tapering towards the edge. This function gets multiplied with the data before the zero-padding occurs, as it should be applied to the data, not the padding. In this way the amplitude of the data gets reduced towards the edge for a more gradual transition. Two functions that are most commonly used are Hann and Hamming windows. These two windows are almost identical, except that Hamming window does not go completely to zero at the edge, while the Hann window does. We have chosen to use a Hamming window for our analysis. For an overview and more in depth treatment of windowing functions, see ref. [210].

The trade-off with using a waveguide 5 micron in width is that allows us to couple more light into it, but supports many more modes. Additionally, the modes are closer together in reciprocal space, making them more difficult or even impossible to resolve, and with

Visible (632.6 nm)			Infrared (1550 nm)		
Mode	k (μm^{-1})	n_{eff}	Mode	k (μm^{-1})	n_{eff}
TE ₀₀	18.98	1.91	TE ₀₀	6.62	1.63
TE ₀₁	18.89	1.90	TM ₀₀	5.97	1.47
TM ₀₀	18.43	1.86			
TM ₀₁	18.31	1.84			
TE ₁₀	15.28	1.54			
TE ₁₁	15.07	1.52			

Table 6.1: Overview of all the modes expected for a waveguide $1.6\mu m$ in width, with their propagation constants k and effective refractive index n_{eff} .

it complicates our ability to separate the in-plane electric field components. In order to reduce the amount of modes present, and allow for better k-space separation, we opt to measure on a smaller waveguide, which is $1.6\mu m$ across instead. For IR this means that we only have the TE₀₀ and TM₀₀ modes, and for the visible we have also greatly reduced the multitude of modes present from 16 down to 6. A full overview of all the modes that we expect to occur in this waveguide can be found in Table 6.1. By utilizing the full scan range of the system along the propagation direction of $180\mu m$, we maximize our k-space resolution in order to attempt to resolve all the individual modes. From here we can then attempt to further reconstruct the E_x and E_y fields based on the known mode profile.

We start with the simplest mode, which is the TE₀₀ mode, for which we know that we mostly have an electric field whose in-plane component is perpendicular to the direction of propagation. Since $\nabla \cdot \mathbf{E} = 0$ for a TE mode, we expect most of the amplitude to be perpendicular to the direction of propagation (E_y), with little light in E_x [211]. In Fig. 6.8 and 6.9 we show the results of rotating the data using the Jones matrix for an elliptical phase retarder (see Section 2.5) for the selected TE₀₀ mode for respectively IR and visible fields. For both the IR and the visible measurements (a & c) we observe that E_y is a uniform band with the maximum of amplitude at the centre of the waveguide, whereas the E_x field has a minimum at the centre of the waveguide instead. Please also note the difference in amplitude between $|E_x|$ and $|E_y|$. As it should be, the majority of amplitude is in the $|E_y|$ channel. A clear difference between IR and visible is their lateral confinement, with visible light being much tighter confined to the waveguide than the IR light is. The simulated data (b & d) is again generated using a 2D mode solver [207], and the values that are plotted are at 20 nm above the surface of the sample, which is roughly the height of the tip during measurements. Keep in mind that in the simulations the waveguide is a perfect step in height, and hence shows a sharp transition between waveguide and slab, while in the experiment the measurement does not allow for such a perfect step. This explains why the amplitude exhibits such a sharp feature at the waveguide edge for the simulations, which is smoother in the measured data.

By showing that we indeed are able to separate the E_x and E_y fields, we show that in addition to amplitude and phase resolution, we have also maintained polarization resolution when performing these dual colour measurements.




Figure 6.10: Comparison of the amplitudes for the TE_{11} mode. Measurement has been performed on a waveguide approximately 1.6 microns across. (a) and (b) show respectively the E_x and E_y fields when the Jones matrix is optimized for this mode. (c) and (d) show respectively the E_x and E_y fields when the Jones matrix is optimized for the TE_{00} mode instead. A clear difference can be observed between the two E_x fields.

6.4.1. LIMITATIONS OF THE SEPARATING THE FIELD COMPONENTS

Even though we have shown that we are able to resolve the in-plane electric field components, we have found that there are limitations to this. We have found the values for the parameters used for the Jones matrix that allows us to split the TE_{00} mode into E_x and E_y as we expect from simulations. However, while it also works well for TE_{01} , using the same parameter values do not lead to proper splitting for the TE_{10} and TE_{11} modes in the visible. It is, however, possible to find a different set of parameter values for the TE_{10} and TE_{11} mode, that do allow us to split those modes nicely into the expected E_x and E_y . If our system was able to separate the fields perfectly, we should find one singular set of values with which we are able to split all modes simultaneously. The fact that this is not the case means that we are not able to resolve the fields perfectly.

Fig. 6.10 shows the difference between the best split we are able to achieve by a Jones matrix operation, and using the same matrix as for the TE_{00} mode for the TE_{11} mode of the visible field. Fig. 6.10 (a) and (b) show the result when using the same parameter values for the Jones matrix used to split the TE_{00} mode optimally, while 6.10 (c) and (d) show the result when using the same parameter values for the Jones matrix used to split the TE_{11} mode instead. While from (b) and (d) little difference can be observed in the field for $|E_y|$, a clear difference can be observed between (a) and (c) for the $|E_x|$ field instead. In (a) we find a stripe of high intensity in the center, as we expect for the TE_{11} mode, which is notably absent in (c). This shows that we are able to separate the electric field components, but are not able to perfectly split them using only one Jones matrix operation for all the modes.

There are multiple factors that could lead to an imperfect split in electric field components. For instance, the aperture of the tip that we use is not perfectly circular, which means the two polarization orientations will not be detected equally. We know that, besides the electric field, we additionally detect the magnetic field, where there is a linear relation between $E_{x(y)}$ and $H_{y(x)}$ [81]. From the performed mode simulations we can additionally extract the constant of proportionality between the components, which turns out to be different for different modes: we find $H_x = 0.005E_y$ for T_{00} , whereas we find $H_x = 0.004E_y$ for T_{10} instead. This difference in proportionality in combination with a slightly elliptical aperture of the near-field probe could offer an explanation for the observed difference in Jones matrix parameters required for the separation of the field components. Furthermore, if we cannot perfectly separate the modes through Fourier filtering, then there can be a little mixing from the TE_{11} into the TE_{10} mode for instance, making it impossible to reconstruct a perfect TE_{10} mode.

6.5. TECHNICAL ANALYSIS

In this section we will be going through some more detailed analysis of the technique we used to measure two colours at the same time in order to probe how well it performs compared to normal near-field microscopy, and where its limitations lie.

6.5.1. COMPARISON OF LOCK-IN AMPLIFIERS

We have used two different types of lock-in amplifiers for the measurements presented in this chapter. We have used three lock-in amplifiers from Stanford Research Instruments (type SR830 DSP) and used the remaining channel on our Zurich Instruments lock-in amplifier (type HF2LI) for our fourth channel. The other channel on the Zurich Instruments lock-in amplifier is used for the tuning fork feedback loop, leaving the other free for the second visible detector. This situation provided a good opportunity to compare the two types of lock-in amplifiers and see if there are differences in the quality of the measured signal. First of all, while the Stanford instruments are dual-phase analog lock-in amplifiers, the Zurich Instruments lock-in amplifier is a digital lock-in amplifier instead. In order to test the difference, we have put the output of the same detector on both a Stanford and Zurich at the same time, so they receive the exact same signal. Furthermore we ensured that the filter strength and their integration time are equal, so any difference in output from the lock-in amplifiers is purely due to the lock-ins themselves.

Fig. 6.11 shows the comparison of the two different lock-in amplifiers. Considering figures 6.11 (a) and (c), we observe that for the IR measurements the difference is negligible, and both types of lock-in amplifier gave virtually the same output. While there is a difference in measured voltage between the two, the results are qualitatively the same. As such, there is no reason to choose one type over the other. However, considering Figures 6.11 (b) and (d) for the visible measurements, we see that there is a significant difference in the quality of the data that was output from the lock-in amplifier, with the Zurich being clearly superior. The higher measured voltage in the Stanford LIA is mainly due to the high intensity events that can be observed as horizontal lines of higher amplitude. We suspect that this occurs due to the HeNe laser introducing unexpected frequencies

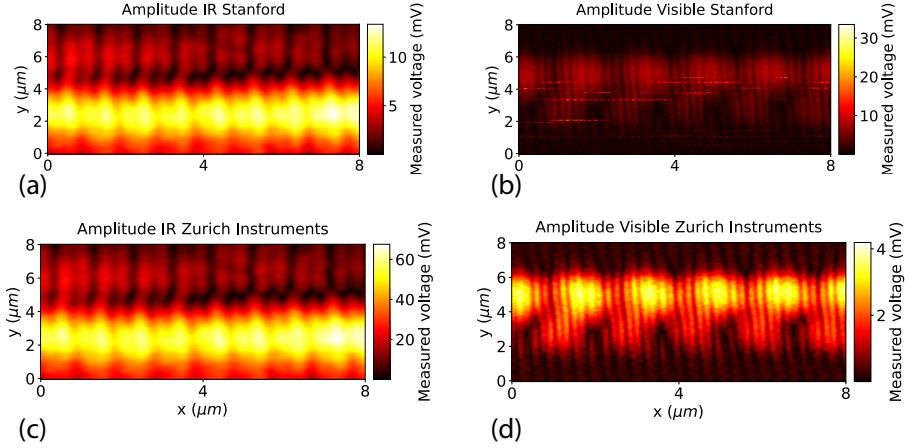


Figure 6.11: Comparison between Stanford instruments and Zurich instruments lock-in amplifiers. Measurement has been performed on a waveguide 5 microns across. (a) and (b) show the measured amplitude fields for respectively the IR and visible fields with the Stanford LIA. (c) and (d) show the same IR and visible fields for the Zurich LIA instead. Comparing (a) to (c) we find little difference. However, (b) and (d) show a significant difference, with the Zurich LIA performing significantly better.

6

around the lock-in frequency, which additionally are not stable in, but rather change frequency. When this noise overlaps with the lock-in frequency, the noise is detected as signal, leading to the high intensity events in the measurement. These events are, due to their unstable nature, not predictable, but on average interfere with the measurement around 5 times for a measurement of half an hour. The digital nature of the Zurich lock-in amplifier is better able to filter out these events. We hypothesise this is the case because there is noise clearly present in the data that gets output from the Stanford that is not present in the data that is output from the Zurich. As such, we recommend to switch to a digital lock-in amplifier for further experiments with this type of HeNe laser. From here on, we have switched out the third analogue lock-in amplifier for another digital one.

6.5.2. SIGNAL-TO-NOISE RATIO & CROSSTALK

We additionally want to investigate the noise present in the performed measurements. Specifically we look at noise from the lasers themselves, as well as noise introduced by the measurements techniques such as background noise on the detectors and electronic noise. Furthermore we investigate if there is any crosstalk as a result of one frequency of light being picked up by the detectors meant for the other frequency.

In order to investigate the noise levels in both colour branches, we have repeated the measurements of the fields, but additionally switched off one of the lasers. This serves the purpose of determining if there is any IR light being picked up by the visible branch and vice versa. Finally we have done one measurement with both lasers off, in order to

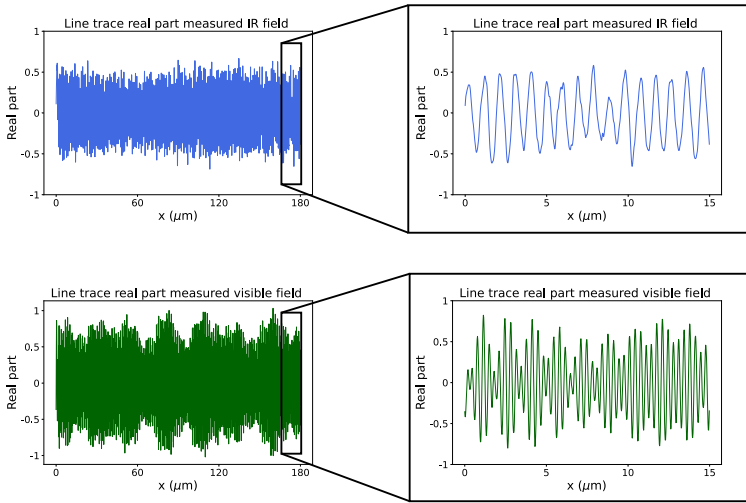


Figure 6.12: Line traces of the real part of the field through the centre along the entire length of the waveguide, and with a zoomed in section. Top shows the trace for the infrared field, and the bottom shows the trace for the visible field. Both clearly show oscillatory behaviour as the dominant contribution to the measured signal.

measure the background noise level without any crosstalk possibilities. Furthermore, out of contact measurements were performed in order to see if any free space light could potentially be picked up. Predictably, some IR light could be detected further away from the surface, while no visible was picked up. This is easily explained by the fact that the evanescent tail of the IR light is much larger than that of the visible light, and hence can be more easily probed from further away.

In order to investigate the possibility of crosstalk between the IR and visible branches, we have compared measurements where both lasers were switched off with a measurement of only the laser on for the other colour. In this way we can compare a measurement of the "wrong" colour, with the measured background. The dichroic mirrors that were used are of the type DMLP950T (950 nm cutoff longpass), which reflect visible light, while transmitting IR light. In the visible regime the dichroic mirror reflects over 99.5% of the light, but also reflects approximately 2% of the IR light in addition, leading to some IR light entering the visible branch. And the visible light that gets transmitted through the dichroic mirror additionally enters the IR branch. For the reference branch of the setup, the AOM's are expected to filter out most of that light due to the difference in angles at which light of different frequencies get modulated. In the signal branch however, the dichroic mirror is what splits the two colours into their respective branches, and bleed-through light goes straight to the detectors instead. Due to the heterodyne detection scheme, we expect that this light will subsequently be mostly filtered by the lock-in amplifiers. Indeed, when we compare the measurements without laser and with the "wrong" laser only, no significant difference in signal or even noise is found. So we can conclude from this that our measurements contain no significant contributions from the opposite colour.

Ideally we would like to put a number on the signal-to-noise ratio, however, it has proven to be difficult to do so in practice. So instead we will consider it in a more qualitative way. By taking a line-trace through the centre of the waveguide, we can investigate the structure of the measured field. For the IR field, we know that we only expect the fundamental TE and TM modes, of which the TE mode is dominant. So when looking at the real part of the measured field, we expect to mainly see a cosine with a certain amplitude and a well defined frequency. For the visible field, we know that there are a lot more modes that interfere, and hence the pattern should be more complex as a result. Still we expect to see mainly cosine behaviour of the fundamental frequency, but with more beating patterns occurring as a result of the amount of modes.

In Fig. 6.12 we have plotted the line trace for both the IR and the visible field for the full length of the waveguide, as well as a zoomed in section of the last 15 microns of the measurement to more clearly show the structure on a smaller scale. For the IR field (top), we clearly see a well-defined cosine behaviour, but with quite a bit of modulation in amplitude. Fitting a cosine to this pattern indeed reveals a main spatial frequency matching the expected TE_{00} mode. All the fluctuations on top of it can be attributed to noise. While the noise is clearly present, the signal is still the most dominant feature present. For the visible (bottom), we clearly also see cosine behaviour, but with the amplitudes fluctuating more rapidly, which again is not unexpected due to the mode structure. Fitting a cosine to the clearest pattern indeed reveals that it matches the frequency for the TE_{10} mode, which we know from the mode analysis to be the dominant mode in the field. Here again we find that the signal is clearly the dominant contribution to the field, perhaps even more so than for the infrared. So even though we did not manage to assign a value to the signal-to-noise directly, it is clear that in both cases the contribution from the bound modes are the most dominant contributions to the measured fields.

While a clear signal-to-noise ratio cannot readily be determined from this, we can make an estimate of the noise. The main contribution to the signal is the oscillations from the mode profile, with the highest frequency being the TE_{00} mode. So any contributions with higher frequencies cannot be from a mode, and as such have to be noise contributions. By considering the fluctuations on top of the oscillation due to this fundamental mode, we can estimate an amount of noise in the system. For IR, we can clearly observe these fluctuations, giving us an estimate of a couple percent. For visible, on the other hand, there is next to no fluctuations detectable on top of the cosine behaviour of the modes, leading to an estimate of noise being far less than a percent. This difference in noise amounts between the two colours is also reflected in Fig. 6.6 and 6.7, where we can consider the contributions to the power spectrum outside of the mode positions: for IR we observe fluctuations at lower values of k_x than expected for modes, which have to be noise contributions. For visible, while some fluctuations are observable, they are much lower in amplitude with respect to the mode peaks.

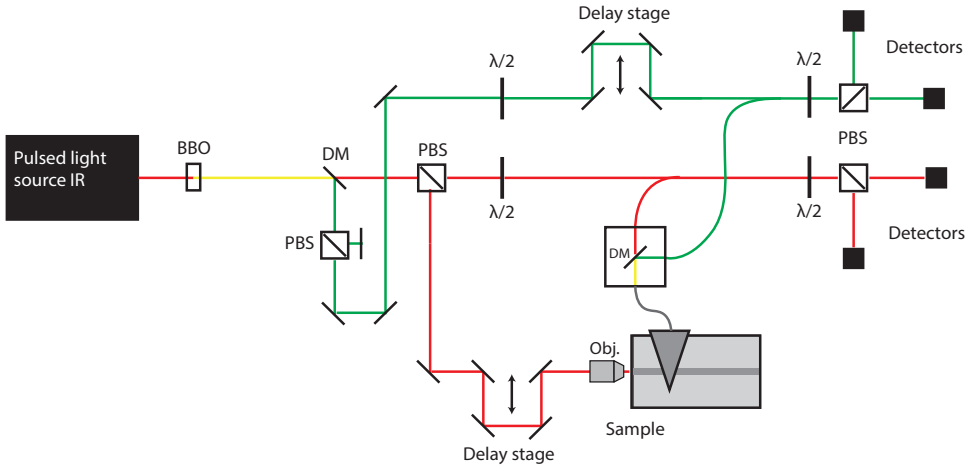


Figure 6.13: Schematic for the suggested setup for a nonlinear experiment. A high-power pulsed laser at telecom frequency is used. Using a BBO a part gets converted to the second harmonic in order to generate a reference beam. The pulse gets sent via a variable delay stage. Afterwards the picked up fundamental and its second harmonic are split again via a dichroic mirror and are recombined with their reference paths, and finally sent to their respective detectors. Green lines represent visible light only, red lines represent infrared light only and yellow lines represent combined visible and infrared light.

6.6. OUTLOOK

As a final part of this chapter we propose follow-up research to be done to put this dual-colour technique to use as described earlier, and provide an experimental guide to perform this experiment. The goal of this proposed experiment is to couple pulsed light with centre frequency at telecom wavelength into the waveguide, while also being able to measure its second harmonic at the same time, allowing us to observe the generation of second harmonic, alongside its fundamental frequency. This allows us to study where the second harmonic light is generated, and correlate the second harmonic fields and their singularities with the optical structure of the fundamental field.

For this experiment, we propose the use of a pulsed laser system with a centre frequency at telecom frequency. We aim at a repetition rate of around 120 kHz in order to have enough peak power to induce non-linearities in the sample. Neutral density filters can be placed to reduce this power if necessary. This pulse is then sent into a simple rib waveguide, where we are then able to measure the fundamental field and the second harmonic that is generated inside of the sample at the same time, with time resolution. For the sample, we prefer to use something that can support both IR and visible frequencies. For this reason we propose a SiN rib waveguide, similar to what was used in the initial experiment. Because it is transparent for both the fundamental and its second harmonic, as well as the ridges of the waveguide locally breaking the inversion symmetry, it makes for a good candidate to perform the experiment. However, as a follow-up to this, one can think of more intricate structures, such as topological structures or plasmonic nanowires for further investigation.

Now in order to detect both colours at the same time, we need to have a reference pulse for both the IR and the visible, that additionally arrive at the same moment as well. In order to make this happen, we send the initial pulse through a barium borate (BBO) crystal, which is able to efficiently convert a part of the pulse into a frequency doubled pulse, which we can use for the visible part of the reference branch. The BBO for dual colour reference generation is another reason that a higher peak power for the pulse is desirable. This allows us to have light in both the IR and visible branches of the reference paths, while retaining the coherence of the pulse, which is needed to employ our heterodyne detection scheme in order to detect light at both frequencies. The second harmonic then gets filtered out of the signal path by simply blocking the path the visible light would take into the signal branch. This can be accomplished by either introducing a physical block in the path between where the visible section is split to reference and where it gets recombined with the IR signal path, or by introducing a filter such as a 1 μm longpass filter in the signal path. A schematic for the proposed setup is shown in Fig. 6.13. The main changes with regard to the previous schematic is the addition of a second delay stage, the BBO crystal and the box after the tip to separate the two colours.

An additional benefit of the pulsed laser setup is that it also enables us to measure with time resolution, allowing us to follow the evolution of the fields in time. In order to do this, one needs to take care that the path lengths of signal and references are equal to within the coherence time of the laser used, including taking into account that the effective path length of free space and fibre are different. This is paramount, as the pulses that travel through reference and signal need to have coherent overlap in order to generate the beating necessary for heterodyne detection. Normally (see Chapter 3) we ensure that we can create overlap between signal and reference via a variable delay stage. However, in order to compare the fundamental field and its second harmonic, we ideally would like to compare them with respect to the same reference time. In other words, at a relative delay of zero with respect to each other. This can only be accomplished if we introduce a tunable path length in either one of the reference branches. Up until now it has been sufficient to tune only the signal path in order to get it equal to the reference path, but since both reference branches are now in use simultaneously, they need to be equal to the other reference branch as well, and for this a tunable length needs to be introduced. We recommend to carefully measure the lengths of both reference branches individually to determine the best candidate for an additional tunable length. In the schematic the delay stage is added to the visible reference path, but that is for illustrative purposes, not meant as a direct recommendation. Alternatively, simple glass blocks can also be introduced into a branch in order to generate more effective path lengths, as needed. Keep in mind that the introduction of the splitter box also introduces more length on the signal side, thus the setup as is will have more length in signal compared to the reference branches. The recommendation is to place the splitter box in a way that the lengths from tip to box, and box to detectors is kept to a minimum.

6.7. CONCLUSION

We have modified our near-field microscope setup to allow for measuring two colours at the same time by utilising a box with a dichroic mirror to split the two different colours that are picked up by our near-field probe. By picking a suitable fibre, which is single mode for both colours, from which to craft our near-field probes, we are able to collect the light of both frequencies with the same probe. With this, we are able to measure both the amplitude and phase of both colours simultaneously. Then by filtering for individual modes in Fourier space, and using a Jones matrix to correct the polarization rotation, we are able to show that polarization resolution is also maintained, to a certain extent. Furthermore, there seems to be little to no crosstalk between the two colours. This means that, at the cost of some throughput, we are able to perform a near-field measurement in the two-colour regime without having to sacrifice information.

7

CONCLUSIONS

As a child I considered such unknowns sinister. Now, though, I understand that they bear no ill will. The universe is, and we are.

Solanum - The Outer Wilds⁷

In this concluding chapter we highlight once more all the main results of the thesis. In addition, we also show the findings that are not yet understood and would make for interesting follow-up research. Finally, for all chapters we make recommendation of what topics could be explored next.

This thesis investigates the world of optics at the sub-wavelength scale, with a focus on understanding the behaviour of singularities that can arise in the electromagnetic field. Chapters 3, 4, and 5 have been fully devoted to deepening our understanding of these infinitesimally small entities, which occur generically in structured 2D light. Chapter 6 does not directly concern itself with singularities, being centered around improving the capabilities of the near-field microscope instead. However, the end game for this novel technique lends itself well for more research into singularities as well. It includes measurements of both varying the frequency of the used light (Chapter 5), as well as time-resolved measurements in order to observe the time dynamics of the singularities (Chapter 3). The experimental studies of singularities are supplemented with a chapter fully devoted to a numerical study (Chapter 4). Here we summarize the main results obtained in this thesis, as well as for each subject highlight some lines of potential follow-up research.

In Chapter 3 we have further explored the analogy between optical phase singularities and real charged particles. By utilising non-monochromatic field, we are able to observe the evolution of a two-dimensional random light field in time, and track the phase singularities in time as well. And with it, we are able to investigate their movement and their diffusive properties. We have revealed that phase singularities in random light do not diffuse linearly, but rather exhibit anomalous diffusion, specifically sub-diffusive behaviour. Additionally, this diffusive behaviour is different when considering faithful or unfaithful singularities. Furthermore, the lifetime of the phase singularities can exhibit a complex decay pattern, which differs from a simple exponential decay. Again this decay is different for faithful and unfaithful singularities. In this thesis we have restricted the investigation to the diffusion and lifetime of phase singularities of the out-of-plane magnetic field component. For future research it would be interesting to extend this investigating to the diffusive behaviour of other types of singularities as well, such as C-points or the Poynting singularities discussed in Chapter 5.

In Chapter 4 we have numerically investigated the melting of optical phase singularity lattices. Having three plane waves interfering in two dimensions will always lead to a trigonal lattice structure (when accounting for singularity charge). By introducing more sources in a continuous way, we have shown that we can melt this lattice structure, while keeping the spatial symmetries intact. We show that this can only happen when the phase of the additional waves is allowed to fluctuate. However, the model that was used has also shown a non-trivial dependence on the number of sources that were introduced, and this dependence is as of yet not well understood. This warrants more research into the model, as well as other potential models that can be used such as the stochastic broadening model discussed briefly in Appendix A. Additionally, we have restricted ourselves to investigating the case of three-fold symmetry only. For future investigations, we recommend considering the case of four through seven initial wave positions, as it fundamentally changes the symmetry inherent to the system.

In Chapter 5 we have investigated the optical singularities that can appear in the transverse Poynting vector field. Since the Poynting vector is comprised of both electric and magnetic field components, the origin of the singularities is more complex than for a complex scalar field. We have shown that the singularities arise from either zeros of

the out-of-plane magnetic field component, or from the in-plane Poynting vector becoming fully imaginary. The magnetic type singularities exhibit a liquid-like correlation, whereas the full ensemble of singularities exhibits a near-isotropic distribution instead. One result that remains unexplained in this study is the occurrence of the dampening behaviour for the $g(r)$ of the polarization-type singularities and warrants further research. As a future objective here it would be interesting to perform the same analysis done for time-resolved data instead. While we do not expect different topological structures in this case, perhaps other unforeseen changes with respect to monochromatic light could occur. In addition, this could be combined with a study of the diffusive behaviour of singularities of the flow field as well, connecting different investigations of this thesis.

In Chapter 6 we have pushed the capabilities of the near-field microscope, by enabling the simultaneous observation of telecom frequency light and visible light in the near field. We have shown that we are capable of performing this measurement without having to sacrifice either phase or polarization resolution. In the future, this will allow us investigate local non linear effects by being able to observe the fundamental field, as well as its second harmonic simultaneously. In this way we can study how and where the second harmonic light is generated inside of a sample. For future research, we proposed an experiment that would allow for investigating this with time-resolution. In this experiment, by using a simple silicon nitride rib waveguide, we want to observe the local generation of second harmonic light. Since the waveguide is mirror symmetric, the global second harmonic light is zero, but not necessarily locally. In addition, this allows us to further study optical singularities by investigating if there is a relation between the singularities in the fundamental field and the singularities in the second harmonic field that are (locally) created.

The main motivation for the performed research is one of curiosity, rather than application driven, with the main aim being to increase our understanding of the behaviour of light and its singularities at the nanoscale. But despite it being fundamental research, there is still potential for applications. Obtaining a better understanding of the spatial distribution of singularities and how to manipulate it allows for the engineering of singularity structures, such as the lattices discussed in Chapter 4. Especially since phase singularities are associated with orbital angular momentum of light. Being able to control the distribution of singularities becomes even more exciting when extending the techniques to for instance C-points. Since C-points carry spin angular momentum [212], they could potentially be using in particle trapping schemes for chiral particles [213]. Having an on-chip way to trap, and possible manipulate these chiral particles by only using light could pave the way for novel biosensing applications [214].

A

SUPPLEMENTARY INFORMATION TO CHAPTER 4

A.1. STOCHASTIC BROADENING

Besides the angular broadening case discussed in Ch. 4, there are more possible ways of introducing this broadening. And in order to further our understanding of the melting of singularity lattices more options should be investigated. One such different way is what we have dubbed stochastic broadening. This method is analogous to the angular broadening case one, but introduces additional variables. The Gaussian envelopes for the phase and amplitudes with a FWHM of respectively $\Delta\theta_a$ and $\Delta\phi$ remain. But now instead of the N sources being equally distributed around the ring in reciprocal space, their positions are additionally randomly selected from a Gaussian distribution with a width of $\Delta\theta_p$. Furthermore, we introduce N_d points per symmetry point. In this way, the likelihood of the additional sources being close in space to the original symmetry point is high, and the chance of finding a point diminishes further away from the symmetry points. As a result, we break the symmetry that is inherent to the grid method of the angular broadening case. If symmetry is not a significant factor for the results that were obtained from the angular broadening case, then we expect that the angular and stochastic broadening give the same results when taking enough samples for certain regions of parameter space.

Analogously to the angular broadening results, we now consider the results for the stochastic broadening case. We look for the different zones that were identified before, and see if the same results as for the angular broadening case can be observed. If the two methods are indeed compatible, we expect to roughly find the same structure. In Fig. A.2 we again show the heat map for the local vortex density, similar to the angular broadening case shown in Fig. 4.12. We show a grid of heat maps generated for the fixed values $\Delta\theta_p = 0.3$ and $N_d = 720$, and variable $\Delta\theta_A$ and $\Delta\phi$. Again we observe a similar structure as with the angular broadening case before, and the three different zones of solid-like,

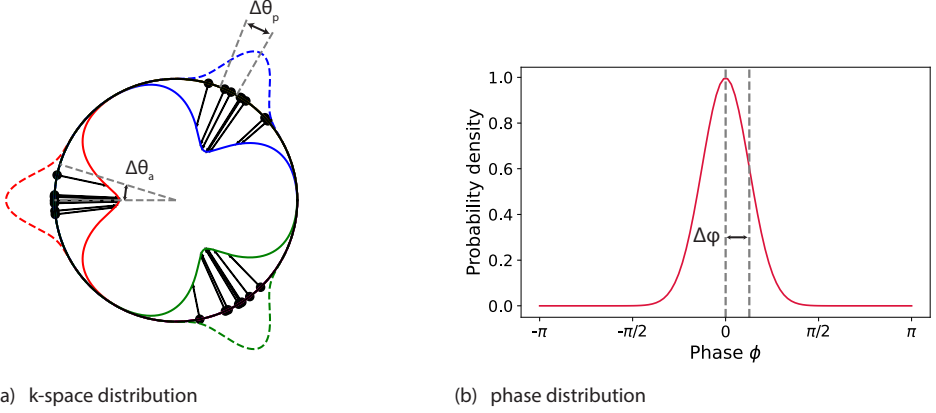


Figure A.1: Schematic illustration of the stochastic broadening case for $N = 3$ symmetry points and $N_d = 8$ waves per symmetry point. The position of these waves is randomly selected from a Gaussian distribution with FWHM $\Delta\theta_p$. This distribution is indicated by the dashed curves. The N_d waves are given an amplitude A_n based on their position on the circle, and is determined by the angular distance to a source point via a Gaussian distribution with a tune able FWHM of $\Delta\theta_a$, which is given by the blue curve in the circle. The phase difference of each of these waves with respect to the closest source point is also drawn from a Gaussian distribution with a FWHM of $\Delta\phi$.

liquid-like and spoke-like appear at roughly the same positions as before: the spoke patterns emerging for low values of phase and angular broadening, the lattice being kept intact without angular broadening, and melting occurring when both values are sufficiently large. As such, we observe qualitatively the same behaviour for both methods of broadening. Furthermore, the spoke pattern is not an artefact of the method, since it appears in both methods.

A.1.1. CONVERGENCE AND DEPENDENCE ON N_d

For angular broadening we have observed that the amount of sources that is used has a non-trivial effect on the state of the system. Furthermore, it seems that for increasing number of sources the state did not converge for the simulated amount of waves. So this also warrants checking with stochastic broadening if the system does seem to converge. As such, it is natural to investigate the effect of increasing number of points N_d on the melting indicators for this method of broadening.

We have fixed the values of $\Delta\theta_a = 0.6$ and $\Delta\theta_p = 0.3$ and determined the melting indicators as a function of $\Delta\phi$ for varying values of N_d in order to observe the influence of this parameter on the melting indicators. In Fig. A.3 the behaviour of the coordination number N_C is depicted. We clearly observe that the melting indicator does appear to converge to a common behaviour where the value of the melting indicator no longer depends on N_d . We see that the behaviour still continuously changes for $N_d \leq 500$, after which it becomes consistent. The only exception here appears to be $N_d = 800$ at the $\Delta\phi = 0.2$ data point.

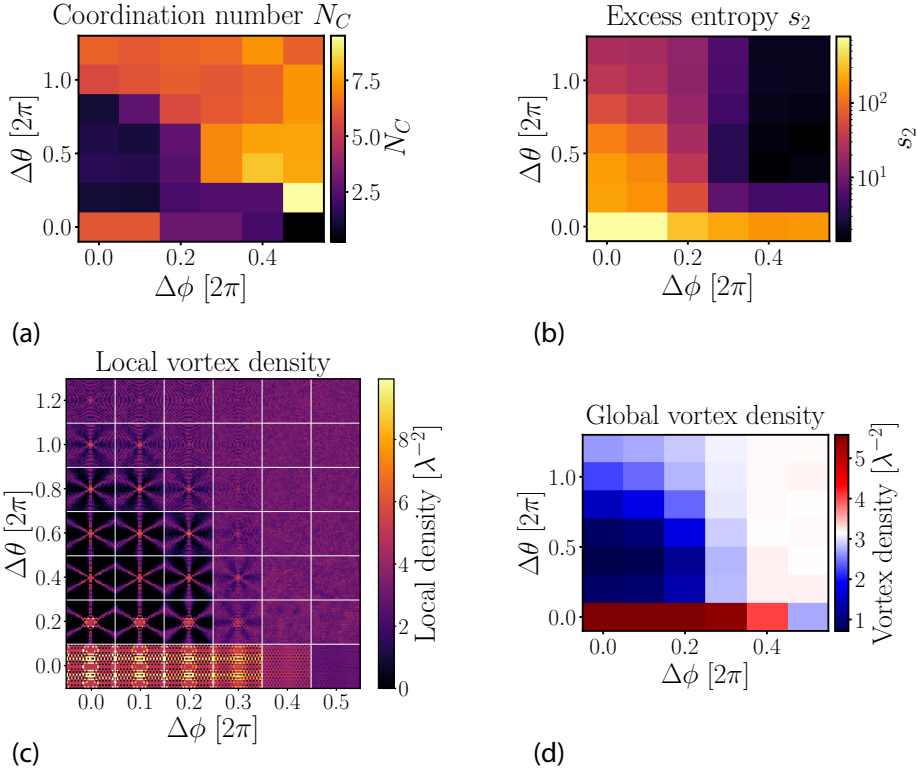


Figure A.2: Heat maps of local vortex densities, separated by white straight lines for the stochastic broadening case. Each block is a separate heat map generated for the parameters. Each heat map has a constant value of $N_d = 720$ waves per source point and a probability width of $\Delta\theta_p = 0.3$ and is averaged over 100 iterations.




Figure A.3: Behaviour of the coordination number N_C as a function of the phase spread $\Delta\phi$ for the stochastic broadening case, with a fixed $\Delta\theta_p = 0.3$ and $\Delta\theta_a = 0.6$. The different coloured lines indicate different amount of waves used to simulate the fields as indicated in the legend. We can see that for increasing number of waves launched per source the transition converges to a certain value of the phase spread, indicating that this system does converge.

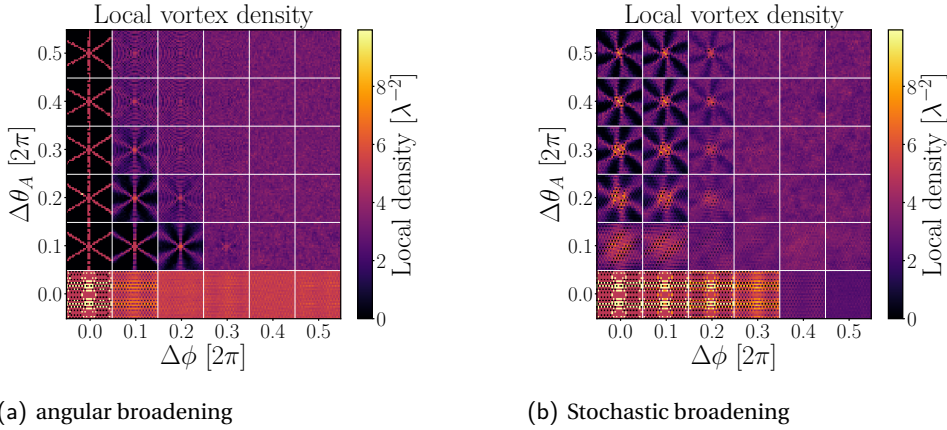


Figure A.4: Comparison of the heat maps for the local vortex density between (a) angular broadening with $N_\theta = 720$ and (b) stochastic broadening with $N_d = 240$ and a uniform distribution signifying the limiting case of $\Delta\theta_p \gg 1$.

A.1.2. CONSISTENCY WITH ANGULAR BROADENING

We consider different models in order to explore different aspects of melting and what seems to induce it. But at their core, they are based on the same principle of broadening a single source in a continuous manner. The assumption is that for certain regions in the parameter space for angular and stochastic broadening, they should converge to the same behaviour. For instance, if we take the amplitude and probability distribution exactly equal for stochastic broadening, we should find the same behaviour for large enough grid points and source points. And if we restrict the amplitude distribution while simultaneously letting the probability distribution converge to a uniform one, we should retrieve angular broadening once again. After all, it is the amplitude distribution that determines which region of k -space is accessible. Here we investigate if the two methods of broadening are indeed compatible when considering these edge cases.

We consider the case where $\Delta\theta_a \gg 1$, which is functionally a uniform distribution. In this case, we expect that the waves will be launched fairly uniformly around the ring. Taken enough iterations, this should average to be the same case as having them equidistributed over the ring. But since we still have the amplitude envelope, we expect that the distribution will be similar to that of the angular broadening case, since that method relies on a uniform distribution around the ring. In order to investigate this case, we have set the distribution to be uniform over the circle.

Since the angular broadening case has been evaluated using 720 sources, and we launch $N * N_D = 3 * N_D$ waves total with the stochastic method, we have opted to run the simulation using $N_D = 240$ waves, as an attempt at reproducing the angular broadening case as closely as possible. In Fig. A.4 we compare the heat maps for the local vortex density for the (a) angular and (b) stochastic cases. There are immediately obvious differences, such as the spoke pattern being much less pronounced in the stochastic case, and the upper left quadrant staying more spoke-like for the stochastic case for slightly higher values

of $\Delta\phi$. Nevertheless, the progression of the zones is qualitatively similar in both cases and the difference between the obvious spoke pattern for continuous and less obvious for stochastic can reasonably be explained through the symmetry that is inherent to the angular broadening case, which is exactly what we attempt to avoid using the stochastic case.

So while the two methods may not be an exact match, the general structure is highly similar, over the same scales of the variables. Furthermore, we also do not expect them to be exactly the same. The conversion between equidistributed grid points for the angular broadening case, and the randomly selected positions for the stochastic case do not immediately mean a one-to-one conversion between the amount of sources in both methods. Furthermore, we have already observed that the amount of sources for the angular broadening case matters. As such we conclude from this that while we cannot readily reproduce the angular broadening case from the stochastic case exactly, we conclude that the two methods are compatible.

A.2. CONVERGENCE OF THE MELTING INDICATORS AS FUNCTION OF AMOUNT OF SOURCES FOR ANGULAR BROADENING

In Section 4.5.1 we have observed that the coordination number does not appear to converge for increasing amount of waves N within the simulated range. For completion, we show here the behaviour for the remaining melting indicators as well.




Figure A.5: Behaviour of the excess entropy s_2 as a function of the phase spread $\Delta\phi$ for the angular broadening case, with a fixed $\Delta\theta = 0.08$.




Figure A.6: Behaviour of the position of the first (non-zero) minimum $r(g_{min})$ as a function of the phase spread $\Delta\phi$ for the angular broadening case, with a fixed $\Delta\theta = 0.08$.




Figure A.7: Behaviour of the ratio R as a function of the phase spread $\Delta\phi$ for the angular broadening case, with a fixed $\Delta\theta = 0.08$.




Figure A.8: Behaviour of the rotational invariants q_n as a function of the phase spread $\Delta\phi$ for the angular broadening case, with a fixed $\Delta\theta = 0.08$.




Figure A.9: Behaviour of the vortex number as a function of the phase spread $\Delta\phi$ for the angular broadening case, with a fixed $\Delta\theta = 0.08$.

BIBLIOGRAPHY

- [1] Gail Higginbottom and Vincent Mom. “Illuminating time: The visibility of temporality in prehistory”. In: *The Oxford Handbook of Light in Archaeology*. Oxford University Press Oxford, 2021, p. 105.
- [2] Norman Lockyer. *The Dawn of Astronomy: A Study of the Temple-worship and Mythology of the Ancient Egyptians*. Macmillan and Company, 1893.
- [3] George Sarton. *Ancient science through the golden age of Greece*. Courier Corporation, 2012.
- [4] Olivier Darrigol. *A history of optics from Greek antiquity to the nineteenth century*. OUP Oxford, 2012.
- [5] Ling-An Wu et al. “Optics in Ancient China.” In: *AAPPS Bulletin* 25.4 (2015).
- [6] Annamalai Manickavasagan and Hemantha Jayasuriya. *Imaging with electromagnetic spectrum: applications in food and agriculture*. Springer, 2014.
- [7] Feng Xia et al. “Internet of things”. In: *International journal of communication systems* 25.9 (2012), p. 1101.
- [8] Stewart Miller. *Optical fiber telecommunications*. Elsevier, 2012.
- [9] Harold J Metcalf and Peter Van der Straten. *Laser cooling and trapping*. Springer Science & Business Media, 1999.
- [10] Nobel Prize Outreach AB 2024. *The Nobel Prize in Physics 1997*.
- [11] Andrew D Ludlow et al. “Optical atomic clocks”. In: *Reviews of Modern Physics* 87.2 (2015), p. 637.
- [12] Patrick Gill. “When should we change the definition of the second?” In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 369.1953 (2011), pp. 4109–4130.
- [13] Kenneth R Brown. *The theory of the GPS composite clock*. Institute of Navigation, 1991.
- [14] Michael A Lombardi et al. “Accurate, traceable, and verifiable time synchronization for world financial markets”. In: *Journal of research of the National Institute of Standards and Technology* 121 (2016), p. 436.
- [15] V Sasikala and K Chitra. “All optical switching and associated technologies: a review”. In: *Journal of Optics* 47.3 (2018), pp. 307–317.
- [16] Arthur Schuster. *An introduction to the theory of optics*. E. Arnold, 1904.
- [17] David J Griffiths. *Introduction to electrodynamics*. Cambridge University Press, 2023.

- [18] Thomas Young. "The Bakerian Lecture. Experiments and Calculations Relative to Physical Optics. [Abstract]". In: *Proceedings of the Royal Society of London Series I* 1 (Jan. 1804), pp. 131–132.
- [19] Niels Bohr et al. *The quantum postulate and the recent development of atomic theory*. Vol. 3. Printed in Great Britain by R. & R. Clarke, Limited, 1928.
- [20] H.Barry Bebb and E.W. Williams. "Chapter 4 Photoluminescence I: Theory". In: ed. by R.K. Willardson and Albert C. Beer. Vol. 8. Semiconductors and Semimetals. Elsevier, 1972, pp. 181–320.
- [21] John David Jackson. *Classical electrodynamics*. John Wiley & Sons, 2012.
- [22] Viktor G Veselago. "Electrodynamics of materials with negative index of refraction". In: *Physics-Uspekhi* 46.7 (July 2003), p. 764.
- [23] Chung Dong Nguyen et al. "Numerical estimation of the complex refractive indexes by the altitude depending on wave frequency in the ionized region of the Earth atmosphere for microwaves information and power transmissions". In: *Progress In Electromagnetics Research M* 52 (2016), pp. 21–31.
- [24] M. V. Berry and M. R. Dennis. "Phase singularities in isotropic random waves". In: *Proc. R. Soc. Lond. A* 456 (2001 2000), pp. 2059–2079.
- [25] Michael Barth and Hans-Jürgen Stöckmann. "Current and vortex statistics in microwave billiards". In: *Phys. Rev. E* 65 (6 2002), p. 066208.
- [26] Isaac Freund, Natalya Shvartsman, and Valentin Freilikher. "Optical dislocation networks in highly random media". In: *Optics Communications* 101.3 (1993), pp. 247–264.
- [27] Isaac Freund and Natalya Shvartsman. "Wave-field phase singularities: the sign principle". In: *Physical Review A* 50.6 (1994), p. 5164.
- [28] Isaac Freund and Valentin Freilikher. "Parameterization of anisotropic vortices". In: *JOSA A* 14.8 (1997), pp. 1902–1910.
- [29] Shihoko Ishii. *Introduction to singularities*. Springer, 2018.
- [30] M. V. Berry and S. Klein. "Colored Diffraction Catastrophes". In: *Proceedings of the National Academy of Sciences of the United States of America* 93.6 (1996), pp. 2614–2619.
- [31] John Lighton Synge. "The gravitational field of a particle". In: *Proceedings of the Royal Irish Academy. Section A: Mathematical and Physical Sciences*. Vol. 53. JSTOR. 1950, pp. 83–114.
- [32] György Szekeres. "On the singularities of a Riemannian manifold". In: *Publicationes Mathematicae Debrecen* 7 7 (1960), p. 285.
- [33] M. D. Kruskal. "Maximal Extension of Schwarzschild Metric". In: *Phys. Rev.* 119 (5 Sept. 1960), pp. 1743–1745.
- [34] Jiro Mizushima, Kazuki Abe, and Naoto Yokoyama. "Bathtub vortex induced by instability". In: *Phys. Rev. E* 90 (4 Oct. 2014), p. 041002.

- [35] A Andersen et al. "Anatomy of a bathtub vortex". In: *Physical review letters* 91.10 (2003), p. 104502.
- [36] AW Marris. "Theory of the bathtub vortex". In: (1967).
- [37] Mark R Dennis, Kevin O'holleran, and Miles J Padgett. "Singular optics: optical vortices and polarization singularities". In: *Progress in optics*. Vol. 53. Elsevier, 2009, pp. 293–363.
- [38] Gregory J. Gbur. *Singular Optics*. CRC Press, 2016.
- [39] Vladimir G. Denisenko et al. "Mapping phases of singular scalar light fields". In: *Opt. Lett.* 33.1 (Jan. 2008), pp. 89–91.
- [40] J.F. Nye and M. V. Berry. "Dislocations in wave trains". In: *Proc. R. Soc. Lond. A* 336 (1605 1974), pp. 165–190.
- [41] MV Berry. "Disruption of wavefronts: statistics of dislocations in incoherent Gaussian random waves". In: *Journal of Physics A: Mathematical and General* 11.1 (1978), p. 27.
- [42] GP Karman et al. "Creation and annihilation of phase singularities in a focal field". In: *Optics letters* 22.19 (1997), pp. 1503–1505.
- [43] L. De Angelis et al. "Spatial Distribution of Phase Singularities in Optical Random Vector Waves". In: *Physical Review Letters* 117.9 (2016), p. 093901.
- [44] L. De Angelis et al. "Persistence and Lifelong Fidelity of Phase Singularities in Optical Random Waves". In: *Phys. Rev. Lett.* 119 (20 Nov. 2017), p. 203903.
- [45] Hugo F Schouten et al. "Creation and annihilation of phase singularities near a sub-wavelength slit". In: *Optics Express* 11.4 (Feb. 2003), p. 371.
- [46] Hugo F Schouten et al. "Phase singularities of the coherence functions in Young's interference pattern". In: *Optics letters* 28.12 (2003), pp. 968–970.
- [47] Bertrand I Halperin. *Statistical mechanics of topological defects*. 1981.
- [48] NB Simpson, L Allen, and MJ Padgett. "Optical tweezers and optical spanners with Laguerre–Gaussian modes". In: *Journal of modern optics* 43.12 (1996), pp. 2485–2491.
- [49] Edward Collett. *Field Guide to Polarization*. SPIE, 2005.
- [50] Mark R. Dennis. "Polarization singularities in paraxial vector fields: Morphology and statistics". In: *Optics Communications* 213.4-6 (2002), pp. 201–221.
- [51] Gianmarc Grazioli. *ggrazioli/LennardJonesFluidLab: Initial Release*. Version 1.0.0. Oct. 2022.
- [52] L De Angelis and L Kuipers. "Effective pair-interaction of phase singularities in random waves". In: *Optics Letters* 46.11 (2021), pp. 2734–2737.
- [53] MR Dennis. "Correlations and screening of topological charges in Gaussian random fields". In: *Journal of Physics A: Mathematical and General* 36.24 (2003), p. 6611.
- [54] BA Van Tiggelen, D Anache, and An Ghysels. "Role of mean free path in spatial phase correlation and nodal screening". In: *Europhysics Letters* 74.6 (2006), p. 999.

- [55] Natalya Shvartsman and Isaac Freund. "Wave-field phase singularities: near-neighbor correlations and anticorrelations". In: *J. Opt. Soc. Am. A* 11.10 (Oct. 1994), pp. 2710–2718.
- [56] Lorenzo De Angelis, Filippo Alpeggiani, and L. Kuipers. "Spatial bunching of same-charge polarization singularities in two-dimensional random vector waves". In: *Physical Review X* 8.4 (2018), p. 41012.
- [57] R. Höhmann et al. "Density and correlation functions of vortex and saddle points in open billiard systems". In: *Physical Review E (Statistical, Nonlinear, and Soft Matter Physics)* 79.1 (Jan. 2009).
- [58] J. A. Barker and D. Henderson. "What is "liquid"? Understanding the states of matter". In: *Rev. Mod. Phys.* 48 (4 Oct. 1976), pp. 587–671.
- [59] JE Lennard and I Jones. "On the determination of molecular fields.—I. From the variation of the viscosity of a gas with temperature". In: *Proceedings of the Royal Society of London. Series A, containing papers of a mathematical and physical character* 106.738 (1924), pp. 441–462.
- [60] John Edward Jones. "On the determination of molecular fields.—II. From the equation of state of a gas". In: *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 106.738 (1924), pp. 463–477.
- [61] Holger Fischer and Olivier J. F. Martin. "Engineering the optical response of plasmonic nanoantennas". In: *Opt. Express* 16.12 (June 2008), pp. 9144–9154.
- [62] M. Burrese et al. "Observation of Polarization Singularities at the Nanoscale". In: *Physical Review Letters* 102.3 (Jan. 2009), p. 033902.
- [63] Martin P. J. Lavery et al. "Detection of a Spinning Object Using Light's Orbital Angular Momentum". In: *Science* 341.6145 (2013), pp. 537–540.
- [64] Jian Feng Li et al. "Shell-isolated nanoparticle-enhanced Raman spectroscopy". In: *Nature* 464.7287 (Mar. 2010), pp. 392–395.
- [65] S. Carretero-Palacios, A. Jiménez-Solano, and H. Míguez. "Plasmonic Nanoparticles as Light-Harvesting Enhancers in Perovskite Solar Cells: A User's Guide". In: *ACS Energy Letters* 1.1 (2016). PMID: 28066822, pp. 323–331.
- [66] S. Arora et al. "Breakdown of Spin-to-Helicity Locking at the Nanoscale in Topological Photonic Crystal Edge States". In: *Phys. Rev. Lett.* 128 (20 May 2022), p. 203903.
- [67] E. Abbe. "Beiträge zur Theorie des Mikroskops und der mikroskopischen Wahrnehmung". In: *Archiv für Mikroskopische Anatomie* 9.1 (Dec. 1873), pp. 413–468.
- [68] E. Hecht. *Optics*. Pearson education. Addison-Wesley, 2002.
- [69] E.H. Synge. "XXXVIII. A suggested method for extending microscopic resolution into the ultra-microscopic region". In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 6.35 (1928), pp. 356–362.
- [70] Wolfgang D. Pohl. *Optical Near-field scanning microscope*. European Patent EP82111974.0 (Dec.27,1982), US Patent 4,604,520 (Dec.20, 1983).

- [71] E. A. ASH and G. NICHOLLS. “Super-resolution Aperture Scanning Microscope”. In: *Nature* 237.5357 (June 1972), pp. 510–512.
- [72] D. W. Pohl, W. Denk, and M. Lanz. “Optical stethoscopy: Image recording with resolution $\lambda/20$ ”. In: *Applied Physics Letters* 44.7 (Apr. 1984), pp. 651–653.
- [73] Nobel Prize Outreach AB 2024. *The Nobel Prize in Physics 1986*.
- [74] Srinivasa M Salapaka and Murti V Salapaka. “Scanning probe microscopy”. In: *IEEE Control Systems Magazine* 28.2 (2008), pp. 65–83.
- [75] R Hillenbrand and F Keilmann. “Material-specific mapping of metal/semiconductor/dielectric nanosystems at 10 nm resolution by backscattering near-field optical microscopy”. In: *Applied Physics Letters* 80.1 (2002), pp. 25–27.
- [76] Yasushi Inouye and Satoshi Kawata. “Near-field scanning optical microscope with a metallic probe tip”. In: *Opt. Lett.* 19.3 (Feb. 1994), pp. 159–161.
- [77] A. Lahrech et al. “Infrared-reflection-mode near-field microscopy using an apertureless probe with a resolution of $\lambda/600$ ”. In: *Opt. Lett.* 21.17 (Sept. 1996), pp. 1315–1317.
- [78] M. L. M. Balistreri et al. “Visualizing the whispering gallery modes in a cylindrical optical microcavity”. In: *Opt. Lett.* 24.24 (Dec. 1999), pp. 1829–1831.
- [79] M. L. M. Balistreri et al. “Local Observations of Phase Singularities in Optical Fields in Waveguide Structures”. In: *Physical Review Letters* 85.2 (July 2000), pp. 294–297.
- [80] M Sandtke et al. “Novel instrument for surface plasmon polariton tracking in space and time”. en. In: *Rev Sci Instrum* 79.1 (Jan. 2008), p. 013704.
- [81] B. Le Feber et al. “Simultaneous measurement of nanoscale electric and magnetic optical fields”. In: *Nature Photonics* 8.1 (Jan. 2014), pp. 43–46.
- [82] D Courjon and C Bainier. “Near field microscopy and near field optics”. In: *Reports on Progress in Physics* 57.10 (Oct. 1994), p. 989.
- [83] Boris le Feber. “Nanoscale electric and magnetic optical vector fields: Mapping & Injection”. PhD thesis. University of Twente, 2015.
- [84] Robin Christian Reddick et al. “Photon scanning tunneling microscopy”. In: *Review of scientific instruments* 61.12 (1990), pp. 3669–3677.
- [85] E. Betzig et al. “Breaking the Diffraction Barrier: Optical Microscopy on a Nanometric Scale”. In: *Science* 251.5000 (1991), pp. 1468–1470.
- [86] J.A. Veerman et al. “High definition aperture probes for near-field optical microscopy fabricated by focused ion beam milling”. Undefined. In: *Applied physics letters* 1998.24 (1998), pp. 3115–3117.
- [87] Sonakshi Arora. “Photonic topological edge states: A nanoscale investigation”. Available at <https://doi.org/10.4233/uuid:f9c4d874-92cc-4a45-bd2f-a37f98f8fb78>. PhD thesis. Delft University of Technology, Nov. 2023.
- [88] L. Novotny and B. Hecht. *Principles of Nano-Optics*. Principles of Nano-optics. Cambridge University Press, 2012.

- [89] EPLJS Betzig, PL Finn, and JS Weiner. "Combined shear force and near-field scanning optical microscopy". In: *Applied physics letters* 60.20 (1992), pp. 2484–2486.
- [90] MLM Balistreri et al. "Tracking femtosecond laser pulses in space and time". In: *Science* 294.5544 (2001), pp. 1080–1082.
- [91] M. L. M. Balisteri. "Coherent imaging of guided optical fields". PhD thesis. University of Twente, 2000.
- [92] Sajeev John. "Strong localization of photons in certain disordered dielectric superlattices". In: *Physical review letters* 58.23 (1987), p. 2486.
- [93] Eli Yablonovitch. "Inhibited Spontaneous Emission in Solid-State Physics and Electronics". In: *Phys. Rev. Lett.* 58 (20 May 1987), pp. 2059–2062.
- [94] EJJY Yablonovitch. "Photonic band-gap structures". In: *JOSA B* 10.2 (1993), pp. 283–295.
- [95] Thomas F Krauss, Richard M De La Rue, and Stuart Brand. "Two-dimensional photonic-bandgap structures operating at near-infrared wavelengths". In: *Nature* 383.6602 (1996), pp. 699–702.
- [96] Eli Yablonovitch, TJ Gmitter, and Kok-Ming Leung. "Photonic band structure: The face-centered-cubic case employing nonspherical atoms". In: *Physical review letters* 67.17 (1991), p. 2295.
- [97] Jean Pol Vigneron and Priscilla Simonis. "Natural photonic crystals". In: *Physica B: Condensed Matter* 407.20 (2012). Proceedings of the conference - Wave Propagation: From Electrons to Photonic Crystals and Metamaterials, pp. 4032–4036.
- [98] Yurii A Vlasov et al. "Active control of slow light on a chip with photonic crystal waveguides". In: *nature* 438.7064 (2005), pp. 65–69.
- [99] Daniela Threm, Yousef Nazirizadeh, and Martina Gerken. "Photonic crystal biosensors towards on-chip integration". In: *Journal of biophotonics* 5.8-9 (2012), pp. 601–616.
- [100] Andreas C Liapis, Zhimin Shi, and Robert W Boyd. "Optimizing photonic crystal waveguides for on-chip spectroscopic applications". In: *Optics Express* 21.8 (2013), pp. 10160–10165.
- [101] Yoshihiro Akahane et al. "High-Q photonic nanocavity in a two-dimensional photonic crystal". In: *nature* 425.6961 (2003), pp. 944–947.
- [102] Philippe Lalanne, Christophe Sauvan, and Jean Paul Hugonin. "Photon confinement in photonic crystal nanocavities". In: *Laser & Photonics Reviews* 2.6 (2008), pp. 514–526.
- [103] C. Liu et al. "Triggering extreme events at the nanoscale in photonic seas". In: *Nature Physics* 11 (Mar. 2015), p. 358.
- [104] Hans-Jürgen Stöckmann. *Quantum Chaos: An Introduction*. Cambridge University Press, 1999.
- [105] H Gersen et al. "Tracking ultrashort pulses through dispersive media: Experiment and theory". In: *Physical Review E* 68.2 (2003), p. 026604.

- [106] Juntao Li et al. “Systematic design of flat band slow light in photonic crystal waveguides”. In: *Optics express* 16.9 (2008), pp. 6227–6232.
- [107] Masaya Notomi et al. “Extremely large group-velocity dispersion of line-defect waveguides in photonic crystal slabs”. In: *Physical review letters* 87.25 (2001), p. 253902.
- [108] Lars H Frandsen et al. “Photonic crystal waveguides with semi-slow light and tailored dispersion properties”. In: *Optics express* 14.20 (2006), pp. 9444–9450.
- [109] R. Clark Jones. “A New Calculus for the Treatment of Optical SystemsI. Description and Discussion of the Calculus”. In: *J. Opt. Soc. Am.* 31.7 (July 1941), pp. 488–493.
- [110] Manfred Hammer. *1-D mode solver for dielectric multilayer slab waveguides*.
- [111] L. De Angelis and L. Kuipers. “Effective pair-interaction of phase singularities in random waves”. In: *Opt. Lett.* 46.11 (2021), pp. 2734–2737.
- [112] Adolf Fick. “Ueber Diffusion”. In: *Annalen der Physik* 170.1 (1855), pp. 59–86.
- [113] Francisco J Buera and Ezra Oberfield. “The global diffusion of ideas”. In: *Econometrica* 88.1 (2020), pp. 83–114.
- [114] Margaret Whitehead. “Diffusion of ideas on social inequalities in health: a European perspective”. In: *The Milbank Quarterly* 76.3 (1998), pp. 469–492.
- [115] Shigeo Kusuoka. “Approximation of expectation of diffusion process and mathematical finance”. In: *Taniguchi Conference on Mathematics Nara’98*. Vol. 31. Mathematical Society of Japan. 2001, pp. 147–166.
- [116] Bjørn Eraker. “MCMC analysis of diffusion models with application to finance”. In: *Journal of Business & Economic Statistics* 19.2 (2001), pp. 177–191.
- [117] Karl Pearson. “The problem of the random walk”. In: *Nature* 72.1867 (1905), pp. 342–342.
- [118] Joseph L Doob. “The Brownian movement and stochastic equations”. In: *Annals of Mathematics* (1942), pp. 351–369.
- [119] Albert Einstein. *Investigations on the Theory of the Brownian Movement*. Courier Corporation, 1956.
- [120] M. von Smoluchowski. “Zur kinetischen Theorie der Brownschen Molekularbewegung und der Suspensionen”. In: *Annalen der Physik* 326.14 (1906), pp. 756–780.
- [121] Igor M Sokolov. “Models of anomalous diffusion in crowded environments”. In: *Soft Matter* 8.35 (2012), pp. 9043–9052.
- [122] L. Bruno et al. “Transition to superdiffusive behavior in intracellular actin-based transport mediated by molecular motors”. In: *Phys. Rev. E* 80 (1 July 2009), p. 011912.
- [123] GM Viswanathan et al. “Lévy flight random searches in biological phenomena”. In: *Physica A: Statistical Mechanics and Its Applications* 314.1-4 (2002), pp. 208–213.

- [124] Andy M Reynolds. “Lévy flight patterns are predicted to be an emergent property of a bumblebees’ foraging strategy”. In: *Behavioral Ecology and Sociobiology* 64 (2009), pp. 19–23.
- [125] Yoav Sagi et al. “Observation of Anomalous Diffusion and Fractional Self-Similarity in One Dimension”. In: *Phys. Rev. Lett.* 108 (9 2012), p. 093002.
- [126] Matthew J. Colbrook et al. “Scaling laws of passive-scalar diffusion in the interstellar medium”. In: *Monthly Notices of the Royal Astronomical Society* 467.2 (Feb. 2017), pp. 2421–2429.
- [127] Liad Levi et al. “Hyper-transport of light and stochastic acceleration by evolving disorder”. In: *Nature Physics* 8.12 (2012), pp. 912–917.
- [128] M Köpf et al. “Anomalous diffusion of water in biological tissues”. In: *Biophysical journal* 70.6 (1996), pp. 2950–2958.
- [129] Benjamin M Regner et al. “Anomalous diffusion of single particles in cytoplasm”. In: *Biophysical journal* 104.8 (2013), pp. 1652–1660.
- [130] Felix Höfling and Thomas Franosch. “Anomalous transport in the crowded world of biological cells”. In: *Reports on Progress in Physics* 76.4 (2013), p. 046602.
- [131] TT Marquez-Lago, Andre Leier, and Kevin Burrage. “Anomalous diffusion and multifractional Brownian motion: simulating molecular crowding and physical obstacles in systems biology”. In: *IET systems biology* 6.4 (2012), pp. 134–142.
- [132] Vasiliki Plerou et al. “Economic fluctuations and anomalous diffusion”. In: *Physical Review E* 62.3 (2000), R3023.
- [133] Enrico Scalas, Rudolf Gorenflo, and Francesco Mainardi. “Fractional calculus and continuous-time finance”. In: *Physica A: Statistical Mechanics and its Applications* 284.1-4 (2000), pp. 376–384.
- [134] Antoine Jacquier and Lorenzo Torricelli. “Anomalous Diffusions in Option Prices: Connecting Trade Duration and the Volatility Term Structure”. In: *SIAM Journal on Financial Mathematics* 11.4 (2020), pp. 1137–1167.
- [135] Isaac Freund. “Saddles, singularities, and extrema in random phase fields”. In: *Phys. Rev. E* 52 (3 Sept. 1995), pp. 2348–2360.
- [136] MV Berry, MR Dennis, and RL Lee. “Polarization singularities in the clear sky”. In: *New Journal of Physics* 6.1 (2004), p. 162.
- [137] Xiaojun Cheng, Yitzchak Lockerman, and Azriel Z. Genack. “Phase singularity diffusion”. In: *Opt. Lett.* 39.11 (2014), pp. 3348–3351.
- [138] Lorenzo de Angelis. “The Singular Optics of Random Light: A 2D vectorial investigation”. Available at <https://doi.org/10.4233/uuid:25dc497c-b218-4d6a-9796-001e9d569975>. PhD thesis. Delft University of Technology, Dec. 2018.
- [139] MV Berry. “Nature’s optics and our understanding of light”. In: *Contemporary Physics* 56.1 (2015), pp. 2–16.
- [140] Andreas Dechant et al. “Continuous-Time Random Walk for a Particle in a Periodic Potential”. In: *Phys. Rev. Lett.* 123 (7 Aug. 2019), p. 070602.

- [141] Daniel Ben-Avraham and Shlomo Havlin. “Diffusion on percolation clusters at criticality”. In: *Journal of Physics A: Mathematical and General* 15.12 (1982), p. L691.
- [142] Ben O’Shaughnessy and Itamar Procaccia. “Diffusion on fractals”. In: *Physical Review A* 32.5 (1985), p. 3073.
- [143] Martin T Barlow. “Diffusions on fractals”. In: *Lectures on Probability Theory and Statistics: Ecole d’Eté de Probabilités de Saint-Flour XXV—1995* (2006), pp. 1–121.
- [144] Gernot Guigas, Claudia Kalla, and Matthias Weiss. “Probing the nanoscale viscoelasticity of intracellular fluids in living cells”. In: *Biophysical journal* 93.1 (2007), pp. 316–323.
- [145] Stephanie C Weber, Andrew J Spakowitz, and Julie A Theriot. “Bacterial chromosomal loci move subdiffusively through a viscoelastic cytoplasm”. In: *Physical review letters* 104.23 (2010), p. 238102.
- [146] Benoit B. Mandelbrot and John W. Van Ness. “Fractional Brownian Motions, Fractional Noises and Applications”. In: *SIAM Review* 10.4 (1968), pp. 422–437.
- [147] Tilmann Gneiting, Hana Ševčíková, and Donald B. Percival. “Estimators of Fractal Dimension: Assessing the Roughness of Time Series and Spatial Data”. In: *Statistical Science* 27.2 (2012), pp. 247–277.
- [148] Steven Orey. “Gaussian sample functions and the Hausdorff dimension of level crossings”. In: *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete* 15.3 (1970), pp. 249–256.
- [149] J. R. M. Hosking. “Modeling persistence in hydrological time series using fractional differencing”. In: *Water Resources Research* 20.12 (1984), pp. 1898–1908.
- [150] Søren Asmussen. *Simulation of stochastic processes*. English. Ed. by Jozef Teugels and Bjørn Sundt. Wiley, 2004, pp. 1570–1572.
- [151] R. B. Davies and D. S. Harte. “Tests for Hurst Effect”. In: *Biometrika* 74.1 (1987), pp. 95–101.
- [152] Stas Burov et al. “Single particle tracking in systems showing anomalous diffusion: the role of weak ergodicity breaking”. In: *Phys. Chem. Chem. Phys.* 13 (5 2011), pp. 1800–1812.
- [153] A. Reyes et al. *yupi: Generation, Tracking and Analysis of Trajectory data in Python*. 2022.
- [154] Valerie J. Anderson and Henk N. W. Lekkerkerker. “Insights into phase transition kinetics from colloid science”. In: *Nature* 416.6883 (Apr. 2002), pp. 811–815.
- [155] Tad Hogg, Bernardo A. Huberman, and Colin P. Williams. “Phase transitions and the search problem”. In: *Artificial Intelligence* 81.1 (1996). Frontiers in Problem Solving: Phase Transitions and Complexity, pp. 1–15.
- [156] Eliezer E. W. Heffern et al. “Phase transitions in biology: from bird flocks to population dynamics”. In: *Proceedings of the Royal Society B: Biological Sciences* 288.1961 (2021), p. 20211111.

- [157] J M Kosterlitz and D J Thouless. "Ordering, metastability and phase transitions in two-dimensional systems". In: *Journal of Physics C: Solid State Physics* 6.7 (Apr. 1973), p. 1181.
- [158] Luuk Crooijmans. "Numerical modelling of melting static optical vortex lattices". Bachelor thesis. Delft University of Technology, Dec. 2021.
- [159] J. B. Clarke et al. "Definitions of terms relating to phase transitions of the solid state (IUPAC Recommendations 1994)". In: *Pure and Applied Chemistry* 66.3 (1994), pp. 577–594.
- [160] Stephen J Blundell and Katherine M. Blundell. *Concepts in Thermal Physics*. Oxford University Press, 2008.
- [161] Rana A. Fine and Frank J. Millero. "Compressibility of water as a function of temperature and pressure". In: *Journal of Chemical Physics* 59.10 (1973), pp. 5529–5536.
- [162] J. J. Neumeier. "Elastic Constants, Bulk Modulus, and Compressibility of H₂O Ice Ih for the Temperature Range 50 K–273 K". In: *Journal of Physical and Chemical Reference Data* 47.3 (July 2018), p. 033101.
- [163] P. W. Anderson. "More Is Different". In: *Science* 177.4047 (1972), pp. 393–396.
- [164] S. Ulrich et al. "Elasticity of highly cross-linked random networks". In: *Europhysics Letters* 76.4 (Oct. 2006), p. 677.
- [165] Christian L. Klix et al. "Glass Elasticity from Particle Trajectories". In: *Phys. Rev. Lett.* 109 (17 Oct. 2012), p. 178301.
- [166] Grzegorz Szamel and Elijah Flenner. "Emergence of Long-Range Correlations and Rigidity at the Dynamic Glass Transition". In: *Phys. Rev. Lett.* 107 (10 Sept. 2011), p. 105505.
- [167] Sergey A. Khrapak. "Lindemann melting criterion in two dimensions". In: *Physical Review Research* 2.1, 012040 (Feb. 2020), p. 012040.
- [168] B. A. Klumov. "On the behavior of indicators of melting: Lennard-Jones system in the vicinity of the phase transition". In: *JETP Letters* 98.5 (Nov. 2013), pp. 259–265.
- [169] Paul J. Steinhardt, David R. Nelson, and Marco Ronchetti. "Bond-orientational order in liquids and glasses". In: *Phys. Rev. B* 28 (2 July 1983), pp. 784–805.
- [170] Natsuda Klongvessa et al. "Aging or DEAD: origin of the non-monotonic response to weak self-propulsion in active glasses". In: (Apr. 2022).
- [171] William B. Streett, Harold J. Raveché, and Raymond D. Mountain. "Monte Carlo studies of the fluid-solid phase transition in the Lennard-Jones system". In: *The Journal of Chemical Physics* 61.5 (Sept. 1974), pp. 1960–1969.
- [172] C. E. Shannon. "A mathematical theory of communication". In: *The Bell System Technical Journal* 27.3 (1948), pp. 379–423.
- [173] E. T. Jaynes. "Information Theory and Statistical Mechanics". In: *Phys. Rev.* 106 (4 May 1957), pp. 620–630.

- [174] N. Jakse and A. Pasturel. “Excess Entropy Scaling Law for Diffusivity in Liquid Metals”. In: *Scientific Reports* 6.1 (Feb. 2016), p. 20689.
- [175] István Borzsák and András Baranyai. “On the convergence of Green’s entropy expansion”. In: *Chemical Physics* 165.2 (1992), pp. 227–230.
- [176] D. Shechtman et al. “Metallic Phase with Long-Range Orientational Order and No Translational Symmetry”. In: *Phys. Rev. Lett.* 53 (20 Nov. 1984), pp. 1951–1953.
- [177] A.Y. Cho and J.R. Arthur. “Molecular beam epitaxy”. In: *Progress in Solid State Chemistry* 10 (1975), pp. 157–191.
- [178] K Stam et al. “Large atom number Bose-Einstein condensate of sodium”. In: *The Review of scientific instruments* 78 (Feb. 2007), p. 013102.
- [179] Floris MH Cromptvoets et al. “Longitudinal focusing and cooling of a molecular beam”. In: *Physical review letters* 89.9 (2002), p. 093004.
- [180] M. C. Gao and M. Widom. “Information Entropy of Liquid Metals”. In: *The Journal of Physical Chemistry B* 122.13 (2018). PMID: 29461826, pp. 3550–3555.
- [181] M. A. van Gogh et al. “Poynting singularities in the transverse flow-field of random vector waves”. In: *Opt. Lett.* 45.9 (2020), pp. 2600–2603.
- [182] John David Jackson. *Classical electrodynamics*. 3rd ed. New York, NY: Wiley, 1999.
- [183] M V Berry. “Optical currents”. In: *Journal of Optics A: Pure and Applied Optics* 11.9 (Aug. 2009), p. 094001.
- [184] J. H. Poynting. “On the Transfer of Energy in the Electromagnetic Field”. In: *Philosophical Transactions of the Royal Society of London* 175 (1884), pp. 343–361.
- [185] A. Boivin, J. Dow, and Emil Wolf. “Energy Flow in the Neighborhood of the Focus of a Coherent Beam*”. In: *Journal of the Optical Society of America* 57.10 (Oct. 1967), p. 1171.
- [186] K. Dholakia and T. Čižmár. “Shaping the future of manipulation”. In: *Nature Photonics* 5.6 (2011), pp. 335–342.
- [187] Isaac Freund. “Poincaré vortices”. In: *Optics Letters* 26.24 (2001), pp. 1996–1998.
- [188] T.W. Gamelin and R.E. Greene. *Introduction to Topology*. Dover books on mathematics. Dover Publications, 1999.
- [189] Lorenzo De Angelis et al. “Index-symmetry breaking of polarization vortices in 2D random vector waves”. In: *Optica* 6.9 (2019), pp. 1237–1243.
- [190] G. G. Stokes. “On the Composition and Resolution of Streams of Polarized Light from different Sources”. In: *Transactions of the Cambridge Philosophical Society* 9 (Jan. 1851), p. 399.
- [191] A.Ya. Bekshaev and M.S. Soskin. “Transverse energy flows in vectorial fields of paraxial beams with singularities”. In: *Optics Communications* 271.2 (2007), pp. 332–348.
- [192] M. R. Dennis. “Phase critical point densities in planar isotropic random waves”. In: *Journal of Physics A: Mathematical and General* 34.20 (2001), pp. L297–L303.

- [193] Y. Kuznetsov. *Elements of Applied Bifurcation Theory*. Springer, New York, NY, 2004.
- [194] J. F. Nye, J. V. Hajnal, and J. H. Hannay. “Phase Saddles and Dislocations in Two-Dimensional Waves Such as the tides”. In: *Proceedings of the Royal Society of London Series A* 417.1852 (May 1988), pp. 7–20.
- [195] Roman Kuchkuda. “An Introduction to Ray Tracing”. In: *Theoretical Foundations of Computer Graphics and CAD*. Ed. by Rae A. Earnshaw. Berlin, Heidelberg: Springer Berlin Heidelberg, 1988, pp. 1039–1060.
- [196] R. P. Drake. “Perspectives on high-energy-density physicsa)”. In: *Physics of Plasmas* 16.5 (Mar. 2009), p. 055501.
- [197] Edmundo A Gutierrez-D, Jamal Deen, and Cor Claeys. *Low temperature electronics: physics, devices, circuits, and applications*. Elsevier, 2000.
- [198] A. Di Piazza et al. “Extremely high-intensity laser interactions with fundamental quantum systems”. In: *Rev. Mod. Phys.* 84 (3 Aug. 2012), pp. 1177–1228.
- [199] C. Giuliano and L. Hess. “Nonlinear absorption of light: Optical saturation of electronic transitions in organic molecules with high intensity laser radiation”. In: *IEEE Journal of Quantum Electronics* 3.8 (1967), pp. 358–367.
- [200] E. Cumberbatch. “Self-focusing in Non-linear Optics”. In: *IMA Journal of Applied Mathematics* 6.3 (Sept. 1970), pp. 250–262.
- [201] Robert W. Boyd. *Nonlinear Optics*. Ed. by Robert W. Boyd. Third Edition. Burlington: Academic Press, 2008, pp. 391–428.
- [202] Theodore H Maiman et al. “Stimulated optical radiation in ruby”. In: (1960).
- [203] Nobel Prize Outreach AB 2024. *The Nobel Prize in Physics 1981*.
- [204] N. Bloembergen et al. “Optical Second-Harmonic Generation in Reflection from Media with Inversion Symmetry”. In: *Phys. Rev.* 174 (3 Oct. 1968), pp. 813–822.
- [205] P. Guyot-Sionnest and Y. R. Shen. “Bulk contribution in surface second-harmonic generation”. In: *Phys. Rev. B* 38 (12 Oct. 1988), pp. 7985–7989.
- [206] Yongmin Jung, Gilberto Brambilla, and David J. Richardson. “Comparative study of the effective single mode operational bandwidth in sub-wavelength optical wires and conventional single-mode fibers”. In: *Opt. Express* 17.19 (Sept. 2009), pp. 16619–16624.
- [207] Manfred Hammer. *Mode solver for 2-D multilayer waveguides, variational effective index approximation*.
- [208] Yiping P. Du et al. “Reduction of partial-volume artifacts with zero-filled interpolation in three-dimensional MR angiography”. In: *Journal of Magnetic Resonance Imaging* 4.5 (1994), pp. 733–741.
- [209] Rotem A Elgavish and Donald B Twieg. “Improved depiction of small anatomic structures in MR images using Gaussian-weighted spirals and zero-filled interpolation”. In: *Magnetic Resonance Imaging* 21.2 (2003), pp. 103–112.

- [210] Thomas Theul, Helwig Hauser, and Eduard Gröller. “Mastering Windows: Improving Reconstruction”. In: (May 2000).
- [211] E. Snitzer. “Cylindrical Dielectric Waveguide Modes*”. In: *J. Opt. Soc. Am.* 51.5 (May 1961), pp. 491–498.
- [212] Sonja Franke-Arnold, Les Allen, and Miles Padgett. “Advances in optical angular momentum”. In: *Laser & Photonics Reviews* 2.4 (2008), pp. 299–313.
- [213] E. Otte and C. Denz. “Optical trapping gets structure: Structured light for advanced optical manipulation”. In: *Applied Physics Reviews* 7.4 (Nov. 2020), p. 041308.
- [214] Wei Ma et al. “Chirality-based biosensors”. In: *Advanced Functional Materials* 29.1 (2019), p. 1805512.

QUOTES BIBLIOGRAPHY

1. Peter Jackson, *The Lord of the Rings: The Fellowship of the Ring* (New Line Cinema/WingNut Films, 2001).
2. The Cinematic Orchestra, *That Home* (Album: Ma Fleur. Publisher: Ninja Tune, London, United Kingdom, 2007).
3. Benoît Mandelbrot, *The Fractal Geometry of Nature* (W. H. Freeman and Co, San Fransisco, United States, 1982).
4. Dream Theater, *Wither* (Album: Black Clouds and Silver Linings. Publisher: Roadrunner, New York City, United States, 2009).
5. Emily, *Stardew Valley* (Windows PC version, ConcernedApe, 2016).
6. Robert Jordan and Brandon Sanderson, *Towers of Midnight* (Book 13 of The Wheel of Time series, Tor Books, New York City, United States, 2010).
7. Solanum, *The Outer Wilds* (Windows PC version, Mobius Digital, United States, 2019).
8. Ghost Brigade, *Breakwater* (Album: Until Fear No Longer Defines Us. Publisher: Season of Mist, Marseille, France, 2011).

CURRICULUM VITÆ

Matthijs Alexander VAN GOGH

28-03-1992 Born in Utrecht, the Netherlands.

EDUCATION

2004–2011 High school
Christelijk Gymnasium Utrecht

2011–2015 Double Bachelor in Physics & Mathematics
Universiteit Utrecht
Thesis: Spontaneous Creation of Solitons by Shock Cooling
Sodium Bose-Einstein Condensates
Supervisor: Prof. dr. P. van der Straten

2015–2017 Master in Theoretical Physics
Universiteit Utrecht
Thesis: Anomalous Magnetic Effects in Three-Dimensional
Weyl Semimetals
Supervisor: Prof. dr. ir. H.T.C. Stoof

2018–2024 PhD. in Physics
Delft University of Technology
Thesis: The Structure of Two-Dimensional Light: Its singular
behaviour
Promotor: Prof. dr. L. Kuipers

AWARDS

2022 Best paper award NFO 2022 conference

ACKNOWLEDGEMENTS

Stay with me, until fear no longer defines us.

Breakwater - Ghost Brigade⁸

This is it. If you are reading this, that means that I have managed to successfully complete my thesis, and what you are reading is the fruits of years of labour. Such a journey is never undertaken alone and as such there are a lot of people that have been part of this ride, without whom I would not have been able to complete this adventure.

First and foremost I would like to thank my promotor **Kobus**. Thank you for all the discussions, "teaching mode" moments during group meetings, and wisdom you have (tried to) impart on me over the years. I am very glad that you are not just a boss who cares about work and productivity, but understands that there is more to life than work, and that life is not all sunshine and roses at all times. Thank you for not giving up on me, when at times I had (nearly) given up on myself. I also want to extend my thanks to my second promotor **Yaroslav**. While we did not interact that much during my time here, besides the greetings when we pass each other in the hallway (which happened a lot), the times that I did step into your office for something you always made time for me. I do at times regret not leaning into theory more for my thesis, perhaps we could have collaborated on a project somewhere.

A massive thanks, of course, goes out to all the amazing people that I have had the pleasure of working with over the years in the nano-optics group. Every one of you brought something unique to the group and I am grateful for all the great times we have had over the years. **Marc**, your endless enthusiasm, positivity, and curiosity never ceased to amaze me, and is something I deeply admire. I enjoyed having you around a lot, and I really enjoyed our conference trip to Canada together. Despite saying that you would not be continuing as a post-doc, you still ended up as one, and I am excited to see where your road is going to lead you. **Sonakshi**, when you first joined the group, I tried to mentor you and teach you what I had learned from my first year of the PhD. But now, as I am close to the end of my PhD, it is you who has provided me with support and mentoring instead. And for that I am extremely grateful. **Daniël**, you have brought some much needed youngster energy to the group, which it was sorely lacking at the time. It's great to see you running all kinds of marathons, keep that up! **Onima**, while you started out quite reserved, it was really nice to see you come out of your shell, especially at Veldhoven last year. It is really great to see you doing well here. **Irina**, you were always the motherly figure of the group. While your calls for the Friday lab cleanings were always met with groans from everyone else, they were quite necessary. The lab has never seen

that level of clean anymore since you finished. I also had a lot of fun at the summer school in Varenna that we attended. **Thomas**, of all people in the group over the years, you have taught me the most. It was you who taught me how to operate Phantom in the first place, and helped me numerous times with the science I was doing. The amount of times I barged into your office because I had more questions are too many to count. You are also the only one I know whose booze collection rivals my own, a rare occurrence. **Aron**, the level of support on both the technical and software front were extremely valuable to the group. And if there were any Friday drinks or social activities being suggested, you were always the one be all for them. Thanks for being a driving force for those. But please refrain from mixing together random drinks for the next group barbecue. **Filippo**, thanks for all being the driving force behind the weird lunchtime discussions that were had, such as if peaches and nectarines are the same or not (they are not). **Lorenzo**, your infectious curiosity and enthusiasm I still remember clearly. Along with the singing in the lab of course, which you hopefully still can do at your current job. **Javi**, always proud to burn stuff with lasers along with Marc. I enjoyed having you as my office mate during your time with us. **Felipe**, the most spiritual of us all, somehow excelled at chucking axes. Is that what they taught you at the silent retreat? Wherever you are, I hope you are living your best life. **Su-Hyun**, I am very happy you managed to secure a permanent position in your home country. I also fondly remember your chiral dog who was trained to give a high-five as well. **Martin**, I could always tell your presence in the lab by the 80's music coming from your side (along with the occasional swearing). It has become quite silent since you left, as no one is openly playing music anymore. **Corné, Freek and Luuk**, who did their MEP/BEP under my supervision, thank you for all the interesting (and sometimes difficult) questions you have asked me in that time. It was a pleasure to supervise you.

A big thanks goes out to all those who have assisted me around the department in the years I spent there, especially **Heleen** and **Etty** who I could rely on for all my administrative and logistical questions and needs. Since then, **Lidewij** has taken over from Etty and has shown herself to be every bit as capable. And Heleen's role been taken over by first **Anita**, and later **Karin**. Both have done a fantastic job at it and I am grateful for all their support. This list would not be complete without **Tino** for all the technical work around the department that needs doing. And of course I am grateful for the rest of the management team for keeping this department up and running on the non-scientific side of it. Finally, I also want to thank the janitorial staff for keeping the building clean. I do not want to image the state of this place without your daily diligence.

While my direct colleagues have been the ones I interacted with the most, I am fortunate to have met many more great people here in Delft. And they are every bit as important to the enjoyment of my time as a PhD. **Sonia**, you have made me feel like almost a part your group, even taking us with you on a group outing as well to the climbing forest in Amsterdam. This is reflected in your group as well, where especially **Abel, Chris and Sabrya** were always there at the lunch table, making for very enjoyable (and sometimes too long) lunch breaks. **Jana**, your sense of humour reminds me of my own, and I could count on you to keep me updated on all the cleanroom drama. **Parsa**, I wish we would have gotten to know each other earlier in our PhD's so we could have made music more

often. It is a rare occurrence for me to find someone with a similar eclectic taste in music. And I am grateful for all the other people that I had the pleasure to interact with over the years. If you are reading this, and you have not been named explicitly, you are certainly amongst this group and my thanks for having met.

But there are many more people in my life than those I met in Delft. I am lucky to have multiple friends that have obtained a PhD as well, and who also have had their fair share of troubles during it, and with whom I could vent about all the frustrations that are part and parcel for a PhD. A constant for the last decade have been my crew from the almanak committee and their partners: **Swinda, Menno, Carolien, Marij, Erik, Ruben, Monique, Niels, Laurens** and **Shawn**. I am happy we have remained friends after everyone went off to do their respective things after graduation. To many more Christmas dinners, summer barbecues and the occasional cocktail parties when I decide to host one.

Wiggert, Jaimy and **Nicolaas**, I am glad to count you amongst my friends, where we went from online gaming in pandemic times to having a super fun LAN party weekend, and going to concerts together. To many more of those. And thank you for offering (well, insisting) to help me move to the new house. To my friends going way way back: **Bram**, it's always fascinating how our paths keep intersecting. We went a while without talking, and I somehow stumbled across you on the campus of TU Eindhoven where you had apparently started a PhD, with me being there being a complete one-off. It will never cease to amaze me how our taste in musics remains almost identical over the years. We can not see each other for years sometimes and continue as if we just spoke last week. **Ruben**, from all the time we spent at our athletics training, the almost all-nighters playing Halo, going to Graspop, and the many nights we spent at the Boothill Saloon, it has been a journey. Thank you both for having experienced life along with me for more than 25 years at this point.

To all of my families (which have become a lot), both my own (the **Voerman, van Gogh** and **Radler** families) and the ones I married into (the **Kleizen** and **Tervoert** families), who were always asking about how it was going with my PhD and when I would be done. The answer is now, finally. My brother **Stef**, while our brotherly bond was perhaps a bit shaky in our younger years, I am happy that it has been mended considerably over the years. We are both grown adults by now after all. Physically that is, mentally on the other hand... Thanks for all the memorable Van Gogh borrels where we grab a beer and chat. Now that I have returned to living closer to Utrecht, I am happy to hold those more frequently in the future. And no, I am still not becoming a programmer. To my mom **Ellen**, thank you for all the love and support you have given my over the years. For all the nice camping trips that we made to the Ardennen and to Cadzand. And for always welcoming me when I need a place to crash, such as when I gambled on the last train home (and lost) after having one of the aforementioned Van Gogh borrels. It probably won't be the last time. To my dad **Paul**, despite living an ocean away, you are the one I chat with the most, every Sunday evening to be precise. Thank you for your continued interest in my work and my life, despite the physical distance. Living far away has its benefits, such as always having a nice vacation home for me. Thanks for all the great roadtrips that we made when we visited in the summer. I have seen much more of the US than many of its residents at this point.

My dear cat **Pien**, your crazy antics have brought a lot of life and joy into the household. I am very happy that you are around to cheer me up and cuddle whenever I feel like it. Your purring is my favourite sound in the world.

Above all I am extremely grateful to my amazing wife **Marinda**. This journey has not been easy, and while we knew going in that it was not going to be, neither of us could foresee what was to come. I deem it very unlikely that I would have seen this through to the end without you by my side. Thank you for being there when life got tough. We have had many great trips together during the 7.5 years we've been together by now, with the highlight of course being our wedding and subsequent honeymoon to Costa Rica. What a magical adventure that was. I look forward to all the amazing adventures that we are going to have in the future, which I am sure are going to be many.

