

Springer Theses

Recognizing Outstanding Ph.D. Research

Alexander John Taylor

Analysis of Quantised Vortex Tangle

 Springer

Springer Theses

Recognizing Outstanding Ph.D. Research

Aims and Scope

The series “Springer Theses” brings together a selection of the very best Ph.D. theses from around the world and across the physical sciences. Nominated and endorsed by two recognized specialists, each published volume has been selected for its scientific excellence and the high impact of its contents for the pertinent field of research. For greater accessibility to non-specialists, the published versions include an extended introduction, as well as a foreword by the student’s supervisor explaining the special relevance of the work for the field. As a whole, the series will provide a valuable resource both for newcomers to the research fields described, and for other scientists seeking detailed background information on special questions. Finally, it provides an accredited documentation of the valuable contributions made by today’s younger generation of scientists.

Theses are accepted into the series by invited nomination only and must fulfill all of the following criteria

- They must be written in good English.
- The topic should fall within the confines of Chemistry, Physics, Earth Sciences, Engineering and related interdisciplinary fields such as Materials, Nanoscience, Chemical Engineering, Complex Systems and Biophysics.
- The work reported in the thesis must represent a significant scientific advance.
- If the thesis includes previously published material, permission to reproduce this must be gained from the respective copyright holder.
- They must have been examined and passed during the 12 months prior to nomination.
- Each thesis should include a foreword by the supervisor outlining the significance of its content.
- The theses should have a clearly defined structure including an introduction accessible to scientists not expert in that particular field.

More information about this series at <http://www.springer.com/series/8790>

Alexander John Taylor

Analysis of Quantised Vortex Tangle

Doctoral Thesis accepted by
the University of Bristol, Bristol, England

 Springer

Author

Dr. Alexander John Taylor
H H Wills Physics Laboratory
University of Bristol
Bristol
UK

Supervisor

Prof. Mark Dennis
H H Wills Physics Laboratory
University of Bristol
Bristol
UK

ISSN 2190-5053

Springer Theses

ISBN 978-3-319-48555-3

DOI 10.1007/978-3-319-48556-0

ISSN 2190-5061 (electronic)

ISBN 978-3-319-48556-0 (eBook)

Library of Congress Control Number: 2016955928

© Springer International Publishing AG 2017

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made.

Printed on acid-free paper

This Springer imprint is published by Springer Nature

The registered company is Springer International Publishing AG

The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Tanto monta cortar como desatar
(Cutting is as good as untying)
Topologically misguided Spanish proverb

Supervisor's Foreword

I am pleased that Springer is publishing Alexander Taylor's exceptional thesis. The quantised vortex tangle he analyzes is found in random waves: mathematical models of interference in three-dimensional space displaying great complexity, which are very different from the simple structure of textbook plane waves. The most significant features of these waves are vortices, which are stationary filaments in the interference pattern, and zeros in the wave's intensity around which the energy flows. As centres of rotation of a surrounding flow (such as in a fluid), vortices hold a fascination for physicists. Vortices in waves are also phase singularities, meaning that the wave field's phase (represented by the argument of the wave's complex amplitude) is not defined on a vortex, and the full cycle of phases occurs on any loop around it. A wave field's topology is thus manifested in its vortex structure. This thesis describes a further topology—the kind of knotting exhibited in the tangle of vortex lines, which is explored using computer models of wave interference in several distinct model systems.

Although vortices were understood as a general phenomenon in wave interference over 40 years ago (by John Nye and Michael Berry, in Bristol), their study has intensified in the last couple of decades, especially in light waves, where the experimental manipulation of laser beams is enabled via holograms, which has led to the fields of singular optics and structured light. Indeed, laser beams can be structured to contain knotted and linked vortices (in theoretical work by Michael Berry, Robert King and myself, and realised experimentally by the group of Miles Padgett in Glasgow).

On the other hand, random waves are a universal and natural feature in many kinds of wave, such as the surface of the sea, in speckle patterns of light reflected from rough surfaces, in acoustic noise in a room, or quantum matter waves in a chaotic billiard cavity. The natural mathematical model for such waves is simply a sum of plane waves, with random directions, amplitudes, and phases, whose structure is surprisingly rich and subtle. Michael Berry and I, in our calculation of the local statistical properties of vortices in this model (including their density and curvatures), speculated that systems of random waves could contain knotted

vortices, although there are no known techniques to answer that question analytically. Subsequent numerical experiments on optical speckle, with Kevin O'Holleran and Miles Padgett in Glasgow, found linked vortex loops to be relatively common, but no examples of knots.

Alexander's thesis addresses this problem by large scale computation, generating random wave fields in volumes much larger than previously: vortex filaments are tracked in the 3D volume, and then calculations are made of topological knot invariants for vortex loops, and fractal properties of the entire tangle. The vortex tracking method is a refinement of previous methods used to follow vortex lines, which resolves the volume of the interference pattern into voxels, and candidate vortices as phase singularities are located in edges of the voxel lattice. In a generalization of the classical bisection method for finding roots, Alexander's procedure re-resolves these regions of the volume, significantly improving the resolution of the vortex pattern. With these improvements, quantitative agreement is possible with the analytically-derived, smooth curvature statistics of random vortices, and the methods were successfully extended to find distributions of other interesting geometric quantities such as persistence length and torsion.

The topological identification algorithm also improves previously published methods: rather than relying on a single discriminating knot invariant, which may be computationally costly to calculate for a complicated, tangled loop, instead several simpler invariants are calculated, whose fingerprint has sufficient discriminatory power for unambiguous identification of most tabulated knots to a reasonable level of complexity (up to around nine crossings). Not only are knots found to occur in random waves, but the statistics of many thousands of random wave simulations in several systems show a diverse range of different knots with varying topological complexity, which is different in the three principal systems (plane waves in a cube with periodic boundary conditions, hyperspherical harmonics in the 3-sphere, and eigenfunctions of the 3D quantum harmonic oscillator). The main results were extended and published after Alexander's Thesis was completed, in Taylor and Dennis, *Nature Communications* vol 7, article 12346 (2016).

The results are of interest to students and scientists working in singular optics, quantum chaos and quantum turbulence, where random wave models and vortex tangles occur. Mathematicians interested in random geometry and topology may find grounds for new conjectures in statistical topology, an area where exact, analytical results are notoriously rare; only a little is known about the topology of these systems at lengthscales beyond a few wavelengths. Knotted random walks, which random tangled vortices exemplify, are very important in polymer physics, and there are tantalizing connections in the kinds of knotting topology exhibited in different systems which are explored in the thesis.

Alexander's work fits into the rich tradition of singular optics and wave physics from Bristol Theoretical Physics going back over four decades. The results described here are only possible with modern high-performance computers, which are necessary for the recognition of highly complex knots. Alexander has made this

subject his own, and, as with the best graduate students, he has changed the way I, as his supervisor, approach the subject. I am especially pleased that Alexander's work answers many questions posed in my own PhD thesis. His thesis provides an excellent summary of the results and provides a background to this intriguing area. We are very grateful to the EPSRC and the Leverhulme Trust for financial support which has made this work possible.

Bristol, UK
September 2016

Prof. Mark Dennis

Abstract

This thesis is an investigation of the tangled vortex lines that arise in the interference of complex waves in three dimensions; they are nodal lines of the intensity where both the real and imaginary components of the wave field cancel out, and are singularities of the complex phase about which it sweeps out a quantised total change. We investigate the behaviour of this tangle as expressed in random degenerate eigenfunctions of the 3-torus, 3-sphere and quantum harmonic oscillator as models for wave chaos, in which many randomly weighted interfering waves produce a statistically characteristic vortex ensemble.

The geometrical and topological nature of these vortex tangles is examined via large-scale numerical simulations of random wave fields; local geometry is recovered with sufficient precision to confirm the connection to analytical random wave models, but we also recover the (high order) torsion that appears analytically inaccessible, and quantify the different length scales along which vortex lines decorrelate. From our simulations we recover statistics also on much larger scales, confirming a fractality of individual vortices consistent with random walks but also comparing and contrasting the scaling of the full vortex ensemble with other models of filamentary tangle.

The nature of the tangling itself is also investigated, geometrically where possible but in particular topologically by testing directly whether vortex curves are knotted or linked with one another. We confirm that knots and links exist, but find their statistics greatly influenced by the nature of the random wave ensemble; vortices in the 3-torus are knotted far less than might be expected from their scales of geometrical decorrelation, but in the 3-sphere and harmonic oscillator exhibit more common and more complex topology. We discuss how this result relates to the construction of each system, and finish with brief discussion of some selected topological observations.

Acknowledgements

I have to start by thanking my supervisor, Mark Dennis. I cannot imagine what Ph.D. life would have been like without his remarkable enthusiasm and dedication, and not least for suggesting such an interesting project in the first place!

I also owe much to the discussions with academics at both Bristol and elsewhere. In particular, thanks to Rami Band for his help and discussions academic and otherwise, to Stu Whittington for allowing me to probe his fount of knowledge and the illuminating insights that always followed, and in Bristol to James Annett, Michael Berry and John Hannay for the useful and interesting suggestions and interactions throughout the last few years. But these are just those whose paths I have crossed the most; thanks to all of the Bristol Physics Theory Group, including those no longer around, for contributing to such a welcoming environment with such diverse goings on.

It is important that I not forget the roles of my fellow students, particularly those I have shared an office with at one time or another. Dare I say, without you it would have been really boring! Oh, and to James Ring, I just noticed your own acknowledgement of me; suffice to say I was as delighted to prattle on as you apparently were to listen.

Outside the office, thanks to everyone with whom I have played Go, juggled, shared a flat (that's you Richard!) or even just a drink, for I am not sustained by science alone. Special mention to the Kivy team, they are pretty cool dudes. And a little outside the box, I am grateful to everyone that contributed to the scientific software ecosystem from which I've drawn heavily—when it comes down to it, this would have been way harder without that stuff.

To my parents, it is cliché to say it but yes; I have finished writing now. I could not have done it without you, so thanks for everything! (actually really now).

Contents

1	Introduction	1
1.1	General Introduction	1
1.2	What Does It Mean to Be Tangled?	2
1.3	Phase Singularities in Two and Three Dimensions	6
1.4	Random Waves and Vortex Tangle	15
1.5	Knotting and Linking as a Statistical Tool	22
1.6	Other Random Tangles and Random Walks	26
1.7	Variations on Random Waves	29
1.7.1	The 3-Torus (Periodic Boundary Conditions)	29
1.7.2	The 3-Sphere	34
1.7.3	The Quantum Harmonic Oscillator	37
1.8	Thesis Outline	39
	References	40
2	Numerical Methods	45
2.1	Numerically Locating Vortices	45
2.2	Tracking Vortices with a Resampling Algorithm	53
2.3	Vortex Tracking in Different Systems	60
2.4	Geometrical Details in Different Systems	65
2.4.1	The 3-Torus	65
2.4.2	The 3-Sphere	65
2.4.3	The Quantum Harmonic Oscillator	67
2.5	Summary of Numerical Results	68
2.5.1	The 3-Torus	68
2.5.2	The 3-Sphere	70
2.5.3	The Quantum Harmonic Oscillator	71
	References	72
3	Geometry and Scaling of Vortex Lines	75
3.1	Introduction to Geometrical Statistics	75
3.2	Local Geometry of Random Wave Vortices	77
3.3	Statistics of Curvature and Torsion	81

3.4	Geometric Correlations	85
3.5	Three Measures of Fractality	87
3.6	Scaling Relations of Random Tangle	94
3.7	Crossing Number and Writhe of Vortex Loops	100
3.8	Discussion	104
	References.	106
4	Topological Methods	109
4.1	A Practical Introduction to Knot Theory.	109
4.2	Calculating Invariants	121
4.2.1	The Alexander Polynomial	122
4.2.2	The Determinant and Roots of the Alexander Polynomial	124
4.2.3	The Jones and HOMFLY-PT Polynomials	126
4.2.4	Hyperbolic Volume.	128
4.2.5	Minimal Crossing Number	128
4.2.6	The Linking Number.	129
4.2.7	The Fary-Milnor Theorem.	130
4.3	Simplifying Space Curves.	130
4.4	Topological Complexity	137
	References.	138
5	Knotting and Linking of Vortex Lines	143
5.1	Finding Knots.	143
5.1.1	The 3-torus	144
5.1.2	The 3-sphere	147
5.1.3	The Quantum Harmonic Oscillator	150
5.2	Probabilities of Knotting.	153
5.3	Comparison of Knot Type Statistics	162
5.4	Finding Links	166
5.4.1	The 3-torus	166
5.4.2	The 3-sphere	168
5.4.3	The Quantum Harmonic Oscillator	169
5.5	Probabilities of Linking	170
5.6	Knotting in Lines with Non-trivial Homology	174
5.7	Postselection Relations	177
5.8	The Smallest Knot	179
5.9	Discussion	184
	References.	186
6	Conclusions	189
	Appendices	193

Chapter 1

Introduction

1.1 General Introduction

It was Leonardo da Vinci who first studied the complex, roiling behaviour of a turbulent fluid, being the first person known to have identified the role of *vortices* as vital to describing the complex swirling patterns he observed [1], and comparing their tangled filamentary structure to a woman's hair. Figure 1.1 reproduces one of his original sketched diagrams, in which one flow of water mixes with another with the swirling vortices picked out via the streamlines of velocity that coil around them.

Da Vinci would not have known at the time that such complex vortex tangles are not just confined to classical fluids, but are even more fundamental to describing spectrally extended flows such as the equivalent turbulence phenomena in quantum fluids, or even to simple linear superpositions of complex scalar waves such as might describe a light beam, the tides, or an eigenfunction of a chaotic cavity. In these latter cases, the field contains vortices of the *complex phase* around which the field may be considered to rotate, forming a complex tangle not so visually dissimilar to the classical turbulence observed centuries ago, but doing so without the complex nonlinear dynamics and interactions of a mechanical fluid; its nature comes entirely from the random nature of the wave ensemble. It is these linear quantised vortex tangles, arising as a generic phenomenon of interacting waves, whose behaviour is the subject of this thesis.

Given some tangle of one-dimensional filaments, it is natural to ask what characterises it; are the vortex tangles observed by da Vinci in some way similar to those of the random wave interference we will investigate? What does it even mean that they are tangled, does this imply some universal properties of the filamentary ensemble, and can we quantify geometrically how individual filaments wind around one another? Such questions have been asked in a range of different tangled systems, and we will use these as a baseline for our own comparison. This thesis presents a numerical investigation by way of simulating large three-dimensional regions of random wave interference, tracing the paths of vortex filaments in order to directly analyse statistics of their large scale tangling.




Fig. 1.1 Turbulence in mixing fluid, drawn by Leonardo da Vinci, image from [2]. *Lines* show the streamlines of velocity

Figure 1.2 gives an example of vortices in the type of *wave chaos* we investigate, the one-dimensional phase singularities in a three-dimensional superposition of 100 randomly oriented plane waves. Each vortex line is coloured by a different hue, and the vortices in the plotted cube of side length 7λ wind about one another in complex ways. It is this characteristic behaviour that we will investigate. The rest of this Chapter introduces the details of this notion of wave chaos, and of how it leads to such a filamentary tangle.

An alternative direction of enquiry is to ask directly whether the vortices of the tangle are *knotted* or *linked* in the mathematical sense of knot theory. This captures a different aspect of behaviour to that of the geometry, relating more fundamentally to how vortex strands wind about themselves, and potentially varying significantly even amongst filamentary tangles that appear geometrically similar. This has been used in other systems to capture aspects of randomness, and we will analyse vortex tangle on this basis; are vortices knotted and linked, do they display the universality of an unbiased random walk, and to what extent is their nature specific to wavefield that they inhabit?

1.2 What Does It Mean to Be Tangled?

We have introduced already the tangle of phase singularities in random waves, with an example depicted in Fig. 1.2. This tangle is geometrically complex and apparently disordered; individual lines wind about themselves and about other lines, bending




Fig. 1.2 The tangle of filamentary phase singularities in wave chaos. Each different line marks the path of a single singularity, obtained by numerical simulation of 100 interfering plane waves of equal wavelength but Gaussian random amplitude and uniformly random direction on the sphere. The plotted cubic cell has side length 7λ . The details of this random wave model are explained in Sect. 1.4, and of the numerical algorithms used to track their paths in Chap. 2. Each singularity line is coloured in a different random hue, and either forms a closed loop within the simulated region or terminates on its boundaries

and twisting along their extent such that they appear well mixed—we might imagine that if these lines were made of string, they would not be easily separated if pulled tight. They are thus tangled in some sense, but what does this mean for their statistics; how should we quantify tangling, and are such differences specific to a given system or in some ways *universal*?

These questions are important in the context of tangles that might appear morphologically similar but arise from very different physical constraints. For instance, the tangle of long molecules in a polymer melt is constrained only mechanically and has many characteristic behaviours [3], whereas the paths of vortices in classical turbulence such as studied by da Vinci are not so well defined because the vortices are not so constrained, being able to merge or part with one another even though the bulk is somehow tangled [4]. This general behaviour is preserved in *quantum turbulence* [5], though now the vortices are quantised as will be explained in the next Section. Despite this, the bulk structure still depends somehow on the entwining of the full ensemble of filamentary strands. This leads back around to our example

of Fig. 1.2, in which there is neither mechanical exclusion nor even time dynamic evolution, and the tangle is of filamentary singularities of the complex phase. We introduce below some of the measures of tangling that we will apply to this system, some of which have previously been investigated in other contexts to which we will compare.

One obvious type of question to ask is about the individual filaments; how exactly do they curve, is their direction persistent or rapidly decorrelating, and are they ordered or disordered on large scales? These questions are known to be able to classify individual lines into different classes of which the most disordered is the *random walk*, the limit at which the line is totally decorrelated beyond a certain scale, implying many statistically characteristic properties on large scales [6]. There has been much investigation of random walks in the literature, and we examine some specific results that we make use of in this thesis in Sect. 1.6.

Classifying individual filaments as random walks has immediate implications for tangled lines, but ignores other geometric properties of the bulk such as how the curves interact not just with themselves but with one another; for instance, self-avoiding tangle might display quite different characteristics to one in which the filaments wind independently. Any large scale properties of the line ensemble might depend on these details, but tangles that otherwise appear morphologically similar can be very different on this level.

One type of measurement that has been used to characterise tangling on the large scale is based directly on the way curves wind around one another, or equally expressed in how frequently they cross in projection [7]. In the simplest case, the *average crossing number* (ACN) integrates a certain property of how one curve winds about another or itself; for any two curves L_1 and L_2 (which may be the same curve) this is defined as [7, 8]

$$\text{ACN} = \frac{1}{4\pi} \oint_{\mathcal{L}_1} \oint_{\mathcal{L}_2} \frac{|(\mathbf{X} - \mathbf{Y}) \cdot d\mathbf{X} \times d\mathbf{Y}|}{|\mathbf{X} - \mathbf{Y}|^3}, \quad (1.1)$$

where \mathbf{X} and \mathbf{Y} denote one point from each curve. In fact the calculation is equivalent to counting the the number of times the curve intersects itself in linear projection to a plane, averaged uniformly across the sphere of directions of projection; the integral counts the contribution to the total solid angle of each intersection. This integral applies to at most two curves, the ACN of an entire tangle would be the sum over all individual curves and all pairs of curves. The ACN naturally captures a sense in which tangles are visibly complex, it is increased by both increased local disorder and by curves winding about one another, and it is not sufficient for the curves to be individually complex but spatially separated as then they would only intersect in projection for a small range of directions and the overall contribution to the ACN would not be large.

This is the crudest measurement of this type, counting nothing more than the number of intersections. If we instead do not take the modulus of the numerator we see that in a given projection there are also two different types of crossing depending on the cross product of the two strands (and additionally normalised by $\mathbf{X} - \mathbf{Y}$). The




Fig. 1.3 The two different orientations of a crossing where one line intersects another in planar projection. **a** Shows a clockwise crossing (we also label this ‘+’ or ‘c’) and **(b)** an anticlockwise crossing (also ‘-’ or ‘a’), though for the calculation of writhe it is only important they are distinguished

two options are drawn in Fig. 1.3 and are labelled as *clockwise* (‘+’) or *anticlockwise* (‘-’) for reasons we will see later, though for this calculation it is only necessary that they are distinguished. The integral over two curves becomes

$$\text{signed crossings} = \frac{1}{4\pi} \oint_{\mathcal{L}_1} \oint_{\mathcal{L}_2} \frac{(\mathbf{X} - \mathbf{Y}) \cdot d\mathbf{X} \times d\mathbf{Y}}{|\mathbf{X} - \mathbf{Y}|^3}, \quad (1.2)$$

which is no longer a simple average number of the crossings in projections but a *signed* average, counting +1 for each clockwise crossing and -1 for each anticlockwise. The name of the operation depends on whether it is applied to a single curve or to two. In the first case, it is the *writhe*, and now characterises a more subtle facet of how a curve may wind about itself; the writhe is increased by crossings of one sign appearing more commonly than the other, which can only occur if the strand coils more consistently in one direction about itself [9]. The writhe thus detects a previously unseen aspect of the large scale tangling of a curve with itself that may be compared between filamentary strands of different origins; on average its expectation would be zero for a random walk, but the modulus of its writhe will have a specific expectation, and its value will vary depending on the local geometry of the curve.

Although it captures a new large scale property of tangled strands, the writhe is still fundamentally geometrical, and we see a very different result if we consider the integral over two *different* strands. This result is thought to have first been considered by Gauss [10, 11], for which reason these integral formulae often bear his name. The difference when considering two different curves is that the integral is *invariant* to changes in geometry, the average signed sum of clockwise and anticlockwise crossings never changes unless one curve actually passes through the other [10]. This sum for two links is called the *linking number* L_n , and is the first *topological* measure we have seen. Its sign is arbitrary up to the choice of orientations along the curves, and so it is $|L_n|$ that is invariant as we have defined it.

The linking number, measured across the entire tangle by considering linking of every loop with every other, captures something quite different to the geometrical quantities discussed so far; it is non-zero only if two curves directly wind around one another, not in a local geometric fashion but globally considering both curves as closed loops. Since it is not affected by the local geometry (excepting changes in which loops pass through one another or otherwise change their topology by joining or splitting), it is not clear how fundamental it should be to different types of tangle. On one side, we might anticipate that tangling on the large scale could lead to statistically stable topology that directly counts the tangling of lines with one another, but it might also be the case that different local physics leads to quite different large scale results without universality of the statistics. The idea of topologically distinct states does not just apply to two or more linked loops but to individual strands, in which case the different topological states are different *knots*. We have already seen that the Gauss linking integral is not an invariant of the knot when applied to a single loop, but in fact other measures *do* exist to distinguish differently knotted curves.

The analysis of topological statistics of our tangle systems makes up a large part of this thesis, including characterising exactly what kinds of knot or link may appear in our tangles arising from wave chaos, measuring exactly how common non-trivial topology is, and comparing these results with both random walks and other random tangle where possible.

1.3 Phase Singularities in Two and Three Dimensions

Having introduced our larger goal of analysing what characterises filamentary tangle, we explain in these next two Sections the origins of our filamentary tangle of phase singularities. We begin here with a brief but more quantified explanation of what it means for a spatially extended wave field to have a *phase singularity*, how these phenomena arise naturally and generically in superpositions of complex scalar waves, and why they are described as *vortices*.

The general concept of a mathematical singularity is a point (or set of points) where some quantity is not well defined or well behaved, changing abruptly (such that limiting values are not consistent) or becoming infinite. A simple example is the question “what time it is at the North or South pole?”; stepping infinitesimally in different directions will send you to completely different time zones and a different answer to the question. Since all the time zones meet at a point, the time limits to different values from different directions, and the point of intersection is a singularity of time. Further, if we walk in a closed path about the pole and count the change in time whenever we pass between time zones (ignoring the $+12 \rightarrow -12$ date line jump), we achieve always the same quantised time change of 24h. This is thus a *quantised* singularity.

An early example of a singularity in wave phenomena was observed by William Whewell mapping the wavefronts of tidal maxima along coastlines [12]. Using measurements along coastlines bordering the North sea and extrapolating his values to

its centre, he realised that the lines must at some point meet, with the tidal lines being rotating wavefronts with a period of roughly 12 hours. The singularity here is in the time of high tide, which is not defined at the intersection of its contours. Equivalently we can think of the *phase* of the tidal waves, the current progression of every point in the oscillation of high to low to high again, which changes with time as the tidal wavefronts move but still exhibits the singularity in its centre.

This tidal phase may be thought of as the direction of a vector field; at every point, an arrow indicates the direction of movement of a wavefront, and the singularity is the point at which this direction is undetermined. This is then a generic property of vector fields. In two dimensions we can look at this as a property of complex scalar fields, each of which has a magnitude and direction via its argument (or phase) on the Argand plane, and a singularity of direction is then a singularity of this complex phase. An example occurs in the function $(x + yi) \exp(-(x^2 + y^2))$, whose intensity and phase are depicted in Fig. 1.4a shows the intensity, normalised to tend to zero at infinity but also dipping to zero at the origin. The intensity pattern alone does not reveal the structure of the phase, shown in Fig. 1.4b to vary continuously around the central point. Again, it is a singularity because the phase has no consistent definition at the point where all phases meet, as its limiting value from any direction would be different. In fact, the integrated phase along any closed path about the core (ignoring other singularities) is precisely 2π and in general can only be $2m\pi$ for some integer m in order that it be continuous; other results could arise only from the line integral passing through another singularity such that the phase is discontinuous. Thus, this vortex is quantised just like that of the time at a pole in our original example. Figure 1.4c, d show the intensity and phase of an example higher order vortex as occurs in a Laguerre-Gauss mode, which might describe the form of a laser beam from a cylindrical cavity [13] (the image would be a two-dimensional slice through the beam). The intensity in this case is concentrated in a bright ring, but dips to zero at the origin at which location a vortex of second order is present; the integrated phase along any closed path about this core would be $\pm 4\pi$ depending on the orientation of the path.




Fig. 1.4 Examples of complex vortices in two dimensions of order one and two. **a** and **b** show the intensity and phase of the function $(x + yi) \exp(-(x^2 + y^2))$, exhibiting a first order complex vortex at the origin, at which the intensity must be zero as the phase is not defined. The integrated phase in any closed path about the origin is $\pm 2\pi$. **c** and **d** show the intensity and phase of a Laguerre-Gauss mode with radial index 0 and azimuthal index 2, with a second order vortex in the center about which the integrated phase along a closed path is $\pm 4\pi$

A phase singularity also may have one of two different *orientations*, with the phase circulating either clockwise or anticlockwise by any consistent definition of direction; both of the above examples have the same orientation, we encounter hues in the same order in both cases as we travel clockwise in a path about the origin.

Since the phase takes values in the range 0 to 2π and is defined modulo 2π , we have represented it via the hues of the colour wheel, as in the colour bar of Fig. 1.4b, which are themselves periodic and so wrap continuously about the phase singularity. This choice of colour scheme is used throughout this document to represent a complex phase, and the same (though less frequently) for the intensity as shown in Fig. 1.4a.

We also note that similar singularities occur in *director* fields specifying a direction at every point (i.e. an unsigned vector), though now multiple types of singularity can exist depending on the local rotation of the directors, and the director need not rotate entirely in a path about a singularity but is invariant under a rotation of just π . Director singularities occur in physical systems including being expressed in many different ways in liquid crystals [14] in which the individual molecules are elongated and can be described as a director field in the bulk, or in polarisation which is a director quantity admitting the same singularities which occur generically in light [15]. We will return to these ideas later.

We will proceed to introduce some of the properties of complex singularities by considering how they may occur naturally in a physical context; we consider an optical slit experiment, such as first proposed and performed by Young in 1804 [16]. The setup begins with a coherent light source placed to shine towards a screen but with a partition between the two. In time independent form (and ignoring polarisation), the propagation of light from such a source obeys the *Helmholtz equation*

$$\nabla^2\Psi(\mathbf{r}) = -k^2\Psi(\mathbf{r}) , \quad (1.3)$$

where ∇ is the *Laplacian* $\nabla = \frac{d^2}{dx^2} + \frac{d^2}{dy^2} + \frac{d^2}{dz^2}$, Ψ is the wavefunction describing the optical field and k is a wave number related to the wavelength λ as $k = 2\pi/\lambda$. This can be realised in practice with a laser beam, in which case the propagation of light is approximately along a single axis and the wavefunction may be approximated as $\Psi(\mathbf{r}) = \Psi_T(\mathbf{r}_T) \exp(ikz)$ where z is the axis of propagation and Ψ_T is the wavefunction in the transverse plane indexed by \mathbf{r}_T . Expanding (1.3) with this form and ignoring the second order derivative in z (negligible compared to radial derivatives) leads to the standard approximation of the *paraxial wave equation* [17].

$$\nabla_T^2\Psi_T(\mathbf{r}_T) = 2ik\frac{\partial}{\partial z}\Psi_T(\mathbf{r}_T) , \quad (1.4)$$

in which ∇_T is the *transverse Laplacian* $\nabla_T = \frac{d}{dx} + \frac{d}{dy}$.

Young's original experiment was equivalent to cutting slits in the partition between source and screen (though Young did the opposite, placing a narrow obstacle in the path of his own beam [16]). We approximate the slits as infinite, vertically oriented line sources in three dimensions unless stated otherwise. With a single slit, assuming

it is narrow enough (i.e. the slit can be approximated as a delta function in light intensity), the resulting light field is (far from the source) a *plane wave* of the form

$$\Psi(x, y) = \exp(ik \cdot \mathbf{r}_T), \quad (1.5)$$

for a two-component *wavevector* \mathbf{k} and $\mathbf{r}_T = (x, y)$ (i.e. in the transverse plane aligned with z). The intensity of this ideal result is shown in Fig. 1.5a, while Fig. 1.5d shows the complex phase (in this optical context, the *optical phase*) of this result; the argument of its complex value at each point. With a single slit, the intensity of the single plane wave is constant everywhere, while its phase varies continuously along its axis (for our single vertical slit, the horizontal direction).

Introducing a second slit recovers the core of Young's famous experiment and exposes the first phase singularities; the two plane waves (one from each slit) *interfere* and cancel out perfectly along dark lines in the intensity of Fig. 1.5b. In the phase pattern of (e) the phase is seen to jump by π , just as would be the case through an order one vortex point, but this singularity is not a vortex but instead extended. In a real experiment this result would be modulated by the incoming intensity of the coherent light source, and the finite slit widths would give rise to an interference pattern better described by a sinc function, but the general principles would be the same. It is this very clear example of wave interference that made Young's experiment so important, providing clear evidence that light should be properly described as a wave phenomenon.

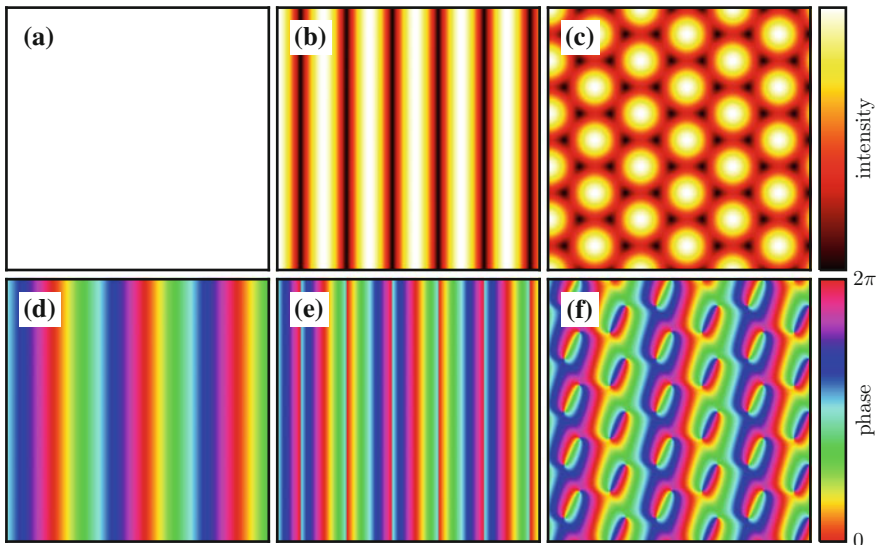


Fig. 1.5 Interference of one, two and three plane waves. **a–c** show the intensity patterns and **d–f** the complex phases across a small region of the wavefunction, in each case with one, two and three plane waves respectively. The three wavevectors occupy three points of an equilateral triangle centered on the origin and have equal magnitudes. Each figure has side length 2.9λ

Although a powerful demonstration of light's wave-like nature, this behaviour is not generic in interfering waves. Figure 1.5c shows the intensity pattern with the addition of a third slit (now horizontal rather than vertical) and therefore *three* interfering plane waves, the last of which has a wavevector pointing vertically. The intensity is now non-zero almost everywhere but vanishes at single points, at which all three waves cancel out (i.e. sum to zero). These points are now phase singularities just as in our original example of a complex function. In fact this behaviour is generic in wave interference; adding additional plane waves to the superposition may change the pattern of singularities positions, especially if the waves have different amplitudes or directions, but phase singularities are still single points with these same properties except in the limit where the amplitude of one wave dominates and no singularities exist.

We have located these points mainly as singularities of the phase, but they are also identified by their other characteristics; they are *nodal points* of the intensity, representing the *zero level set* of the wavefunction. Their phase circulation makes the points *quantised vortices* of the phase—quantised because the phase integrated along a closed path around a single singularity cannot take an arbitrary value but instead can only return (modulo 2π) to its starting value, unless it passes through another singularity. We will use these terms interchangeably throughout this thesis as appropriate, also often referring to the quantised phase singularities in complex fields as simply *vortices*.

Our examples in Figs. 1.4 and 1.5f include only vortices that are perfectly regular, but in general the anatomy of a vortex has other characteristics. We have already noted that although their phase circulation must be continuous, its integrated value is not confined to be $\pm 2\pi$ but instead may be any multiple $2\pi m$ for positive or negative integer m . However, although such higher order vortices are possible in principle, they do not occur in generic wave superpositions, as such an interference pattern is tremendously unlikely to occur at random [15]. The vortex is unstable even in constructed functions; in these cases even tiny imperfections will actually collapse the vortex to two nearby vortices of order one (an integrated phase of 2π each), which in practice will always happen if performed experimentally [18]. The local phase structure also admits other characteristic patterns, such as different degrees of anisotropy of the phase about the vortex.

Figure 1.6 shows two examples of more complex plane wave superpositions with more than just three components, demonstrating how vortices may have a more complex local structure than those in previous simple demonstrations. (a) and (b) show sections of the intensity and phase respectively from a sum of 20 different plane waves with equal wavelengths but now with random, independent, identically distributed (IID) complex Gaussian amplitude coefficients. (c) and (d) show the same for 40 plane waves with *different* wavelengths, and random Gaussian IID amplitudes additionally modulated by a Gaussian based on their wavelength. These two different choices of field construction give visibly very different results, including quite different vortex patterns with the monochromatic field exhibiting a more regular vortex grid. These differences will manifest in different correlation functions and local





Fig. 1.6 More complex superpositions of plane waves. **a** and **b** show the intensity and phase respectively of a sum of 20 plane waves with random wavevector \mathbf{k} directions and random amplitudes (complex circular Gaussian with mean zero), but the same wavelengths. The square area has side length 3.85λ . **c** and **d** show the same plots for 40 plane waves with random wavevector directions and Gaussian amplitudes, but different wavelengths with an additional amplitude modulation by a Gaussian $\exp(-|\mathbf{k}|/15)$. The side length of the plotted cell is 7.2Λ , where Λ is a wavelength analogue corrected appropriately for the Gaussian spectrum [19]

vortex structures, which are also clearly different to the rigidly uniform vortices of Fig. 1.5f.

We have so far established the existence of complex phase singularities as line or (generically) point phenomena in two dimensional wave interference, but two dimensions is itself a special case, and none of our previous examples are physically limited to the plane (unlike the tides or time zones on a two-dimensional globe). For instance, in the optical beam context each interference pattern would be only a two-dimensional slice through a propagating three dimensional field, which must have structure of its own along the axis of propagation (we label this the z -direction). Alternatively, considering directly the interference of plane waves, a three-dimensional wavefunction would be obtained by considering component waves with wavevectors oriented in three dimensions. In this sense, the two-wave interference of Fig. 1.5b, e gives rise to three-dimensional phase singularity *sheets*, while the vortex

Fig. 1.7 Vortex lines in a small three-dimensional volume, plotted in two different ways. The wavefunction is a sum of 20 plane waves of equal wavelengths but uniformly random three-dimensional directions on the sphere, and the side length of the plotted cubic cell is 1.73λ , marked by the enclosing *black lines*. In **a**, the one-dimensional vortex lines are shown in blue, the translucent *red blobs* show contours where the intensity is half its maximum over the volume, and on one side of the cell is plotted the complex phase in this plane. In **b**, the zero contours of the real and imaginary components of the field are shown in *green* and *red* respectively, with vortices again shown in *blue* at their lines of intersection. In both plots, contours are picked out with a marching cells algorithm and are accurate only up to the initial sampling resolution of 0.017λ



points of (c) and (f) are positions on straight, one dimension vortex *lines* through three dimensions, which would also appear in planes at different distances from the slit apparatus. However, just as in the two-dimensional slice this does not expose the full range of vortex behaviour, the lines may be curved and twisted in superpositions of more plane waves or even three waves with different amplitudes.

Figure 1.7a shows a small cubic region of a Gaussian random monochromatic wave superposition of 20 different plane waves, chosen with wavevectors uniformly

randomly distributed on the direction sphere. The Figure shows not the value of the superposition function but the paths of its (one-dimensional) vortices, picked out as blue lines. As expected, they are not straight but *curved*. The plotted red surfaces mark contours of 70 % of the maximum intensity over the plotted volume and, since they are nodal lines, the vortices wind around these bright patches; they exist instead at the cores of ‘dark’ regions (those with low intensity, which in e.g. a light field would be literally dark). One face of this cubic cell shows the two-dimensional phase pattern across its surface, nodal lines meeting this face can be seen to match the singularities in this cross-section. These vortex paths were obtained numerically by simulating the wavefield and tracking vortex lines using the methods explained and discussed in Chap. 2.

These one-dimensional vortex lines are the generic manifestation of phase singularities in three dimensions, and other structures are unstable. For instance, the singularity sheets of the ideal two-wave superposition would collapse immediately under a small perturbation (by adding a third wave or deforming one of the initial plane waves) to a highly anisotropic vortex line somewhere near to the original sheet. However, they would not in general collapse further to true point singularities in three dimensions, which are also a very special case; it is normal that the component waves cancel out along one dimensional lines, which also cannot simply terminate at points but must ultimately close with themselves to form *loops* (or in principle travel forever without closing, or in a physical scenario such as a light beam terminate at some interface).

The stability of vortex loops may be understood by reference to their nature as nodal lines of the complex wavefunction, which must necessarily mean that both the real and imaginary components are zero. The nodal sets of the real and imaginary components form two-dimensional surfaces which generically intersect along lines, though they might non-typically touch at a single point in an appropriate limit (on one side of which lies a tiny vortex loop, and on the other side no vortex at all), or even less easily coincide completely along sheets. As with the case of a higher order vortex, for which multiple zero-surfaces of the real and imaginary components must simultaneously intersect along a line, this is unstable and does not occur generically. Considering these nodal lines as intersections of real and imaginary level sets also makes their vortex nature clear; excepting the unstable special case of the surfaces intersecting without crossing, the nodal lines must clearly always exhibit a 2π phase circulation as their surrounding region is naturally divided into (smooth connected) volumes where the real and imaginary parts are both positive, both negative, or one of each (not in this order).

The intersection of real and imaginary zero surfaces is shown in Fig. 1.7b; the vortices shown in (a) match their lines of intersection, as expected, and in fact the zero surfaces are themselves very complicated. They form continuous sheets with high genus rather than simple almost-ellipsoidal surfaces on the wavelength scale. Following this description, we see that vortices do occur naturally in complex wave superpositions as one dimensional curves that form closed loops, or extend forever without closing, or in a real system potentially terminate on its boundaries. They are not generically straight lines as in our initial examples of two- or three-wave




Fig. 1.8 Example phase structures about a vortex line, with different self-winding numbers. In **a** the self-winding is zero, but in **b** it is one, as the local phase structure rotates once about the vortex core before the vortex closes

superpositions, but curved in a way that already shows a certain level of tangling, not being confined to a single overall direction but looping back on themselves freely. Although we have noted that vortex lines must be continuous and so either extend infinitely or eventually form a closed loop, they do not necessarily have to do so on a small lengthscale, and there are no closed loops enclosed within these example superposition volumes. Whether this occurs commonly in an arbitrary wave superposition will depend on the parameters of its construction, and this is one of the statistics we examine later.

Although Fig. 1.7 does not show the full three-dimensional phase structure, the isosurface for any given phase value would also be a continuous two-dimensional sheet, terminating only along vortex lines, with all the space between these lines being filled with non-intersecting phase sheets of different values. The phase surfaces are a natural *fibration* of the complement of the vortex lines (i.e. the manifold of \mathbb{R}^3 with their paths removed), with such a structure existing for the vortices of any complex field. This gives vortices an extra parameter, a *self-winding* counting the number of times the local phase rotates around the vortex core before it closes [19]. Figure 1.8 shows example local phase structures that might surround a vortex line with different self-winding; at any point around the tube in (a) or (b) every phase appears once, but in (b) the phase structure itself has rotated one time about the vortex core by the time the vortex closes. In real fields vortex lines may have different winding numbers.

This phase structure also has physical implications; if the complex field described a wavefunction such as that of a light beam under the Helmholtz equation, it would in a sense describe the local ‘flow’ of the light, which truly does rotate around the vortex lines. In the optical example, a small particle absorbing this light would be pushed around following such a path [20]. In this sense we have returned to something similar to our original example of Leonardo da Vinci’s classical turbulence; where he saw womens’ hair in the turbulent vortices of fluidic mixing, we see something visually similar and yet very different as a natural linear wave interference phenomenon. Were we able to see the dark lines in a light field, we would see our own tangled pattern of vortices, this time even more clearly thanks to their quantised nature; the change of phase does not indicate some arbitrary rotation that could continuously slow down until the vortex disappears, but is locked to the clear presence or absence of a vortex core. In a real physical field not made simply of plane waves with equal

wavelengths (non-monochromatic and therefore time-dependent), we would even see these vortices move around. In the next sections we provide a more quantitative analysis of what it means for a ‘tangle’ to manifest physically, different systems in which such a phenomenon may arise, previous analysis, and the vortex tangles that will be the focus of our own work.

Before doing so, we return to some of the other types of phase singularity that we will refer to elsewhere in this thesis. While we have shown the progression of the ideas behind complex phase singularities from two to three dimensions, other singularities have analogous (although in some ways very different) behaviour. For instance, the singularities of a director field in two dimensions can extend to three dimensions, though as in two dimensions there are then multiple types of singularity, and in fact the type of a singularity line may change along its length according to the local polarisation geometry [21], and in optical speckle may display conformations not unlike those of the optical phase [22]. The same is true in liquid crystals considered as a director field, though much of the interest in singularities is then in their interaction with boundary conditions [14].

1.4 Random Waves and Vortex Tangle

It has long been observed that naturally interfering physical wavesystems may be described by Gaussian random superpositions of interfering waves, those where each wave component has uniformly random direction and IID Gaussian weighting. This includes phenomena as diverse as waves on the ocean [23–25], the acoustics of cavities based on random wave interference [26, 27], the analysis of random noise such as occurs in physical devices [28] and the interpretation of statistical significance in PET scan images [29] based on the mathematical foundation of [30]. The same is true of the phenomenon of optical speckle in which light from a laser beam carries a distinctive, temporally static distribution of bright and dark regions, in this case modelled as a direct result of the many interfering wave components of reflection from any surface—since almost any surface is rough and disordered on the scale of the wavelength [31]. Perhaps most surprisingly amongst these, the eigenfunctions in ergodic cavities have a particular random character at high energies, and Berry’s *ergodic mode hypothesis* proposes that infinite sums of plane waves with Gaussian random coefficients and random phases are a good statistical model for generic high-energy eigenfunctions of such a system [32]. This captures properties including the uniform spatial distribution of such eigenfunction in the semiclassical limit of energy $E \rightarrow \infty$ (Shnirelman’s theorem [33]), although Berry’s hypothesis suggests that a generic high energy eigenfunction will be well described by the model at high finite energies, not just in the limit. In effect, this is semiclassically a suitably random sampling of an otherwise unbiased classical phase space. Returning to acoustics, a similar idea is expressed in the behaviour of regularly vs irregularly shaped rooms [26], the latter exhibiting a quite different statistical behaviour when ergodic to those that are regular.

Formally, we write such a random wave model (RWM) as [27]

$$\Psi(\mathbf{r}) = \sum_{n=0}^N a_k \exp(i\mathbf{k} \cdot \mathbf{r}), \quad (1.6)$$

where $\Psi(\mathbf{r})$ is the wavefunction resulting from the sum of plane waves with wavevectors \mathbf{k} uniformly randomly distributed on the direction sphere, each with IID complex Gaussian random amplitude a_k such that the complex phase of each wave is also uniformly random on the interval $(0, 2\pi)$. The width of the Gaussian distribution from which a_k is drawn may additionally depend on k , discussed further below. The isotropic RWM is the limit of $N \rightarrow \infty$ (ignoring issues of convergence for which an additional normalisation factor could be applied). This equation describes the ensemble of complex random waves; in the case of real waves we could simply take only the real or imaginary component (which are independent and statistically identical to one another), or replace the exponential with a sine wave of random phase. We discuss real random waves where stated explicitly below, but throughout the rest of this thesis mainly refer to the complex model.

This model fixes the wave ensemble only up to a choice of the weighting of the random amplitudes a_k ; under an appropriate normalisation these satisfy $\frac{1}{2} \sum_k a_k = 1$, and following the convention of [19, 34, 35] this is in the limit of infinite wavevectors equivalent to

$$1 = \int d^n \mathbf{k} \Pi(k), \quad (1.7)$$

where $\Pi(k)$ is the *radial power spectrum* (i.e. the width of the IID Gaussian distribution from which a_k is drawn depends only on $|\mathbf{k}|$). This definition also imposes a distinction between two- and three-dimensional power spectra as a two-dimensional distribution is not isotropic in three dimensions, discussed further below. Statistics of the random wave ensemble depend very strongly on the choice of $\Pi(k)$; we have already seen a qualitative two-dimensional example in Fig. 1.6. In (a) the spectrum is monochromatic, $\Pi(\mathbf{k}) = \delta(|\mathbf{k}| - \sqrt{E})$, while in (b) wavevectors are chosen from a square grid in \mathbf{k} -space and have their IID Gaussian amplitudes modulated additionally by $\exp(-|\mathbf{k}|/15)$. The resulting pattern of vortices differs greatly between the two examples. The effect of different spectra on vortex statistics has been investigated in detail in the literature [19, 34].

Not all of our original random wave examples use the same spectra; for instance, the analysis motivated by sea waves of [23, 24] and that of random noise in [28] apply to general power spectrums as is appropriate given the ways these systems arise. Other physical examples require a more specific power spectrum; the optical speckle model of [36, 37] is an anisotropic propagation of a two-dimensional initial plane of Gaussian spectrum, intended as a direct model for random waves in a laser beam (though in this case, even the power spectrum is not truly isotropic and so there may be other deviations from the isotropic RWM). In contrast, the analysis of quantum chaos [32] and the previously-referenced work on acoustics [26] use

random wave models of constant wavevector magnitude (i.e. monochromatic in the optical sense) since they relate to eigenfunctions of cavities; all waves have the same fixed magnitude of momentum. We will use the general term wave chaos to describe all these models, but stress that there is only a direct connection with quantum chaos in the monochromatic regime.

The analysis of this thesis will investigate only eigenfunctions, modelled by wave superpositions of constant energy. This is a choice, not a limitation of the wave chaos systems explained later, and all of our methods could be generalised to other power spectra by using appropriately weighted wave superpositions with different wavevectors. We discuss later the implications of this choice.

Figure 1.9 shows two examples of wavefields described by (1.6), monochromatic and in two dimensions, numerically calculated. (a) and (b) show the intensity and phase in an example with 100 interfering waves, and (c) and (d) the same with 10000. Although these values for N are very different, there are no statistical differences between the images that are obvious by visual inspection; given only the images, it would be difficult if not impossible to tell by eye which is closer to the infinite wave limit.

The motivation for a random wave model is that it makes accessible a wide range of statistical properties that cannot be obtained by direct calculation on a specific system. This includes in real waves (amongst other quantities) the distributions of maxima and lengths of nodal lines [23–25], density and velocities of critical points [23], the correlation functions of critical points [35, 38], distributions of nodal line lengths and their correlations [39], the curvature of nodal lines [40], and statistics of nodal domain sizes and shapes [41]; several of these examples test directly the predictions of RWMs against the actual statistics of ‘real’ systems, with such comparisons including experimental tests such as in experimentally pumped ergodic cavities [42].

In complex waves, many similar statistics have been investigated in the literature. In two dimensions this includes the densities of vortex points [43] and their correlation functions [35, 38], as well as other analysis of their geometrical statistics. In three dimensions we are interested in the behaviour of the nodal lines, and have already shown in Fig. 1.2 how they behave in a sum of 100 monochromatic plane waves; the vortices bend and tangle very similarly to our earlier example with just 20 superposing waves. The properties of these line tangles have been examined in the literature; [34] investigates many of their statistics in three dimensions, including mean densities, speeds of vortex lines in non-monochromatic superpositions, and the local geometric characteristics of the ellipse describing phase line anisotropy about a vortex core. This work also includes some two-dimensional correlation functions of complex phase singularities. [44] observes how short vortex loops can appear from nothing (or the opposite) under a perturbation of wavevectors, with a genesis at a single point before rapid growth through an elliptical form after which higher order effects begin to dominate. Similar mathematics is able to describe *reconnection* of vortex lines, the geometry by which two different vortex segments may meet at a single point before, under a further perturbation of the field, breaking apart again with the segment end points now connected in a different way.




Fig. 1.9 Intensity and phase in the random wave model with 100 or 10000 component waves. **a** and **b** show the intensity and phase respectively of a wave constructed following (1.6) with k selected uniformly on the circle in two dimensions and with $N = 100$. **c** and **d** show the same with $N = 10000$, in all cases from a monochromatic spectrum. All images show a square region of side length 7λ

Many of these statistical results on vortex lines are obtained from integration on the Gaussian random wave model, in which the values of field quantities are written in terms of correlated Gaussian random variables which can be linearised to an independent integral [19]. We can then write the expectation value of some quantity f using angle brackets

$$P(f) = \langle f(\mathbf{V})\delta(f(\mathbf{V}) - f)P(\mathbf{V}) \rangle_{\mathbf{V}}, \quad (1.8)$$

where $P(f)$ is the probability density of some function f of the field, and $P(\mathbf{V})$ is the probability of simultaneous occurrence of each of the field quantities \mathbf{V} (i.e. ζ , ζ_x , ζ_{xy} etc. where the subscripts indicate differentials) since these quantities are not independent; in practice the calculation can be simplified by linearising their matrix of correlations to obtain the same number of now-independent Gaussian-distributed quantities. Most of the vortex line properties in [34] are obtained by performing this integral analytically, with the condition of being on a vortex line guaranteed by requiring that the real and imaginary field components are zero.

Although a powerful model, not all quantities of random wave vortices appear practically computable via analytic calculation with the Gaussian random model. [34] discuss the *torsion* of dislocation lines but its averages not calculated due to its complex dependence on third order derivatives of the field, and we have found that similar calculations for (for instance) the density of intensity critical points (discussed in [44]) appear intractable even to numerical integration. For other quantities such as the large scale scaling behaviour of vortex lines, or the question of whether they knot or link with one another, it is not even clear how the answer might be obtained analytically from local considerations [34]. For these reasons, we will investigate the vortex tangle of wave chaos by simulation, numerically tracking vortex lines in sample random waves and obtaining geometrical and topological results by direct statistical analysis of the recovered vortex curves, using methodology similar to that of [36, 37] in an optical speckle model but augmented with our own methods described in Chap. 2. We will compare the results with those predicted analytically, but also find accessible this way quantities that are otherwise difficult to compute.

A limiting factor in such numerical simulation is that it is clearly impractical to simulate an infinite superposition of waves; the superposition must be finite. We have already hinted in Fig. 1.9 that this does not apparently pose a problem, as even relatively small numbers of wavevectors in superposition closely approximate the truly isotropic limit. It may even be the case that the zeros of the field are a better statistical approximation of the isotropic limit than other features such as the maxima of intensity, since the latter may depend more subtly on the normalisation of the power spectrum. We can test that a given system is a suitable approximation by checking the values of known analytical results on the geometry and density of vortex curves, and so numerical analysis can in principle proceed by simply sampling a relatively small N in (1.6).

Under such a model it is possible to simulate some volume of a chaotic wavefield, to track the vortex lines within it (we discuss this in detail in Chap. 2), and to analyse properties of their geometry and topology. However, this simulation model puts significant restrictions on the length of a simulated vortex curve, which will eventually close on a relatively short lengthscale or (more commonly) terminate on the boundaries of the simulated region and so is statistically limited in total extent. Since the vortex line will not otherwise ‘feel’ the simulated boundaries (unlike in e.g. eigenfunctions of a finite cavity [45]), it will not avoid them, and a single simulated volume will tend to contain many relatively short vortex segments. This is very visible in Fig. 1.2, which includes 22 closed loops but 310 different vortex segments that end on the boundaries of the cell. Some of these segments might be identified as

regions of the same vortex curve if simulated in a larger volume, but this is unknowable from a simple volume simulation. For this reason, we will investigate other models than truncated sums of plane waves in Euclidean space. We investigate wave chaos in degenerate eigenfunctions of different systems, whose limiting behaviour is the same but in which the behaviour of individual vortices is constrained so as to facilitate the calculation of large scale quantities.

The main emphasis here is that if a system supports degenerate eigenfunctions of the Laplacian, the most natural unbiased sampling of phase space results from random IID Gaussian weighted superpositions of these solutions [46]. At a given finite energy this clearly does not perfectly reproduce the RWM of (1.6), as the sum of degenerate eigenfunctions may not produce the same statistically unbiased wavefield as the isotropic model. However, with a suitable system we can hope that an eigenfunction of some reasonable energy will closely approximate the isotropic RWM, or in the limit of infinite energy be precisely equivalent. Phrased in this way, the monochromatic manifestation of (1.6) may be considered as the Gaussian random sum of eigenfunctions of the Helmholtz equation in \mathbb{R}^3 , isotropic due to the totally unbiased directions of wavevectors amongst individual plane wave eigenfunctions.

A natural example of such an alternative wave superposition in two dimensions are the spherical harmonics Y_l^m , the eigenfunctions of the Laplacian on the sphere which have energy $-l(l+1)$ but also a second quantum number m taking values $|m| \leq l$, with each resulting eigenfunction having the same degenerate energy. A Gaussian random superposition of these degenerate components is an appropriate random wave on the surface of the sphere. Figure 1.10a shows examples of the real spherical harmonics with $l = 7$ and $m \geq 0$, demonstrating how together they capture nodal structure across the entire sphere. Figure 1.10b, c illustrate the intensity and phase respectively of such a Gaussian random sum of the complex spherical harmonics with $l = 15$, pictured on the surface of the sphere. The result is visually similar to our previous depictions of random waves, but a statistical analysis of such superpositions would reveal that its statistics are modified by the nature of the sphere; in particular,




Fig. 1.10 Random waves via spherical harmonics. **a** shows the real part of each spherical harmonic of index $l = 7$ labelled with each of $m = 0 \rightarrow 7$; a complex, Gaussian random sum would be over all m with random phase coefficients. **b** and **c** show the intensity and phase respectively of such a random spherical harmonic with $l = 15$

the eigenfunctions are not an infinite superposition but constructed over a finite area which imposes certain periodicity constraints, and the sphere is *curved* which affects the behaviour of local geometry. Both of these effects tend to zero in the limit of high energies as $l \rightarrow \infty$, and in fact the sum of spherical harmonics can then be demonstrated to be equivalent to the RWM [46], though it is less clear exactly how statistical quantities behave before the limit is approached.

This example is just two dimensional; in this thesis we will consider examples of three-dimensional, complex, Gaussian random eigenfunctions in systems chosen such that the characteristics of their eigenfunctions impose some useful condition on the three-dimensional vortex tangle such that we can analyse and contrast closed vortex loops of significant length both as an investigation of the RWM and of how its properties may break down in the non-isotropic limit. To this end, we will consider vortex tangle in three Gaussian random eigenfunction systems. We list here the only qualitative behaviour of each such system that we investigate, and give in Sect. 1.7 the full mathematical details and examples of each one. We have already stated that our investigation will be only of pure eigenfunctions, IID Gaussian weighted wave functions of degenerate states, though we note again that this is not a limitation of this model; we could equally investigate different spectra via appropriately weighted degenerate eigenfunctions of different energies regardless of the chosen manifold.

Before introducing these systems, we note that the behaviour of vortex curves in three-dimensional manifolds is not simply an extension of the two-dimensional case, as the statistics of two-dimensional spectra are not in general the same as those of two-dimensional slices through a spectrum of wavevectors in three dimensions; for instance, the correlations of magnitude of each of the real and imaginary components depend on the number of dimensions. The *autocorrelation* function $c(\mathbf{r})$ of each of the real and imaginary components of the field is given by $C(\mathbf{r}) = \langle \xi(\mathbf{r}_0)\xi(\mathbf{r}_0 + \mathbf{r}) \rangle = \langle \eta(\mathbf{r}_0)\eta(\mathbf{r}_0 + \mathbf{r}) \rangle$ where ξ and η are the real and imaginary field components respectively, the average is over all positions \mathbf{r}_0 and, by the isotropy of the field, only the magnitude r of the displacement \mathbf{r} is important. For monochromatic waves in two dimensions $C(\mathbf{r})$ is a Bessel function $J_0(KR)$ for two-dimensional wavevector $K = |\mathbf{K}|$ and two-dimensional displacement R , whereas in three dimensions it is a spherical Bessel function $\sin(kr)/kr$ [19], leading to (for instance) half the vortex point density in a slice through three-dimensional random waves as compared to the two-dimensional spectrum [19]. This does not drastically change the nature of the model, but the examples of planar random waves in this Section come universally from planar spectra and it is important to recognise that these will be modified when considering three dimensions. True two-dimensional slices of the three-dimensional structure in each of our systems are given in Sect. 1.7.

The first system of our investigation is the *3-torus*, geometrically equivalent to a unit cube with periodic boundary conditions along all three of its Cartesian axes. The advantage here is that vortices will not terminate on the boundaries of the unit cell but can be traced *through* them as they pass through the boundary and enter the opposite face, so each short vortex segment terminating on the boundary can be identified as part of a much longer vortex loop. This makes accessible far longer vortex segments than would be the case in an equivalently sized region of the RWM.

Gaussian random eigenfunctions of the 3-torus are equivalent to the RWM in the limit of infinite volume since the effect of its periodicity becomes negligible.

Our second system will be the round *3-sphere*, the three-dimensional analogue of the 2-spherical eigenfunctions of Fig. 1.2. The advantage of this manifold is that we can trace vortex tangle across its entire volume, over which all vortices must necessarily be closed. However, just as in two dimensions the resulting tangle will be affected by both this finite volume constraint and the background curvature of the manifold, both of which may be significant at accessible energies. Our analysis will investigate this by contrast with our other manifolds.

The third and final system of our investigation will be the 3D isotropic *quantum harmonic oscillator* (QHO), in which a parabolic potential proportional to the square of the radius spatially limits the behaviour of the eigenfunctions. The difference to the previous two manifolds is that vortices occupy a region of effective finite volume, not directly confined but only rarely entering the region where the potential exceeds the eigenfunction energy, a very different type of constraint to the finite volume of the 3-torus or 3-sphere. We will see later how this affects the behaviour of the vortex tangle. The eigenfunctions of the QHO are equivalent to the RWM in the limit of the energy becoming infinite, in which case the value of the potential is negligible across an arbitrarily large volume, and the wave superposition becomes locally indistinguishable from the isotropic random wave model.

1.5 Knotting and Linking as a Statistical Tool

We introduced in Sect. 1.2 the idea of topological properties of a loop, depending not on its geometry but only on how it ultimately winds around itself or around other loops, as initially picked out by the linking number of two loops that is invariant to geometrical changes unless one loop actually passes through the other. We revisit this idea here with some of the core ideas of knotting and linking and their place in the physical context. Chapter 4 gives a deeper and more formal introduction to the relevant theory of knotting and linking, and we apply these methods later in Chap. 5.

A perhaps more intuitive way to present our original topological introduction is to ask, given a closed loop of string or rope, what different conformations the whole loop can make so long as we do not at any point cut it or otherwise pass it through itself. By twisting, bending and shaking (assuming there is no limit on the curvature of the rope) we may always obtain some tangled and geometrically complex conformation, or the opposite, removing complexity to make the loop into a smooth curve. However, there will be other results we can never achieve, and the set of possible conformations fixes the type of the knot [47]. Mathematically, these physical ideas of continuous physical manipulation without allowing intersection are contained in the operation of *ambient isotopy* which formally defines what operations do not modify the type of the knot [47]; topological invariants such as the linking number introduced in Sect. 1.2 are not modified by this operation. We define this in a more mathematical context in Sect. 4.1.





Fig. 1.11 Four examples of knots. Each image represents a curve through three-dimensional space, with **a** showing the so-called unknot while **b–d** show geometrically different representations of the so-called trefoil knot. Each curve is coloured with the colour wheel along its length only to visually distinguish the different regions of the curve

To give a visual example of knotting, Fig. 1.11a–d shows images of four loops in three dimensions. Each has been coloured along its length, but this colour has no meaning other than to distinguish different sections of the curve. (a) and (b) are both geometrically quite simple, but they have different knot types; if we had a closed piece of string in the shape of (a), it would be impossible to manipulate it to the precise shape of (b) without somehow passing it through itself or cutting it. In fact, these images show the so-called *unknot* and *trefoil knot* respectively, and are the two simplest knots in terms of having the smallest number of crossings in a projection like the one above. This means that it is not possible to draw diagrams of *any* other knot with so few crossings.

Although no other knot permits such a simple diagram (by crossing count), the opposite is not true, and even a simple knot also has arbitrarily complex geometrical representations. For instance, Fig. 1.11c, d are both also trefoil knots, topologically the same as (b), but it is very difficult to determine this visually from these images alone. Being the same knot means only that, if given all three in loops of string, it would be possible to manually contort them to the same shape. It is possible to enumerate knots in a range of different ways, placing them in broad families or classes with different properties, though we leave the details of such classification to later sections.

Linking has a very similar definition, involving not one closed loop but two or more. As before, the meaning can be understood naturally by considering a physical loop of string; two or more such loops are linked if they cannot be separated without (as before) passing one strand through the other. The simplest link with two component loops is shown in Fig. 1.12, the so-called *Hopf link*, while (b) shows the

Fig. 1.12 Two examples of links. Each image represents a pair of curves through three-dimensional space, one coloured in *green* and the other in *red*. **a** is the so-called Hopf link, and **b** the Whitehead link



Whitehead link. As with the example knots, neither representation may be contorted to the other, so they are topologically distinct. Links may be different by having individual components wind more about each other, or by adding more loops that are also topologically entangled with one or more of the others, or by changing the knot type of one or more individual loops. We have already seen the linking number as an example of a quantity that can distinguishes different links, but it is far from a perfect tool for this; it is easy to construct topologically distinct links that have the same linking number, and in fact the Whitehead link has linking number 0 just as would be the case for two unlinked curves.

None of these definitions apply to open curves. In the language of our original definition, this is because any conformation of such an object may be transformed to any other by some appropriate rearrangement even *without* passing the string through itself (any tangle may be removed by simply pulling an open end back through it, just as one undoes a ‘knot’ in a piece of real rope or string). However in a wave superposition we know already that the vortices must form either closed loops, or extend infinitely, or terminate on a system boundary. The wave superpositions of the degenerate eigenfunctions we study do not include solid boundaries, and only the QHO admits vortices travelling infinitely in which case they become geometrically trivial and can be closed at infinity without topological ambiguity (discussed further in Sect. 5.1. We may thus use topological measures of knotting and linking to classify our loop ensembles.

In physical systems, knotting and linking are important quantities in many different contexts and in many different ways. In fact some of the first popular recognition of the subject is down to the suggestion that atoms might be explained as knots in some perfect fluid, an idea due to Lord Kelvin in the 1860s [48], arising from the observation of Helmholtz that vortex filaments in an ideal classical fluid (infinitesimal lines about which the fluid circulates [49]) are topologically stable (in the sense that the type of the knot they form will not change, though the vortex line geometry may) [50], even though the vortex is not quantised as in our wave interference. Kelvin’s idea was that each atom might be identified as a specific knot in some pervasive fluid, with their properties related to the type of the knot. Although a beautiful idea, the hypothesis ultimately failed because an inspection of knots revealed that their characteristic properties could not be matched with those of the atoms. The investigation did inspire some of the earliest works of modern knot theory, in particular the first major knot table of Tait which catalogued knots with up to nine self-intersections (in three dimensions, crossings of one strand over the other) in their minimal projections [51].

In modern physics, knots have reappeared in many different contexts. In molecular physics, atoms themselves may have been ruled out as knots but molecular chains permit a mechanical knotting that can be very important to their behaviour. One of the earliest examples determined experimentally was in DNA, in which the interaction of geometrical and topological constraints is a vital part of how a cell can pack many metres of the extended molecule (itself a twisted double helix under mechanical constraints) into its nucleus, while also extracting segments of the chain to be read [52]. More directly, knotted circular DNA has been examined in experiment,

with implications for its behaviour [53]. Similar concepts appear in proteins, which are generally considered as open curves but may be considered as knotted under some reasonable definition [54]. The mechanics of topological constraints also play a role in the behaviour of polymer melts, in which many long polymer strands are tangled around one another with behaviour strongly related to so-called ‘topological’ interactions with other strands [3]. These are generally not true topological constraints in the sense of knotting, though [55] suggest a model where knotting is applicable.


Other examples appear in the very different context of fields, in which the field lines may be knotted. [56] construct solutions of Maxwell’s equations (i.e. electromagnetic fields) which are everywhere knotted, and in which time evolution preserves the topological structure. These methods are ultimately based on the *Hopf fibration*, a map from the 3-sphere to the 2-sphere such that points on the latter are identified with circles on the 3-sphere, every one of which is topologically linked with every other, and which appears in a variety of physical contexts [57]. These are examples of *knotted fields*, where the knot is not in some filamentary object such as our vortex strands in wavefields, but instead the streamlines of the entire field are knotted.

Finally, knotting also occurs in the phase singularities of different systems such as were discussed in Sect. 1.3. [58] constructs a knotted (classical) vortex in water, demonstrating how in a non-ideal fluid the topology is not conserved as the knot rapidly decays to unknotted components via vortex reconnection. Similar behaviour has been studied in the quantised complex phase vortices of a Bose-Einstein condensate, with qualitatively similar results [59], though knots have not been observed experimentally in such a system. Knots have also been constructed in the director singularities of liquid crystals by inducing them via particular boundary conditions (i.e. by placing objects in the crystal); colloids [60], twisted strips [61], or directly knotted microparticles [62]. A modern suggestion reminiscent of Kelvin’s original motivation is that ‘glueballs’, hypothetical gluon-only particles, might be formed of knotted flux tubes [63]. Or to finish, knots and links have been constructed theoretically and experimentally in the complex phase singularities of a light beam [64], and analytically in superpositions of electron eigenfunctions in the hydrogen atom [65].

A common thread through many of these examples is that these are *synthesised* knots, coerced into existence by choices of field construction, experimental techniques, or boundary conditions. We will use topology quite differently, asking not whether we can construct a knot, but whether knots occur naturally in vortices in our generic random model, and how the statistics of their presence vary between different physical and mathematical systems. One may reasonably expect that knots likely exist in the tangle bulk, since vortices may be extremely long with (naively) easily enough total arclength to tangle with themselves. Figure 1.13 shows a sample vortex loop from our simulations of 3-torus eigenfunctions with energy 1875 as described in Sect. 2.5.1, the curve is very clearly long enough that it may be knotted, and other vortices of the field wind in and out of its path in a way that might be expected to form links. We will later show that such behaviour is common.

Understanding the characteristic behaviour of knotting allows different systems to be directly compared on this footing, not geometrically (quantities such as the curvature and density of filaments) but topologically, looking at how the lengthscales

Fig. 1.13 A sample vortex curve from simulations of Gaussian random eigenfunctions of the 3-torus with energy 675. The vortex is a closed loop with total length 7000λ



of knotting relate to the geometrical lengthscales of the curve at the short scale, and how quantities such as the populations of different types of knot compare. Some analysis of this type has been performed for our earlier examples of (closed) polymer melts [55], and for Gaussian random light beams where a Gaussian random initial condition is propagated [37]. The latter finds only links and no knots, and we discuss later how and why this compares to our own results. We will also make significant comparisons to the behaviour of *random walks*, lines that are statistically random in an appropriate sense of scaling and for which much is known about the generic behaviour. We discuss the literature on this subject in Sect. 1.6.

1.6 Other Random Tangles and Random Walks

We have already emphasised the question of how different geometrical and topological statistics may vary amongst random tangles of different origins, and we return here to note some particular examples to which we will later compare the vortex tangle of wave chaos; in the degenerate random waves, the randomness comes entirely from the (approximately) uniform statistical measure of the degenerate eigenspace, but this is does not have to be the case in other physical and mathematical scenarios. This will also lead to considering more of the behaviour of a random walk.

Our very first physical example of vortex tangle was in the vortices of classical turbulence. These are not such a well defined object as the vortices of a complex field because they are not quantised and so do not have the same conservation laws (except in an ideal fluid), so it is harder to directly measure their structure; while some investigations have been made [4], we do not use them as a major system of comparison.

Classical turbulence does instead have a quantum analogue in wavefunctions evolving under the nonlinear Schrödinger equation. The vortices are again quantised phase singularities of the wavefunction [5]. The turbulent state is characterised by a particular energy scaling which requires a particular type of vortex tangle, since

information about the dissipation is contained entirely within the vortices [66], but this tangle is now a dynamic, changing structures with the geometry of individual vortex curves now having a direct effect on the evolution of the field. Much investigation has been made of the statistics and behaviour of this vortex tangle [5, 7, 67], and we compare to these results in later Chapters. However, it is not clear what amongst these statistics should be generic, and questions about large scale statistical topology do not seem to have been investigated.

Returning to linear fields, we have mentioned the existence of vortex tangle in optical volume speckle such as in a laser beam after reflection from a rough surface, modelled directly as a superposition of random waves due to the rough nature of almost any surface on the wavelength scale [31]. This has been quite well studied both analytically and numerically, including aspects of the geometry [36] and topology [37] of the vortex tangle within it. These works in particular analyse the tangle of vortices in the optical wavefield by simulation in a $2 + 1$ paraxial propagation model with an initial Gaussian spectrum in the plane modelled in three dimensions through numerical propagation and vortex detection as explained in Sect. 2.1. This provides an interesting baseline for a type of random wave behaviour, both as a baseline for the behaviour of random waves but also with the details of the model contrasting with our own isotropic, monochromatic power spectrum. In particular, it provides one of very few searches for knotting in random tangle, with which we compare later.

These examples are all from specific physical systems. While they provide a convenient basis for contrast and comparison, it is a separate question to ask what would be the necessary *universal statistics* to characterise vortex tangle. Such a question is approached by [68, 69] who examine a certain class of generic tangle, with free parameters including whether component filaments are oriented and a fugacity parameter controlling how loops may reconnect. These reconnections (or disconnections of one loop into two) must take place with some probability based on the length of the loop. This model maps statistically to a wide range of physical and mathematical systems, not necessarily demanding specific local geometry but making broad claims about some large scale statistics. This is expected to apply to systems including vortex tangle such as ours, and we describe the details of this comparison in Sect. 3.6. Of all these comparisons, this may provide the best basis for the study of universal behaviour that would encompass not just the monochromatic waves we focus on but the broader class of random wave vortex tangle.

Since all of these tangles including those in random waves appear random in the sense of being unpredictable (the path of a single line does not obviously depend on its path at an earlier point, barring periodicity), it is a different (albeit natural) question to ask if vortices behave as *random walks*, the mathematical formalisation of paths consisting of a series of uncorrelated random steps, and for which in the limit of very small steps one achieves perfectly random Brownian trajectory [70]. If our vortex tangle is truly random, we may expect the component vortex filaments to approach this generic random walk behaviour in an appropriate large lengthscale limit. The principal signature of this phenomenon will be in the *fractality* of each of the long vortex lines, since a truly random curve demands a certain kind of scale

invariance in the statistics of the motion. We will investigate if such a scaling holds, and other related quantities, in Chaps. 3 and 5.

There has been much interest in the topology of random walks, particularly arising from their status as models for the behaviour of different kinds of polymers [3, 6]; in such systems, the large scale behaviour of polymers may be strongly affected by their topology, and so it is important to understand how it behaves as a function of their local geometry and physical effects such as an excluded volume about other molecules. In particular, it is well established that closed random walks exhibit an exponential decrease in the probability of unknotting as a function of their length [71, 72], with curves being knotted with probability 1 in the limit of infinite length [73–75]. Although this applies equally to all ensembles of random walks in this limit, their specific behaviour at small lengthscales may not be the same, with knots appearing in different orders and with different probabilities as a function of length [76]. The topology of individual random walks also has an interesting effect on their fractal scaling, with the fractal dimension of individual knot types apparently appearing like that of self-avoiding walks even as the scaling of all knots considered together is that of a normal unbiased walk [77–79]. It is also possible to analyse linking of random curves, but this is not so well investigated in the literature as linking is not defined for a single random walk, and it is not clear how to naturally create multiple curves that are appropriately random with respect to one another.

These facts about the behaviour of random walks are the basis for much of our later analysis of vortices in wave chaos tangle, particularly the discussion of scaling in Sects. 3.5 and 3.6, and much of our analysis of topology in Chap. 5. Although we may expect that all of our random wave systems have the same limiting behaviour at large lengthscales and energies, it may be that their differences are better expressed in different small scale statistics. In particular, it is unclear if topological properties will display universal behaviour across different systems in the same way as geometrical features, and how each of these different properties is modified by the different constraints of random waves.

As part of this investigation, we will sometimes compare our results to those of the literature on random walks. In some cases it is convenient to compare to a specific result, to which purpose we perform our own numerical investigation of a specific random walk in order to directly compare specific quantities. We cannot choose a generic random walk but must select a model for a closed walk, otherwise it is impossible to define topology or to understand how it affects the bulk.

There are many numerical models of random walks that close but exhibit appropriate random walk behaviour on large scales [72, 78, 80–85]. We choose the method of [80], which lets us select a random polygon of N segments in length and of specific total length although the edge lengths may differ, returning a set of randomly oriented edges that may be joined together to construct the closed polygonal curve. Unlike some of the models above, the lengths of the polygon edges are not uniform but instead take a beta distribution [80], though this does not seem to significantly affect statistics. This method involves an unusual construction, selecting a random walk from the space of all N -sided polygons by a map from the $4N$ -sphere to \mathbb{R}^3 , with closed random walks achieved by isolating the areas of the $4N$ -sphere in which




Fig. 1.14 Three example closed random walks. Each is constructed via the method of [80], explained algorithmically in Appendix A.2. **a**, **b** and **c** have lengths of 100, 1000 and 10000 segments respectively, each coloured along their length only to distinguish different regions of the curves

the sum of edge segments happens to close, and selecting the starting position from within these regions. The mathematical details of this procedure are given in Appendix A.2. These details are only necessary to recreate the precise statistics of this random walk model; the quantities we investigate are largely universal to random walks and so invariant to the model chosen.

We have generated a total of 8798500 random polygons via this method, with N between 25 and 2000, normalised such that the average segment length is constant. Figure 1.14 shows examples of the resulting polygons. The *persistence length* is a standard geometric quantifier of the lengthscale on which the tangential direction of a line decorrelates, and for this random walk type is 1.82 times the average step length (the details of this calculation are fully explained in Sect. 3.4), so the longest samples are around 500 steps in these units; this is much longer than any vortex curve we will find in our results.

1.7 Variations on Random Waves

We have introduced qualitatively the idea that Gaussian random superpositions of degenerate eigenfunctions are an appropriate equivalent to the random wave model in a given system, and that we will use eigenfunctions of the 3-torus, 3-sphere and quantum harmonic oscillator in particular. We give here the quantitative details of the eigenfunctions of each, with examples of the resulting vortex tangle and how it differs from the isotropic random wave limit.

1.7.1 The 3-Torus (Periodic Boundary Conditions)

Our first system is the 3-torus, a flat manifold that may be considered as a cubic unit cell with periodic boundary conditions along all three Cartesian axes of the cell.

The eigenfunctions of the 3-torus are solutions of the Helmholtz equation, just as in the generic RWM;

$$\nabla^2 \psi \equiv \left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2} \right) \psi \equiv -k^2 \psi = -E \psi, \quad (1.9)$$

where the solutions are still simply plane waves. Unlike the isotropic RWM, their spectrum is limited to those meeting the periodicity condition $\Psi(\mathbf{r}) = \Psi(\mathbf{r} + \mathbf{\Delta r})$, for any integer components of $\mathbf{\Delta r} = (l \Delta x, m \Delta y, n \Delta z)$. We thus write individual eigenfunctions of the form

$$\psi_{\mathbf{k}}(\mathbf{r}) = \exp\left(\frac{2\pi i}{L} \mathbf{k} \cdot \mathbf{r}\right), \quad (1.10)$$

where $L = |\mathbf{k}|$, and each component of the vector \mathbf{k} is an integer. This guarantees that the value of $\psi_{\mathbf{k}}(\mathbf{r})$ is unchanged under a translation of distance L along any Cartesian axis. It is convenient to rewrite this as

$$\psi_{\mathbf{k}}(\mathbf{r}) = \exp(i\mathbf{k} \cdot \mathbf{R}), \quad (1.11)$$

where $\mathbf{R} = 2\pi \mathbf{r}/L$ and ranges from 0 to 2π . Under this form for the cell size, the 3-torus periodicity condition becomes

$$\psi_{\mathbf{k}}(\mathbf{r}) = \psi\left(\mathbf{r} + \begin{pmatrix} 2\pi a \\ 2\pi b \\ 2\pi c \end{pmatrix}\right) \quad (1.12)$$

$$= \exp\left(i\mathbf{k} \cdot \mathbf{R} + \begin{pmatrix} 2\pi a k_x \\ 2\pi b k_y \\ 2\pi c k_z \end{pmatrix}\right) \quad (1.13)$$


$$= \psi_{\mathbf{k}}(\mathbf{r}); \quad (1.14)$$

the periodicity condition is guaranteed for any integers a, b, c . Thus, Gaussian random eigenfunctions of the 3-torus may be written as

$$\Psi(\mathbf{r}) = \sum_{\mathbf{k}} a_{\mathbf{k}} \psi_{\mathbf{k}}(\mathbf{r}), \quad (1.15)$$

with complex, independent, identically distributed gaussian-random coefficients a_n that randomise both the amplitudes and phases of the composite waves. The sum is over all possible wavevectors whose integer coefficients give the same constant magnitude (i.e. of the same wavelength), with the degeneracy depending on the energy $|\mathbf{k}|^2$; the wavevectors are sampled not isotropically but from a cubic lattice, with their possible values corresponding to the lattice points that intersect a given energy shell. Figure 1.15 shows an example of this in two dimensions, marking the

Fig. 1.15 The intersection of the shell with energy 65 with the grid of integer valued wavevector components. The *blue line* marks the energy shell, and each *purple line* shows a wavevector meeting the condition $|k| = \sqrt{E}$; these wavevectors are permutations (including in sign) of (1, 8) and (4, 7)



grid of integer valued k_x and k_y in grey and with a shell of energy 65 superimposed in blue. This intersects the integer lattice only at sign and order permutations of the (1, 8) and (4, 7) lattice points, marked in purple, and so only plane waves with these wavevectors would contribute to degenerate eigenfunctions of the 2-torus at this energy. It is clear that the selection cannot be truly isotropic, as the lattice intersects the energy shell only in a small number of places; this problem is exaggerated in the 2-torus compared to the 3-torus since a three-dimensional energy shell naturally exposes more options, but is still a core difference between eigenfunctions of the 3-torus and the isotropic RWM.


Figure 1.16 shows how the number of possible wavevectors (i.e. number of intersections of the energy shell with the integer lattice in three dimensions) behaves as a function of the energy. The degeneracy at a given energy (not counting plane waves equivalent under reversal of the wavevector as this is already accounted for by the fact that IID Gaussian coefficients may be negative) is shown in blue. The maximum possible wavevectors grows approximately as the square root of the energy, but varies significantly; even some relatively high energies have few or no available wavevectors, where it happens that the energy shell is of a value that does not pass through many points of the integer lattice.

This wavefunction form is clearly directly related to the random wave model of (1.6), but with the sum over wavevectors rendered finite by the periodicity condition of the 3-torus. The isotropic random wave model would naturally be recovered in the limit of high energy; the larger the size of the torus in wavelengths, the more wavevectors are valid in the sum.

Within these degenerate eigenfunctions, Ψ is periodic up to a choice of sign through any translation along half the diagonal of a unit cell, as in (1.12) with $a = b = c$;

$$\Psi_k \left(\mathbf{r} + a\pi \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \right) = \pm \Psi_k(\mathbf{r}), \quad (1.16)$$

Fig. 1.16 Wavevector degeneracy in eigenfunctions of the 3-torus at different energies. *Blue* points show the full degeneracy, the number of integer valued (up to sign inversion) k_x, k_y, k_z such that $|\mathbf{k}|^2 = E$, while *red* points show the degeneracy only of possible wavevectors whose entries k_x, k_y, k_z are all odd or all even



where the factor of ± 1 on the right hand side depends on whether a, b, c are all odd or all even. We are largely interested in the zero lines of the superposition, and this relation means that the vortex pattern recurs with twice the periodicity of the full eigenfunction; a translation by π in each axis. This means that the unit cell of the zero line tangle is not the periodic cube of the full 3-torus, but a *truncated octahedron*.

Unfortunately, this form is highly inconvenient for numerical simulation, as the periodicity of the truncated octahedron is more complex than that of the cube. Ψ would be periodic via translation through any of its 12 sides, which is inconvenient to numerically stitch together. This problem can be avoided by selecting the subset of wavevectors where k_x, k_y, k_z are all even or all odd. Under translation by π , these satisfy

$$\Psi \left(\mathbf{r} + \begin{pmatrix} \pi a \\ \pi b \\ \pi c \end{pmatrix} \right) = \exp \left(\mathbf{k} \cdot \mathbf{R} + \begin{pmatrix} \pi a k_x \\ \pi b k_y \\ \pi c k_z \end{pmatrix} \right), \quad (1.17)$$

$$= \pm \Psi(\mathbf{r}), \quad (1.18)$$

for any integers a, b, c (i.e. $k_i - k_j = 0 \pmod{2}$ for any $i, j = x, y, z$). Since the components of \mathbf{k} are all odd or all even, they all transform identically under a translation of π along any Cartesian axis; the sign of the eigenfunction may change, but the vortex pattern is now eight-fold periodic, recurring identically in unit cells that are the regular cubic octants of the full 3-torus. This imposes a significant restriction on the available wavevectors, and Fig. 1.16 shows the number satisfying this extra condition at each energy in red. There are naturally always fewer than would be allowed in the unbiased torus with truncated octahedron vortex periodicity, but there still exist a large number at reasonable energies, similar to or greater than in the examples of the previous Section which seemed to have statistically stabilised. We confirm in Chap. 3 that they are appropriately similar to the isotropic RWM, via comparison of their statistics with known analytic results on the isotropic limit.

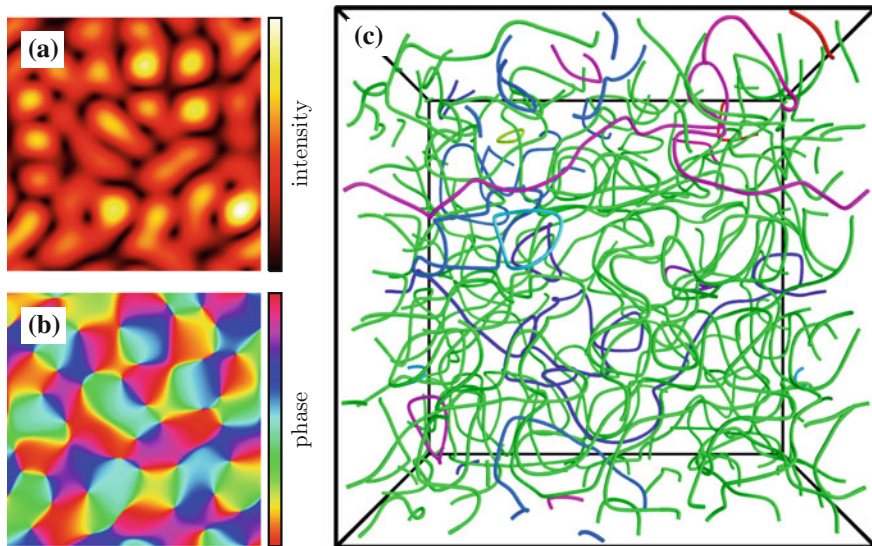


Fig. 1.17 A random degenerate energy 75 eigenfunction of the 3-torus. **a** and **b** show the intensity and phase respectively in a two-dimensional slice through the 3-torus, demonstrating a typical vortex pattern, while **c** shows vortex lines in three dimensions; each vortex is given a distinct colour, and much of the total vortex arclength is in the single *green* example. The side length of the unit cell is 3.44λ . The vortex pattern is recovered using the RRCG algorithm of Sect. 2.2

It is these individual smaller octants of vortex periodicity that we will implicitly consider throughout this work whenever we refer to the 3-torus. In Sect. 2 we only choose energies where the wavevectors have only odd components as this maximises the number of possible wavevectors in a given energy range, and all discussion of torus lengthscales refers to the sidelength and volume of the octant of vortex periodicity, one half or one eighth respectively of the full 3-torus.

Figure 1.17c shows an example energy 75 eigenfunction of the 3-torus. (a) and (b) demonstrate the intensity and phase patterns in a two-dimensional slice through one Cartesian axis, while (c) shows vortices with each different vortex loop coloured with a different random hue. The tangle appears similar to that of the isotropic example in Fig. 1.2, but now almost all of the different segments terminating numerically on the boundary are identified as a single long vortex loop pictured in green.

A result of the 3-torus periodicity is that, although all vortex loops must eventually close, they may wrap around the periodic boundaries of the torus a non-zero total number of times before doing so. This imposes no extra constraints on their local geometry, and they are still closed loops taking into account the connectivity of the manifold. Figure 1.18a, b show two examples from an eigenfunction of energy 243, the first a relatively short line of length approximately 300λ , and the second longer with length around 800λ ; there is nothing locally peculiar about their conformations. We show here a single *periodic segment* of the line; considering the 3-torus as




Fig. 1.18 Vortex curves with non-trivial homology on the 3-torus. **a** and **b** show vortex segments of length 330λ and 800λ respectively, which have been unwrapped from the single numerically sampled torus cell (with periodic boundaries) and joined together as if space were tiled with adjacent cells through which the lines are continuous

representing a unit cell with periodic boundary conditions, we may imagine all of space being tiled by identical unit cells, in which case the vortex curves are continuous lines that, upon intersection with the boundary, pass through to the adjacent cell. Eventually they return to their geometric starting position within a given cell. In this context, loops that wrap around any of the boundaries of the torus a non-zero total number of times are those that reach their starting point relative to a *different* cell, and are both preceded and succeeded by an infinite number of geometrically identical repeats of their length; they represent infinite lines rather than closed loops. The minimal periodic segments are the shortest regions that do not repeat, and are what we refer to in any measure of arclength of these lines. All regions of a single periodic segment appear exactly once within a single unit cell.

Considered this way, we note that these lines have *non-trivial homology*; considering the manifold of the 3-torus in four dimensions, passage through the periodic boundary circles a hole in the manifold. We refer to these as NTH lines. Each may be indexed by the vector of the number of times it passes through the periodic boundary in each Cartesian direction, which in the periodic tiling model would correspond to the total number of cells traversed by a minimal segment before repeating. The distribution of this homology vector is discussed further in Sects. 3.6 and 5.1. We have already noted that the homology does not limit the local geometry, but it does have topological implications, considered in Sect. 5.6.

1.7.2 The 3-Sphere

Our second manifold is the 3-sphere (or *hypersphere*), the curved manifold satisfying the equation

$$x^2 + y^2 + z^2 + t^2 = r^2, \quad (1.19)$$

for spatial dimensions x, y, z, t , and radius r that we will always set to be 1. Although this equation makes use of four spatial dimensions x, y, z, t , only three are independent variables necessary to parameterise its (three-dimensional) volume, just as we often draw the two-dimensional surface of the 2-sphere as the surface of a three-dimensional ball.

There are many coordinate systems that parameterise the volume of the 3-sphere. We will find it most convenient to use the *hyperspherical angles* ψ, θ, ϕ satisfying the following equations in accordance with (1.19),

$$t = r \cos \psi , \quad (1.20)$$

$$z = r \sin \psi \cos \theta , \quad (1.21)$$

$$x = r \sin \psi \sin \theta \cos \phi , \quad (1.22)$$

$$y = r \sin \psi \sin \theta \sin \phi , \quad (1.23)$$

where ψ and θ range from 0 to π , and ϕ from 0 to 2π . This coordinate system is an obvious extension of standard angular coordinates on the 2-sphere. Unlike the 3-torus or harmonic oscillator, the 3-sphere has a non-trivial *metric*, in this case a constant positive local curvature at every point, which must be taken into account in any measures of local geometry. Its general effect is similar to that on lines on the 2-sphere, in that lines that are locally straight follow geodesics that eventually return to their starting point. The total volume of the 3-sphere is finite but without any corporeal boundary, and every vortex curve within an eigenfunction of the 3-sphere has trivial homology. The relevant geometrical details are explained in Sect. 2.4.2, and we deal here only with the random harmonics of the 3-sphere and the basic details of its visualisation.

Eigenfunctions of the 3-sphere are solutions of the equation,

$$\Delta \Psi = -k^2 \Psi = -E \Psi , \quad (1.24)$$

where the Laplace-Beltrami operator Δ extends the Laplacian ∇^2 to the curved manifold. In hyperspherical coordinates this takes the form

$$\Delta = \frac{1}{\sin^2 \psi} \frac{\partial}{\partial \psi} \left(\sin^2 \psi \frac{\partial}{\partial \psi} \right) + \frac{1}{\sin^2 \psi} \left(\frac{1}{\sin \theta} \frac{\partial}{\partial \theta} \left(\sin \theta \frac{\partial}{\partial \theta} \right) + \frac{1}{\sin^2 \theta} \frac{\partial^2}{\partial \phi^2} \right) . \quad (1.25)$$

The final term is simply the usual Laplacian on the 2-sphere, multiplied by $1/\sin^2 \psi$.

The eigenfunctions of (1.24) are known as the *hyperspherical harmonics* which we label $\mathcal{Y}_{Nlm}(\psi, \theta, \phi)$. In angular form, they are given by

$$\mathcal{Y}_{Nlm}(\psi, \theta, \phi) = \sqrt{\frac{2^{2l+1}(N-l)!(1+N)}{\pi(1+l+N)!}} l! \sin^l \psi C_{N-l}^{l+1}(\cos \psi) Y_l^m(\theta, \phi) , \quad (1.26)$$

in which Y_l^m are the spherical harmonics of the 2-sphere, and C_{N-l}^{l+1} are the *Gegenbauer polynomials* [13], also known as ultraspherical polynomials for their role in defining spherical harmonics in higher dimensions. These harmonics have Laplace-Beltrami eigenvalue $N(N+2)$ and permit any $0 \leq l \leq N$ and $-l \leq m \leq l$ for a $(1+N)^2$ -fold degeneracy in l and m . Thus, Gaussian random waves on the 3-sphere are of the form

$$\Psi_N(\psi, \theta, \phi) = \sum_{l=0}^N \sum_{m=-l}^l a_{l,m} \mathcal{Y}_{Nlm}(\psi, \theta, \phi) \quad (1.27)$$

for complex IID Gaussian random variables $a_{l,m}$.

Visualising these functions requires representing the three-dimensional surface of the 3-sphere in \mathbb{R}^4 via projection to the three dimensions of \mathbb{R}^3 . There are many possible projections preserving different aspects of the geometry of the sphere (i.e. lengths, areas, angles, volumes), but it is not possible to preserve all of these relations across all of space; the same is true when representing the surface of a normal sphere on a plane, hence the many choices of projections used in cartographic maps.

When representing the 3-sphere, we make use mainly of *stereographic projection*, a continuous, conformal (angle-preserving) map that takes all but one point on the 3-sphere to a unique point in \mathbb{R}^3 . This projection does not preserve lengths or areas, some regions on the sphere become arbitrarily highly distorted, but all regions of the sphere can be represented save for a single point. The mathematics of this projection in terms of the hyperspherical angles are that

$$r = \tan \frac{\psi}{2}, \quad (1.28)$$

$$z = r \cos \theta, \quad (1.29)$$

$$x = r \sin \theta \cos \phi, \quad (1.30)$$

$$y = r \sin \theta \sin \phi. \quad (1.31)$$

Figure 1.19 shows the random waves of (1.27); (a) and (b) reiterate for reference the nature of a random degenerate spherical harmonic of energy 20 on the 2-sphere, whose statistics will be similar to those of a spherical cross-section of the 3-sphere though with a scaling factor. (c) shows vortices of a 3-sphere eigenstate with $N = 15$ (energy 255) in stereographic projection; near the centre of the projection the random tangle is geometrically similar to that of Fig. 1.17, while vortices that pass close to $\psi = \pi$ are highly distorted by the projection.

Although it is not necessarily clear in Fig. 1.19c, the vortex pattern includes an inversion symmetry of the eigenfunctions under the rotation $\psi \rightarrow \pi - \psi$, $\theta \rightarrow \pi - \theta$ and $\phi \rightarrow \phi + \pi$. Under such a transformation in (1.27), $Y_l^m(\pi - \theta, \phi + \pi) = (-1)^l Y_l^m$ and $C_{N-l}^{l+1}(\cos(\psi + \pi)) = (-1)^{N-l} C_{N-l}^{l+1}(\cos \psi)$, so $\mathcal{Y}_{Nlm}(\pi - \psi, \pi - \theta, \phi + \pi) = (-1)^N \mathcal{Y}_{Nlm}(\psi, \theta, \phi)$. The nodal line structure is invariant to this parity and so is simply symmetric under this rotation; any vortex structure appears twice in any degenerate eigenfunction. This has implications for certain statistics, discussed further in Sect. 5.1. This symmetry is the hyperspherical equivalent of the truncated

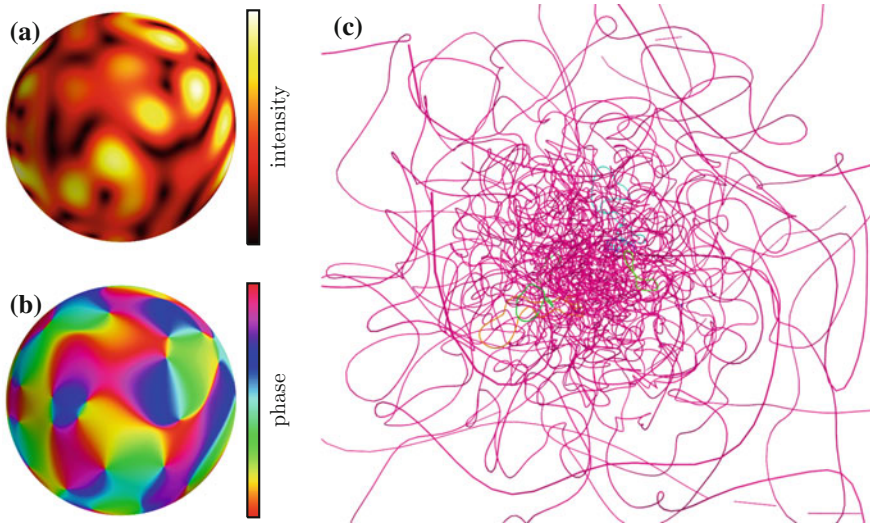


Fig. 1.19 Random degenerate eigenfunctions of the 2-sphere and 3-sphere. **a** and **b** depict the intensity and complex phase respectively in a random superposition of degenerate energy 20 spherical harmonics of the 2-sphere, a pattern that will match (up to a statistical scaling) the pattern on a great-sphere cross section of the 3-sphere. **c** shows vortices in an energy 120 Gaussian random eigenfunction of the 3-sphere, as in (1.27), via stereographic projection; vortices are cut off where they pass close to $\psi = \pi$, and appear much denser in the centre where ψ is near to 0. Each vortex is coloured differently but, although there are 11 different vortex lines in the eigenfunction, the single pink line consumes the majority of arclength and so is the main visible component

octahedron vortex periodicity in generic eigenfunctions of the 3-torus, but in this case poses no numerical problem since it is convenient to locate vortices in the entire volume of the 3-sphere and no improvement would be made by attempting to limit vortex periodicity further.

Beyond projection alone, working with the geometry of the 3-sphere requires corrections for the local metric in order to calculate distances and angles appropriately. The details of the relevant calculations are explored in Sect. 2.4. We also utilise a different type of projection in our numerics, explained in Sect. 2.3, and draw our results from a range of different energies and degeneracies enumerated in Sect. 2.5.

1.7.3 The Quantum Harmonic Oscillator

Our third system of random waves is the *quantum harmonic oscillator*, the set of eigenfunctions of the equation

$$-\nabla^2\Psi + \frac{1}{2}(x^2 + y^2 + z^2)\Psi = E\Psi . \quad (1.32)$$

This may be considered as the standard Schrödinger equation with a harmonic potential, but with units scaled out such that $\hbar = m = \omega = 1$. This system is *superintegrable*; it classically possess more constants of its motion than it has degrees of freedom [86], which in quantum terms tends to mean that different closed-form solutions may arise by separation of variables in different coordinate systems. For instance, we later make use of the eigenfunctions obtained by separation in spherical polar coordinates, but through most of this work use those resulting from separating the equation in the Cartesian directions

$$\Psi(x, y, z) = X(x)Y(y)Z(z), \quad (1.33)$$

in which case the solutions in each of X , Y and Z are given by

$$\Phi_n(x) = \frac{1}{\sqrt{2^n n!}} \left(\frac{1}{\pi}\right)^{1/4} \exp\left(-\frac{x^2}{2}\right) H_n(x), \quad (1.34)$$

where H_n are the *Hermite polynomials* [13]. This eigenfunction has energy $E_n = n + 1/2$, and so the degenerate solutions of (1.32) are given by superpositions $\Phi_p \Phi_q \Phi_r$ with constant $N = p + q + r$;

$$\Psi_N = \sum_{p=0}^N \sum_{q=0}^{N-p} \sum_{r=0}^{N-p-q} a_{p,q,r} \Phi_p(x) \Phi_q(y) \Phi_r(z), \quad (1.35)$$

with $a_{p,q,r}$ being complex IID random variables with Gaussian distribution. The energy of the degenerate eigenfunction under these scalings is $E = p + q + r + 3/2 \equiv N + 3/2$, and the total degeneracy $(N + 1)(N + 2)/2$.

Figure 1.20 shows examples of degenerate eigenfunctions of the QHO in both two and three dimensions; in 2D the degeneracy is a sum over constant $p + q$ with $r = 0$. In both cases, the degenerate eigenfunction exhibits symmetry as $x \rightarrow -x$, $y \rightarrow -y$ and (in three dimensions) $z \rightarrow -z$; this is a result of the symmetry of the Hermite polynomials $H_n(x) = (-1)^n H_n(-x)$ [13]. For a given energy index N , the component eigenfunctions must have the same total parity, giving a total factor $(-1)^N$, just as in both previously discussed systems. For the nodal lines that we are largely interested in, this means that every pattern must occur twice; the implications of this symmetry for our own investigation are explored later where they become relevant, particularly in Sect. 5.1.

The interaction of the eigenfunction with the radius where the potential exceeds the energy, $E = \frac{1}{2}(x^2 + y^2 + z^2)$ (we call this the *classical radius*), is also striking. Almost all of the vortex tangle is confined within the classical region where the energy exceeds the potential, with only short segments entering the classically forbidden region outside these bounds, of which those not heading directly away from the bulk do nothing more complex than reversing direction and re-entering the classical region. Likewise, in two dimensions the point vortices are almost (but not quite) confined to this region. However, a small number of 3D vortex lines do leave the classical




Fig. 1.20 Random degenerate eigenfunctions of the QHO in two and three dimensions. **a** Shows the real component of a 2D eigenfunction with $N = 20$ (i.e. $E = 21$), with the *blue circle* marking the radius at which $E = 1/2(x^2 + y^2)$. **b** Shows vortices in the equivalent in three dimensions with $E = 25 + 3/2$, with the drawn cell marking a cube of side length 1.13 times the classical cutoff radius. Each vortex is given a different colour

region and exit the simulated volume in straight lines; this behaviour is expected of nodal lines, and would also be visible in the component real and imaginary parts of the two-dimensional function, in which case, the number of such lines grows as $N^{1/2}$ [87]. Many of our later measurements will depend on investigating closed loops, but it is unclear whether these lines heading to infinity can be considered as such in a meaningful sense. As with the symmetry of the eigenfunction, we investigate the implications of this in Sect. 5.1.

Although the QHO does not have a curved metric like the 3-sphere, the radius-dependent potential exerts its own effect on the local geometry. We discuss the geometrical implications of this effect for length measurement in Sect. 2.4. The numerical details of the eigenfunction datasets we choose to investigate are given in Sect. 2.5.3.

1.8 Thesis Outline

The material of this thesis is arranged as follows:

In Chap. 2 we explain the numerical and practical details of tracking singularities of the complex phase in three-dimensional complex fields, covering both the methods that have been used in the literature and our own extensions to allow efficient but geometrically and topologically accurate vortex recovery even from large-scale

vortex tangle. We also cover here the analytical and numerical details of geometrical calculation taking account of the local conditions on each of the 3-torus, 3-sphere and QHO. This Chapter concludes with a summary of the data we have collected and analysed with these methods, from which we draw almost all of our later analysis.

In Chap. 3 we investigate quantities of the local geometry, large scale fractality and global scaling in vortex tangle, considering vortices both as individual curves that are expected to behave as random walks and as members of the tangle ensemble that imposes extra constraints on their expected behaviour. We compare to previous analytic results where possible, as well as expectations from other systems. These results include analysis of local geometry, fractality and scaling that have previously been published in [88].

Chapter 4 contains the further analytical and numerical methods of topological analysis, beginning with a broad overview of all the relevant theory of knotting and linking. We continue by explaining the specific details of the topological calculations that we make use of in our own analysis. This includes the numerical details and limitations that this entails and, since vortex curves are often long and prohibitively geometrically complex, methods to simplify them before topological analysis is performed.

In the final Chap. 5, we discuss the results of applying topological measures to our curves. As with their geometry this includes comparison with analytic results and other systems where possible, though there are far fewer known results on vortex tangle topology. We take particular interest in the differences between our systems of wave chaos, in which knotting and linking may be expected to be strongly affected by their different constraints, but whose differences are not directly predictable.

References

1. U Frisch, *Turbulence: The Legacy of A N Kolmogorov* (Cambridge University Press, 1995)
2. J.J. Monaghan, J.B. Kajtar, Leonardo da Vinci's turbulent tank in two dimensions. *Eur. J. Mech. B-Fluid*. **44**, 1–9 (2014)
3. T.C.B. McLeish, Tube theory of entangled polymer dynamics. *Adv. Phys.* **51**(6), 1379–527 (2002)
4. M. Farge, K. Schneider, G. Pellegrino, A.A. Wray, R.S. Rogallo, Coherent vortex extraction in three-dimensional homogeneous turbulence: comparison between CVS-wavelet and POD-Fourier decompositions. *Phys. Fluids* **15**(10), 2886–2896 (2003)
5. S.K. Nemirovskii, Quantum turbulence: theoretical and numerical problems. *Phys. Rep.* **524**, 85–202 (2013)
6. E. Orlandini, S.G. Whittington, Statistical topology of closed curves: some applications in polymer physics. *Rev. Mod. Phys.* **79**, 611–642 (2007)
7. C.F. Barenghi, R.L. Ricca, D.C. Samuels, How tangled is a tangle? *Physica D*, **157**(3), 197–206 (2001)
8. M.H. Freedman, Z.-X. He, Divergence-free fields: Energy and asymptotic crossing number. *Ann. Math.* **134**, 189 (1991)
9. H.K. Moffatt, R.L. Ricca, Helicity and the Călugăreanu invariant. *Proc. R. Soc. A* **439**, 411–429 (1992)
10. R.L. Ricca, B. Nipoti, Gauss' linking number revisited. *J. Knot. Theor. Ramif.* **20**, 1325–1343 (2011)

11. M. Epple, Geometric aspects in the development of knot theory, in *History of Topology*, ed. by I.M. James (Elsevier Science B V, 1999), pp. 301–357
12. M.V. Berry, Making waves in physics: three wave singularities from the miraculous 1830s. *Nature* **403**, 21 (2000)
13. M. Abramowitz, A.I. Stegun, eds., *Handbook of Mathematical Functions with Formulas, Graphs and Mathematical Tables*, Chap. 22 (Dover, 1965), pp. 773
14. G.P. Alexander, B.G. Chen, E.A. Matsumoto, R.D. Kamien, Colloquium: Disclination loops, point defects, and all that in nematic liquid crystals. *Rev. Mod. Phys.* **84**, 497 (2012)
15. M.R. Dennis, K. O'Holleran, M.J. Padgett, Singular optics: optical vortices and polarization singularities. *Prog. Optics* **53**, 293–363 (2009)
16. T. Young, Experiments and calculations relative to physical optics. *Philos. T R. Soc. Lond.* **94**, 1–16 (1804)
17. M. Born, E. Wolf, *Principles of Optics* (Pergamon Press, 1980)
18. F. Ricci, W. Löffler, M.P. van Exter, Instability of higher-order optical vortices analysed with a multi-pinhole interferometer. *Opt. Express* **20**, 22961–22975 (2012)
19. M.R. Dennis, Topological singularities in wave fields. PhD thesis, University of Bristol, 2001
20. M.J. Padgett, L. Allen, The angular momentum of light: optical spanners and the rotational frequency shift. *Opt. Quant. Electron.* **31**, 1–12 (1999)
21. J.F. Nye, Lines of circular-polarization in electromagnetic-wave fields. *Proc. R. Soc. Lond.* **389**, 279–290 (1983)
22. F. Flossmann, K. O'Holleran, M.R. Dennis, M.J. Padgett, Polarization singularities in 2D and 3D speckle fields. *Phys. Rev. Lett.* **100**, 203902 (2008)
23. M.S. Longuet-Higgins, The statistical analysis of a random, moving surface. *Philos. Trans. R. Soc. A* **249**, 321–387 (1957a)
24. M.S. Longuet-Higgins, Statistical properties of an isotropic random surface. *Philos. Trans. R. Soc. A* **250**, 157–174 (1957b)
25. D.E. Cartwright, M.S. Longuet-Higgins, The statistical distribution of the maxima of a random function. *Proc. R. Soc. A* **237**(1209), 212–232 (1956)
26. G.M. Morse, R.H. Bolt, Sound waves in rooms. *Rev. Mod. Phys.* **16**, 70–147 (1944)
27. M.R. Dennis, Gaussian random wavefields and the ergodic mode hypothesis, in *New Directions in Linear Acoustics and Vibration*, ed. by M. Wright, R. Weaver (Cambridge University Press, 2010), pp. 59–76
28. S.O. Rice, Mathematical analysis of random noise. *BSTJ* **23**, 282–332 (1944)
29. K.J. Worsley, A.C. Evans, S. Marrett, P. Neelin, A three-dimensional statistical analysis for CBF activation studies in human brain. *J. Cereb. Blood Flow* **12**, 900–918 (1992)
30. R.J. Adler, A.M. Hasofer, Level crossings for random fields. *Ann. Probab.* **4**, 1–12 (1976)
31. J.W. Goodman, Some fundamental properties of speckle. *JOSA* **66**(11), 1145–1150 (1976)
32. M.V. Berry, Regular and irregular semiclassical wavefunctions. *J. Phys. A* **10**, 2083–2091 (1977)
33. A.I. Shnirelman, Ergodic properties of eigenfunctions. *Usp Mat Nauk* **29**, 181–182 (1974)
34. M.V. Berry, M.R. Dennis, Phase singularities in isotropic random waves. *Proc. R. Soc. A.* **456**, 2059–2079 (2000)
35. M.R. Dennis, Nodal densities of planar Gaussian random waves. *Eur. Phys. J.: Spec. Topics* **145**, 191–210 (2007) (Conference on Nodal Patterns in Physics and Mathematics, Wittenbrg, Germany, JUL 24–28, 2006)
36. K. O'Holleran, M.R. Dennis, F. Flossmann, M.J. Padgett, Fractality of light's darkness. *Phys. Rev. Lett.* **100**, 053902 (2008)
37. K. O'Holleran, M.R. Dennis, M.J. Padgett, Topology of light's darkness. *Phys. Rev. Lett.* **102**, 143902 (2009)
38. G. Foltin, Signed zeros of Gaussian vector fields—density, correlation functions and curvature. *J. Phys. A.* **36**, 1729 (2003)
39. G. Foltin, S. Gnutzmann, U. Smilansky, The morphology of nodal lines: random waves versus percolation. *J. Phys. A* **37**, 11363–11372 (2004)

40. M.V. Berry, Statistics of nodal lines and points in chaotic quantum billiards: perimeter corrections, fluctuations, curvature. *J. Phys. A* **35**, 3025–3038 (2002)
41. A.G. Monastra, U. Smilansky, S. Gnutzmann, Avoided intersections of nodal lines. *J. Phys. A* **36**, 1845–1853 (2003)
42. R. Höhmann, U. Kuhl, H.-J. Stöckmann, J.D. Urbina, M.R. Dennis, Density and correlation functions of vortex and saddle points in open billiard systems. *Phys. Rev.* **79**, 016203 (2009)
43. M.V. Berry, Disruption of wavefronts: statistics of dislocations in incoherent Gaussian random waves. *J. Phys. A* **11**, 27–37 (1978)
44. M.V. Berry, M.R. Dennis, Topological events on wave dislocation lines: birth and death of loops, and reconnection. *J. Phys. A* **40**, 65–74 (2007)
45. M.V. Berry, H. Ishio, Nodal-line densities of chaotic quantum billiard models satisfying mixed boundary conditions. *J. Phys. A* **38**, L513–L518 (2005)
46. F. Nazarov, M. Sodin, On the number of nodal domains of random spherical harmonics. *Am. J. Math.* **131**(5), 1337–1357 (2009)
47. C.C. Adams, *The Knot Book* (American Mathematical Society, 1994)
48. Lord Kelvin (Sir William Thomson). On vortex atoms. *Proc. R. Soc. Edinb.* **VI**, 94–105 (1867)
49. K. Moffatt, Vortex dynamics: the legacy of Helmholtz and Kelvin, in *Symposium on Hamiltonian Dynamics*, Vortex Structures, Turbulence, volume 6 of IUTAM Bookseries, ed. by A.V. Borisov, V.V. Kozlov, I.S. Mamaev, M.A. Sokolovskiy (Springer, Netherlands, 2008), pp. 1–10
50. P.G. Tait, J. Für die reine und Angewandte Mathematik Translation of H von Helmholtz. **55**(25–55), 1858 (1867)
51. P.G. Tait, On knots I, II, III. In *Scientific Papers*, vol. 1 (Cambridge University Press, 1900)
52. D. Buck, DNA Topology, in *Applications of Knot Theory*, ed. by D. Buck, E. Flapan, vol. 66 of *Proceedings of Symposia in Applied Mathematics* (American Mathematical Society 2008)
53. A. Stasiak, V. Katrich, J. Bednar, B. Michoud, J. Dubochet, Electrophoretic mobility of dna knots. *Nature* **384**, 122 (1996)
54. K.C. Millett, Tying down open knots: a statistical method for identifying open knots with applications to proteins. *Ser Knots*, 203–217 (2005)
55. J. Qin, S.T. Milner, Counting polymer knots to find the entanglement length. *Soft Matter* **7**(22), 10676–10693 (2011)
56. H. Kedia, I. Bialynicki-Birula, D. Peralta-Sals, W.T.M. Irvine, Tying knots in light fields. *Phys. Rev. Lett.* **111**, 150404 (2013)
57. H.K. Urbantke, The hopf fibration—seven times in physics. *J. Geom. Phys.* **46**, 125–150 (2003)
58. D. Kleckner, W.T.M. Irvine, Creation and dynamics of knotted vortices. *Nat. Phys.* **9**, 253–258 (2012)
59. D. Proment, M. Onorato, C.F. Barenghi, Vortex knots in a bose-einstein condensate. *Phys. Rev. E* **85**, 036306 (2012)
60. U. Tkalec, M. Ravnik, S. Čopar, S. Žumer, I. Muševic, Reconfigurable knots and links in chiral nematic colloids. *Science* **333**, 62–65 (2011)
61. T. Machon, G.P. Alexander, Knots and nonorientable surfaces in chiral nematics. *PNAS* **110**, 14174–14179 (2012)
62. W.T.M. Irvine, D. Kleckner, Tangled loops and knots. *Nat. Mater.* **13**, 229–231 (2014)
63. R.V. Buniy, T.W. Kephart, Glueballs and the universal energy spectrum of tight knots and links. *Int. J. Mod. Phys. A* **A20**, 1252–1259 (2005)
64. M.R. Dennis, R.P. King, J. Barry, K. O’Holleran, M.J. Padgett, Isolated optical vortex knots. *Nat. Phys.* **6**, 118–121 (2010)
65. M.V. Berry, Knotted zeros in the quantum states of hydrogen. *Found. Phys.* **31**, 659–667 (2001)
66. T. Araki, M. Tsubota, S.K. Nemirovskii, Energy spectrum of superfluid turbulence with no normal-fluid component. *Phys. Rev. Lett.* **89** (2002)
67. D. Kivotides, C.F. Barenghi, D.C. Samuels, Fractal dimension of superfluid turbulence. *Phys. Rev. Lett.* **87** (2001)
68. A. Nahum, J.T. Chalker, Universal statistics of vortex lines. *Phys. Rev. E* **85**, 031141 (2012)
69. A. Nahum, J.T. Chalker, P. Serna, M. Ortuno, A.M. Somoza, Length distributions in loop soups. *Phys. Rev. Lett.* **111**, 100601 (2013)

70. K. Falconer, *Fractal Geometry: Mathematical Foundations and Applications*, Chap. 3 (Wiley, 1997)
71. M.L. Mansfield, Knots in hamilton cycles. *Macromolecules* **27**, 5924–6 (1994)
72. K. Koniaris, M. Muthukumar, Knottedness in ring polymers. *Phys. Rev. Lett.* **66**(17), 2211–2214 (1991)
73. D.W. Sumners, S.G. Whittington, Knots in self-avoiding walks. *J. Phys. A* **21**, 1689–1694 (1988)
74. N. Pippenger, Knots in random walks. *Disc. Appl. Math.* **392**, 273 (1989)
75. Y. Diao, The knotting of equilateral polygons in \mathbb{R}^3 . *J. Knot. Theor. Ramif.* **04**(02) (1995)
76. T. Deguchi, K. Tsurusaki, Numerical application of knot invariants and universality of random knotting. *Banach Center Publ.* **42**, 77–85 (1998)
77. A. Dobay, J. Dubochet, K. Millett, P.E. Sottas, A. Stasiak, Scaling behavior of random knots. *PNAS* **100**, 5611–5615 (2003)
78. N.T. Moore, R.C. Lua, A.Y. Grosberg, Topologically driven swelling of a polymer loop. *PNAS* **101**, 13431–35 (2004)
79. J. des Cloizeaux, Ring polymers in solution: topological effects. *J. Phys. Lett.*, **42**, L433 (1981)
80. J. Cantarella, T. Deguchi, C. Shonkwiler, Probability theory of random polygons from the quaternionic viewpoint. *Commun. Pure Appl. Anal.* **67**, 1658–1699 (2014)
81. J.P.J. Michels, F.W. Wiegel, On the topology of a polymer ring. *Proc. R. Soc. A* **403**, 269–284 (1986)
82. T. Deguchi, K. Tsurusaki, Topology of closed random polygons. *J. Phys. Soc. Jpn.* **62**, 1411–1414 (1993)
83. J. des Cloizeaux, M.L. Mehta, Topological constraints on polymer rings and critical indices. *J. Phys.* **40**, 665–670 (1979)
84. R. Lua, A. Borovinskiy, A.Y. Grosberg, Fractal and statistical properties of large compact polymers: a computational study. *Polymer* **45**, 717 (2004)
85. S.M. Rappaport, Y. Rabin, A.Y. Grosberg, Worm-like polymer loops and Fourier knots. *J. Phys. A* **39**, L507–L513 (2006)
86. M.A. Rodríguez, P. Tempesta, P. Winternitz, Reduction of superintegrable systems: the anisotropic harmonic oscillator. *Phys. Rev. E* **78**, 046608 (2008)
87. W.E. Bies, E.J. Heller, Nodal structure of chaotic eigenfunctions. *J. Phys. A* **35**, 5673–5685 (2002)
88. A.J. Taylor, M.R. Dennis, Geometry and scaling of tangled vortex lines in three-dimensional random wave fields. *J. Phys. A* **47**(46), 465101 (2014)

Chapter 2

Numerical Methods

In our numerical analysis of quantised vortex tangle in random waves, we must be able to locate vortices numerically in three-dimensional complex fields, recovering both their local geometry well enough for comparison with known geometrical results and their global topology such that no information about how they wind about each other is missed. In this chapter we describe and explain the numerical techniques that we use to do so.

We begin in Sect. 2.1 by discussing the various techniques used in the literature to locate vortices in both two and three dimensions via numerical discretisation of the field in which they are present. Section 2.2 describes our own algorithm, extending these basic techniques with our own additions to create an algorithm that can efficiently trace vortex lines on large lengthscales by recursively focusing numerical effort on their local surroundings. In Sect. 2.3 we describe the necessary adaptations to this algorithm in order to best capture the nuances of the three different systems in which we will investigate random waves; the 3-torus, 3-sphere and quantum harmonic oscillator. Section 2.4 follows with the analytical and numerical details of recovering geometrical statistics of these vortices, accounting for the local details of each system, and in Sect. 2.5 we list all the details of the numerical datasets from which most of our later results are drawn.

2.1 Numerically Locating Vortices

A central problem in this thesis is the numerical detection and tracking of quantised vortices in random waves as they wind through three-dimensional space, accurately reconstructing their local space curve geometry, correctly resolving the way nearby vortex segments wind around one another, and doing so efficiently enough to sample large volumes of random vortex tangle. Similar problems in resolving properties of vortex filaments occur in many different areas of physics or mathematics, both related

to our random waves and in entirely different contexts. We describe and compare here existing methods to do so in two and three dimensions, including an analysis of which techniques are most suited to locating vortices in our own complex scalar fields. This establishes the basic principles for our own tracking algorithm which extends some of these techniques, explained in Sect. 2.2.

The most directly similar methods in the literature include tracing nodal lines in random wave models [1, 2] or other numerically simulated fields [3], tracked vortices in optical beams [4–7], and certain simple discrete models of cosmic strings [8], as well as the detection of vortices in fields that are still complex scalar but with more complicated driving physics, such as vortices in Bose-Einstein condensates through simulations of the Gross-Pitaevskii equation [9, 10]. The same ideas appear more widely in understanding the unquantised vortices of classical turbulence [11] and in locating the more varied types of director field singularity in liquid crystals [12] or polarisation fields [13]. Detecting phase singularities is even important either physically or algorithmically to a range of physical image analysis problems, including understanding voltage patterns in the heart [14–16], interpreting MRI data [17, 18], or in reliably ‘phase-unwrapping’ noisy data that has been measured or stored modulo some constant [19]. Not all of these are directly relevant to our own work, but most make use of the same core ideas, which we will explore further below.

The simplest relevant problem is that of detecting vortices in a two-dimensional complex scalar field, where the vortices are point defects in the phase; if the field is a planar slice of a 3D system, these points are where one-dimensional vortex curves pass through the plane. Figure 2.1 shows an example area of a monochromatic Gaussian random field of the type discussed in Sect. 1.4; the side length of each image is 3.8λ , covered with a numerical grid of side length 350, easily enough to show all




Fig. 2.1 A region of an example two-dimensional Gaussian random wavefunction. The side length of the given square is 3.8λ , and **a**, **b** show the intensity and phase respectively. The numerical problem discussed in this section is how to locate the vortices in the phase efficiently and reliably

features of the field in high detail. The task of this Section is to locate its vortices with far fewer numerical sample points.

As discussed in Chap. 1, vortex lines are characterised both by their being nodal lines such that the value of the field is zero along their length, and by the complex phase winding around them some integer multiple of 2π times (in practice always once). It is possible in principle to use either of these properties to locate a vortex numerically, but in practice the changing phase is far more useful as in comparison to the intensity it can be more easily detected; the intensity may be near-zero across a wide area and can have a complex local structure, making gradient descent unreliable, but the phase change can be detected even far from the vortex as the total change remains the same about any path that encircles the vortex core, as long as no other vortices are involved. The detection of vortices is thus reduced to a phase unwinding problem, in which a closed path is sampled through the field via a set of sampling points, and the phase interpolated continuously between these samples. A vortex is present if the sum of phase differences between each consecutive pair of sampled lattice points is $2\pi m$, with any m in principle being resolvable with enough samples.

Figure 2.2 shows an example of this process using a circular path around a vortex in a Gaussian random field, with the path shown in (a) for the intensity and (b) for the phase of the field, and in (c) the integrated path beginning from the marked point and proceeding clockwise; we discuss the blue plot here, which uses 1000 sample points about the given contour, easily enough to observe variation at small distances. The phase along the path has a fairly complicated structure, changing sometimes extremely rapidly while in other places being relatively flat, but once it has been ‘unwrapped’ to remove the jump accounting for measurement modulo 2π the full phase change is equal to 2π and clearly detects the vortex. Although this example uses a circular sampling path about the vortex, any other shape would reveal the same result as long as only one vortex is present in the region it encloses (or multiple vortices with the same total phase sum). The vortex could in principle also be located by searching the intensity of the wavefunction for the nodal point at the vortex core, but this clearly cannot be so easily detected far from its real location—if we imagine for instance a gradient descent algorithm from some point on the contour, we might easily find only a local minimum or a different vortex line entirely.

Although robust, this initial method requires that very many sample points be selected in a path around the vortex, which is not practical for efficiently locating vortices whose position is otherwise unknown, such as in our random fields; it is hard to pick an appropriate integration contour without already knowing the position of a vortex. An ideal algorithm should instead locate vortices using the fewest possible sample points in the field. It is natural to base the procedure on a Cartesian grid, with the closed paths that may surround a vortex being the border of each square grid plaquette, and the sampled field values on this path being only the four grid points of the plaquette corners. The phase unwinding is reduced to just two rules; the phase change between each pair of adjacent sampling points is taken to be in the direction that minimises its absolute change modulo 2π , and a vortex is present if all four phase jumps have the same sign. This condition is not the most general, a vortex could be present if there are three large phase jumps in one direction even if the final




Fig. 2.2 Phase unwinding around a vortex. **a, b** Show the intensity and phase respectively of a $(1.27 \lambda)^2$ area of a complex scalar Gaussian random wave with 12 component wavevectors, with a marked sampling contour proceeding clockwise from the marked point. **c** Shows the phase around this contour, sampled with 1000 points (*blue*) or just four points (*red*). The integrated change along both curves is 2π (though in local regions it may both increase or decrease) after unwrapping in both cases through the single large jump with magnitude greater than π that indicates a traversal of the phase periodicity from 2π to 0. The contour thus encircles a vortex. I_{\max} is the maximum intensity within the visible region, not the entire field

jump is backwards, but this only occurs for an extremely anisotropic vortex and does not occur often. We discuss in the next Section how our own algorithm guarantees vortex continuity, which removes this omission as a concern.

Figure 2.2c shows in red the result of sampling the integration contour in (a) and (b) with only four (equidistant) points, linearly interpolating the phase between them. Although this discards almost all information about the local phase structure of the vortex, and despite the sample points being as far as 0.35λ from the vortex core, applying the rules of the previous paragraph to these four points correctly detects that a vortex exists. It is no longer so visually clear where the phase jumps through 2π to 0, but the change is still detected following the rule that the total phase change (modulo 2π) is assumed to be less than π , and so the vortex is properly detected. This is only an example of how a very sparse sampling can detect vortices, we note that in plausible situations the sparse sampling would fail even though the dense sampling of 1000 points does not, and we discuss such situations further shortly.

In some models of vortex tangle, there is no more to the simulated field than this initial discretisation. For instance, in the \mathbb{Z}_3 model each vertex of the lattice is identified only by one of three values, which in this case we can consider as equidistant phases modulo 2π (for instance, $0, 2\pi/3$ and $4\pi/3$) [20], a notion related to the Potts model of statistical physics [21]. Assigning each vertex a random choice of these values creates a discrete phase field, around each plaquette of which we could search for vortices using the above algorithm based on just its four corners. In two dimensions this returns an unambiguous result (each plaquette simply does or




Fig. 2.3 Phase and vortex points of a two-dimensional random wave sampled at three different resolutions. **a–c** Show the same random wave sampled at Cartesian lattice spacings of 0.01λ , 0.1λ and 0.15λ respectively, with the total grid side length 2.1λ . Vortices are identified at each point via the phase change around the four vertices of a Cartesian sampling grid, with clockwise vortices coloured *blue* and anticlockwise vortices coloured *red*

does not contain a vortex), but we will return to this model later as an example of how further information about the vortices cannot be found in such a discrete model.

Returning to the question of how low resolution sampling of the field may introduce errors in the detected vortices, Fig. 2.3 shows an example of a function sampled on a Cartesian grid at resolutions of 0.01λ , 0.1λ and 0.15λ . In all cases, vortices are detected solely from the phases of the four lattice points of each grid plaquette, as discussed above. At the highest resolution case, this procedure is extremely reliable and it is clear that every vortex has been detected correctly. This results from the smoothness of the field; when so far below the wavelength scale it cannot change too rapidly and four sample points are almost guaranteed to properly sample its path.

At a lattice spacing of 0.1λ we begin to see how this algorithm is not always reliable, with the numerical search at this resolution shown in Fig. 2.3b but failing to detect precisely one vortex (ignoring the one example where vortex was closer to the edge than the lattice spacing). This is because two vortices approach closely, and the lower resolution is insufficient to distinguish them. Such an occurrence will sometimes lead to no vortex being detected as they have opposite signs and their phases cancel out, or in this case the phase anisotropy and sampling points are such that the double vortex appears instead as just one. It can also happen that the two vortices are both detected even though very close to one another, if the lattice sampling points happen to be placed correctly and the phase anisotropy is not too high, but there is no example of this in the Figure. Although the presence of the other vortices is detected, their apparent positions may be incorrect by up to approximately the lengthscale of the numerical lattice spacing; no smaller lengthscale is numerically sampled, so vortices can only be detected near to (but not exactly at) their real positions in the continuous field.

Decreasing the resolution slightly more to 0.15λ lattice spacing leads to far more significant errors; almost all of the vortex points from the first two choices of resolution are no longer detected correctly, mostly being missing entirely. This appears to be because the scale has become large enough that the phase is commonly

anisotropic around the lattice plaquette, so it becomes increasingly a question of chance as to whether its four sampling points happen to have the correct position to see the presence of a vortex (though they are also now more likely to do so even if no vortex is present). Vortices could still be detected if the four lattice points were replaced with more samples in a path around the plaquette, the phase change still sums to 2π , but four points alone are no longer sufficient.

These overall constraints lead to our choosing a resolution around 0.1λ or smaller in our later numerical analysis of vortex lines in our random wave systems. The choice is somewhat heuristic rather than based on a specific property of the vortex recovery, but we are less concerned about missing vortices for reasons that will be explained later and so this basic resolution is sufficient. For similar reasons, similar Cartesian lattice methods are common throughout the literature [5, 10, 15, 17, 19].

The new algorithm presented in Sect. 2.2 resolves the detection failure problem through vortex continuity in the third dimension, but other modifications have been suggested in the literature. A first option is to use not a Cartesian grid but a triangular one, so that just three points encircle each singularity [22]; the local resolution is still low, but a regular triangle permits less ambiguity about the nature of the phase change, as it becomes totally unambiguous that all three triangle sides must have a phase changing in the same direction.

The quality of triangle-based sampling suggests another potential optimisation; since the phase of random waves varies only slowly below a certain scale, it is reasonable to *guess* that the point in the centre of each grid plaquette has a value that is the average of all the corner points. The plaquette can then be broken into four triangles, each individually searched for phase vortices. This method has the advantage that the spatial location of the vortex can be guessed on a scale lower than the sampling resolution, and in preliminary numerical experiments we have verified its stability. However, its accuracy is still not guaranteed and in practice our own three-dimensional method renders it unnecessary for fields where the full phase function is known, so we do not make use of it.

Using just four sampling points also means that vortices of higher orders, or two or more nearby vortices below the sampling distance, cannot be detected due to aliasing; the phase unwinding rules can only detect a 0 or $\pm 2\pi$ phase change. In random fields, nearby vortices are uncommon (the vortex-vortex correlation function goes to 0 at the origin [23]), and vortices of higher order generically do not occur in random waves.

These methods all relate to locating point vortices in two-dimensional phase fields, but our interest is in the three-dimensional tangle of one-dimensional vortex curves. The ubiquity of grid-based data and the robustness of phase unwrapping in 2D makes it desirable to extend these techniques; the 3D grid considered as an array of many cubic *grid cells*, with every face of every cell considered as a 2D phase unwrapping problem based on its four corners. A grid cell is marked as containing a vortex if a 2π phase change is detected around any of these faces, and since the vortex line must both enter and leave the cell, a local segment may be reconstructed by joining the two faces around which the phase circulation is found. Since the direction of




Fig. 2.4 Vortices in random wave fields, tracked in three dimensions through a basic phase unwinding method. In (a), Cartesian grid cells are marked in *green* if a vortex is detected through one or more of their sides, and in *red* if vortices are detected through three or more of their sides. In (b), the local line segments of vortices are approximated based on the cell sides through which they pass. The plotted cubic region has side length 4.33λ .

phase circulation must be consistent along the vortex, this orients the local segment according to the direction of phase change around these faces.

Figure 2.4 shows examples of this basic technique; Fig. 2.4a marks cells detecting a vortex through at least one face in green, and through at least three faces in red, denoting cells where the algorithm must be insufficient (other cells also incorrectly detect vortices, but not so obviously). Figure 2.4b shows the vortex segments from each cell, straight lines between the centres of the cell faces through which a vortex was detected. These cannot all be joined to continuous lines, as segments are missing where the basic algorithm failed to detect vortices. This can happen for all of the same reasons that apply to vortices in two dimensions (though the lengthscales are different between two-dimensional random waves and slices of three-dimensional volume tangle), but also for new reasons. For a vortex segment to be presented it must be detected not just once but twice, through both the two sides of the cell that it passes through. If it is not, we do not know its path. Vortices also cannot be resolved if more than one enters the same cell, in which case more than two cell sides have vortices detected and it is impossible to know which way the vortices are joined without more information. Slightly different mistakes arise if multiple vortices pass through the same face, in which case they will often not be detected properly (just as in two dimensions). This can even lead to previously-detected vortex paths being shown to be incorrect; even if two adjacent cells contain vortex segments that seem to be part of continuous chains, it is still possible that both vortices pass through the adjoining cell face and connect differently, but such that their passage is invisible in the discretised lattice.

This basic algorithm is the basis for three-dimensional vortex tracking in many of the systems mentioned above, where the vortices need only be approximately located [10, 19] or where experimental limitations do not permit a higher resolution [4]. However, the issues listed above mean it is not in general sufficient to precisely locate and distinguish vortices. This problem is well recognised and multiple solutions have been used; in a 2 + 1D optical beam model, it was possible to ‘guess’ the correct local connectivity based on the smoothness of the field, a procedure applicable to our own work though such a statistical technique cannot always give the correct result [4, 5]. The issue also appears directly in the \mathbb{Z}_3 discrete phase lattice model introduced earlier, in which connectivity can only be resolved randomly if using a Cartesian lattice.

This ambiguity has also been resolved by using a tetrahedral grid [24] which directly clarifies vortex paths since tetrahedra with random phase vertices will never lead to multiple vortices passing through their volume, but applying this idea to continuous fields adds complexity and does not avoid problems when vortices approach closely or where numerical error means their passage through a grid cell wall is not detected. As with two dimensions, the use of tetrahedra may also be extended to split each cubic cell of a Cartesian grid, though rather than guessing the central phase (now an average of 8 corner points) it requires less interpolation to take advantage of the third dimension and split along four new planes of the cubic cell to give 4 individual tetrahedra in a sublattice of the Cartesian sampling; these choice of planes to split the cell is arbitrary, though a cell-by-cell comparison of each choice may help make the method robust. [25] similarly uses the centre point of the cell and of each of its faces to construct 24 different tetrahedra through whose sides the vortex is tracked with the same purpose. Again, these methods might improve the ability to distinguish and resolve vortices below the sampling resolution, but does not guarantee doing so in all otherwise-ambiguous cases.

This section has focused on locating vortices by a direct extension of two-dimensional phase unwrapping, but this is not the only method that has been employed in the literature. A somewhat different method focuses not directly on the phase, but on the *intersection* of two surfaces; the independent zero level surfaces of the real and imaginary components [2, 3] (discussed in Chap. 1. These are located by inspecting the corners of each grid cell in a Cartesian grid, as with the vortex phases above, but with a more involved scheme accounting for the triangulation of the zero surfaces rather than the location of one-dimensional line segments passing through the plane. Locating vortices reduces to testing the intersection of triangles in the surface reconstructions, a process which may be surprisingly efficient. Conveniently, there is no numerical ambiguity regarding vortex line closure, as the line is defined via isosurfaces whose triangulation is guaranteed to be continuous by the nature of the algorithms used—unlike the phase unwrapping above which may fail to detect a vortex under some conditions. However, this does not guarantee that the closure is physically correct if the resolution is low, or that all vortex loops are detected.

The main disadvantage of these techniques is that the surface triangulation is unnecessary to extract the location of vortex lines; if the same reconstruction can be

obtained via the phase alone, many of the same optimisations will be available, but the surface triangulation and intersection calculations may be omitted. For instance, our own algorithm in Sect. 2.2 could be modified to locate vortices via isosurface intersection, but we would anticipate identical results at the cost of a greater computation time.

A different method employing isosurfaces simply looks for contours of low intensity (i.e. for complex Ψ , the intensity $|\Psi|^2$). In principle, at some sufficiently low value these contours would only form tubes around vortex lines. In practice this performs very poorly for the same reason that the intensity is not useful for finding vortices in two dimensions; a very high resolution is necessary for the contour tube to closely surround the vortex, and if two vortices approach closely their contour tubes will merge.

A final technique from the literature is to try to avoid searching a full grid for vortices (in which many grid cells will be empty), instead locating an initial point on a vortex and iteratively marching along the curve [1, 3]. This has the advantage of efficiency, but the disadvantage that the location of vortices becomes statistical; it is impossible to guarantee finding at least one point on every vortex without searching almost all grid cells (a problem our method partly shares, but minimises by always searching the entire numerical grid for vortices). Additionally, walking along the vortex may be difficult or unstable, particularly where it approaches close to another (possibly unidentified) vortex curve. For this reason, we do not trace vortices this way; these techniques might be more suitable for approximately tracing curves such as in illustrating gradient flow [1], or if trying to follow a single vortex curve over a very large arclength rather than also tracking the whole local ensemble of curves.

2.2 Tracking Vortices with a Resampling Algorithm

None of the methods of Sect. 2.1 fully resolve the problem of three-dimensional vortex tracking, relying on assumptions and approximations that may not always hold, or which are insufficient in continuous fields. For these reasons, we introduce our own extension to this core algorithm, based instead on recursively *resampling* the field on smaller Cartesian grids of higher resolution wherever the vortex path is ambiguous. This allows us to locate and trace vortices efficiently with a relatively low initial sampling resolution but without sacrificing topological correctness, and returns piecewise linear representations of vortex curves with up to hundreds of vertices per wavelength. We will refer to this technique as the recursively resampled Cartesian grid method (RRCG). In this Section we give the details of the resampling algorithm, and of our geometrical optimisations to make precise vortex tracing possible.

The core new feature of the tracking algorithm implemented here is that a new, local sampling lattice is automatically generated in regions where the spatial details of the tangle are not resolved by the initial resolution of the sampling grid. This occurs mainly when two vortices enter the same grid cell, but also when an unusual vortex position and anisotropy means it is only detected to pass through one side of




Fig. 2.5 Demonstration of the supersampling vortex tracking algorithm. Each sphere at a vertex of the sampling grid represents a phase voxel, with phase indicated by *greyscale*. Lines between the voxels are coloured following a linear interpolation of phase between them (following the shortest change mod 2π). The patterned lines represent vortex lines, the increase *red-white-blue* denoting the direction of the vortex line, from a right-handed increase of phase. These lines enter or leave the cell through faces around which the total phase change is 2π . This Figure originally appeared in [26]

the cell, while in reality we know that it must both enter and leave. In both cases, the topology within the grid cell is ambiguous, and cannot be perfectly resolved using any of the methods in the previous section.

Figure 2.5 shows an example grid cell where this resampling is necessary. In Fig. 2.5a, vortices are detected to pass through four of the cell faces, oriented by the direction of phase change (an arbitrary choice; we orient vortices such that the phase circulates clockwise around them). Without more information it is impossible to know which incoming and outgoing sides should be joined. In Fig. 2.5b, the 8 points of the local grid cell are replaced by a new $3 \times 3 \times 3$ grid; 8 new grid cells, each with side length one half of the original. Each cell of this new grid is examined for vortices in exactly the same way, this time fully distinguishing the paths of incoming and outgoing vortices as they do not approach closely on the scale of the new grid cells. The procedure is also *recursive*, if vortex topology is ambiguous on the scale of the new grid the resampling procedure is performed again, at an even higher local resolution, in the region of this specific smaller grid cell that still cannot be resolved.

To fully resolve all vortex filaments, we must both repeat this procedure and allow it to spread through the sampling lattice. The algorithm is as follows:

1. Check for vortex segments in all cells of the initial grid (i.e. looking for 2π phase changes around the border of each plaquette face). If there is no ambiguity in a grid cell, join the centres of the two faces where a phase change is detected with an oriented vortex line segment.
2. Resample all grid cells in which vortices are detected but the topology is ambiguous. The local grid cell sidelength of the resampling grid is changed to the

- (initial sidelength)/ N for integer N_{cell} , initially 2 but stored and incremented by one each time the procedure is run in a given cell.
3. If the local topology is still ambiguous, repeat step (2), further incrementing the local resolution.
 4. Spread the new local resolution to all neighbouring cells bordering faces of the initial cell through which a vortex is detected. Each of these is also resampled to the same local resolution.
 5. If step (4) changes the detected vortex configuration in any of the neighbouring cells, repeat (4) with neighbours of this new cell. Thus, the new resolution spreads until no ambiguity remains.
 6. If the detected vortex configuration of *any* cell changed through steps (2–5), jump back to step (2).

In this fashion, volumes of higher resolution are generated where necessary, and spread through the initial grid until each ambiguity is resolved. Step (6) repeats the process, catching and correcting any regions through which a vortex was initially detected, but which resampling has revealed to be a numerical artefact rather than a real line segment. After following this procedure, every vortex is an unbroken chain of piecewise linear segments, forming either a closed loop within the simulated grid cell or terminating on its boundaries. Due to the recursive procedure, the topology of these curves can be correctly resolved even where they approach (in principle) arbitrarily closely, where multiple (even $\gg 2$) vortices enter the same cell, or even where multiple vortices pass through the same initial grid cell face and would normally be undetectable.

It is important to note that this algorithm does not just increase the resolution of the small number of cells where vortex detection errors can be seen, but also in all surrounding cells. This is vital because the initial vortex detection may simply be wrong, with particular phase anisotropy patterns and unfortunate grid spacing leading to the apparent presence of a vortex in the initial lattice even though a vortex does not pass through this region the continuous field being sampled. Allowing the resampled region to spread entirely resolves this problem, as it means that once a cell has been resampled (and its detected vortex contents possibly changed) its neighbours will always be checked for consistency with this new result, and in turn resampled themselves until the detected vortices are consistent everywhere. This also resolves most of the issues explained in the previous Section when two vortices are connected in an unseen way via the passage of both through a single grid plaque. In practice this phenomenon tends to lead to high phase anisotropy and numerical failures that force a resampling, which in turn detects the true local geometry and topology. We are also able to optionally force this procedure every time different vortices occupy adjacent cells, but find that doing so almost never reveals new problems. Alternatively, the opposite may happen with the resampling grid containing new numerical errors, but in this case the resampling is simply repeated at a higher resolution and the local smoothness of the field means that all ambiguities are eventually resolved.

As a benchmark of standard performance under real use, several of our core datasets with 3-torus eigenfunctions use an initial grid cell sidelength of 0.087λ ;

at this resolution, around 4 % of grid cells contain vortices of which roughly 5 % must be re-resolved at least once and approximately 0.1 % re-resolved at five or more times the initial resolution. Sampling the initial grid at even twice the initial resolution of the original grid would have an 8-fold impact on calculation time, so the resampling procedure is essential to efficiently resolve the local topology. Full details of the different parameters used and numerical statistics are given in Sect. 2.5.

These options are also visualised in Fig. 2.6; (a) shows resampled regions in the same example cell as Fig. 2.4, with boxes showing regions resolved at two to seven times the initial resolution, the ratio denoted by the intensity of the colour. Only a small number of these seeded the resampling, and most of the cases of several adjacent resampled cells are due to the resampled region spreading by necessity as it reveals errors in adjacent cells and is necessary to ensure connectivity. This gives some insight into exactly how important the resampling procedure is to reconstructing the tangle. Figure 2.6b shows the same vortices but without the resampling indicator boxes; the steps of our algorithm guarantee that these vortices are now all continuous, unbroken space curves terminating on the boundaries of the volume.

Although this method topologically resolves vortex lines very precisely on the sub-wavelength scale, the individual linear segments do not accurately reproduce the spatial conformation of the vortex filament, as it is still assumed that each vortex line segment passes through the centre of the face where it are detected. Thus, even if all the sampled vertices are close to the true curve, the local tangents cannot be accurate. We add three further optional steps to improve the local fit to the curve, which are applied mainly in Chap. 3 where accurately resolving local geometry is most important.

First, once the vortex lines have been located, the resampling procedure is applied again to all the cells through which the vortex line passes, even if the topology is already well distinguished. This accurately traces the shape of the vortex curves without expending further numerical resources on the voids between. In principle, extremely small vortex loops below the sampling resolution may not be detected, but these are extremely rare and contribute a negligible amount to the total vortex arclength (Sect. 3.6). Figure 2.6d illustrates the result of this procedure applied to our previous example function, with the initial result of our algorithm in (b). Although the vortices in both cases have the same overall topology, the enforced resampling of the entire vortices in (d) can be seen to recover much more of their precise geometry. Applying such a locally high resolution to the initial grid would also track the vortices well even without the RRCG procedure, illustrated in Fig. 2.6c which shows the grid cubes in which vortices were detected even without the RRCG algorithm after searching the same function as in (a) but with four times its initial resolution applied uniformly across the lattice. Even viewing only the cells containing vortices it is now possible to follow the paths of individual vortex lines, but this calculation takes significantly longer than the RRCG resampling. Using this method to sample high energy eigenfunctions with an even higher resolution would quickly extend the computation time to days or weeks. Our algorithm focused on the vortices takes only hours at most (or more normally minutes, details in Sect. 2.5) and allows a local resolution along vortices of several hundred sample points per wavelength.




Fig. 2.6 Vortices detected in a 3-torus eigenfunction, showing the results of different algorithmic optimisations. The torus eigenfunction has energy 75, sampled on a grid of size 25^3 in (a), (b) and (d), an initial resolution of 0.1λ . In (a), the *green lines* denote recovered vortex segments and the *red boxes* show the degree of local resampling—up to 4 times the initial resolution according to the intensity of the colour. **b** Shows the same vortex cell, but displaying only the vortices. **c** Shows the same vortex cell sampled at four times the original resolution, a lattice of size 80^3 , showing only cells detected to contain one or more vortex segments. **d** Shows vortices detected with the lower initial resolution, but with an enforced resampling along the entire length of vortex curve to a local grid size of 0.025λ , equivalent to (c) but applied only to vortex-containing cells rather than the entire volume. Of these cells, the method in (d) is used mainly in Chap. 3, while the simpler method of (b) is sufficient in Chap. 5. These cells all have side length 4.33λ

Second, rather than assuming the vortex line penetrates the centre of the grid face through which it is detected, we instead treat the face plane as a new, two-dimensional vortex location problem. The intersection point is located by resampling the face with a 2D square grid of some multiple of the local 3D resampling resolution, and locating the square plaquette of this new grid around which there is still a 2π phase change. The centres of these new plaquettes in the 2D face are joined by straight lines in the 3D grid to approximate the real vortex curve. This optimisation is illustrated in Fig. 2.7, in which the initial four phase points of (a) reveal only that a vortex probably passes between them, while the 49 times higher resolution in (b) isolates the vortex well away from the point where the phases meet. Although this higher resolution is only necessary for the recovery of local geometrical quantities, the numerical cost of the face 2D resampling is relatively low (clearly lower than that of resampling an entire grid cell by a factor of the cube of the resolution multiplier), so we apply it frequently in datasets where only topology is relevant.

Finally, even with both of these optimisations the recovered curve is noisy on some small scale, thanks to remnants of the original lattice and the discretisation of the tangent vector. This poses problems for geometrical recovery as quantities such as the torsion (Sect. 3.3) are extremely sensitive to even small distortions. This remaining noise can be removed through smoothing and downsampling the recovered curve; since the recovered vertices cluster around the true curve, it will be better reconstructed by an appropriate *average* of their positions.




Fig. 2.7 Resampling of a face of a grid cell, gaining improved geometric information about the vortex position. In (a), the sampling resolution is 0.1λ and there is only enough information to show that a vortex is probably present. In (b) the resolution is seven times higher (typical in our simulations), and the vortex can be located more precisely in the plane. The *red dot* denotes the (clockwise) vortex, numerically located via the phase changes between the adjoining four samples

We smooth the curves by keeping only every n th vertex for some integer n , and fitting with a piecewise polynomial spline curve interpolation. A forced perfect fit would introduce high frequency noise to the polynomial thanks to remnants of the original lattice, so we require only that the spline curve fit closely approximates the original vertices. We enforce that the sum of square distances between original and interpolated points does not exceed some small value; $\sum_v (\mathbf{r}_v - \mathbf{r}_i)^2 < \delta L$, where \mathbf{r}_v represents the vertices of the original piecewise linear curve, \mathbf{r}_i is the interpolated spline curve at each of these points, δ is a fitting parameter, L is total length of the current vortex curve and the sum is over the vertices. This procedure resembles that of [27] in tracing superfluid vortices, though here we are concerned with accurately recovering the small scale geometry rather than removing noise on larger scales. We also optionally perform a direct smoothing by of the curve by convolving by a weighting distribution, though we find that this is largely redundant if combined with the polynomial spline fitting.

Figure 2.8 compares different choices of downsampling and δ parameters, along with their direct effect on the position of the peak of the curvature distribution for a larger line ensemble with these parameters applied (see Sect. 3.3, but used here only to test stability). If the polynomial fit is forced to pass perfectly through every point or no downsampling is performed, the peak is significantly higher (and the distribution would be highly distorted), but as both types of smoothing are increased the peak settles to approximately 0.15 ± 0.05 . For this reason, we use $\delta = 0.004$, but the choice is somewhat arbitrary and we accept an inherent error on the order of 3% of values recovered this way. Likewise, we downsample curves to every fifth




Fig. 2.8 The effect of different smoothing parameters on recovered vortex geometry. The vertical axis in both (a) and (b) shows the maximal value of the curvature PDF of the vortex curves (described fully in Sect. 3.3), i.e. the probability of the most probable curvature, an example of a geometrical parameter whose value depends on the details of smoothing. **a** Shows how this is affected by the δ parameter in our polynomial fit, and **b** the downsampling of the initial vortex curve

point before applying the polynomial fit. Although both patterns appear to be settling to a lower value of curvature, in fact this is due to the increased smoothing further distorting the curve.

2.3 Vortex Tracking in Different Systems

The algorithms described in Sects. 2.1 and 2.2 are easily adapted to vortex tracking in different systems by choosing parameters such that the area of interest is covered by a Cartesian grid. We cover here the details of doing so in the systems covered by this thesis; the 3-torus (periodic boundary conditions), the 3-sphere, and the quantum harmonic oscillator, as in Sect. 1.7.

For two of these systems this extension is trivial; with the quantum harmonic oscillator, the vortex tangle covers some statistically small region of real space, and the grid must only cover this. In principle, ‘hairpin’ vortices may extend arbitrarily far outside the classical radius where the potential is equal to the energy, but these are rare events that do not significantly affect topological statistics (Sect. 5.1), and the grid may easily be large enough to cover the interesting region. Our practical choices of parameters are discussed in Sect. 2.5.

When the eigenfunctions being studied are those of the 3-torus and so have Cartesian periodicity, the grid must itself must be made periodic, with plaquettes on one side identified with plaquettes on the opposite side, and lines joined continuously through these boundaries. We extend the algorithms above by allowing the ‘spreading’ stage to resample cells through periodic walls, ensuring a perfect vortex recovery. With this adjustment, the full extent of the manifold is appropriately sampled contingent only on an appropriate resolution according to the wavelength; our choices of parameters are again discussed in Sect. 2.5. Although lines may pass through the periodic sides many times, the full curve may be recovered by simply ‘unwrapping’ these passages, as in Fig. 2.9; each arc of the curve, terminating on the boundary of the simulation cell, is shifted by an appropriate multiple of the periodicity to form a single continuous curve. The closed loops may wrap a non-zero number of times around the torus and so do not form closed loops following this procedure, but this poses no numerical problem as long as the jump between their start and end points is ignored in geometrical measures. Such lines are topologically more complex discussed in Sect. 5.6.

Our third system is the degenerate eigenfunctions of the 3-sphere, S^3 . These demand a further extension to the algorithm, as the 3-sphere cannot be tiled by regular quadrilaterals of arbitrary size. Other regular tilings are available, but are undesirable since (as discussed in Sect. 2.1) they add practical difficulty. It is more desirable to adapt our Cartesian grid method.

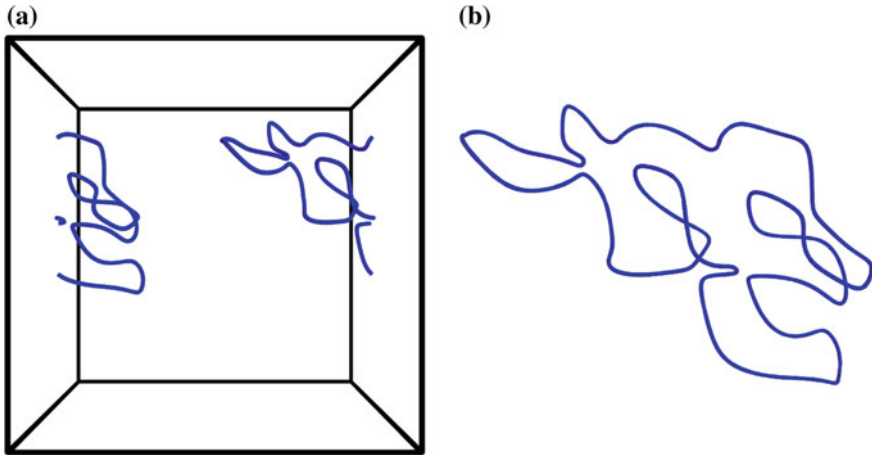


Fig. 2.9 A short vortex loop passing through the 3-torus boundaries in a simulated cell. **a** Shows its initial representation, with multiple components cut along the simulation boundary, while **b** shows the result of ‘unwrapping’ the line by shifting components by multiples of the periodicity in order to recover a trivial loop

We simulate the 3-sphere via the eight *octants* of the 3-cube (the *hypercube*), the shape in \mathbb{R}^4 with edges of equal lengths along each of four orthonormal directions. Its eight octants are the analogue of the six faces of a normal cube in \mathbb{R}^3 . We cover each octant with a Cartesian grid on which the metric is not uniform. This is most clearly visualised by the analogous construction in two dimensions; consider a cube lying within a sphere, with both having their centre of mass at the origin. Any point of the cube may be projected to a unique point on the sphere along the straight line that begins at the origin and passes through that point. Each face of the cube can be assigned a natural Cartesian grid across its area, and any process on the surface of the sphere modelled by performing calculations on this grid while retrieving distances from the projection to the sphere. Although the metric is not uniform across each face, every face is identical, and the geometry of the sphere may be recovered by joining faces along appropriate edges. In this respect, the unfolded faces form a *net* of the 2-sphere.

In our case of the 3-sphere, the process is identical, but with the projection now via a straight line from the origin of \mathbb{R}^4 and the six 2D faces of the cube replaced by three-dimensional octants of the hypercube. Other properties remain the same; geometrical positions in each octant are identified with coordinates on S^3 by projection along a straight line passing through the origin in \mathbb{R}^4 , and the geometrical connectivity of the 3-sphere is recovered by joining appropriate faces of each octant lattice.

Following the discussion in Sect. 1.7.2, the 3-sphere manifold satisfies

$$x^2 + y^2 + z^2 + t^2 = r^2. \quad (2.1)$$

and may be understood via three angular coordinates ψ, θ, ϕ ranging from $0 \rightarrow \pi$, $0 \rightarrow \pi$ and $0 \rightarrow 2\pi$ respectively. The projection to the 3-sphere from the hypercube requires inverting (1.20)–(1.23) to give

$$\psi = \arccos\left(\frac{t}{r}\right), \quad (2.2)$$

$$\theta = \arccos\left(\frac{z}{r \sin \psi}\right) = \arccos\left(\frac{z}{\sqrt{r^2 - t^2}}\right), \quad (2.3)$$

$$\phi = \arctan\left(\frac{y}{x}\right), \quad (2.4)$$

which project any point in \mathbb{R}^4 radially onto the 3-sphere of radius r .

By analogy with the 2-spherical example above, we simulate the 3-sphere by running our numerical algorithm from the previous sections once for each of the octants of the hypercube, covering each octant with a Cartesian lattice whose points are mapped to the hyperspherical angles via (2.2)–(2.4), allowing vortices to be traced without further modification of our numerical algorithm. Each face of each octant is associated with a face of another, which corresponds to the same plane on the 3-sphere. Vortices within each individual octant may enter or leave through one of its faces, but the full continuous vortex line is traced by joining it with the equivalent line leaving or entering the associated face of an ‘adjacent’ octant.

The main advantage of this decomposition is that octants are locally quite flat, the metric does not vary significantly across their volume, so our algorithm’s assumptions about vortex density still hold to a good approximation. Other possible choices of simulation method do not perform so well; for instance, we could impose a Cartesian lattice on the stereographic projection of the 3-sphere and track vortices via the inverse map from stereographic coordinates to hyperspherical angles, but stereographic projection significantly distorts lengths and so vortex density varies significantly in the projected space. This latter choice makes the RRCG algorithm much less efficient, as it is difficult to concentrate numerical effort appropriately when large areas of the Cartesian grid correspond only to small hyperspherical volumes. Although this relatively low metric is sufficient for the RRCG algorithm to perform well, further corrections are necessary to recover local geometry of the resulting curves, discussed in Sect. 2.4.2.

Figure 2.10 demonstrates vortices traced numerically in each of our eight octants. The octants are placed adjacent to one another along identified faces where possible, but are not distorted from the Cartesian lattice recovered from the simulation. Since they represent a net of the 3-sphere, they could be folded together in \mathbb{R}^4 to construct the hypercube. The effect of the minimally distorted metric can be seen in the visually consistent vortex density within each octant; lengths of vortices measured in this discontinuous projection will correspond approximately (but not precisely) to distances on the 3-sphere, discussed further below.




Fig. 2.10 Vortices traced in a hyperspherical eigenfunction via our octant method. The eigenfunction is a random hyperspherical harmonic of energy 80. Vortices in each octant are given a different colour, and octants positioned so that (where possible) moving through an octant wall maps to a continuous path in hyperspherical angles

Figure 2.11 shows the same vortices, coloured in the same way, but represented by stereographic projection. Most octants are now much more distorted, particularly the outermost in which the central point of the octant maps to the point at infinity; vortices passing close to this point may travel very far from the origin the projection. Despite this distortion, all vortices are now continuous, closed loops.

Although we have stressed the relative flatness of each octant, it is important to note that the metric does vary across their extent, and arclength as seen in Fig. 2.10 does not directly correspond to arclength as calculated properly in hyperspherical coordinates. Figure 2.12 demonstrates the ratio of metric distance to Cartesian distance (without metric correction) across the volume of each octant, via two planar slices; in the centre of the octant the hyperspherical distance is twice the local Cartesian distance, while at the edges this factor is reduced to approximately 0.88. This anisotropic distortion is not strongly visible in Fig. 2.10, but must be taken into account when calculating the true length. The details of this calculation are given in Sect. 2.4.


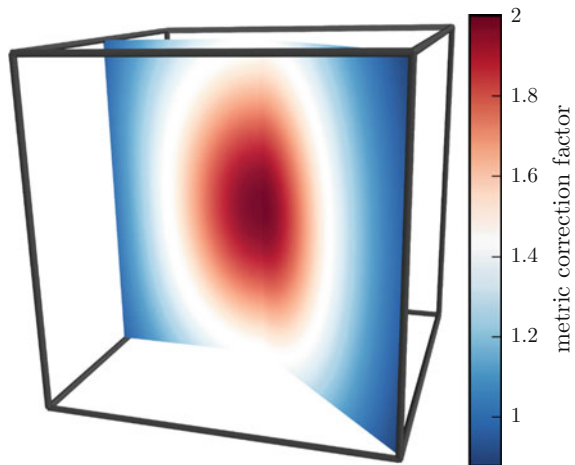


Fig. 2.11 Vortices traced in a hyperspherical eigenfunction via the octant method, projected stereographically. The eigenfunction is the same random eigenfunction as in Fig. 2.10, and vortices in each octant coloured identically, but now projected stereographically in their hyperspherical coordinates such that all vortices are continuous. The black lines show the azimuthal octant cell boundaries, also stereographically projected via the 3-sphere

Fig. 2.12 Distance on 3-sphere as fraction of local distance on the Cartesian grid of a numerical octant. The distance is shown through two planes within the octant, one halfway through one axis, and the other following a diagonal. The distance factor ranges from minimum 0.8776 in the corners to maximum 2.0 in the centre



2.4 Geometrical Details in Different Systems

Performing different kinds of geometrical and topological analysis necessitates a consistent notion of lengths. We give here the details of calculating these quantities for vortices in the context of each of our three systems, accounting for specific properties including their periodicity, curved metric and varying background potential.

2.4.1 The 3-Torus

The 3-torus is flat (it admits trivial metric) and includes no potential, and so angles between vectors may be calculated just as in \mathbb{R}^3 . The periodicity of the manifold does impose an extra condition on shortest distances between points; distances may not always be calculated directly between the two points, but (if relating to a larger section of vortex curve) must correctly take account of the periodic boundary conditions.

In practice, it is almost-always sufficient to ‘unwrap’ a vortex curve through the periodic walls, shifting each segment within the periodic cell by a multiple of the periodic side length corresponding to the number of times the boundaries have been traversed since an arbitrarily chosen starting point along the line; it is by this procedure that the vortex segments within a single cell of the 3-torus may be reconstructed to a continuous line as shown in Sect. 1.7.1. For loops that close with trivial homology, this map to \mathbb{R}^3 loses no information and lets the curves be analysed as normal. For loops that wrap some non-trivial number of times around the boundaries, the unwrapped curve will not close but instead remain open by some (potentially large) multiple of the cell side length according to its homology vector. The loop is still periodic across this gap, but this poses no practical problem as long as analysis does not treat this as a geometrical jump but only a representation of periodicity.

2.4.2 The 3-Sphere

The 3-sphere is unique among the systems studied here in exhibiting inherent background curvature; any calculation of distances or angles must take account of the path taken around the manifold.

The round metric on the 3-sphere in hyperspherical angular coordinates is given by

$$ds^2 = dx^2 + dy^2 + dz^2 = r^2 (d\psi^2 + \sin^2 \psi (d\theta^2 + \sin^2 \theta d\phi^2)), \quad (2.5)$$

and distances along a given path $\Phi(s)$ may be obtained by integrating along its length with this metric. For a vortex curve L , the total length is the integral $\int_L ds$

In practice, the calculation is made easier by considering vortex curves as *piecewise-geodesic* on S^3 , the generalisation of piecewise linear curves to arbitrary manifolds. Such a parameterisation is a natural result of the numerical algorithms in the previous sections, and the vortices of the sphere are specified in terms of a series of hyperspherical angles joined by minimal arcs. These minimal arcs (which are the geodesics between the points) are segments of *great circles* of the 3-sphere, defined as intersections of S^3 with a 2-plane passing through the origin in \mathbb{R}^4 [28]. Any two points of the 3-sphere sit on precisely one such great circle (unless they are antipodal in which case they sit on an infinite number of equivalent great circles), and they divide the circle into two segments, one of which is guaranteed to be the shortest distance between the points along the spherical metric.


Figure 2.13 demonstrates three such great circle segments on the 2-sphere (this could be a cross-sectional 2-sphere of the 3-sphere). The distance between each pair of points reduces to

$$d = r\kappa, \tag{2.6}$$

$$= r \arccos(\mathbf{v}_1 \cdot \mathbf{v}_2), \tag{2.7}$$

where r is the radius of the sphere and κ is the angle between the positions of the points in \mathbb{R}^{n+1} , given by the normalised vectors \mathbf{v}_1 and \mathbf{v}_2 ; in the case of the 3-sphere, these vectors are their positions in \mathbb{R}^4 as given by (1.20)–(1.23). Obtaining the lengths of vortices reduces to the sum of these geodesic distances along the vortex curve.

Fig. 2.13 Distances and angles between arcs on the 2-sphere. The 3-vectors \mathbf{v}_1 , \mathbf{v}_2 and \mathbf{v}_3 denote points on the sphere, joined by geodesic arcs a and b which are segments of the great circles marked in blue. The length of each arc is given by (2.6), and the angle χ between the arcs is given by $\mathbf{t}_a \cdot \mathbf{t}_b$ in (2.8)



The calculation of angles must be treated similarly; since each geodesic segment is not simply a vector but an arc, the angle between each pair of segments is given by the angle between the arcs at the point of their intersection. This is illustrated at the intersection of arcs a and b in Fig. 2.13, they are separated by the angle

$$\chi = \arccos(\mathbf{t}_a \cdot \mathbf{t}_b) = \arccos\left(\frac{\mathbf{v}_2 - \mathbf{v}_1(\mathbf{v}_1 \cdot \mathbf{v}_2)}{|\mathbf{v}_2 - \mathbf{v}_1(\mathbf{v}_1 \cdot \mathbf{v}_2)|} \cdot \frac{\mathbf{v}_3 - \mathbf{v}_2(\mathbf{v}_2 \cdot \mathbf{v}_3)}{|\mathbf{v}_3 - \mathbf{v}_2(\mathbf{v}_2 \cdot \mathbf{v}_3)|}\right). \quad (2.8)$$

2.4.3 The Quantum Harmonic Oscillator

The quantum harmonic oscillator is outwardly most similar to the 3-torus; the eigenfunctions live in \mathbb{R}^3 without projection. However, the notion of distance still depends on the radius due to the changing potential in (1.32), which may be rewritten as

$$\nabla^2 \Psi = (E - \frac{1}{2}(x^2 + y^2 + z^2))\Psi, \quad (2.9)$$

$$= k_{\text{eff}}^2 \Psi, \quad (2.10)$$

where k_{eff} is an effective local wavevector, with an effective local wavelength $\lambda_{\text{eff}} = 2\pi/k_{\text{eff}}$. In this form the equation matches the random wave model of (1.6), but with a position-dependent wavelength that tends to infinity as the radius nears the boundary where $E = 1/2(x^2 + y^2 + z^2)$. Distances along vortex curves in eigenfunctions of the QHO may be calculated by integrating along the curve including a factor of this local wavelength.

This description only applies within the radius where the energy exceeds the potential. Outside this region, λ_{eff} becomes imaginary and is no longer a good descriptor of the local distance. When measuring the length of vortex curves we simply exclude these regions as not contributing to the geometry of the tangle, as vortices are confined to very simple behaviour that does not reflect the random wave model, being able only to travel straight to infinity (Sect. 1.7.3), or doing nothing more complex than reversing direction and re-entering the main tangle in what we will call hairpins.

Measuring arclength in this way is only an approximate descriptor of vortex arclength, taking account of the varying lengthscale of correlation in regions of different potential. However, it is not a complete description as it does not account for the anisotropic directional correlations near the classical radius. It is this phenomenon that encourages hairpins, with vortices rapidly changing direction despite the low local wavelength. A full parameterisation of arclength would need to instead take account of a local metric controlling this difference, but we do not do so here.

2.5 Summary of Numerical Results

We give here the details of all our core numerical datasets whose analysis comprises the rest of this research, including details of input parameters, examples of results, information on the practical efficiency of our algorithms, and basic statistics of the output vortices.

All statistics of timing refer to implementations of our algorithms via Python and Cython, running on a typical laptop machine at the time of writing (Intel Core i5-3320M CPU). It is likely that significant speed improvements could be made to the core algorithm, but our implementations are optimised only so far as necessary to retrieve appropriate results.

2.5.1 The 3-Torus

Our core datasets on the 3-torus have energies 75, 243, 675 and 1875 in (1.11). The numbers of simulated cells in each case, along with basic statistics, are given in Table 2.1.

The initial grid spacing in all cases is 0.0173λ . Local regions are resampled at a higher resolution following Sect. 2.2. In the case of energy 75, an extra resampling stage is performed along the vortex lengths, forcing a local grid spacing of no greater than 0.00247λ . In general, around 0.19 % of cells in the initial sampling grid contain a vortex segment of which about 0.15 % of these cases must be re-resolved. After this resampling and any further smoothing as in Sect. 2.2, each line is resolved with at least 50 sample points per wavelength, or at least 140 sample points per wavelength in the data with energy 75.

All other energies use a lower initial resolution, leading to around 3.9 % of grid cells containing vortices, with a further 4.6 % of these being re-resolved to guarantee

Table 2.1 Details of the simulations we have carried out on eigenfunctions of the 3-torus. The initial and resampled resolutions are only the minimal resolution per grid cell, and the local resolution may be higher if required by the algorithm. Errors given are standard deviations, reflecting the overall distribution of line lengths

Energy E	Number of cells simulated	Approx time per cell	Number of loops		Number of NTH lines	
			Total	Average	Total	Average
75	2628	2 h	39395	14.5 ± 11.1	6915	2.55 ± 1.54
243	19845	80 s	1211553	61.1 ± 27.4	66633	3.36 ± 1.56
675	9514	6 min	2580009	271.2 ± 83.3		
1875	51	20 min	54200	1062.7 ± 112.9	191	3.75 ± 0.83

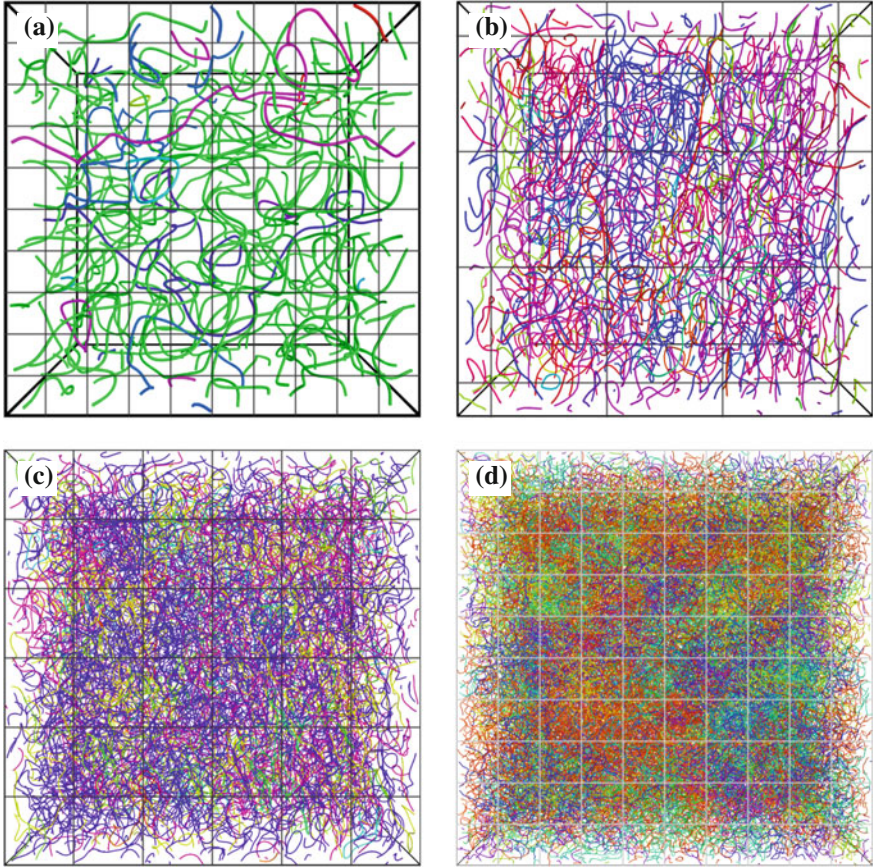


Fig. 2.14 Numerically tracked vortices in simulated cells of the 3-torus. **a–d** correspond to energies 75, 243, 675 and 1875 respectively. In each figure, every vortex is coloured by a different (random) hue. The grid on each image shows a spacing $25\times$ the grid spacing of the initial numerical sampling grid

local continuity; the lower initial resolution successfully shifts computational effort to the resampling procedure, rather than on closely examining empty regions of the grid.

Figure 2.14 demonstrates vortices in example simulation cells at each energy.

As a test of the recovered geometrical statistics at energies 243, 675 and 1875 in which the numerical resolution is relatively low, we have also simulated twice several eigenfunctions of energy 243. The first simulation was performed as above, and the second with the origin of the numerical grid shifted by $(\sqrt{2}, \sqrt{2}/2, 2\sqrt{2})$. The choice of perturbation is arbitrary, but intended to shift each axis by a non-trivial factor of the grid spacing. The recovered vortex curves should be topologically the same in both simulations, but the offset of the origin will expose geometrical

differences as the RRCG algorithm will not necessarily resample the same local regions of the vortices. The difference in the total recovered arclength (measured over the numerically recovered piecewise-linear vortex curves) will give an estimate of the error on the numerical recovery.

Over 49 eigenfunctions of energy 243, the average total arclength recovered was 2220λ , and the average difference between the results of simulations on the same eigenfunction just 2.4λ with a standard deviation of 1.8λ ; both on the order of just 0.1 % of the total arclength. The maximum difference found was 7.5λ , and the minimum 0.009λ . We conclude that the RRCG procedure gives consistent results even at the numerical resolutions we make use of, and although we must account for these errors in any later measurements of length, they are not likely to be significant. These numbers should also generalise without modification to our simulations at energies 675 and 1875, as although the total vortex arclength in a single cell will change, the fractional error of the RRCG algorithm will not.

2.5.2 The 3-Sphere

Our core datasets of the 3-sphere have energies 120 and 255 in (1.24). The numbers of simulated cells in each case, along with basic statistics, are given in Table 2.2.

All of these simulations are performed via a lattice of 50^3 points for each octant of the tesseract, as described in Sect. 2.3. This corresponds to a varying wavelength resolution depending on the energy of the eigenfunction and position in the cell due to the metric described in Sect. 2.4. We are not concerned with replicating the precise local geometry and this resolution is always sufficient for our algorithm to fairly efficiently distinguish vortices.

Figure 2.15 shows examples of the vortex tangle in 3-sphere simulations at each of these energies. In each case, only the regions surrounding the central seven cells of the tesseract (in stereographic projection) are plotted, with vortices outside this region being topologically and geometrically trivial in the projection but travelling far from this central volume.

Table 2.2 Details of the simulations we have carried out on eigenfunctions of the 3-sphere. Errors given are standard deviations, reflecting the overall distribution of line lengths

Energy E	Number of cells simulated	Approx time per cell/min	Number of loops	
			Total	Average
120	2949	5	30910	10.5 ± 3.8
255	5697	11	107856	18.9 ± 5.50

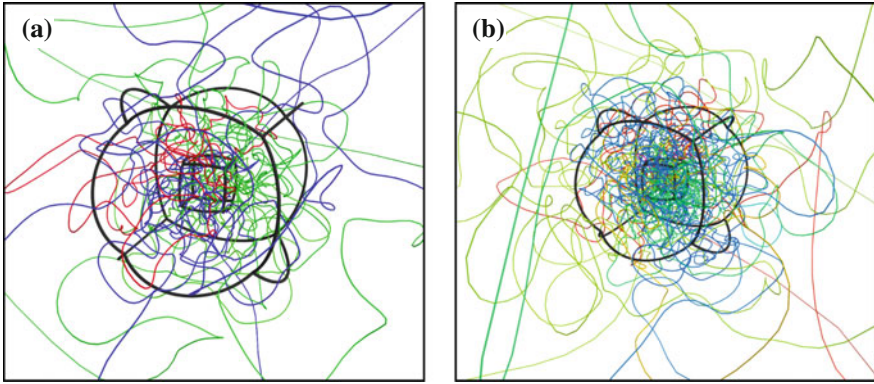


Fig. 2.15 Numerically tracked vortices in simulated degenerate eigenfunctions of the 3-sphere. **a, b** correspond to energies 120 and 255 respectively. The vortices are drawn in stereographic projection, with black lines showing some edges of the tesseract whose eight cells have been individually sampled, also plotted via projection to the 3-sphere. In each figure, every vortex is coloured by a different (random) hue. Only the area immediately surrounding the seven tesseract cells closest to the origin in projection is shown, with vortices outside this area corresponding to only a small true distance

2.5.3 The Quantum Harmonic Oscillator

Our core datasets of the QHO have energies $10 + \frac{3}{2}$, $15 + \frac{3}{2}$, $20 + \frac{3}{2}$ and $25 + \frac{3}{2}$ in (1.32). The numbers of simulated cells in each case, along with basic statistics, are given in Table 2.3.

The numerical grid size in all cases is identical, with vortices sampled in the region $-15 < x, y, z < 15$ at a resolution of 0.2 along every axis. The number of wavelengths that this resolution corresponds to varies both with position and with

Table 2.3 Details of the simulations we have carried out on eigenfunctions of the QHO. The ‘number of infinite’ column gives the number of vortices that extend to infinity (though becoming geometrically trivial outside the classical radius), while the ‘number of loops’ gives the number of vortices that close normally, largely within the classical radius. Errors given are standard deviations, reflecting the overall distribution of line lengths

Energy $E - \frac{3}{2}$	Number of cells simulated	Approx time per cell/min	Number of infinite		Number of loops	
			Total	Average	Total	Average
10	1032	2.5	10671	10.3 ± 2.9	3135	3.89 ± 2.15
15	296	3.5	4660	15.74 ± 3.84	2005	7.06 ± 3.65
20	5279	5.1	112384	21.29 ± 4.41	65318	12.40 ± 4.90
25	1413	7.5	37964	26.9 ± 4.9	22062	15.6 ± 6.3




Fig. 2.16 Numerically tracked vortices in simulated degenerate eigenfunctions of the QHO. **a–d** correspond to energies E of 10, 15, 20 and 25 respectively. In each figure, every vortex is coloured by a different (random) hue. The drawn cell has total width 2.2 times the radius in the eigenfunction at which $E = 1/2(x^2 + y^2 + z^2)$, though the data tracks vortices over a larger volume to more reliably detect hairpins where vortices leaving this region return to the bulk of the tangle

energy (Sect. 2.4), but we are not concerned with geometrical correctness in these datasets and so the difference is not important.

Figure 2.16 shows examples of vortices in QHO cells at each of these energies.

References

1. K.-F. Berggren, P. Ljung, Nature of streamlines for Berry-type wave functions in open 3D cavities, in *Mathematical Modeling of Wave Phenomena, AIP Conference Proceedings*, vol. 1106 (2009), pp. 253–259

2. P. Ljung, A. Ynnerman, Extraction of intersection curves from iso-surfaces on co-located 3d grids, in *SIGRAD* (2003), pp. 23–28
3. J.-P. Thirion, A. Gourdon, The 3D marching lines algorithm. *Graph. Model. Im. Proc. The 3D marching lines algorithm*. **58**(6), 503–509 (1996)
4. K. O’Holleran, *Fractality and topology of optical singularities*. PhD thesis, Department of Physics and Astronomy, Faculty of Physical Sciences, University of Glasgow (2008)
5. K. O’Holleran, M.R. Dennis, F. Flossmann, M.J. Padgett, Fractality of light’s darkness. *Phys. Rev. Lett.* **100**, 053902 (2008)
6. J. Leach, M.R. Dennis, J. Courtial, M. Padgett, Vortex knots in light. *New J. Phys.* **7** (2005)
7. K. O’Holleran, M.J. Padgett, M.R. Dennis, Topology of optical vortex lines formed by the interference of three, four and five plane waves. *Opt. Express* **14**, 3039–3044 (2006)
8. I.B. Halperin, Statistical mechanics of topological defects, in *Les Houches Session XXV–Physics of Defects*, ed. by R. Balian, K.I.éman, J.-P. Poirier (1981)
9. B.M. Caradoc-Davies, R.J. Ballagh, K. Burnett, Coherent dynamics of vortex formation in trapped bose-einstein condensates. *Phys. Rev. Lett.* **83**(5), 895–898 (1999)
10. B.M. Caradoc-Davies, *Vortex Dynamics in Bose-Einstein Condensates*. PhD thesis, University of Otago, Dunedin, New Zealand (2000)
11. M. Farge, K. Schneider, G. Pellegrino, A.A. Wray, R.S. Rogallo, Coherent vortex extraction in three-dimensional homogeneous turbulence: comparison between CVS-wavelet and POD-Fourier decompositions. *Phys. Fluids* **15**(10), 2886–2896 (2003)
12. G.P. Alexander, B.G. Chen, E.A. Matsumoto, R.D. Kamien, Colloquium: Disclination loops, point defects, and all that in nematic liquid crystals. *Rev. Mod. Phys.* **84**, 497 (2012)
13. F. Flossmann, K. O’Holleran, M.R. Dennis, M.J. Padgett, Polarization singularities in 2D and 3D speckle fields. *Phys. Rev. Lett.* **100**, 203902 (2008)
14. E. Matene, V. Jacquemet, Fully automated initiation of simulated episodes of atrial arrhythmias. *Europace* **14**, v17–v24 (2012)
15. K. Umopathy, K. Nair, S. Masse, S. Krishnan, J. Rogers, M. Nash, K. Nanthakumar, Phase mapping of cardiac fibrillation. *Circ. Arrhythmia Electrophysiol.* **3**, pp. 105–114 (2010)
16. J.M. Rogers, Combined phase singularity and wavefront analysis for optical maps of ventricular fibrillation. *IEEE T Bio-Med. Eng.* **51**(1), 56–65 (2004)
17. S. Chavez, Q.S. Xiang, L. An, Understanding phase maps in mri: a new outline phase unwrapping method. *IEEE Trans. Med. Imaging* **21**(8), 966–977 (2002)
18. R. Yamaki, A. Hirose, Singularity-spreading phase unwrapping. *IEEE T Geosci. Remote* **45**(10), 3240–3251 (2007)
19. H.S. Abdul-Rahman, M.A. Gdeisat, D.R. Burton, M.J. Lalor, F. Lilley, C.J. Moore, Fast and robust three-dimensional best path phase unwrapping algorithm. *Appl. Opt.* **46**(26), 6623–6635 (2007)
20. T. Vachaspati, A. Vilenkin, Formation and evolution of cosmic strings. *Phys. Rev. D* **30**, 2036–2045 (1984)
21. F.Y. Wu, The Potts model. *Rev. Mod. Phys.* **54**, 235 (1982)
22. L.J. Rantner, L. Wieser, M.C. Stühlinger, F. Hintringer, B. Tilg, G. Fischer, Detection of phase singularities in triangular meshes. *Method Inform. Med.* **46**, 646–654 (2007)
23. M.R. Dennis, Nodal densities of planar Gaussian random waves. *Eur. Phys. J. Spec. Top.* **145**, 191–210 (2007). Conference on Nodal Patterns in Physics and Mathematics, Wittenberg, Germany, July 24–28, 2006
24. M. Hindmarsh, K. Strobl, Statistical properties of strings. *Nucl. Phys.* **B437**, 471–488 (1995)
25. J.F. Nye, Evolution of the hyperbolic umbilic diffraction pattern from airy rings. *J. Opt. A* **8**, 304–314 (2006)
26. A.J. Taylor, M.R. Dennis, Geometry and scaling of tangled vortex lines in three-dimensional random wave fields. *J. Phys. A* **47**(46), 465101 (2014)
27. A.W. Baggaley, C.F. Barenghi, Spectrum of turbulent Kelvin-waves cascade in superfluid helium. *Phys. Rev. E* **83**, 134509 (2011)
28. H. Manning, *Geometry of Four Dimensions*. Applewood Books (2011)

Chapter 3

Geometry and Scaling of Vortex Lines

In this Chapter we examine the local geometric structure and statistics of vortex lines in the Gaussian random wave model. We compare the results of our simulations to analytic results on the ensemble and with similar systems that have been studied numerically, as well as calculating quantities that have not previously been studied.

Some of the contents of this section have been published previously as [1]. This includes in particular Sects. 3.1–3.6 and 3.8, though interspersed with new results and discussion.

Section 3.1 introduces the details of how we will carry out this investigation, which will largely involve our results from simulations of the 3-torus. Sections 3.2 and 3.3 look at the curvature and torsion of vortices in this system, considering them as geometrical space-curves. This concept is continued in Sect. 3.4 in which we observe how geometrical correlations determine the shape of the curve, but it emerges that correlations exist only on small scales and so we take a different approach to examining large scale quantities. Section 3.5 looks at vortex lines and the tangle itself as fractals, directly checking if they behave consistently across different large length-scales. Section 3.6 covers the similar concept of how the nature of the vortex tangle may lead to specific scaling relations in quantities such as the distribution of vortex lengths. Section 3.7 moves back to large scale behaviour, looking at some simple geometrical measures of how the vortex curves wind about themselves, which may fall into characteristic classes depending on the nature of their random construction. We conclude with brief discussion in Sect. 3.8, covering the implications of our observations and some possibilities for further investigation.

3.1 Introduction to Geometrical Statistics

The study of the geometry of vortex lines revolves around the behaviour of vortex curves on small lengthscales ($\lesssim \lambda$), at which the amplitude of a Gaussian random wavefield is a smooth complex function of position so the nodal lines are themselves smooth space curves. Previous works in the literature use this fact to calculate

statistical quantities such as the density of vortex lines per unit volume [2–4] and the probability distribution of the vortex lines' curvature [3], treating the wave field as an isotropic Gaussian random function as described in Sect. 1.4. Some quantities, such as the statistical distribution of the torsion of nodal lines, seems analytically intractable by these methods [5]. We address some of these problems in later sections.

At larger lengthscales, analysis of local geometry is not sufficient to determine analytically most quantities, and it is necessary to use numerical experiment. For instance, previous simulations of a nodal lines in random optical fields have indicated that at distances much larger than the correlation length (comparable to λ), a typical random nodal line looks like a Brownian random walk [6]. For this reason we will explore numerically the crossover in the geometric behaviour of the tangle of vortices between the local regime where the vortices are smooth space curves, and longer lengths where the lines and tangles display various measures of statistical self-similarity.

Although all three of our random wave systems are examples of Gaussian random modes whose limiting behaviour is the isotropic Gaussian random wave model, they are not all so equivalent at the energy scales we investigate (as summarised in Sect. 2.5). We anticipate that the 3-torus will be closest in behaviour to the true isotropic RWM, since its eigenfunctions arise purely from a constraint on the available wavevectors and are not additionally affected by the curvature of the 3-sphere or potential of the QHO; these not only push the system further from the limiting behaviour, but make numerical analysis of the local geometry more difficult since it is necessary to account for the metric in any calculation that is not highly local. The 3-torus is also the system with the longest vortex lines, since multiple segments within a periodic volume join at the boundaries to form one large line, and so is the only system in which analysis of truly large lengthscales is possible (thousands of wavelengths rather than hundreds at most for the other systems). This property of making large vortex loops accessible has also been the reason for using periodic boundary conditions in previous work on random optical fields [6, 7]. For these reasons, we will largely present results from simulations of the 3-torus throughout this Chapter, though we discuss our other random wave systems where it makes sense to do so.

Although we focus on eigenfunctions of the 3-torus, we utilise random eigenfunctions at different energies, including contrasting them where appropriate. In our study of smooth local geometry we use in particular simulations at energy 75 at a much higher resolution than any others, with at least 140 recovered vertices per wavelength in the piecewise-linear representation of the vortex curve. Such a high resolution is vital to accurately recovering subtle local quantities. Later Sections use higher energies sampled at much lower resolutions of more than 0.1λ per vertex-vertex step. The lower resolution is sufficient for the calculation of larger scale quantities.

Although the 3-torus is likely closest to the true random wave model following these arguments, the limited choice of wavevectors to enforce periodicity must still move the system away from the isotropic random limit. One way to gauge this effect is to compare the vortex densities in two and three dimensions with their

counterparts calculated analytically from the isotropic random wave model [3]: the ideal isotropic wave vortex tangle penetrates an arbitrary 2D plane at points with a density $2\pi/3\lambda^2 \approx 2.09/\lambda^2$, whilst in 3D the density of vortex *arclength* per unit volume is twice this, $4\pi/3\lambda^2 \approx 4.19/\lambda^2$. In our simulations with energy 1875, the 2D density is $(2.12 \pm 0.03)/\lambda^2$ in planes perpendicular to the axes of cubic periodicity, and $(2.14 \pm 0.03)/\lambda^2$ through planes passing through the diagonal of a periodic cell face. Both are consistent with the theoretical result, and isotropic within the error bounds. The 3D vortex densities from simulation are $(4.45 \pm 0.03)/\lambda^2$, $(4.61 \pm 0.04)/\lambda^2$ and $(4.49 \pm 0.03)/\lambda^2$ in cells with energies 75, 243, 1875 respectively, now not quite fully consistent with the theoretical result. We suspect this reflects the periodicity constraint, which forces vortex lines into loops more often than in the fully isotropic model (discussed further in Sect. 3.6), but does not appear significantly to affect our main results. All of these arclength densities are higher than the theoretical result, but do not strictly tend toward it as the energy is increased—it may be that other factors such as the distribution of wavevectors available are also important (though the strict number of wavevectors in our simulations increases with energy); we see a similar effect in the distribution of NTH lines, discussed in Sect. 3.6. We will make similar observations later when comparing geometrical quantities to analytical results from the literature, but the deviation is consistently small and so we believe the 3-torus well represents the isotropic random wave model.

We will also compare the properties of wave vortex tangles in 3-torus eigenfunctions with tangled filaments in other physical systems such as vortex filaments in turbulent superfluids, and polymer melts. Due to the ergodic mode hypothesis described in Sect. 1.4 they are also of interest by comparison with the high energy complex eigenfunctions of chaotic 3 Dcavities whose limiting behaviour appears to be the isotropic RWM, and for which the vortices in non-time-reversal-symmetric eigenfunctions [8] and resonances [9] have been extensively studied in two dimensions. The wave vortices are the 3 Dnodal counterpart to the nodal lines bounding nodal domains in the 2Dreal chaotic case, whose behaviour has also been well characterised [10–14]. Finally, they will be considered in the sense of true random walks and purely statistical random tangle as was introduced in Sect. 1.6.

3.2 Local Geometry of Random Wave Vortices

At the scale of a wavelength or less, we approach the geometric structure of the tangled melange of vortex filaments by considering the statistical shape of a single vortex line, ignoring the line’s global topology and the conformations of other lines nearby. On these scales, the shape of a vortex line is limited by the equation of which the wavefunction is an eigenfunction. In the case of the 3-torus, this is the Helmholtz equation. Although superoscillatory regions may in principle vary arbitrarily quickly [15] (in fact the phase structure sufficiently close to a vortex is already superoscillatory), such behaviour is unlikely to significantly affect a typical point on

a vortex. The vortices are thus smooth space curves, and as such can be parameterised by a *curvature* κ and *torsion* τ which uniquely determine its shape [16]. We define these formally below, but they may be thought of initially as how much the vortex is curving at a point, and how the plane of this curvature changes from point to point.

As natural measures of the geometrical tangling of a single space curve, κ and τ have been studied in other systems including superfluid vortices [17], the defect lines of liquid crystals [18], molecular systems in which strings of molecules may be approximated as 1D curves, and even generalisations to discrete lattice models [19]. In these systems the vortex geometry has direct physical implications; in superfluids, the local curvature and torsion are directly related to the dynamic conformation of a vortex filament, which in the so-called *local induction approximation* fully describes the motion of the vortex core [20]. In liquid crystals the curvature and torsion can be measured for singularities of the director field and can also be used here as components of a dynamical model [21], while in molecular systems such as polymers they are limited by the physical characteristics of the molecules but vary depending on the system [18]. This is not the case for our linear wave superpositions, in which vortices do not carry energy and are temporally static, and furthermore do not depend on the choice of any physical parameter except the wavelength.

Given a curve $f(t)$ for some parameter t , the formal definition of its curvature and torsion requires first reparameterising the curve by *arclength*, s . This is defined via an integral along its length,

$$s = \int_a^b \sqrt{|\dot{f}(t)|^2} dt, \quad (3.1)$$

where a and b are two points on the curve (or the same point if integrating all around a loop), and the arclength is parameterised by the function $s(t)$ for some parameter t [16]. For piecewise linear curves, this simplifies to the sum of piecewise linear segment lengths between the start and end points. We can then define the curvature and torsion, which occur as part of the *Frenet-Serret* formalism of space curves. This describes an orthonormal reference frame at every point on the curve, made up of the *tangent* \mathbf{T} , the orthogonal *normal* \mathbf{N} which points in the direction of the rate of change of tangent, and the *binormal* \mathbf{B} that completes the frame. The orthonormality of these vectors follows naturally from their definitions, beginning with

$$\frac{d\mathbf{T}}{ds} = \kappa \mathbf{N}, \quad (3.2)$$

where κ is defined to be the curvature and \mathbf{N} is the unit vector giving the direction of its change. \mathbf{N} must be orthonormal to the tangent \mathbf{T} since $\mathbf{T} \cdot \frac{d}{ds} \mathbf{T} = 0 = \kappa \mathbf{T} \cdot \mathbf{N}$.

The curvature thus naturally parameterises the rate of change of the tangent vector. The binormal \mathbf{B} arises naturally as $\mathbf{T} \times \mathbf{N}$, which must complete the orthonormal frame. These orthonormal components fulfil the relations

$$\frac{d\mathbf{T}}{ds} = \kappa\mathbf{N} , \quad (3.3)$$

$$\frac{d\mathbf{N}}{ds} = -\kappa\mathbf{T} + \tau\mathbf{B} , \quad (3.4)$$

$$\frac{d\mathbf{B}}{ds} = -\tau\mathbf{N} , \quad (3.5)$$

where τ is the torsion which gives the rate of change of the plane in which the vortex curves.

As written, these relations all relate to space curves parameterised by arclength s . In practice, our vortices come from simulation and are parameterised as a series of piecewise linear edges of unequal lengths. If we label the space curves as $\boldsymbol{\gamma}(t)$, where t is this parameterisation, we obtain appropriate derivatives via the chain rule,

$$\dot{\boldsymbol{\gamma}} = v^2\kappa\mathbf{N} + \dot{v}\mathbf{T} \quad (3.6)$$

$$\ddot{\boldsymbol{\gamma}} = v^3\kappa\tau\mathbf{B} + (2v\dot{v}\kappa + v^2\dot{\kappa})\mathbf{N} + (v\dot{v} + \ddot{v})\mathbf{T} , \quad (3.7)$$

where $\dot{\boldsymbol{\gamma}}$ denotes a derivative in t and v is a *velocity* arising from $\dot{\boldsymbol{\gamma}} = v\mathbf{T}$ and so $v = |\dot{\boldsymbol{\gamma}}|$. Using the orthogonality properties of the vectors \mathbf{T} , \mathbf{N} and \mathbf{B} , it is straightforward to show that

$$\kappa = \frac{|v\ddot{\boldsymbol{\gamma}} - \dot{v}\dot{\boldsymbol{\gamma}}|}{v^3} \quad (3.8)$$

$$\tau = \frac{\dot{\boldsymbol{\gamma}} \times \ddot{\boldsymbol{\gamma}} \cdot \ddot{\boldsymbol{\gamma}}}{|v\ddot{\boldsymbol{\gamma}} - \dot{v}\dot{\boldsymbol{\gamma}}|^2} . \quad (3.9)$$

It is these relations that we use to recover κ and τ from our numerical simulations, with derivatives of $\boldsymbol{\gamma}$ being replaced by appropriate discrete forms based on vertex indices.

We will also calculate both κ and τ directly via Monte Carlo simulation on the Gaussian model. For these calculations it is necessary to obtain the Frenet-Serret framing from the values of the field, for which we use the vorticity $\boldsymbol{\omega} \equiv \frac{1}{2} \text{Im} \nabla\psi^* \times \nabla\psi$. This vector points along the direction of the vortex line, i.e. $\boldsymbol{\omega} = \omega\mathbf{T}$, and so with an appropriate normalisation it gives all the information necessary to calculate the Frenet frame; in (3.6)–(3.9), $\dot{\boldsymbol{\gamma}} = \boldsymbol{\omega}$, $v = \omega$ and $d/dt = \boldsymbol{\omega} \cdot \nabla$. Thus, the following expressions for curvature and torsion are obtained by direct substitution into (3.8) and (3.9) (evaluated along nodal lines) [5]

$$\kappa = \frac{|\boldsymbol{\omega} \times (\boldsymbol{\omega} \cdot \nabla)\boldsymbol{\omega}|}{\omega^3} , \quad (3.10)$$


$$\tau = \frac{\boldsymbol{\omega} \times (\boldsymbol{\omega} \cdot \nabla)\boldsymbol{\omega} \cdot (\boldsymbol{\omega} \cdot \nabla)^2\boldsymbol{\omega}}{|\boldsymbol{\omega} \times (\boldsymbol{\omega} \cdot \nabla)\boldsymbol{\omega}|^2} . \quad (3.11)$$

The curvature therefore depends on second derivatives of the field, and the torsion on third derivatives. The numerator of κ occurs in the denominator of τ , suggesting a statistical anticorrelation discussed further below.

The curvature and torsion may be informally understood in terms of a local fit to a *helix* of radius r and pitch c : $\mathbf{a}(a) = (r \cos(t), r \sin(t), ct)$ (a helical vortex may be realised in a wave of the form of a perturbed screw dislocation [5, 22]). An example is shown in Fig. 3.1, with the Frenet frame marked. The curvature and torsion of the helix are the same at all points: $\kappa = r/(r^2 + c^2)$ and $\tau = c/(r^2 + c^2)$. The local radius of curvature is $1/\kappa > r$; in the limit $c \rightarrow 0$, $\kappa \rightarrow 1/r$ and the curve is locally a planar circular arc. In the opposite limit $r \rightarrow 0$, the helix becomes a straight line and $\tau \rightarrow 1/c$. This limit is therefore singular (the straight line has an arbitrary value of torsion, which is technically defined only when $\kappa > 0$). This suggests that when physical considerations do not bound τ , it can have very large fluctuations and is numerically ill-behaved when the curve is locally almost straight. Such a situation can certainly occur in vortices, and we will see the consequences of this in the next section.

Figure 3.2 shows a simulated volume of vortex tangle taken from a region of an energy 75 3-torus eigenfunction, in which the vortex lines are coloured according to κ and τ . Both quantities change smoothly, but exhibit rapid changes on the sub-wavelength scale. By eye the curvature is usually small, but with localised regions where the curving is much sharper, particularly where pairs of vortex lines approach closely. The torsion is also often small with strongly concentrated regions, but these tend to occur in the less sharply curved segments of the vortex lines, indicating that the plane of curvature often rotates about the vortex axis when the local curvature is small. This is consistent with the anticipated anticorrelation of (3.11).

Fig. 3.1 The Frenet Serret framing of \mathbf{T} , \mathbf{N} , \mathbf{B} illustrated via the fit to a helix of radius r . The curvature is a local fit to the *circle* with the radius of the helix, while the torsion relates to the out-of-plane movement






Fig. 3.2 Vortex lines in random wave tangle, coloured by their curvatures and torsions. The vortex lines are from a simulation of the 3-torus with energy 75 and are coloured by by curvature in **(a)** and torsion in **(b)**, determined numerically. Both **a**, **b** show the same the same $(2.17 \lambda)^3$ volume of one of the cells described in Sect. 2.5, with the axis ticks marking 0.5λ increments. Figure from [1]

3.3 Statistics of Curvature and Torsion

We address here the statistical distributions of curvature and torsion, comparing our numerical model to analytic expectations where available. The curvature probability density function (PDF) for vortex lines in the continuum random wave model was found in [3] to have the Cauchy-like distribution

$$P(\kappa) = \frac{3^{5/2} \kappa \kappa_c^3}{(3\kappa_c^2 + \kappa^2)^{5/2}}. \quad (3.12)$$

The *characteristic curvature* κ_c is a constant depending on moments of the wavevector power spectrum, with k_n denoting the n th moment. It is defined as [5]

$$\kappa_c = \sqrt{\frac{9k_4 - 5k_2^2}{45k_2}}. \quad (3.13)$$

As we are considering eigenfuntions of the 3-torus, all wavevectors have the same magnitude and the power spectrum is a radial delta function. In this special case, $k_n = k^n$ such that $\kappa_c = 4\pi/\sqrt{45}\lambda \approx 1.87/\lambda$ (but this would be different in e.g. a polychromatic Gaussian random light field [5]). Thus, the peak of the distribution is at $2\pi/\sqrt{15}\lambda \approx 1.62/\lambda$, and its first two moments are

$$\langle \kappa \rangle = \sqrt{3}\kappa_c \approx 3.24/\lambda, \quad (3.14)$$

$$\langle \kappa^2 \rangle = 2\langle \kappa \rangle^2 = 6\kappa_c^2 \approx 21.1/\lambda^2, \quad (3.15)$$

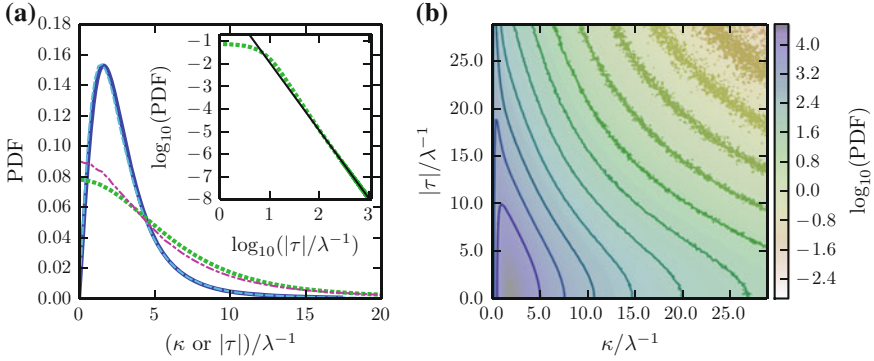


Fig. 3.3 Curvature and torsion PDFs from the Gaussian random wave model and numerical simulations of random waves. **a** Shows PDFs for curvature and torsion from simulations of the 3-torus with energy 75, as well as from analytic results and Monte Carlo integration. **b** Shows the joint PDF $P(\kappa, \tau)$ drawn from Monte Carlo integration. In **a**, $P(\kappa)$ is determined by (3.12) (—), and from the random wave simulations (---). The torsion PDF is found by Monte Carlo integration (⋯), and retrieved simulations (---). The inset to **(a)** is a double-log plot indicating how τ scales at larger values. All Monte Carlo results come from 6×10^{10} points. Figure from [1]

whereas higher moments diverge. The values of these moments indicate that the geometrical structure lengthscales associated with vortex curvature are certainly sub-wavelength.

The curvature PDF can also be extracted from our numerical simulations, following the geometrical recovery methods of Sect. 2.2 and using a discrete approximation for the Frenet-Serret parameterisation as in the previous section. The distribution recovered this way is shown in Fig. 3.3a, along with the analytical result of (3.12). The visual agreement is very good; the main discrepancies are an oversampling at small κ , a small shift of the peak to lower κ , and an undersampling at high κ . All three of these are remnants of the smoothing and downsampling procedure on the linear piecewise representation of the curves, which will tend to towards underestimating the curvature as the price for correctly recovering the bulk of the distribution.

More quantitatively, the peak of the numerically recovered distribution is at $\kappa = (2.71 \pm 0.05)/\lambda$, a shift of 3.4% from the analytic result. Its first two moments are

$$\langle \kappa \rangle = (3.18 \pm 0.06)/\lambda, \quad (3.16)$$

$$\langle \kappa^2 \rangle = (18.9 \pm 0.4)/\lambda^2 = 1.87 \langle \kappa \rangle^2, \quad (3.17)$$

$$1/\sqrt{\langle \kappa^2 \rangle} = (0.230 \pm 0.002)\lambda, \quad (3.18)$$

all comparable to the analytic results but consistently shifted to slightly lower curvatures, with the surprising exception that the expected relationship $\langle \kappa^2 \rangle = 2 \langle \kappa \rangle^2$ is relatively strongly affected by the shift. This result is consistent with our checking in Sect. 3.1 that the 3-torus well represents the isotropic random model. It is also

provides further support for our choice of smoothing parameters in Sect. 2.2, the deviation from the analytical result is consistent with the difference in recovered curvature as a result of different smoothing parameter choices.

We also observe that curvature is not statistically equally distributed amongst all loops, the shortest have a very specific puckered ellipse shape in which the curvature is higher than average because the loop has a fairly uniform high curvature in order to close on a small scale while satisfying the smoothness of the field. These loops are anticipated in [23] as the small scale limiting behaviour of any vortex curve, though they are few in number and the statistics of curvature would be barely changed if they were ignored. The probability distribution of loop lengths is discussed in Sect. 3.6.

The torsion of a nodal line in a random wave field is less easy to interpret than its curvature. As can be seen in Fig. 3.2, τ varies rapidly over lengthscales comparable to or shorter than those of κ and, as discussed above, may be large even where a vortex filament is almost straight—it describes instead the local twisting of the Frenet frame about the tangent axis, which (unlike a change in tangent) is not clear to the eye.

The statistical properties of the torsion of vortex lines have not explicitly been covered previously (except for brief comments in [3, 5]), as the analytic form of τ in (3.11) is significantly more complicated than that of the curvature κ . It depends on cubic combinations of fields and their derivatives, rather than quadratic combinations as for curvature, which poses challenges for analytic evaluation by Gaussian techniques. Furthermore, the torsion depends on the first, second and third derivatives of field quantities. Since the third derivatives are correlated with first derivatives following the Helmholtz equation, it is not possible to use the same simplifications as in the curvature calculation—for which a choice of symmetry and coordinate system may eliminate many of the terms [3].

We instead consider the torsion PDF via Monte Carlo integration on the ensemble of the Gaussian model. The method of integration follows [3, 5] and was introduced in Sect. 1.4; for torsion specifically, the result is that the PDF is given by

$$P(\tau) = \langle \tau(\mathbf{V}) \delta(\tau(\mathbf{V}) - \tau) P(\mathbf{V}) \rangle_{\mathbf{V}}, \quad (3.19)$$

where the average is a dislocation average resulting from selecting points on vortex lines such that the real and imaginary field components are zero and the vector of Gaussian random field quantities \mathbf{V} consists of all relevant field derivatives from (3.11), in this case everything up to the third derivatives. The probability $P(\mathbf{V})$ includes all of their correlations and correctly-weighted variances which account for sampling only along vortex lines; in this sense, the primary random variables are the independent Gaussian distributed quantities derived from the vector of field derivatives \mathbf{V} , not the derivatives themselves. Evaluating this quantity by Monte Carlo integration means numerically selecting points at random from this multidimensional Gaussian distribution, evaluating the vorticity ω which is a function of the zeroth and first derivatives, and passing this result to (3.11) to retrieve the resulting τ . We repeat this procedure many times, and evaluate probability density distributions as histogram averages of these results.

Such a Monte Carlo torsion PDF is shown alongside that of curvature in Fig. 3.3a, calculated from 6×10^{10} numerical sample points. Its shape is distinctly different to that of the curvature as it is a signed quantity, symmetric about $\tau = 0$, which is also the peak of the distribution. The first moment of the unsigned distribution is

$$\langle |\tau| \rangle = (6.19 \pm 0.09)/\lambda . \quad (3.20)$$

The torsion PDF decays much more slowly than that of the curvature, reflecting its quantitative instability when $\kappa \approx 0$, in which regions of the curve the plane of curvature may rotate rapidly without affecting the local smoothness of the curve. The high- τ scaling of the distribution is also shown in Fig. 3.3a, inset, with linear fit $\log_{10} P(\tau) = (-3.01 \pm 0.01) \log_{10}(\tau)$, suggesting that as $\tau \rightarrow \infty$, the probability density scales as τ^{-3} . We note that this limit is reached only when $\tau \gg \lambda$; the fit becomes stable only for $\tau \gtrsim 100/\lambda$. If by analogy with κ_c in (3.13) there is some characteristic torsion τ_c determining the characteristic scaling of the analytic distribution, then $\tau_c \gg \kappa_c$.

The τ^{-3} scaling means that the second and higher moments of τ diverge, as suggested to be likely in [5]. It is hard to identify by eye what the analytic form of $P(\tau)$ might be; a range of PDFs with different analytic forms (including Cauchy-like distributions) occur as the distribution of various measures of vortex core twist and twirl [5, 24]), but we have not been able to directly fit any of these families with *ad hoc* parameter choices to the Monte Carlo torsion distribution.

We also recover the torsion from our simulations, with the resulting PDF shown in Fig. 3.3a. The fit to the Monte Carlo result is again reasonable; as with the curvature, it is shifted towards $\tau = 0$ as information is lost in the smoothing and resampling procedure. This shift is more pronounced here, reflecting the numerical instability in measuring the torsion; since it may vary significantly when the line is almost straight, it is highly sensitive to numerical noise of any kind. In practice it is not feasible to recover the torsion PDF more accurately.

The first moment of the recovered distribution is

$$\langle |\tau| \rangle = (6.00 \pm 0.06)/\lambda . \quad (3.21)$$

Unsurprisingly, the visual shift to low torsions translates to a lower moment, and the match to the Monte Carlo integral is less good than that of the curvature to the analytic result. However, despite this relatively large discrepancy between the Monte Carlo and numerical distributions, the shift represents a change of only around 3%. This is because the tail of the distribution is overrepresented in the recovered numerical distribution, a remnant of the high torsions that result from the initial lattice structure of the numerics, even though these are largely smoothed away.

The Monte Carlo integration also makes accessible the joint curvature-torsion PDF, shown in Fig. 3.3b. It is consistent with the apparent distributions of Fig. 3.2 and with the anticipated anticorrelation of κ and τ . In fact, the curvature PDF is almost recovered on the symmetry line $\tau = 0$, with only a small shift to the higher

curvatures that make low τ statistically likely. In contrast, the torsion distribution is very poorly recovered when $\kappa = 0$, since κ occurs directly in the denominator of (3.11).

3.4 Geometric Correlations

The probability density functions of Sect. 3.3 are visibly consistent with what is seen by eye in Fig. 3.2, but they do not describe how curvature and torsion vary locally *along* a vortex line. We consider here some simple statistical lengthscales that quantify this local behaviour, describing the curve lengths over which geometry quantities become decorrelated and the new local behaviour becomes unpredictable.

A natural lengthscales measurement for random curves is the *persistence length* L_p , which quantifies the distance along the curve over which correlations in the direction of the tangent are lost [25, 26]. It is a standard quantity in the physical analysis of the segmented chains of polymer physics, in which case it may be analytically accessible through physical properties of the polymer chain [26], though it is also used more generically in classifying the behaviour of random walks [27]. Our vortex curves are confirmed as random walks later in Sect. 3.5, though unlike in the polymer case their persistence relates to the smoothly changing tangent direction rather than the inherently discrete bonds between molecules; instead of physical constraints, the persistence length is limited by the smoothness of the field.

The persistence length is defined from the approximately exponential decay of the *direction cosine correlation* in random walks; the correlation of tangent vector directions $\mathbf{T}(s)$ with respect to arclength s , with

$$\langle \mathbf{T}(s) \cdot \mathbf{T}(s + \Delta s) \rangle \overset{\infty}{\sim} \exp(-\Delta s/L_p), \quad (3.22)$$

where $\mathbf{T}(s')$ is the tangent to the curve at the point labelled by s' , parameterised by arclength. The relevant Δs scale here is typically longer than any relevant lengths that can be extracted from the curve analytically by Taylor expansion, unlike curvature and torsion which are derived at a single point via high derivatives.

This direction cosine correlation is plotted in Fig. 3.4. It has approximately the exponential shape of (3.22), but with deviations at both ends. First, as $\Delta s \rightarrow 0$ the correlation function is not exponential but instead determined by the mean square curvature through application of the Frenet-Serret relations of (3.3), giving

$$\langle \mathbf{T}(s) \cdot \mathbf{T}(s + \Delta s) \rangle = 1 - \frac{1}{2} \langle \kappa^2 \rangle (\Delta s)^2 + O(\Delta s)^4. \quad (3.23)$$

Extracting such a fit from Fig. 3.4 suggests $\langle \kappa \rangle = (3.2 \pm 0.1)/\lambda$, which is consistent with analytic result of (3.16). This parabolic fit persists to $\Delta(s) \approx 0.35 \lambda$, after which the fit begins to better approximate an exponential.

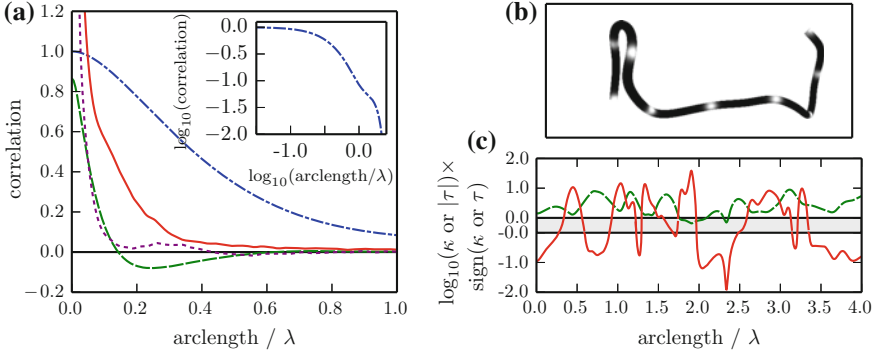


Fig. 3.4 Variation of direction cosine, curvature and torsion with respect to arclength s along random vortex lines. **a** Shows correlation functions of direction cosine (as in (3.22)) (—), curvature correlation minus one (3.24) (---), signed torsion correlation (3.25) (····) and unsigned torsion correlation minus one (—). **b** Shows a typical short vortex segment taken from our numerical experiments, coloured with a white dot at every 0.5λ arclength distance. **c** Shows the varying curvature and torsion along the segment of (**b**). The scale of (**c**) is the signed logarithm $\log_{10}(\kappa \text{ or } |\tau|) \times \text{sign}(\kappa \text{ or } \tau)$, except in the shaded area which tracks the raw κ or τ between $\pm 1/\lambda$. Figure from [1]

The direction cosine correlation also lacks a simple exponential shape at higher arclengths beyond $\Delta s \approx \lambda$, with the correlation instead decaying less rapidly than anticipated in (3.22). We explain this as arising from the vortex curve *not* being a true random walk on the scale of a wavelength, with the smoothness of the field enforcing small but persistent correlations at larger distances; in a sense, the curve is random but still strongly feeling the effect of the rest of the field, which means that though its tangent would be expected to decorrelate following the local smoothness of the field, this assumption does not necessarily hold at intermediate scales. This effect is not too strong, the curve remains roughly exponential and appears to totally decorrelate at larger scales. The width of the decay function still gives an estimate for the correlation length scale similar to the ideal persistence length; its second moment is $(0.590 \pm 0.002)\lambda$, and so tangents become totally decorrelated over roughly the wavelength scale.

The persistence length is not the only correlation distance to characterise random curves. We can also compare it directly against that of the curvature and torsion, i.e. the correlation functions

$$C_{\kappa}(\Delta s) \equiv \frac{\langle \kappa(s)\kappa(s + \Delta s) \rangle}{\langle \kappa(s) \rangle^2}, \quad (3.24)$$

$$C_{\tau}(\Delta s) \equiv \frac{\langle \tau(s)\tau(s + \Delta s) \rangle}{\langle |\tau(s)| \rangle^2}. \quad (3.25)$$

These correlation functions are also shown in Fig. 3.4a. The curvature correlation is well-behaved, with vertical intercept matching (3.17). The torsion is noisier, and diverges as $\Delta s \rightarrow 0$ since its denominator is not defined in this limit.

As with the direction cosine correlation, the decay lengths of (3.24) and (3.25) give an indication of the length of vortex segment over which curvature and torsion become decorrelated. Both decay over a distance much shorter than a wavelength or the tangent correlation lengthscale; in the case of the curvature, the second moment of $(0.055 \pm 0.001)\lambda$ sets a rough lengthscale of correlation. Since the torsion correlation diverges, we instead fit its shape beyond $\Delta s = 0.1$ to an exponential with decay length $(0.24 \pm 0.03)\lambda$.

An alternative measure of the torsion correlation distance comes from the *unsigned* torsion correlation, with τ replaced by $|\tau|$ in (3.25), which is also shown in Fig. 3.4a. As with the torsion, it diverges as $\Delta(s) \rightarrow 0$, but beyond $\Delta(s) \cong 0.04 \lambda$ it fits well to an exponential with decay length $(0.08 \pm 0.01)\lambda$. This sets a correlation lengthscale for torsion similar to that of the curvature.

These correlation curves and their associated lengthscales describe different features of the random filament, matching what is seen in Fig. 3.2. The persistence length quantifies the arclength interval of approximately fixed tangent direction along the curve as it undulates; this is the length along which a typical vortex line is (almost) straight. In between these sections are shorter intervals, whose length is the curvature correlation length, at which this direction changes rapidly. There are other potentially even shorter intervals of high torsion (often when the curve is nearly straight) where the plane of curvature varies rapidly—including jumping rapidly from positive to negative.

Figure 3.4b–c show a segment of a vortex curve randomly sampled from our numerical experiments, along with the curvature and torsion along its length, demonstrating how the correlation functions are expressed. Again, both exhibit small scale peaks, with widths following the correlation lengths above. These peaks are often strongly related to regions of hyperbolic interchange where two vortices approach closely, the hyperbolic analogues of the elliptic small loops previously mentioned to occur at small scales [23, 28, 29]. As with the loops, the smoothness of the field demands that these regions be highly curved but almost planar.

3.5 Three Measures of Fractality

In the previous section, we considered the geometry of individual vortex lines in random waves at a scale of a wavelength or less, and found good agreement at this local level between our numerical simulations of the 3-torus and the isotropic Gaussian random wave model. However, such an approach cannot be directly extended to describe the geometry of the vortex tangle at scales much beyond the wavelength. Local methods based on Taylor expansion require too many terms to determine what quantities should be calculated, and even the probability distribution of torsion seemed analytically intractable. We instead investigate the scaling of the tangle on large scales; is it self-similar, and what is the *fractal dimension* of the vortex curves?

Intuitively, the idea of a fractal dimension quantifies how the complexity of a pattern changes through different lengthscales. If a pattern has a consistent behaviour through different lengthscales then it may be called *self-similar*, and may have a fractal dimension different from its topological dimension. A well-known example is the *Koch snowflake*, formed by beginning with an equilateral triangle and progressively replacing the centre of each side with a further triangle, progressing infinitely such that the surface exhibits identical protrusions at all scales. Although the boundary of the shape is topologically one-dimensional, it has fractal dimension $\log(4)/\log(3)$ [30]; intuitively, the fractal may be thought of as ‘more space-filling’ than a normal line. Most importantly for us, the concept can extend to random systems. For instance, an unbiased random walk in the plane has discrete fractal dimension 2; in the fractal sense, it fills both dimensions [31].

Fractality has been well investigated in physical models involving tangles of random filaments. For instance, amongst some of our previous examples the discrete vortices in the \mathbb{Z}_3 Cartesian lattice model behave as Brownian random walks [32], and similar results are found for vortex tangles in random optical waves [6], including Brownian fractal scaling in the radius of gyration of vortex loops [33] (defined later). Such behaviour is also characteristic in other physical systems such as polymer melts, in which it has long been suggested that a chain in a polymer melt scales as a random walk (the *Flory ideality hypothesis*) [34, 35], though in this case the tangle is of molecular chains rather than field defects (and in reality, other physical considerations may break this assumption [36]). These results all apply to individual filamentary strands, but it is also possible to determine the fractal dimension of the full vortex tangle. One simple quantity is the box-counting fractality (described later in this section), as measured in simulations of evolving superfluid turbulence [37], itself an extension of calculations for vortices in a classically turbulent field [38]. We compare all of these quantities with those of the vortex tangle in eigenfunctions of the 3-torus, verifying that the fractal scalings of our vortices at long lengthscales are consistent with random walk behaviour, and potentially with global similarity to other tangles. These results are thus the 3D counterparts to the numerous random walk and percolation properties discussed for nodal lines in real 2D random waves as a model for quantum chaos [12].

These fractal measures cover a wide range of length scales, and so the nature of the boundary conditions may affect the result. Following Sect. 1.7.1, the total vortex line length per simulated cell is finite and individual filaments are always closed, but they may wrap around the periodic boundaries one or more times before doing so, potentially having non-trivial homology. Both these lines and long loops may span many cells if we consider the periodic boundary conditions as relating to an infinite tiling of identical copies. Taking a single period as the total arclength of a NTH line, the total length of individual vortex lines varies from less than a wavelength in the smallest loops to the longest lines whose length is of comparable to the total arclength in the cell. These longest lines can have total arclengths of over 40000 λ in 3-torus eigenfunctions with energy 1875, a side length of just 21.7 λ . We anticipate that closed loops and NTH lines will have identical fractal characteristics on the scale of the periodic cell, but it is unclear if this should necessarily persist at the largest

scales, or if even closed loops may be affected beyond the periodic scale. For these reasons, we will be careful in our analysis to compare fractalities both below and above the scale of a single periodic cell, as a further test of whether the periodicity affects this result.

The self-similarity of a single vortex curve can be determined by comparing, for one point on the curve fixed and the other varying along the curve, the arclength s between the pairs of points to the (unwrapped) Pythagorean straight line distance R [6, 32],

$$\langle R \rangle \sim C s^n, \tag{3.26}$$

where C is a constant and n^{-1} is the fractal dimension of the line: by this measure, $n = 1$ for a smooth line, $n = 1/2$ for a Brownian random walk, and $n = 0.588$ for a self-avoiding random walk [31, 39]. The range at which (3.26) applies should be larger than the persistence length of the curve (it is smooth at smaller distances), and less than the overall length L of the curve (beyond which it repeats, either as a loop or a periodic line).

Figure 3.5 shows this scaling for our simulated vortex filaments from 3-torus eigenfunctions of energy 1875. The result for each vortex line is an average over 1500 starting points on each curve. This gives a characteristic scaling relationship for each curve, producing a straight line fit as expected even for each individual

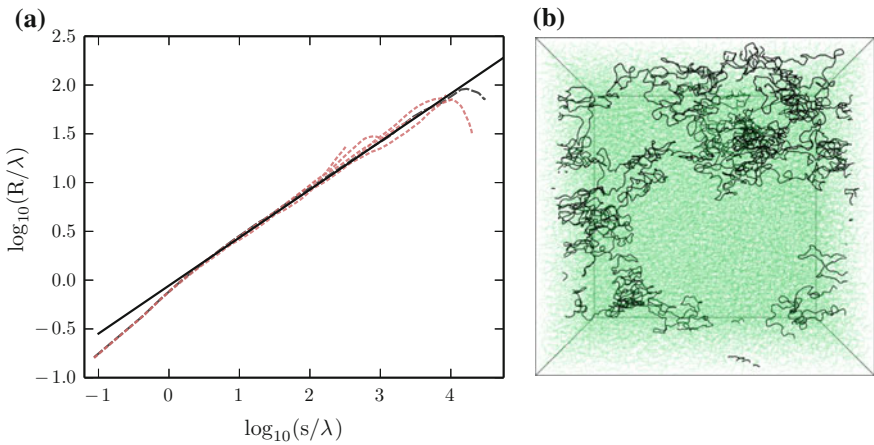


Fig. 3.5 The scaling of straight line distance between vortex points against the arclength distance between them. **a** The log of the straight line distance R , averaged over 1500 pairs of points on the vortex curve separated by arclength s , against $\log_{10}(s)$. Between $0.5 \leq \log_{10}(s) \leq 1.3$ the curves are fit to the same straight line (—) arising as the average gradient from 3000 individual curves, shown as the continuous *black line* with gradient 0.504 ± 0.002 . Also shown are 5 examples of these curves (---), alongside a single much longer curve (---) almost three orders of magnitude longer than the fitted region but with the same fractal scaling persisting at larger scales. In **b**, a sample periodic line of length 401λ (in the middle of the length distribution), wrapped within its periodic cell, with several other lines and hundreds of loops (*green*). Figure from [1]

vortex line, since they are long enough that the average over starting points on their length properly samples all different local behaviours.

To limit contamination from (periodic) boundary effects, we initially consider a range of R constrained by the periodic side length $0.5 \leq s \lesssim 21.7\lambda$, and averaging over 3000 curves with a range of overall lengths. The sample in Fig. 3.5b shows an example of length 400λ ; this is relatively short amongst our samples, but clearly is already long and complex. Over this scale we find $n = 0.504 \pm 0.002$, close to the expected 0.5 and indicative of Brownian fractality, similar to the results of [6]. The lowest lengthscale for the Brownian scaling is around $s \approx 3.5\lambda$ (somewhat higher than the persistence length $L_p = 0.590\lambda$ of Sect. 3.4), indicating that even when $s > L_p$, the smoothness of the field affects the line's statistical self-similarity and matching the non-random behaviour of the tangent correlation in this region. The scaling at shorter lengthscales gives $n \approx 0.63 \pm 0.02$, although it is not clear what determines this exponent; our sampling resolution within these cells does not permit a resolution lower than about 0.1λ .

We also fit to 10 of the longest lines (arclength greater than 10000λ) over their full length, not just within the span of a single cell; these vortex curves have lengths of at least $10^4\lambda$. The scaling exponent for these lines is $n = 0.493 \pm 0.03$ over three decades (significantly longer than the side length of the cell), with a single example shown in Fig. 3.5a. This scaling behaviour continues the initial trend and appears highly typical in periodic cells of different sizes, limited only by the periodic side length.

The characteristic hooked shape at the longest scales is a consequence of the periodic boundary conditions and the fact that cells tend to contain only one such very long line. This may be understood heuristically as follows. Each periodic cell must have the same vortex flux entering and leaving opposite faces. This constrains the paths of lines with non-trivial homology; for every passage through the periodic boundaries, the same line or a different one must make an equal but opposite passage for the total flux to sum to 0. An NTH line thus cannot typically roam far from any reference point, as the other lines in the cell must counter its flux, but are typically much shorter and so could not compensate if the line had travelled too far. The long NTH line therefore cannot venture far and instead behaves almost like a loop, with a given periodic section typically returning to a cell near its starting point. This manifests as the observed hook shape in both loops and NTH lines; in both cases, the line ultimately closes with a gap of length 0, or of negligible size on the scale of the NTH line length. This hook is actually not quite characteristic of all NTH lines; in cells with several shorter lines, none are strongly constrained in this way, and may display random walk fractality over their full length. These observations are illustrated in Fig. 3.6, plotting for each NTH line its overall Pythagorean distance separation between the start and end points of its periodic length against this length. As expected, the results roughly match the expected square root due to their random walk behaviour, except that at the highest scales NTH lines are increasingly forced to almost close and so no longer match the theoretical expectation.

Overall, this measure random walk fractality seems generic over a range of scales including well above the side length of a periodic volume, but breaks down at the very

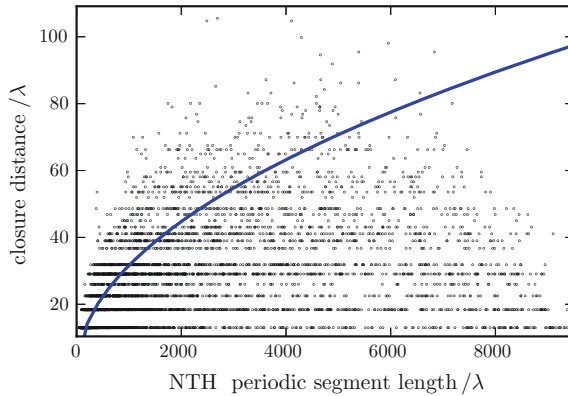


Fig. 3.6 Straight-line closure distances of vortex lines with non-trivial homology plotted against their lengths. Each point shows this result for a single NTH line, with 5372 samples from 2145 eigenfunctions. The discretisation is due to the discrete number of periodic cells a line may have traversed. Lengths refer to the length of a single periodic section of the line, and the *blue curve* shows the square root of this

largest lengthscales due to the constraints imposed by periodicity. Since this imposes a termination condition on individual lines, it seems that in the unbounded case every vortex line will behave as a random walk between the initial cutoff at around 3.5λ and the scale of its own maximum extent. The distribution of line lengths in the general case is examined in Sect. 3.6.

An alternative method of extracting the fractal self-similarity of *closed* loops is to compare the loop's overall length L against its radius of gyration (ROG) r_g , the root mean square distance of all points on the loop from its centroid. We have already noted that at the smallest scale, an analysis based on local Taylor expansion shows that small vortex loops typically resemble small puckered ellipses; this implies that the radius of gyration scales approximately linearly with length, i.e. $r_g \propto L$, as previously mentioned to be an expected result of the small scale smoothness of the field [23]. However, for longer loops to have the characteristics of Brownian walks we expect that $r_g \propto L^{1/2}$ [33]. A log-log plot of loop length against radius of gyration for simulated vortex loops is shown in Fig. 3.7, clearly demonstrating these two regimes, again with a rather sharp transition at $L \approx 3.5 \lambda$. These loops are taken from cells of side length 7.8λ ; the smaller size is numerically optimal for tracing vortices in many cells in parallel, necessary since the radius of gyration cannot be averaged across a single loop, so many data points are necessary to check the scaling.

The fit beyond $L \approx 3.5 \lambda$ has gradient 0.52 ± 0.01 , consistent with the 0.5 expected for random walks, though there is wide variation amongst individual loops. Below the transition the fit is instead to a gradient of 0.98 ± 0.01 , consistent with the scaling of simple ellipses. Figure 3.7b–g shows examples of randomly selected loops across all these scales; small loops such as (b) demonstrate the small scale ellipse shape required by the smoothness of the field, while approaching the gradient cutoff

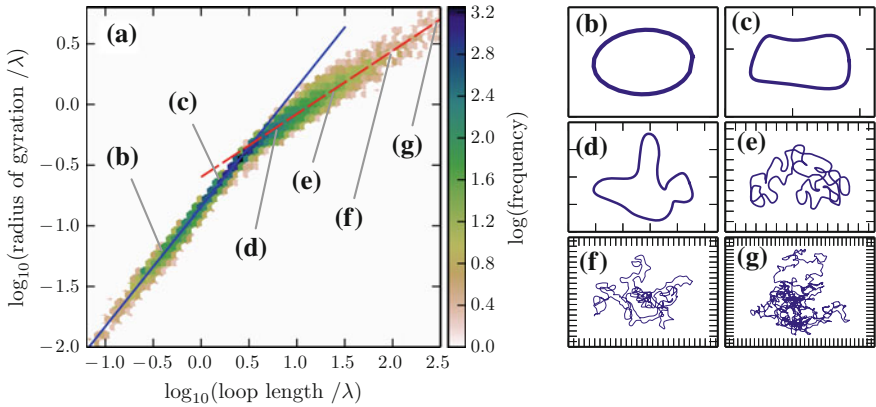


Fig. 3.7 Scaling of random vortex loops. **a** Log-log plot of radius of gyration r_g against loop length L , for 39487 loops from 650 simulated 3-torus eigenfunctions of energy 243. The fit for the lower values (—) has gradient 0.98 ± 0.01 , while the fit for higher values (—) has gradient 0.52 ± 0.01 , with the cutoff between regimes at $(3.5 \pm 0.1)\lambda$. **b–c** Show examples of loops randomly selected at different length scales. Figure from [1]

introduces higher order terms that disrupt this shape as in (c). Beyond the cutoff the loops become tangled, no longer limited by the local smoothness of the field.

This ROG fit applies at almost all lengthscales, but it is not clear if the small number of the very longest loops also fit the pattern. We do not investigate this here, as the number of such loops is not significant and their geometry may be relatively strongly affected by the imposition of periodic boundaries (through which they pass many times). However, the distinction is relevant and discussed further in Sect. 5.7. The overall result strongly matches that of the Pythagorean distance; vortex loops behave as random walks beyond a lengthscale of 3.5λ .


The preceding scaling measures apply to vortex curves considered in isolation. It is also possible to examine the self-similarity of the tangling *neighbourhood* of any point in the vortex field by means of the Minkowski dimension of the tangle, otherwise known as the box-counting dimension, and closely related to the mathematical Hausdorff dimension [30]. This is obtained by tiling the cell with *boxes* of side length δ , and considering the number of boxes $N_b(\delta)$ containing a vortex line segment as a function of δ . As above, we consider this in the range of large δ , as the tangle is smooth in the range $\delta \approx \lambda$.

Following [38], the box counting fractal dimension n_b is defined as the scaling exponent in the expression

$$N_b(\delta) \propto \delta^{-n_b} \quad (3.27)$$

over some range of δ (which we expect to be larger than λ); in the limit $\delta \rightarrow 0$, N_b approaches the number of vortex sampling points; when $\delta \lesssim \lambda$, $n_b \approx 1$, since the vortex tangle consists of lines. However, when δ is very large (perhaps around the

Fig. 3.8 Box counting fractal dimension of vortices in 3-torus eigenfunctions of energy 1875. The fit is to the gradient 2.90 ± 0.01 , averaging over 10 simulation cells with $N = 25$. Figure from [1]



average vortex-vortex spacing), n_b must simply equal 3 as every box will include at least one vortex segment. Any fractal region, if it exists, occurs between these two extremes.

The box-counting dimension is therefore a measure of the proportion of space (at a scale given by δ) which is filled by the curves. The box counting scaling for our numerical tangles is shown in Fig. 3.8; the fit is to $n_b = 2.90 \pm 0.01$ across a full decade of δ : clearly $n_b = 3$ indicates space is homogeneously filled at these scales. This space-filling regime in fact begins at around $\delta \approx 0.6 \lambda$, well below the fractality scale of the other measurements and roughly on the scale of the persistence length—thus, the range in which the box-counting dimension is unity appears negligible.

In fact, this lower limit of 0.6λ is close to the reciprocal square root density of points where vortex lines cross an arbitrary plane, of $\sqrt{3/2\pi}\lambda \approx 0.69\lambda$; this is a measure of the mean spacing of vortex points. The fact that the space-filling regime begins at this scale seems to be an indicator of the rigidity and regularity of random fields satisfying the Helmholtz equation [14, 40]. This property has been related in 2D to the infinite screening length of the vortex points as topological charges [3, 41]; the vortex points do not in general approach closely, and so this lengthscale limits their small-scale fractality in three dimensions. Given the screening length of distributions of wavenumbers (such as a Gaussian), we anticipate that vortices in these fields would display $n_b = 3$ only at larger values of δ .

The gradient simply becomes 3 as box size approaches and exceeds the average inter-vortex spacing, with every box containing at least one vortex point. Since this is well below the scale of the periodic boundaries, this measure is not sensitive to periodic behaviour during the passage between periodic cells, though any statistical effect on the vortex bulk may still affect the box counting dimension.

Box counting fractality has also been investigated numerically in the vortex tangle of superfluid turbulence, where the scaling is directly affected by the dynamics of the flow and normal fluid interaction [37]. The fractal dimension depends on the vortex line density, but ranges only between 1.4 and 1.7 over a wide range of physical parameters. This is a strong contrast to vortices in random waves; although the vortex bulk is fractal, the lower dimension implies the line density is not distributed locally

isotropically, instead surrounding vortex voids at all length scales. This is not evident in the self-similarity scaling of a single vortex, but is revealed when considering the scaling of the tangle as a whole.

3.6 Scaling Relations of Random Tangle

A natural counterpart of fractal scaling measurements are more general scaling properties of the *loop ensemble*. That vortices behave as random walks has direct implications for the scaling of loop frequency with length [32, 42], but the ideal distribution may not hold in a system with periodic boundaries and particularly with lines of non-trivial homology which cannot be appropriately scale invariant [6, 32]. More recent work has approached the generic problem of vortex tangle in disordered three-dimensional fields via mappings to supersymmetric models that permit analytic results using field theory including the full line length distribution accounting for the boundary condition [42, 43], but it is not clear to what extent these results apply to our random wave systems, or what tunable parameters are appropriate. We address these questions; to what extent does the vortex tangle form a self-similar loop soup, how do the periodic boundaries and NTH lines affect the ideal scaling relationships of the statistical ensemble, and is the system instead better described by loop soup models that take into account the periodic boundaries?

A first scaling relation for the probability density function of loop lengths is provided by the assumption of scale invariance of the loop ensemble at all scales (and with no boundaries). Following [32, 44], this would mean that smoothing irregularities and removing all loops below a scale ξ gives rise to a new loop ensemble statistically indistinguishable to the original ('viewing the system at a lower resolution'). We label loops with a 'size' S given by the trace of their bounding box ($\Delta x + \Delta y + \Delta z$) according to an arbitrary Cartesian reference frame, then from dimensionality the number of closed loops n_l with size from S to $S + dS$ within a given volume must have the form

$$dn_l \propto \frac{dS}{S^4} f \left[\frac{S}{\xi} \right], \quad (3.28)$$

for f some dimensionless function of the relative lengthscale. Following the assumption of scale invariance, f must be a constant independent of this relative scale, and so

$$dn_l \sim S^{-4} dS \quad (3.29)$$

The size S is itself a fractal quantity, which for random walks exhibits that total length $L \sim S^2/\xi$ [32], similar to (3.26). Substituting gives

$$dn_l \sim \xi^{-3/2} L^{-5/2} dL \quad (3.30)$$

$$n_l \sim \xi^{-3/2} L^{-3/2} \quad (3.31)$$

This scaling relation must be obeyed by any scale invariant system of random loops. Following Sect. 3.5 this should include our own random vortices, but cannot persist over all scales, with behaviour below about 3.5λ no longer being fractal, and loops on the longest scales potentially being subtly affected by the boundary conditions. The ideal scale invariant distribution also does not permit lines with non-trivial homology. In (3.29), the probability of a randomly selected line being of this type is zero, but such lines do exist in our ensemble.

Figure 3.9a shows the loop and NTH line length scalings from our simulations of eigenfunctions of the 3-torus with energies 243 and 675, using these datasets as they maximise the number of vortex lines recovered. The loop length distribution shows three regimes; first, below around 3.5λ , the distribution is limited by the smoothness of the field and longer loops become progressively more common. Beyond this region, loops become fractal, and the self-similarity assumption holds with $N(L)$ decaying as $L^{-(1.447 \pm 0.007)}$. This is not fully consistent with (3.31), which is likely an artefact of periodicity but is not a major effect; compared to the ideal isotropic case, lines must always close on relatively short scales, rendering the ideal continuum distribution of loops inaccessible. Finally, the scaling breaks down entirely as loops become more frequent at the longest arclengths. In the continuum model, arbitrarily long loops would occur (though with decreasing probability), but our system is instead dominated by statistical constraints on the total vortex length and the conservation




Fig. 3.9 Probability distributions for vortex curve lengths in the 3-torus. **a** Shows a log-log plot of number of loops at each length for loops with trivial homology in eigenfunctions of energy 243 (blue) and energy 675 (green), and for lines with non-trivial homology of energy 243 (red) or 675 (purple). The given line of best fit has gradient -1.447 ± 0.007 , from 1245899 loops in 4616 simulation cells with energy 243. **b** Shows the same statistics, but with loops and lines with non-trivial homology combined, and lengths given as a fraction of total arclength in the eigenfunction with the vortex. **a** Is an extended version of a Figure in [1]

of phase through each periodic wall. The very longest lines cannot have non-trivial homology because (as in Sect. 3.5) each NTH line must have at least one partner travelling in the opposite direction to conserve phase through the walls of the periodic cell, meaning that the maximum NTH length even in principle is the total arclength within the cell minus one cell sidelength, with the rest of the arclength made up by another NTH line with opposite homology. It seems that cells with vortex length concentrated into a single line are still quite common, but as in Fig. 3.9a this is preferentially expressed in a single large loop regardless of (3.31).

Figure 3.9a also shows separately the distribution of lines with non-trivial homology, which are entirely absent in the self-similarity model of loop scalings. These first appear almost as soon as the available arclength is enough to support them (i.e. the periodic sidelength), but their distribution beyond this point quickly becomes almost flat, with all lengths from shortly beyond their minimum being almost equally likely except at the greatest possible lengths where (as above) loops become more common again.

It is natural to ask if NTH lines are simply direct artefacts of the periodic boundaries, sections of longer lines that would join to form loops in a larger cell, or whether they instead correspond to infinite lines that might be typical even in the isotropic RWM. If they were identified as segments of loops, NTH lines might be expected to occupy a progressively lower fraction of the total arclength as cell size increases. This fraction is shown in Table 3.1 for each of our numerical cell sizes; although there is some variance in the results, the NTH line fraction does not exhibit a simple trend towards 0% (or even necessarily an overall decrease); even considering that this must be the case in an appropriate infinite size limit, it seems that other effects dominate the distribution at the energies we investigate.

These fractions are also generally consistent with the 73% NTH lines in similar numerical experiments on an optical propagation model [6] and with 81% NTH lines in the simple \mathbb{Z}_3 model [32] (a figure consistent when the model is adjusted for different lattices [45, 46], and tunable if the system is biased [47]). It seems that NTH lines are fundamental to the nature of these systems; although they are products of the periodic boundary conditions, they are not sensitive solely to the periodic radius.

The full nature of disordered three-dimensional tangle in random systems with small scale correlations has recently come under scrutiny, with their universal properties related to field theories which can yield analytic results [43], as introduced

Table 3.1 Fraction of vortex length in 3-torus eigenfunctions consumed by lines with non-trivial homology. The fraction at energy 1875 is particularly unreliable as it is recovered from only a small number of numerical simulations, of which one contained no NTH lines

Energy	NTH line length fraction	Standard deviation	Standard error
75	0.66	0.36	0.01
243	0.792	0.241	0.004
675	0.814	0.2	0.009
1875	0.58	0.41	0.17

in Sect. 1.6. A major result of relevance is that the length distribution of all loops (ignoring homology) is universally Poisson-Dirichlet across a range of systems [42, 48–50], even being observed in earlier results where it had not originally been recognised [42], and conjectured to apply to many other physical systems including the aforementioned models of optical vortices and cosmic strings [42, 43]. It is likely therefore to apply to our own model, and we describe here the result of this comparison.

[42] predicts three regimes. First, below a small-scale correlation length cutoff, behaviour is model dependent. Next, the Brownian character of loops still implies a dimension-dependent scaling identical to (3.31). Finally, the size of the periodic sidelength D sets a lengthscale D^2 beyond which loops are sensitive to system boundaries and the universal Poisson-Dirichlet form appears,

$$P(L) = \theta \mathcal{L}^{-1} \left(1 - \frac{L}{f \mathcal{L}} \right)^{\theta-1}, \quad (3.32)$$

where \mathcal{L} is the total arclength of all lines in the system (with no distinction by homology) and f is a normalisation constant depending on parameters of the system. θ relates to a model-dependent *fugacity* (for directed loops such as vortices, θ is equal to the fugacity), which may significantly affect the shape of the distribution, and whose value in our system must be found by inspection.

Figure 3.9b shows the vortex length fraction PDF for all loops in our 3-torus eigenfunctions of energy 243 and 675. Its shape very closely matches the form described above; the small scale model-dependent regime is in our system the growing number of small loops below the 3.5λ lengthscale at which loops become fractal. The second regime has already been observed in Fig. 3.9a, and discussed above. The third regime was not previously described, but its flat shape perfectly matches (3.32) when the fugacity $\theta = 1$, as expected for vortices (other systems permitting more ‘colourings’ of loops would have other fugacities). As a check on this fit, the change between a self-similar loop soup and the extended Poisson-Dirichlet tail should occur at lengths well above the square of the sidelength [42]. This value is marked for each system in Fig. 3.9b for both energies, and does appropriately separate the two regimes. Overall, these statistical predictions match our model extremely well, and we conclude that the vortices of Gaussian random waves are indeed an expression of generic vortex tangle as considered in [43].

Although the form of (3.32) matches our model well, the Poisson-Dirichlet result suggests that the probability of non-trivial homology lines must ultimately decrease with the size of the system, as the scale-invariant loop regime with gradient $\frac{3}{2}$ grows to dominate behaviour at all scales. This conflicts with our previous observation that the fraction of NTH lines does not clearly decrease with energy in the eigenfunctions we have studied. This suggests that the NTH line fraction is significantly affected by effects other than the energy alone; for instance, it may be related to the power spectrum distribution of plane wave modes making up the Gaussian random eigenfunctions.


Although this explains the statistical commonality of NTH lines in our results, it also reinforces the expectation that their arclength fraction should decrease with the size of the system, which we have seen to not clearly be the case. This is perhaps because the NTH line distribution depends on the precise distribution of degenerate wavevectors, which is not fully isotropic in our systems and so may perhaps affect the NTH lines in a consistent way. We would still expect that this effect becomes negligible as the energy of the eigenfunctions increases, but it may be that our sampled energies are low enough that this is not yet the case. Similar conclusions would apply to the simulations of a cosmic string model [2] and optical speckle [6] that we have already discussed, with the particular constraints of each system (including e.g. the discrete nature of the cosmic string model and the anisotropy of the optical speckle) explaining their varying NTH line fractions.

The literature provides a statistical viewpoint on the origin of the Poisson-Dirichlet distribution, which may give insight into its appearance in our model. The numerical processes used in [42, 50, 51] may all be described by variations on a *split-merge* algorithm, in which an iterative procedure chooses two loops at each stage with probabilities proportional to their lengths, joining them if they are different curves, or splitting them if the same line is chosen twice. The Poisson-Dirichlet distribution is the stationary limit of this process [42]. Our own system lacks natural dynamics, but the analogy of this procedure might be to take a random walk through the configuration space of random wave parameters. The changing wave field would include a series of moving interaction nodes mediating *reconnection events* between vortex strands [23]. The Poisson-Dirichlet distribution is a natural result if the (continuous) distribution of these reconnections samples each line in an appropriately random fashion, such that the probability of a given line being involved is proportional to its length. Thus, the Poisson-Dirichlet distribution would be a natural result of randomly selecting the system parameters.

Although the analysis of this section has focused on the statistics of the 3-torus, these statistics of loop length are accessible for vortex tangle in eigenfunctions of both the 3-sphere and QHO, and may give insight into how they differ from one another. In the limit of the RWM, the periodic sidelength of the 3-torus will tend to infinity and the Poisson-Dirichlet region reduce to nothing with the fractal loop scaling instead continuing indefinitely. This same pattern must be recovered in both the 3-sphere and QHO in their own large-size limits, but it is less clear how well the Poisson-Dirichlet pattern will hold in the presence of curvature or a position-dependent potential which may dominate behaviour on some scales.

Figure 3.10 shows the vortex loop length fraction PDF for 3-sphere eigenfunctions of energy 120 and 255. The arclength parameterisation of vortex curves now follows the piecewise-geodesic form of Sect. 2.4.2, though we expect that this should be appropriate. The form of the length fraction distribution is more complex, but seems to roughly match that of (3.32) and Fig. 3.9b; below a fraction of around 0.003 of the total length we see the system-specific regime which is more complex than that of the 3-torus, with a separate peak in length probability. After this is a polynomial decay (i.e. linear on the log-log plot) that seems to correspond to the fractal loops regime whose gradient magnitude is close to $\frac{3}{2}$ in the 3-torus, though in this case a fit to the

Fig. 3.10 Vortex length fraction distribution in degenerate eigenfunctions of the 3-sphere. The blue line shows energy 120, and the green line energy 255. The purple line shows a fit to part of the energy 120 distribution, with gradient -0.65 ± 0.01 , and the red line to part of the energy 255 distribution with gradient -0.70 ± 0.01



gradient gives 0.65 ± 0.01 or 0.70 ± 0.01 at energies 120 and 255 respectively. It seems that the curvature of the manifold is high enough on the wavelength scale to distort the result from this expectation. The gradient is higher at the higher of energy 255, which would be consistent with tending toward the $\frac{3}{2}$ limit as energy increases. Finally, beyond a length fraction of around 3 % of the total the distribution of lengths is fairly flat as anticipated by the Poisson-Dirichlet distribution, though the fit is not so good as in the case of the 3-torus; it may again be the case that the geometry of the 3-sphere at these energies distorts the result from this ideal form.

The same plot for QHO eigenfunctions of energy $20 + \frac{3}{2}$ is shown in Fig. 3.11, including in (a) the result separated by whether lines leave to infinity or close within the classical radius, and in (b) the joint result. The pattern is quite different to that




Fig. 3.11 Vortex length fraction distribution in degenerate eigenfunctions of the QHO with energy 20. **a** Shows individually normalised PDFs for lines that close at infinity in blue (though their length is calculated only in adjusted wavelengths as in Sect. 2.4.3), and for loops that close within the bulk tangle in red. **b** Shows the combined PDF considering all lines simultaneously

of the 3-torus and 3-sphere, and it is not possible to separate the distribution into the three regimes previously discussed. It may be that our measure of local wavelength distorts the distribution and that it would better match if accounting properly for an anisotropic local metric, but it seems unlikely that this would completely reshape the distribution, particularly considering that the potential is relatively small in much of the classical volume of the QHO. The difference from the theoretical result is likely instead largely due to the potential exerting a strong effect on the length distribution; while the $\frac{3}{2}$ fractal scaling and Poisson-Dirichlet regimes will be recovered in the large energy limit, the energy needed for boundary effects to become negligible may be relatively large.

3.7 Crossing Number and Writhe of Vortex Loops

The measures in previous sections quantify in a certain sense the scales at which geometry becomes random, unpredictable both in the sense that knowledge of local geometry does not parameterise the curve beyond a small scale, and in behaving like a random walk at larger scales. We now turn to analysis of some of the large scale geometrical measures originally introduced in Sect. 1.2, quantifying how curves wind about themselves on the large scales by looking at the average crossing number and the writhe, by comparison to those of random walks.

Integral formulae for ACN and writhe were given in (1.1) and (1.2), but in both cases the calculation comes to a sum over the number of self-intersections in planar projections of the curve, averaged across all projections uniformly distributed on the direction sphere. Although these quantities are strongly related to the linking number from which their integrals originate, we note again that neither is a topological invariant, as when the crossings of a single loop are considered both are sensitive to changes in its local geometry regardless of whether the loop passes through itself. Previous work in the literature has applied these measures to including DNA molecules [52] or other polymers [53], and in vortices including fluids [54] and superfluids [55].

The writhe may also be understood in the context of so-called *Călugăreanu's theorem* [56, 57], with notable similar and extended work due to Fuller [58], White [59] and Pohl [60], of whom the former two may also lend their names to statements of the result. This applies to a framed curve at which every point has a consistent orientation about the curve axis, or alternatively to the sheet of a *ribbon* as its radius tends to 0. The theorem states that

$$\text{linking number of framing} = \text{twist} + \text{writhe} \quad (3.33)$$

where the writhe is the average sum of signed crossings of the curve with itself (as above), the twist is the local rate of rotation of the ribbon axis, and the linking number of the framing is the linking number of the two edges of the ribbon with one another [57, 61, 62]. The expression may be considered as an invariant quantity for a given curve; with a set self-linking number, any conformation of the curve

may alter the writhe or the twist only at the expense of the other. This can be of great significance in physics, where physical constraints might encourage a curve to ‘convert’ dynamically one quantity for the other; for instance, a highly twisted molecule relieving stress by coiling up, reducing local twisting at the expense of a greater writhe. Although the local phase structure around a vortex line gives a framing that must obey (3.33), we focus on the writhe alone as a measure of coiling.

It is impractical to totally isotropically average the writhe or ACN over all projection axes in the analysis of our curves, so we instead approximate both with estimated quantities based on an average of a limited number of projection directions selected (approximately) uniformly on the sphere. There are many ways of selecting such sets of points, with perfect tilings leaving each point the same local neighbourhood and including (for instance) lattices based on the Fibonacci numbers [63]. We use a method of ‘generalised spirals’ suggested in [64], arising from an approximate solution to the problem of minimising energies for points on the sphere interacting through a power law potential. The primary advantage of this method is that it provides tilings for any number of points, a convenience that outweighs local imperfections (discussed below). Figure 3.12 shows example point sets with different total numbers of points; the generalised spirals are mostly uniform but are most locally non-uniform (if anywhere) for a small number of points near their poles.

It is still necessary to choose an appropriate number of sample points to accurately recover the ACN and writhe, as they may have a large variance between projections. Figure 3.13 shows how the numerically recovered ACN and writhe behave as a function of the number of sample directions in the random walk model described in Sect. A.2. Both quantities exhibit large fluctuations when the number of sampled directions is small, but quickly settle down; we thus take our later measurements with 50 sample points, enough that error in the quantities has generally dropped below 10 % for a given curve (anomalous curves are over-represented in the visually distinct curves of the graphs), but before the averaging over more points gives significantly diminishing returns. Other methods have been suggested for optimising the behaviour of the writhe calculation, possibly with superior limiting behaviour as a function




Fig. 3.12 Sets of points on the sphere, selected to be approximately uniform. The selection of angles is via the method of [64]. Figures a–c show 10, 50 and 100 points respectively, with colours from blue to red showing the order in which the points are selected




Fig. 3.13 Behaviour of writhe and crossing number in a standard random walk model. **a** and **b** Show behaviour of the average crossing number and writhe respectively. In each case, the plot shows the fraction of the final result (with 100 sampling points) at every lower number of samples. Both plots show 50 different curves in different colours in order to demonstrate the variance between different geometries

of curve length [65], but simply enumerating the crossings of vectors in projection is sufficiently efficient in practice.

Some aspects of the ACN and writhe quantities have been examined in random walks. The ACN has been shown slightly faster than linearly in a closed random walk of n fixed length steps [66], but this special case is certainly different in self avoiding walks ([67] suggest $n^{1.122 \pm 0.005}$ in numerical experiments), and is further different when comparing the inter-crossing number of two different walks [68] or when considering confinement [69]. This sensitivity may make a non-standard scaling visible in the crossing numbers of our own curves, given their previously discussed interaction between small scale non-random and large scale random behaviour.

The writhe has been similarly investigated, though in this case it is its modulus $|Wr|$ that is interesting as the writhe is a signed quantity that will always average to zero. [70] suggest it should have the form $Wr \propto n^\eta$ for self-avoiding random walks with n steps, and fit to $\eta = 0.522 \pm 0.004$ from Monte Carlo lattice simulations. [67] obtain a similar result but fit to 0.500 ± 0.005 in open curves. It is not clear what should be expected for random walks that are not self-avoiding, such as our vortex lines appear to be according to their fractal scaling. If anything, it is likely that the writhe will grow a little faster.

Figure 3.14a shows how the ACN behaves as a function of vortex length in 3-torus eigenfunctions of energy 675. The variance is initially very large, but quickly

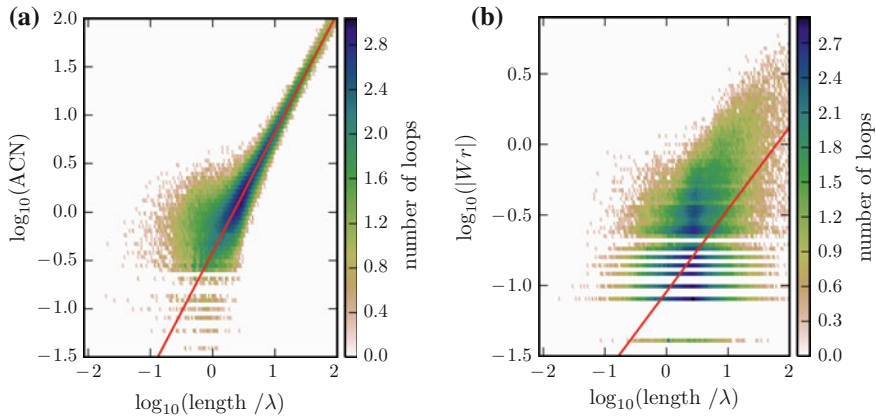


Fig. 3.14 Histograms of ACN and writhe for 270649 vortex curves from 1000 eigenfunctions of the 3-torus with energy 675. The scaling of ACN is shown in (a) with a best fit of gradient 1.233 ± 0.003 , and the scaling of mod writhe $|Wr|$ in (b) with a best fit of gradient 0.585 ± 0.008 . The latter fit is more robust than the spread of writhe may suggest, being fairly stable even when fitting only to particular arclength ranges

settles down to an apparent polynomial dependence. The cutoff for this change is around 3.5λ , consistent with all our previous results of this being the point at which the curve begins to behave as a random walk; below this point, it may be that the requirement for vortex loops to close very locally increases the ACN. We fit the distribution beyond this cutoff to $\langle \text{ACN} \rangle \propto s^{1.233 \pm 0.003}$ for arclength s . This is quite close to the exponent of 1.122 obtained for self-avoiding walks [67]. Although we do not know if an unbiased random walk should have precisely the same scaling, this result is not inconsistent with our expectations.

The same results for the writhe are shown in Fig. 3.14b, though now the variance is proportionally much larger as the writhe grows more slowly. The writhe is also visibly strongly discretised at low value, which arises from the sampling of 50 directions and means that the lowest possible numerical results appear in multiples of $1/50$. Although both these effects affect the appearance of the graph, it displays a strong power law dependence even if the log-log plot is fit to a straight line over different regions, with fit to a gradient of 0.585 ± 0.008 . As expected, this is a less rapid increase than that of the ACN, but it is more different to previous results for a self-avoiding walk than our result for the ACN. This may arise from our walks not truly having self-avoiding statistics, and may also be affected by the particular geometry of vortices. As we have seen in previous sections, there is an apparent intermediate regime between the decorrelation of local geometry and the advent of random walk behaviour; it may be that tangling in this regime increases the rate at which the writhe increases.

3.8 Discussion

We have described various geometric features of the vortex curves in random wave tangle, using eigenfunctions of the 3-torus as a model for isotropic Gaussian random waves. Along a single vortex curve at scales of around a wavelength, we have seen that the probability distributions of curvature and torsion extracted from the simulations agrees well with the results of the isotropic Gaussian random wave model, and have identified numerically the different lengths of order λ at which various scalings occur. At larger length scales, we have verified that the tangle homogeneously fills space, but individual lines look like Brownian random walks and the overall tangle is an extremely good match to a certain general model of random tangle.

The geometric properties of vortex lines arise purely from interference, and the probability distributions of curvature and torsion are not difficult to extract (at least numerically) using the random wave model. Vortex lines in isotropic random functions with different power spectra were found to have similar properties to random Helmholtz waves in [3] (such as the same curvature distribution (3.12) but with different κ_c); it is not clear whether similar PDFs lead to similar conformations of vortex lines, or whether they depend more subtly on correlations along vortex lines. According to the fundamental theorem of space curves [16], a curve's shape is determined uniquely by the curvature and torsion as functions of arclength. One might expect that the random walk behaviour arises from the vanishing of the various correlation functions along the vortex lines, but do these have additional subtle features which characterise them as random wave vortex curves? For instance, preliminary study of a different class of closed curves, specified by random finite Fourier series in Cartesian directions [71], look rather different to random vortex curves on visual inspection, yet have qualitatively similar curvature and torsion PDFs, and direction cosine, curvature and torsion correlation functions. Therefore, what geometric properties of random vortex lines characterise their geometry in particular?

We investigated in Sect. 3.7 a particular type of geometrical entanglement, expressed via the average crossing number and writhe of our vortex curves. This concept of tangling naturally leads to questions of topology; while the writhe and crossing number capture a certain intuitive sense of how the curve interacts with itself, they do not actually put very strong bounds on how the curve really winds about itself. It is easy to construct, for instance, a curve with large average crossing number that in fact barely winds about itself at all, merely by perturbing its local structure. Alternatively, a curve might wind about itself a great deal but otherwise be quite geometrically simple, such that it does not rapidly accumulate crossings. In these cases the knotting of the curve would capture a kind of true entanglement irrespective of local geometry, and this is part of the motivation for our studying its behaviour for the remainder of this thesis.

A persistent theme in our study of the 3-torus has been the existence of lines with non-trivial homology. By comparison with the isotropic random wave model these seem like strange artifacts of the connectivity of the 3-torus, but they occur in many periodic systems of random vortices, and turn out to be a natural phenomenon

in generic periodic tangle. In fact, we find that each vortex cell typically contains one very long line of nontrivial homology, somewhat analogous to the existence of a percolating cluster in regular percolation theory [72, 73], and a small number of shorter ones. A similar observation has been the basis for a percolation model of nodal domains in two dimensions [12], and it is possible that a deeper connection may be drawn with nodal lines (which may be an appropriate analogue of nodal domains in three dimensions). The two-dimensional analysis is based on *avoided crossings* where nodal lines might reconnect under a perturbation of conditions, and the same phenomenon has been mentioned to rely on critical points of the intensity in the case of vortex lines [23]. However, the differences in three dimensions may complicate any comparison; the distribution of reconnection points is not well understood and nodal lines are not in every way equivalent to those in two dimensions since they do not bound entire nodal domains.

Our comparisons to other physical systems have largely focused on a few particular examples of complex vortex tangle, such as in the optical context, the discrete \mathbb{Z}_3 model and quantum turbulence. Another type of geometric similarity might be drawn between vortex tangle and the physical tangle of a polymer melt, a system of many tangled molecular strands. The physics of such a system are very different to those of complex vortices, entanglement is dynamic and mechanical with molecules achieving overall movement by reptation, and there are many subtleties relating to the specifics of the molecules. Overall, the system may be described by a tube model [74], in which an individual molecule effectively performs its semi-random walk within a tube of some radius, being a random walk on large scales with respect to the tube diameter, but with random walk behaviour curtailed on shorter scales by interaction with other molecules [74]. This gives rise to a geometric notion of entanglement length, relating the behaviour of the tube strand to the lengthscale on which the polymer molecule interacts with others.

Such a model seems to qualitatively capture some of the behaviour we observed through measures of local geometry and fractality, in particular the existence of a lengthscale in between that where the line decorrelates (for us, around a wavelength) and that on which it is a random walk. However, the physics of vortex lines is so different that the existence of a deeper connection is unclear; the polymers of a melt are physically tangled by their inability to pass through one another, a very different constraint to that of vortex lines that are not truly confined (they could reconnect with one another under a change of field parameters). The properties of a polymer melt also depend on whether its polymers are closed loops or open chains, a further physical distinction that does not apply to vortices. Even ignoring these differences, it might be the case that random wave tangle is simply not well described even qualitatively by representing surrounding vortices as a tubular constraint. Nevertheless, we could analyse a given stationary vortex state in the terms of a tube model, and it would be interesting to investigate whether random wave vortices are well described in this way. Perhaps this very mechanically founded notion of geometric entanglement is more general than it first appears.

References

1. A.J. Taylor, M.R. Dennis. Geometry and scaling of tangled vortex lines in three-dimensional random wave fields. *J. Phys. A* **470** (46), 0465101 (2014)
2. I.B. Halperin, Statistical mechanics of topological defects, eds. by R. Balian, Kléman, J.-P. Poirier *Les Houches Session XXV—Physics of Defects* (1981)
3. M.V. Berry, M.R. Dennis, Phase singularities in isotropic random waves. *Proc. Roy. Soc. S* **456**, 2059–79 (2000). Including erratum
4. H. Ishio, A.I. Saichev, A.F. Sadreev, K.-F. Berggren, Wave function statistics for ballistic quantum transport through chaotic open billiards: statistical crossover and coexistence of regular and chaotic waves. *Phys. Rev. E* **64**, 056208 (2001)
5. M.R. Dennis, *Topological singularities in wave fields*. PhD thesis, University of Bristol (2001)
6. K. O’Holleran, M.R. Dennis, F. Flossmann, M.J. Padgett, Fractality of light’s darkness. *Phys. Rev. Lett.* **100**, 053902 (2008)
7. K. O’Holleran, *Fractality and topology of optical singularities*. PhD thesis, Department of Physics and Astronomy, Faculty of Physical Sciences, University of Glasgow (2008)
8. M.V. Berry, M. Robnik, Quantum states without time-reversal symmetry: wavefront dislocations in a non-integrable Aharonov-Bohm billiard. *Proc. Roy. Soc. A* **19**, 1365–1372 (1986)
9. R. Höhmann, U. Kuhl, H.-J. Stöckmann, J.D. Urbina, M.R. Dennis. Density and correlation functions of vortex and saddle points in open billiard systems. *Phys. Rev. E* **79**, 016203 (2009)
10. M.V. Berry, Regular and irregular semiclassical wavefunctions. *J. Phys. A* **10**, 2083–2091 (1977)
11. G. Blum, S. Gnutzmann, U. Smilansky. Nodal domains statistics—a criterion for quantum chaos. *Phys. Rev. Lett.* **88**, 114101 (2002)
12. E. Bogomolny, C. Schmit, Percolation model for nodal domains of chaotic wave functions. *Phys. Rev. Lett.* **88**, 114102 (2002)
13. A.G. Monastra, U. Smilansky, S. Gnutzmann, Avoided intersections of nodal lines. *J. Phys. A* **36**, 1845–1853 (2003)
14. E. Bogomolny, C. Schmit, Random wave functions and percolation. *J. Phys. A* **40**, 14033 (2007)
15. M.V. Berry, M.R. Dennis, Natural superoscillations in monochromatic waves in d dimensions. *J. Phys. A* **42**, 022003 (2009)
16. A. Gray, *Modern Differential Geometry of Curves and Surfaces with Mathematica* (CRC Press, 1997)
17. A.W. Baggaley, C.F. Barenghi, Spectrum of turbulent Kelvin-waves cascade in superfluid helium. *Phys. Rev. E* **83**, 134509 (2011)
18. R.D. Kamien, The geometry of soft materials: a primer. *Rev. Mod. Phys.* **74**, 953 (2002)
19. M.G. Bickis, The torsion of a three-dimensional random walk, eds. by S.G. Whittington, D.W. Summers, T. Lodge *Topology and Geometry in Polymer Science*, pp. 23–28. (Springer, 1998). Chapter 5
20. B.K. Shivamoggi, G.J.F. van Heijst, Motion of a vortex filament in the local induction approximation: reformulation of the Da Rios-Betchov equations in the extrinsic filament coordinate space. *Theor. Comp. Fluid Dyn.* (2009)
21. O.D. Lavrentovich, P. Pasini, C. Zannoni, *Defects in Liquid Crystals: Computer Simulations, Theory and Experiments* (Springer Science and Business Media, 2001)
22. J.F. Nye, M.V. Berry, Dislocations in wave trains. *Proc. Roy. Soc. A* **336**, 611–654 (1974)
23. M.V. Berry, M.R. Dennis, Topological events on wave dislocation lines: birth and death of loops, and reconnection. *J. Phys. A* **40**, 65–74 (2007)
24. M.R. Dennis, Local phase structure of wave dislocation lines: twist and twirl. *J. Opt. A* **6**, S202–S208 (2004)
25. E. Orlandini, S.G. Whittington, Statistical topology of closed curves: some applications in polymer physics. *Rev. Mod. Phys.* **79**, 611–642 (2007)
26. P.J. Flory, *Statistical Mechanics of Chain Molecules* (Interscience Publishers, 1969)

27. P. Grassberger, On persistency in self-avoiding random walks. *Phys. Rev. Lett.* **89A**(8), 382–4 (1982)
28. J.F. Nye, Local solutions for the interaction of wave dislocations. *J. Phys. A* **6**, S251–S254 (2004)
29. M.R. Dennis, K. O’Holleran, M.J. Padgett, The fractal shape of speckled darkness. *SPIE Proc.* **6905**, 69050C (2008)
30. K. Falconer, *Fractal Geometry: Mathematical Foundations and Applications*, chapter 3. (Wiley, 1997)
31. J.C. Le Guillou, J. Zinn-Justin, Critical exponents for the n -vector model in three dimensions from field theory. *Phys. Rev. Lett.* **39**, 95–98 (1977)
32. T. Vachaspati, A. Vilenkin, Formation and evolution of cosmic strings. *Phys. Rev. D* **30**, 2036–2045 (1984)
33. K. O’Holleran, M.R. Dennis, M.J. Padgett, Topology of light’s darkness. *Phys. Rev. Lett.* **102**, 143902 (2009)
34. P.J. Flory, *Statistical Mechanics of Chain Molecules* (Oxford University Press, New York, 1988)
35. P.G. de Gennes, *Scaling Concepts in Polymer Physics* (Cornell University Press, 1979)
36. J.P. Wittmer, P. Beckrich, A. Johner, A.N. Semenov, S.P. Obukhov, H. Meyer, J. Baschnagel, Why polymer chains in a melt are not random walks. *EPL* **77**(5), 56003 (2007)
37. D. Kivotides, C.F. Barenghi, D.C. Samuels, Fractal dimension of superfluid turbulence. *Phys. Rev. Lett.* **87** (2001)
38. J.C. Vassilicos, J.G. Brasseur, Self-similar spiral flow structure in low Reynolds number isotropic and decaying turbulence. *Phys. Rev. E* **54**, 467–485 (1996)
39. J. Zinn-Justin, *Quantum Field Theory and Critical Phenomena*. International series of monographs on physics, 2002 edn. (Oxford Science Publications, 2002)
40. G. Foltin, S. Gnutzmann, U. Smilansky, The morphology of nodal lines: random waves versus percolation. *J. Phys. A* **37**, 11363–11372 (2004)
41. D.A. Kessler, I. Freund, Short- and long-range screening of optical phase singularities and C points. *Opt. Commun.* **281**, 4194–4204 (2008)
42. A. Nahum, J.T. Chalker, P. Serna, M. Ortuno, A.M. Somoza, Length distributions in loop soups. *Phys. Rev. Lett.* **111**, 100601 (2013a)
43. A. Nahum, J.T. Chalker, Universal statistics of vortex lines. *Phys. Rev. E* **85**, 031141 (2012)
44. A. Vilenkin, Cosmic strings, eds. by G. Gibbons, S.W. Hawking, S.T.C. Siklos *The Very Early Universe* (Cambridge University Press, 1983)
45. M. Hindmarsh, K. Strobl, Statistical properties of strings. *Nucl. Phys.* **B437**, 471–488 (1995)
46. M.B. Hindmarsh, T.W.B. Kibble, Cosmic strings. *Rep. Prog. Phys.* **58**, 477–562 (1995)
47. J. Robinson, A. Yates, Cosmic string formation and the power spectrum of field configurations. *Phys. Rev. D* **54**(8), 5211–5216 (1996)
48. C. Goldschmidt, D. Ueltschi, P. Windridge, Quantum heisenberg models and their probabilistic representations. *Contemp. Math.* **552**, 177 (2011)
49. S. Grosskinsky, A.A. Lovisololo, D. Ueltschi, Lattice permutations and Poisson-Dirichlet distribution of cycle lengths. *J. Stat. Phys.* **146**, 1105–1121 (2012)
50. A. Nahum, J.T. Chalker, P. Serna, M. Ortuno, A.M. Somoza, 3D loop models and the $C P^{n-1}$ sigma model. *Phys. Rev. Lett.* **117**, 110601 (2011)
51. A. Nahum, J.T. Chalker, P. Serna, M. Ortuno, A.M. Somoza, Phase transitions in 3d loop models and the cp^{n-1} σ model. *Phys. Rev. B* **88**, 134411 (2013b)
52. J. Arsuaga, Y. Diao, M. Vazquez, Mathematical methods in dna topology: applications to chromosome organization and site-specific recombination, eds. C.J. Benham, S Harvey, W.K. Olson, D.W. Sumners, D. Swigon *Mathematics of DNA Structure, Functions and Interactions*, pp. 7–36 (Springer Science + Business Media, 2009a)
53. T. Edvinsson, C. Elvingson, G. Artega, Variations in molecular compactness and chain entanglement during the compression of grafted polymers. *Macromol. Theory Simul.* **9**, 398–406 (2000)

54. F. Maggioni, S. Alamri, C.F. Barenghi, R.L. Ricca, Velocity, energy and helicity of vortex knots and unknots. *Phys. Rev. E* **82**, 026309 (2010)
55. C.F. Barenghi, R.L. Ricca, D.C. Samuels, How tangled is a tangle? *Physica D* **157**(3), 197–206 (2001)
56. G. Călugăreanu, L'intégral de Gauss et l'analyse des noeuds tridimensionnels. *Rev. Math. Pures Appl.* **4**, 5–20 (1959)
57. G. Călugăreanu, Sur les classes d'isotopie des noeuds tridimensionnels et leurs invariants. *Czech Math. J.* **11**, 588–625 (1961)
58. F.B. Fuller, The writhing number of a space curve. *Proc. Nat. Acad. Sci.* **68**(4), 815–819 (1971)
59. J. White, Self-linking and the Gauss integral in higher dimensions. *Am. J. Math.* **XCI**, 693–728 (1969)
60. W.F. Pohl, The self-linking number of a closed space curve. *J. Math. Mech.* **17**, 975–985 (1968)
61. M.R. Dennis, J.H. Hannay, Geometry of Călugăreanu's theorem. *P. Roy. Soc. A* **461**(2062), 3245–3254 (2005)
62. H.K. Moffatt, R.L. Ricca, Helicity and the Călugăreanu invariant. *P. Roy. Soc. A* **439**, 411–429 (1992)
63. J.F. Hannay, J.F. Nye, Fibonacci numerical integration on a sphere. *J. Phys. A* **37**, 11591 (2004)
64. E.A. Rakhmanov, E.B. Saff, Y.M. Zhou, Minimal discrete energy on the sphere. *Math. Res. Lett.* **1**, 647–662 (1994)
65. P.K. Agarwal, H. Edelsbrunner, Y. Wang, Computing the writhing number of a polygonal knot. In *SODA '02 Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 791–9 (2002)
66. Y. Diao, A. Dobay, R.B. Kusner, K. Millett, A. Stasiak, The average crossing number of equilateral random polygons. *J. Phys. A* **36**, 11561 (2003)
67. E. Orlandini, M.C. Tesi, S.G. Whittington, D.W. Sumners, E.J. Janse van Rensburg, The writhe of a self-avoiding walk. *J. Phys. A* **27**, L333 (1994)
68. Y. Diao, A. Dobay, A. Stasiak, The average inter-crossing number of equilateral random alks and polygons. *J. Phys. A* **38**, 7601–7616 (2005)
69. J. Arsuaga, B. Borgo, Y. Diao, R. Sharein, The growth of the mean average crossing number of equilateral polygons in confinement. *J. Phys. A* **42**, 465202 (2009b)
70. E.J. Janse van Rensburg, E. Orlandini, D.W. Sumners, M.C. Tesi, S.G. Whittington, The writhe of a self-avoiding polygon. *J. Phys. A* **26**, L981 (1993)
71. S.M. Rappaport, Y. Rabin, A.Y. Grosberg, Worm-like polymer loops and Fourier knots. *J. Phys. A* **39**, L507–L513 (2006)
72. D. Stauffer, A. Aharony, *Introduction to Percolation Theory* (CRC Press, 1994)
73. G. Grimmett, *Percolation*, Grundlehren der mathematischen (Springer, Berlin Heidelberg, 2010)
74. T.C.B. McLeish, Tube theory of entangled polymer dynamics. *Adv. Phys.* **51**(6), 1379–527 (2002)

Chapter 4

Topological Methods

We introduced in Chap. 1 the notions of knotting and linking, and the idea that they may provide a way to characterise tangles in a way distinct from the geometrical measures covered in the previous chapter. We explain here more formally the mathematical, numerical and practical details of detecting knotting and linking, classifying their different types, and applying these algorithms in practice to vortex curves from our simulations of wave chaos.

In Sect. 4.1 we walk through all the basic knot theoretic ideas of which we will make use. Section 4.2 gives the further practical and numerical details of calculating specific topological invariants, a core part of our later analysis. Section 4.3 moves from knot theory to numerical simplifying space curves while preserving knotting and linking, which will be vital to our own analysis. We finish in Sect. 4.4 by comparing different notions of topological complexity that we will use to judge our own vortex curves.

4.1 A Practical Introduction to Knot Theory

We address here the general details of how we may practically determine whether one or more curves are knotted or linked, particularly through the use of *topological invariants* that map all topologically identical curves (i.e. of the same knot or link type) to the same result. Later sections expand on the specific tools and methods we will use, including the practical details of our own topological calculations and algorithms.

As introduced in Sect. 1.5, the core concept of knotting is intuitive; if we take two closed loops of string in the shape of (for example) the three space-curves in Fig. 4.1a–c, it may or may not be possible to smoothly deform one loop into the shape of another without passing the string through itself. If such a deformation is not possible, the loops are two different types of knot, otherwise they are the same [1].




Fig. 4.1 Images of three-dimensional space curves and their projective diagrams. **a–c** Show images of three dimensional space curves coloured by hue along their length, and topologically the unknot (0_1), trefoil (3_1) and figure-eight (4_1) knots respectively. **d–f** Are projective of each of these curves, with the line broken at crossing points to denote the arc passing beneath the other. Each crossing is indexed by a number, and labelled by a letter a (anticlockwise) or c (clockwise) denoting the orientation of the crossing as in Fig. 1.3. The curves are also oriented according to the marked arrows, but this is an arbitrary choice and does not follow directly from the original space curve

We have already mentioned that the simplest loop may be arranged so as to have no crossings in projection, as in Fig. 4.1a, and is known as the *unknot*, while Fig. 4.1b has at least three crossings in any projection and is called the *trefoil knot*, and Fig. 4.1c at least four crossings with the name *figure-eight*. The methods of distinguishing these are explained below, but we will often refer to these three simplest knots by name. These examples are the three simplest knots by crossing number; no other knot admits a projection with so few crossings.

More rigorously stating this intuitive fact requires some stricter mathematical definitions. In mathematical terms, a knot is an embedding of a circle (S^1) into Euclidean 3-space (\mathbb{R}^3) or the 3-sphere (S^3) [2], though we will only ever consider knotting in a space that is locally homeomorphic to \mathbb{R}^3 even if this involves a projection from the 3-sphere. As in Sect. 1.5, the knot type of the embedding is unchanged under an ambient isotopy, which we can now define as a mapping between manifolds that is a homeomorphism of the complement of the curve and so cannot change its knot type; it cannot pass the curve through itself, and so fulfils the role of rearranging the string in the physical picture [1, 2]. These definitions apply to smooth curves and so admit the concept of a *wild knot*, in which some topological pattern may recur infinitely in line sections of decreasing length in the limit as that length approaches zero [1, 2]. Wild knots are difficult to examine topologically, but do not occur in the vortices of wave chaos and so we simply ignore them.

These definitions become particularly clear if we consider polygonal embeddings of S^1 , finite sets of vertices joined by straight edges that close in a loop. Such curves are not smooth and so wild knots are not possible (except in the limit of zero-length segments). The notion of an ambient isotopy between polygons is the same, but any such ambient isotopy can now be constructed as a sequence of *elementary deformations* [3]. The first of these takes an edge of the polygon and replaces it with two new edges, with the condition that no other edge passes through the triangle formed of the original and two new edges. The second does the opposite, taking two consecutive edges of the polygon and replacing them with the straight line between their preceding and succeeding vertices, again with the condition that no edge segment passed through the triangle originally formed of the two edges and the new replacement edge. Under these definitions it is clear that an elementary deformation cannot pass the curve through itself, and so the rules of ambient isotopy are fulfilled. In practice all of our numerical topology is based on piecewise-linear representations of vortex loops, which can thus be manipulated as polygons rather than smooth curves.

It is important to note that these mathematical definitions are a little different to the everyday idea of knotted string, in which the knot is physically constrained by the inability of the string to pass through itself; the mathematical presence of knotting in some curve need not be physical, being a function only of how the curve winds around itself and not of what would happen if we actually tried to provoke a self-intersection. This allows us to meaningfully talk about knotting in vortex systems; although vortex lines may reconnect under a change of parameters, knotting is well defined for a given stationary state as long as it does not include a self-intersection. The definitions also make clear that knotting is not well defined for non-closed curves, as under an appropriate ambient isotopy *all* such open curves are identical, even if they might be intuitively tangled in the physical picture. In reality, the everyday knotting of open curves is not topological but a practical property based on their friction and behaviour under stress—both irrelevant here [1].

These definitions deal with knots as fully three-dimensional embeddings of S^1 , but it is impractical to work directly with these objects. Instead, all the topological information is contained in the *projective diagram* of the space curve, obtained by projecting the curve onto a plane while preserving knowledge of whether the curves passes *over* or *under* itself at points of self-intersection (the knot is compressed to a *thickened 2-sphere* [4], where the thickness refers to the retention of over-under information at apparent self-intersections). The projection must be *regular* (no more than two arcs cross at a single point, and no intersections are on vertices of the piecewise linear curve [2]), but this is generically the case for a random projection of a real curve. Figure 4.1d–f shows three examples of such projections for the respective space-curves depicted in Fig. 4.1a–c, with the information about which curve’s segment passes under or over the other shown by a short gap in the lower strand as is standard.

A projective diagram may be *oriented* by giving the curve a consistent direction along its length, and information about the topology of the curve is contained solely in the ordered set of these self-intersections including over/under information and

also the orientation (or *sign*) for each crossing as described in Sect. 1.2 and depicted in Fig. 1.3. We label the clockwise type ‘c’ and the anticlockwise type ‘a’ in later discussion. This sign is necessary information to fully describe the topology of a curve beyond just the ordered list of crossings and over-under information. Figure 4.1d–f includes the orientations and example crossing labels for each of the space curves in Fig. 4.1a–c.

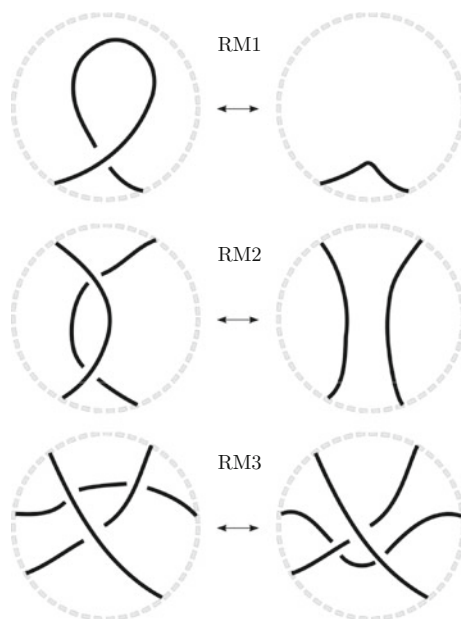
Since the list of crossings contains all the topological information about a curve, many different notations have been advanced to describe a given diagram [5, 6]. The differences are not directly relevant to our analysis and so we do not discuss them here. Where we occasionally refer to a diagram in this way, we make use of the *Gauss code*, simply the list of crossings in the order they are encountered as one walks along a curve from an arbitrary base point, along with information in each case of whether the current arc travels over or under the crossing, and the sign. Figure 4.1d–f are labelled in this way. For instance, the unknot projection in Fig. 4.1d has Gauss code $1+c, 1-c$ (first travelling *over* the clockwise crossing labelled 1, then *under* the same crossing). Similarly, the trefoil in Fig. 4.1e has $1-a, 2+a, 3-a, 1+a, 2-a, 3+a$, and the figure-eight in Fig. 4.1f is $1-a, 2+c, 3-c, 1+a, 4-a, 3+c, 2-c, 4+a$. The same information can be written more concisely, but doing so is not a major concern here.

Just as knots each have many spatial representations via ambient isotopies, a knot may be represented by an infinity of different projective diagrams (including with different crossings), and it is important to be able to transform between topologically equivalent projections. To this purpose, a *planar isotopy* may deform the projection plane, stretching or shrinking the diagram to any other configuration that does not change the set of crossings [1, 3]. Three so-called *Reidemeister moves* may modify the crossings, demonstrated in Fig. 4.2 and labelled by the identifiers RM1, RM2 and RM3. These simple moves, applied repeatedly, can transform any diagram representing a given knot into any other [2, 7]. Each move applies to the configuration of lines within some bounded, simply-connected region which may have any size or shape, as long as no other line segment of the diagram passes through it.

It is important that although only Reidemeister moves 1 and 2 directly change the number of crossings in the diagram, they are not always sufficient to simplify the tangle [1, 3]; it is not difficult to invent a diagram of a given knot that requires the application of Reidemeister 3 to achieve a minimal crossing representation, even though this move does not itself reduce the number of crossings. These moves are also not the only crossing operations that may be safely performed, though any others can be expressed as combinations of Reidemeister moves; we make use of some shortcuts to remove many crossings at once in Sect. 4.3 as this is algorithmically convenient [8, 9].

These operations on projective diagrams also provide a simple way to begin to tabulate different knots, based on the *minimal crossing number*, the smallest number of crossings in any of the projective diagrams representing a given knot. This measure relates intuitively to the complexity of the knot; if the knot may be reduced to almost no crossings, it is in some sense ‘less complex’ than a knot that may never be simplified so far. For instance, we noted earlier that the unknot of Fig. 4.1a has

Fig. 4.2 The three Reidemeister moves for projective diagrams. The moves are labelled as RM1, RM2 and RM3 respectively. Any regions of the knot diagram outside the marked circle are unchanged during the application of the move. The third Reidemeister move could also be drawn with the moving strand passing over the other two



minimal crossing number 0, which in these terms means that any diagram of the unknot may be reduced to a diagram with 0 crossings via Reidemeister moves. In contrast, the trefoil knot of Fig. 4.1b may not be reduced below 3 crossings—in fact, it is the only knot with a minimum of three crossings.

Before describing further the tabulation of knots, we must address the difference between knots that are *prime* and those that are *composite*. A composite knot is one that may be constructed by joining together other knots through a ‘knot sum’ or ‘composition’ operation [1]. Figure 4.3 demonstrates this procedure via projective diagrams; representations of the trefoil knot and figure-eight knot are each cut in a single place, and the now-open ends joined together in a way that does not change the previously existing crossings. This creates a new knot that could be further modified through planar isotopies and Reidemeister moves. As a topological procedure this does not need to preserve any orientation with which the curves have been labelled, though physical filaments that are naturally oriented may have such restrictions. A




Fig. 4.3 The knot sum, denoted by #, that joins (or in reverse, separates) two component knots. In this case, a trefoil knot in red and figure-eight knot in green are cut at the marked dots and joined to a single knot denoted trefoil # figure-eight

prime knot is one in which the reverse of this procedure is not possible, the knot cannot be factored into simpler components [1] (except the trivial factorisation into the same knot composed with the unknot). Thus, every composite knot must be composed ultimately of prime knots, and the process of cataloguing knots requires counting only those that are prime, and this is implicit in the following discussion of knot tables; all discussion of numbers of knots refers to those that are prime, and the (much larger) number of composite knots is not stated. We will use the symbol # to denote knot composition, for instance the composition of an unknot and trefoil knot is $\text{unknot}\#\text{trefoil}$, and is in fact equal to the trefoil since composition with the unknot is the identity of knot summation [1]. We also use exponentiation to denote multiple components of the same knot, such as trefoil^2 for a composite double-trefoil. To judge the structure of knots it is important that the unknot itself is not composite [1] (it is impossible to compose two non-unknots to produce an unknot), and that there may be distinct compositions of two knots (i.e. cutting at different arcs gives topologically distinct results).

These ideas have been a basis for knot tabulation since the earliest knot tables, of which the first major work was an enumeration of the knots with 10 or fewer crossings due to Tait and Little [10–12]. Subsequent work corrected mistakes and extended the original work [13, 14], leading to the milestone of Rolfsen’s revised table in 1976 [15] which included all knots with minimal crossing number up to and including 10. Famously, Rolfsen’s version of the table still contained one incorrect addition, listing 256 diagrams of which two were the same knot; the now well known ‘Perko pair’ [16] (named after their discoverer). Such a pair of diagrams is shown in Fig. 4.4 in order to emphasise the difficulty in rigorously constructing such tables and in identifying knots in general; it is essentially impossible by visual inspection alone to identify that the diagrams represent the same knot type. The problem is so insidious that a similar mistake persists even in some modern resources [17, 18]—the original diagrams of the Perko pair were labelled as ten-crossing knots number 161 and 162, but in later tables the incorrect number 162 is often removed entirely with the labels of the four remaining ten-crossing knots shifted down by one. The mistake is to report number 161 and the *new* number 162 as the identical Perko pair—easy to do since complex knots are so hard to distinguish, but quite wrong.

We use the conventions of the online *Knot Atlas* [19] in identifying and labelling knots, in which knots with up to 10 crossings correspond to those of Rolfsen with




Fig. 4.4 The Perko pair, diagrams 10_{161} and 10_{162} in the tabulation of Rolfsen. Although visually very different, both diagrams are equivalent under Reidemeister moves and planar isotopies

the Perko pair corrected, and which also extend the catalogue to knots with minimal crossing number up to 15. Specific knots are referenced by their crossing number and a mostly-arbitrary numeric identifier, such as 3_1 for the trefoil knot, as it is the first and only 3-crossing knot, or 7_2 for the second knot that happens to have minimal crossing number seven, or 10_{161} for the Perko knot from the previous paragraph. Knots with 11 or more minimal crossings may also have an additional letter identifier, with references such as 11_{n107} . This convention matches most other sources of knot tabulations, though we reference mainly shorter knots whose names date back to Rolfsen or earlier. These labels do not demonstrate chirality without additional information—knots that are not ambient isotopic to their mirror image are given the same label—but this difference is not important to our results.

Figure 4.5 shows a depiction of Rolfsen’s knot table from [20], showing diagrams of all knots with 9 or fewer crossings, in order to inform our later discussion of knot complexity. The first few knots may readily be distinguished by eye, but even at 9 crossings the projections display many varied features. If the projections were made more complex by applying Reidemeister moves, they would also quickly become difficult to match with their initial states.

More recent work has drastically expanded these early tables, with the assistance of numerical tools and more advanced techniques. By 1998, [21] had fully catalogued the first 1701936 knots, all those with a minimal crossing number of 16 or below [21]. More recent work has efficiently enumerated the many millions of *alternating* knots (those with a projection that has crossings alternating between over and under as one walks along the knot in a fixed direction [1]) well into 20 crossings [22, 23], but these are thought to make up only an exponentially small proportion of knots with given crossing number [21] (a conjecture proven for links [24], discussed later). The growth of these numbers suggests a rapid scaling of the number of possible knots with minimal crossing number, a relation in fact shown to be exponential [25].

To focus on these tables is to ignore the remarkable wealth of other properties that knots may have—many of which are essential to constructing tabulations and for any analysis that we perform later, since distinguishing projections only by sequential application of Reidemeister moves is generally impractical. There are few known limits on the number of Reidemeister moves that may be necessary to transform between two diagrams—recent upper bounds are that the number of moves needed to transform between two diagrams with a total of n crossings is at most $\exp^{(c^*)}(n)$ with $c = 10^{1,000,000}$ [26], or the polynomial bound $(231n)^{11}$ if one diagram is of the unknot [27]; neither bound is of practical use. The problem of identifying if a knot is the unknot was famously resolved by Haken [28], later followed by the basics of a procedure for any two diagrams [29], but these algorithms are slow (potentially time- and memory-exponential in crossing number) and again not practical for the comparison of arbitrary curves.

A more practical means of distinguishing and identifying knots is provided by the study of topological invariants such as the linking number that we have already introduced; quantities that may be computed based on the knot (e.g. encoded in its crossings in projection), whose value is always the same no matter the actual geometry of the knot or link and so does not change under the application of Reidemeister




Fig. 4.5 Knots with 9 or fewer crossings. Figure from [20]; the knots match those in the table of Rolfsen [15], with a different presentation. Each knot is labelled in the upper-left by its now-standard identification label, and below by its Alexander polynomial listing only the second half of the coefficients (the first half may be recovered by its symmetry)

moves or ambient isotopies. Not all invariants apply to both knots and links, including the linking number which (in the guise of writhe) is clearly not geometrically invariant. Another simple example of a knot invariant is the minimal crossing number discussed earlier [1], since it relates to minimal projections of the knot (following an arbitrary number of Reidemeister moves or planar isotopies), and so is invariant to changes that make the projection more complex. However, this shifts the burden of knot recognition to the application of Reidemeister moves, which we have seen is generally not simple. It is not even clear if the minimal crossing number of a composite knot is the same as the sum of its components' [1]. An exception to this difficulty is the knots that are alternating. For these, any alternating *reduced diagram* (one where every *isthmus* has been removed, a generalisation of the first Reidemeister move to untwist any crossing that can be isolated by two continuous non-intersecting infinite lines, illustrated in Fig. 4.6), has minimal crossing number [30]; this result includes that the composition of two alternating knots is itself alternating, and has minimal crossing number equal to the sum of its component knots. Unfortunately this fact about alternating knots is not useful in our study of random curves, as they are sufficiently complex that it is difficult to know if an alternating projection exists. Despite the general difficulty in calculation, finding the minimal crossing number is still desirable as a way to compare directly to knot tabulations and the crossing-based understanding of complexity, and we discuss its calculation further in Sect. 4.4.

There are other invariants with the same property of being unchanging only by applying to knot representations that are somehow 'minimal', such as the *bridge number* (the minimal number of arcs with at least one over-crossing in any diagram of a given knot) [1], the observation of the Fary-Milnor theorem that the integrated curvature along a smooth line must be less than or equal to 4π [31, 32] (the sum of angles between segments for piecewise linear curves [33]), or minimal 'energies' of knots under some energy functional [34]. These generally share the difficulty of calculation of the minimal crossing number, and give very little information about the




Fig. 4.6 An 'isthmus' in a knot diagram. Each grey circle denotes an arbitrary knot diagram, connected to the exterior only via the black lines shown. As long as the diagram may be partitioned this way, the crossings in the centre can be added removed without changing the knot type, though actually doing so may require many Reidemeister moves


knot in any case, partitioning knots only into large sets. We make different use of the Fary-Milnor theorem in Sect. 4.3, but will need other tools to efficiently distinguish knots.

Most useful to us are the *knot polynomial* type invariants, which are surprisingly able to associate a polynomial with any given knot, invariant to ambient isotopies or (in projection) to Reidemeister moves and planar isotopies. Since the polynomial is a function of the crossings in any knot diagram, not of some hard-to-find minimal representation, it can be relatively easily calculated even from a particularly tangled and complex representation of a knot. We leave details of the calculation of these polynomials to Sect. 4.2, but cover the basic types here.

The earliest (and perhaps most famous) polynomial invariant is the *Alexander polynomial*, discovered in 1928, and the only known polynomial invariant of knots for 60 years. It may be efficiently calculated (see Sect. 4.2), but is far from a perfect invariant; there are infinitely many knots with the same Alexander polynomial, including those that are indistinguishable in this way from the unknot. Figure 4.7 shows the first example, the 11 crossing knot labelled $K11n34$ in the standard tables. There are also eight other 13-crossing knots with the same Alexander polynomial as the trefoil, and many more with higher minimal crossing numbers, so the Alexander polynomial may not be relied upon to distinguish knots beyond the very simple. In addition, the Alexander polynomial does not detect chirality; knots always have the same Alexander polynomial as their mirror image, even where these are actually different knots.

We also make more limited use of two more recent such invariants, the Jones polynomial [1, 35] and the HOMFLY-PT polynomial [1, 36] whose acronym reflects its simultaneous discovery by several different researchers. Both are more complex to calculate but distinguish more knots; the Jones polynomial distinguishes some chiral pairs of knots, and may possibly perfectly detect the unknot (if no other knots with trivial Jones polynomial exist), though this result is not proven [37], and there are still infinitely many other knots with the same Jones polynomial [38]. The HOMFLY-PT polynomial is more general again, being a polynomial of two variables that generalises both the Jones and Alexander polynomials and distinguishes even more knots, but is still not a perfect invariant that gives all knots a different polynomial [39].

Fig. 4.7 Topologically distinct knots with the same Alexander polynomial invariants. **a** Is simply the unknot, while **b** shows the 11-crossing $K11n34$, the next knot (by crossing number) to also have Alexander polynomial 1



The present analysis makes use of one other type of invariant, the *hyperbolic volume*, defined as the volume of the complement of the knot in S^3 [40–42]. This volume only exists if the complement admits a hyperbolic metric, which is not the case for all knots, in which cases the hyperbolic volume is defined to be zero. In fact, following the famous result of Thurston, knots are one (and only one) of *torus knots* (those that can be drawn on the surface of a torus without passing through themselves), *satellite knots* (those containing an incompressible, non boundary-parallel torus in their complement [1]), and *hyperbolic knots* as above. The advantage of the hyperbolic volume is both that it is convenient to calculate numerically (Sect. 4.2.4) and that it is a powerful identifier of prime knots as all composite knots are satellite knots, and so have hyperbolic volume zero. Identifying prime knots thus reduces to checking the hyperbolic volume is non-zero, or identifying if a knot is a torus knot. Considering the manifold of the knot complement is also the basis for many other invariants, such as the *fundamental group* of the manifold [1, 15], but we do not consider these here.

The combination of these invariants is a fairly good knot detector, and as discussed Sect. 4.2 form part of the core of our own analysis. Still, the combination is not perfect, and other invariants would provide more information at the cost of a more complex calculation. Modern knot theory has generalised these tools in multiple remarkable ways; each of these polynomials is one of a much larger family of *quantum invariants* via quantum groups [43], and has coefficients that may be obtained by considering ‘knots’ with self-intersection to obtain a larger class of *finite type* or *Vassiliev invariants* [44, 45]. The Alexander polynomial is categorified by Floer homology [46, 47], and the Jones polynomial by Khovanov homology [48, 49], in which the polynomials become characteristics of a more complex structure that contains more information. In all of these cases the calculation of the invariant is more difficult, but the results are better knot discriminators than the original polynomials. For instance, Khovanov homology detects the unknot [50], and the full class of Vassiliev invariants gives strictly more information than the individual polynomial invariants, possibly including terms that distinguish all knots [51]. Unfortunately, the more powerful of these tools are far more difficult to evaluate for a given knot, and are generally impractical for our use in numerical classification.

This introduction to topology has focused so far entirely on knotting, but most of the details we need for the analysis of linking follow directly from these. The only difference is that instead of considering the topology of a single curve we consider that of multiple loops that cannot pass through one another. Figure 4.8 shows three examples in planar diagram form; Fig. 4.8a is the simplest link, often called the *Hopf link*, while Fig. 4.8b shows the *Whitehead link* whose significance is discussed below and Fig. 4.8c shows the *Borromean rings*, an example of a link with three components. Figure 4.8d shows a final example that is similar to the Hopf link but in which one of the components is a trefoil knot. As with planar diagrams of knots, all the topological information needed to ascertain the knot is in the list of crossings, which we could enumerate just as before.

Just as with knots, there exists an infinite family of links, distinguished by having different numbers of components, different numbers of windings about one another,




Fig. 4.8 Planar diagrams of some example links. Each separate loop of each link is given in a different colour. **a** Shows the so-called Hopf link, **b** the Whitehead link, **c** the Borromean rings and **d** a Hopf-like link with no standard name of its own. Each of these diagrams is a minimal projection of the given link; it would not be possible to reduce the total number of crossings through Reidemeister moves and planar isotopies

or different knot topologies of the individual loops. For instance, Fig. 4.8d is topologically distinct from Fig. 4.8a through the knotting of one of the component curves, even though the crossings between the two different curves have not changed. Links also share the properties of knots in being prime or composite. It is possible to catalogue links by crossing number [14, 24], but our later analysis will not focus on the direct identification of links so we will not elaborate on these.

The concept of knot invariants also extends to links, but not all invariants of knots are also linking invariants. For instance, the Alexander polynomial as we describe in Sect. 4.2.1 does not work for links, though the concept can be extended to do so [15]. Some other invariants do work without modification, such as the Jones and HOMFLY polynomials, but these are harder to calculate. In our later analysis we instead focus on the linking number, first introduced in Sect. 1.2 and intuitively related to the number of times one curve winds around another. This is a relatively poor invariant, it is trivial to construct links with linking number equal to that of the unlink (for instance, the Whitehead link of Fig. 4.8b), but is numerically highly convenient as explained in Sect. 4.2.6. The linking number is very similar to the writhe of Sect. 3.7, but unlike this quantity it is invariant to ambient isotopies of the link; it is an invariant of links but not of knots.

4.2 Calculating Invariants

We have explained in Sect. 4.1 that there are many invariants of knots and links, some more practically useful than others, and explain here the basics of calculating those that we will make use of for identifying and classifying knotted vortex curves. Before doing so, we briefly describe the procedure by which we convert a piecewise linear three-dimensional space curve to a list of crossings (in fact a representation of the Gauss code), which will be the first stage in calculating any of the quantities discussed later. Numerically retrieving the Gauss code is not difficult and similar procedures are standard throughout the literature.

Given a space curve represented as a polygon in three dimensions (i.e. a set of vertices joined by edges), the procedure begins by projecting to a plane (in practice, one of the Cartesian axes of the piecewise linear representation by ignoring the ‘depth’ information about the curve). This choice is arbitrary and unimportant as long as the projection is ‘good’ in the sense of Sect. 4.1. We also select a starting vertex on the curve and an orientation along its length, this choice is also arbitrary.

To find the crossings, we compare every edge joining two vertices with every other in the two dimensions of projection. This reduces to a simple check for the intersection of two straight lines, starting at the previous vertex and ending at the next. If the edges intersect, the indices of intersection (parameterised by arclength) are stored, along with the information of which edge passes over the other by interpolation of their depthwise coordinates. We also store the orientation of the crossing, which maps to whether the cross product of the projected over-strand edge with the projected under-strand edge points upwards or downwards from the projection plane. After comparing all edges with all others, the list of recovered crossings is sorted by the arclength index at which each crossing occurs. This basic procedure clearly recovers a Gauss code representation of the curve; the ordered crossings, over/under information and orientation.

This basic procedure is actually very inefficient for general curves, in which many edges are not nearby in projection. We make two main optimisations; first, every edge of the polygon is compared only with every later edge, since clearly both crossings (from the perspective of the under-strand and of the over-strand) can be recovered by detecting the intersection a single time. Second, we keep a list of all edge lengths in the piecewise-linear polygon, along with the distance from the edge with which we are currently searching for crossings; at any point as we walk around the curve, we do not have to check for intersection if we are so far from the starting point that the next edge cannot reach it. This allows many intersection tests to be skipped, dramatically speeding up the procedure.

The resulting numerically recovered Gauss code is our input to any of the procedures explained in the subsections below. Our implementation of all of these procedures uses Python and Cython, interfacing with external tools where stated. We apply basic optimisations as above, but these algorithms are not otherwise so time consuming as to require special discussion.

4.2.1 The Alexander Polynomial


The Alexander polynomial is generally one of the *least* discriminating of the polynomial invariants; there are many knots with the same Alexander polynomial such as the example in Fig. 4.7. However, it has the advantage of being easily calculated in polynomial time (and with a small prefactor).

The Alexander polynomial may be calculated via a presentation matrix based on the crossings of a projective diagram. Each crossing involves three arcs; the incoming strand passing underneath, the outgoing strand from underneath, and the strand that passes over (and divides) the others. These strands are labelled with polynomials according to the sign of each crossing, clockwise or anticlockwise, with the labels shown in Fig. 4.9a or b [52].

From this labelling, one constructs an *Alexander matrix*. Picking a start point anywhere on the diagram and an orientation along the curve, we walk along the line labelling each new crossing with an index that is incremented each time; these indices mark the row of the Alexander matrix. We also label each new arc with a number that is incremented every time we pass *underneath* a crossing (modulo the total number of arcs), and these numbers mark the columns of the matrix. Figure 4.10 demonstrates this process for the figure-eight knot; the curve is oriented by the arrows and our labelling begins from the marked dot. The arc indices are given in red, the crossing indices in blue, and the orientation of the crossing (+ for clockwise or – for anticlockwise) in green.

With this procedure complete, every crossing is associated with three arcs, of which each arc has its own column index, and the crossing has its own row index. For each one, we populate the given row and column in the Alexander matrix with the polynomial label of that arc. All other entries of the matrix are set to zero. Thus, the Alexander matrix includes three non-zero entries in each row (the three arcs of the crossing indexed by that row), and at least two entries in each column marking the crossings at which the arc begins and ends, plus any crossings it passes over on the way. The result of this procedure for the example in Fig. 4.10 is

Fig. 4.9 Clockwise and anticlockwise crossings labelled for calculation of the Alexander polynomial. **a** Shows a clockwise crossing, and **b** an anticlockwise crossing






Fig. 4.10 A diagram of the figure eight knot, labelled for calculation of the Alexander polynomial. The curve is given an arbitrary starting point marked by the dot and an arbitrary choice of orientation. Proceeding around the curve from here, each arc of the diagram is given a label starting from 1 and incremented when passing under a crossing, marked in red. Each crossing is given a label the same way, marked in blue. Each crossing is also labelled in green with its orientation, + for clockwise and - for anticlockwise

$$\begin{pmatrix} 1-t & -1 & t & 0 \\ -1 & \frac{1}{t} & 0 & 1-\frac{1}{t} \\ 0 & 1-\frac{1}{t} & -1 & \frac{1}{t} \\ t & 0 & 1-t & -1 \end{pmatrix}, \tag{4.1}$$

in which every column contains three non-zero entries, indicating that every arc passes over one crossing before terminating, and matching the diagram of Fig. 4.10.

The Alexander polynomial is given by the determinant of any principal minor of this matrix, for instance removing the lowest row and rightmost column of (4.1)

$$\begin{vmatrix} 1-t & -1 & t \\ -1 & \frac{1}{t} & 0 \\ 0 & 1-\frac{1}{t} & 1 \end{vmatrix} = -t + 3 - \frac{1}{t}. \tag{4.2}$$

The result is the Alexander polynomial up to multiplication by $\pm t^n$ for some integer n (it is really the ordered list of monomial coefficients that is invariant). In this case, the result matches that expected for the figure-eight knot. This procedure follows [52], and is similar to the original method of Alexander [53]. It may clearly be calculated from the Gauss code explained in Sect. 4.1, as these preserve all the necessary information about the order and signs of the crossings.

An addendum to this core method is that the matrix may only be properly constructed if three *different* arcs meet at each crossing. If this is not the case (e.g. for the crossing of Reidemeister move 1, where a single arc loops back to the same crossing), the algorithm would place two different algebraic crossings in the same row and column of the matrix. To avoid this, it is necessary to first perform an extended version of Reidemeister move one as in Fig. 4.11, which may remove many crossings at once where a single arc passes only over or under other arcs before returning to a single origin point of another crossing. The thin lines in the diagram represent these other




Fig. 4.11 An example of Reidemeister move 1 extended to remove many crossings at once. This is permissible as long as the Reidemeister arc passes only over or only under other arcs before returning to the Reidemeister crossing. This diagram could also be written as a combination of Reidemeister move one and applications of moves two and three, but this would require many more steps

arcs, which could have any other geometry and set of crossings between themselves as long as they pass only over or only under the Reidemeister one arc. Only the over-arc version is necessary to resolve ambiguity in constructing the Alexander matrix, but this concept of generalised Reidemeister moves is a powerful way to simplify curves beyond what is easily possible with the standard Reidemeister moves alone [8], and so in practice we always apply both variants.

All of this algorithm clearly maps easily to our numerical Gauss code whose recovery was described in the introduction of this section. It is necessary only to convert the arbitrary arclength parameterisation of crossing labels to successive integers, and the labels for arcs can be recovered simply by looking at the distance between crossings in this representation. Since the orientation is also stored, the correct crossing-arc labels can be directly entered into the Alexander matrix.

The relative ease of this matrix calculation contributes to the popularity of the Alexander polynomial, even while more powerful invariants are nowadays available. For this reason, it is a core component of our knot identification in later sections, with further optimisations as discussed in the next subsection. However, this algorithm does not work when applied to links with more than one intersecting component. We use other invariants are used to classify these, discussed in later sections.

4.2.2 *The Determinant and Roots of the Alexander Polynomial*

A less powerful but even more easily calculated invariant is the *determinant* of the knot, defined as the modulus of the Alexander polynomial $\Delta(t)$ evaluated at -1 [15], for instance in the case of the trefoil knot

$$\Delta(t) = 1 - t + t^2 \quad (4.3)$$

$$|\Delta(-1)| = |1 + 1 + 1| = 3. \quad (4.4)$$

This invariant is clearly less discriminating than the Alexander polynomial itself, as many different polynomials may evaluate to the same determinant. For instance, the Alexander polynomial of the trefoil knot is $1 - t + t^2$ and the first knot with the same full Alexander polynomial is the 13-crossing 13_{n3663} , but the 8-crossing 8_{19} also has $|\Delta(-1)| = 3$ despite its Alexander polynomial being the quite different $\Delta(t) = 1 - t + t^3 - t^5 + t^6$. Other clashes occur even sooner, for instance the figure-eight knot and one of the two five-crossing knots 5_1 both have determinant 5.

The advantage of this invariant is that it is particularly easy to compute numerically using the algorithm of Sect. 4.2.1; it is simple to calculate the determinant of a integer-valued matrix even with many thousands of entries, while symbolic algebra packages struggle to do so beyond a few hundred. This is particularly important in our investigation of complex knots, whose diagrams often may not easily be reduced to only hundreds of crossings.

This technique can be made more discriminating by calculating values at roots other than -1 , but other integers are not very useful because the algorithm of Sect. 4.2.1 guarantees the result only up to an unknown factor of t . This is not a problem if the full polynomial is known, in which case it can be normalised (e.g. such that the leading term is in t^0 and the coefficient positive), after which replacing t with any number will yield a consistent result. However, the value of the determinant is that it can be calculated without ever obtaining the full Alexander polynomial, so the correct normalisation factor of t^n is not known. Although this means it is not useful to substitute other integers, we instead note that $|\Delta(t)|$ will be consistent evaluated at any root of unity $t = \exp(2\pi i/n)$.

In fact, we find that the first three non-trivial roots ($n = 1, 2, 3$, evaluating $|\Delta(t)|$ at $t = -1$, $t = \exp(2\pi i/3)$ and $t = i$) are sufficient to discriminate knots almost as consistently as the full Alexander polynomial. Table 4.1 lists the knots below 12 minimal crossings that are identified less well by these values than the full Alexander polynomial, in a lookup of Alexander polynomials for knots with up to 15 crossings. Of the 802 such prime knots, only 15 are identified less well by by these evaluations than by the full Alexander polynomial, and there are at most 2 extra misidentifications each within the sampled knot table. Extending to knots of 12 crossings, a further 70 of the 2176 such knots are identified less well, but these are not shown.

The knot determinant has a long history of use for knot identification, thanks to its ease of calculation even with limited computational resources [54–56]. The use of roots of unity does not seem common in the literature, excepting occasionally the similar idea of evaluating the polynomial at 2 or other integers [54, 56] (and accounting for the unknown factor of t in the result), or different normalisations of the Alexander polynomial to allow its useful calculation at different values [57]. Perhaps the practical availability of more powerful invariants made this choice of optimisation unnecessary—Vassiliev invariants in particular can be efficiently calculated (in polynomial time) and were used by [58] to augment the determinant. In our

Table 4.1 All the 15 knots with 11 or fewer crossings whose combination of Alexander polynomial evaluated at the first 3 roots of unity is less discriminating than the full Alexander polynomial, when compared against all knots with 15 or fewer crossings. The second two columns show the number of knots in the dataset with the same Alexander polynomial, and the number with only the same values when evaluated at these roots

Knot identifier	Same $\Delta(t)$	Same $ \Delta(-1) , \Delta(i) $ and $ \Delta(\exp(2\pi i/3)) $
9_4	8	9
$K11a188$	8	9
$K11a367$	2	3
$K11n14$	16	17
$K11n13$	5	6
$K11a144$	8	9
$K11n137$	6	7
$K11a310$	7	9
$K11n121$	16	17
$K11n109$	6	7
$K11n153$	6	7
$K11a342$	1	3
$K11n92$	39	40
$K11a207$	4	5
$K11n70$	31	32

later analysis, we make heavy use of the Alexander polynomial evaluated at roots of unity for the (relatively) loose identification of highly complex curves with many hundreds or thousands of crossings, for which rapid calculation of more powerful invariants remains impractical.

4.2.3 The Jones and HOMFLY-PT Polynomials

The Jones and HOMFLY-PT (hereafter just HOMFLY) polynomials are two of the most well known of the larger class of polynomial invariants. Unlike the Alexander polynomial, they may not be realised as a simple matrix determinant. Instead, they are most commonly described in terms of *skein relations*, giving the polynomial of some knot diagram in terms of the polynomial of multiple others differing by one or more crossings. The skein relations of the Jones polynomial $V(q)$ are

$$V(\bigcirc) = 1 , \tag{4.5}$$

$$(q^{1/2} - q^{-1/2})V(\text{X}) = q^{-1}V(\text{Y}) - qV(\text{Z}) , \tag{4.6}$$

where the contents of each bracket represent a knot diagram or a local area of a diagram (i.e. with the loose ends of each strand connected to the rest of the diagram), and the equation relates the Jones polynomial of different knot diagrams under replacement of local crossings with a crossing of the opposite sign or a local configuration that no longer contains a crossing (but preserves the original arbitrary orientation of the line sections). The full Jones polynomial can always be recovered by recursively replacing crossings in this way until each diagram eventually represents the unknot [1]. This is not the only way to write a skein relation for the Jones polynomial, for instance it may also be obtained from a *Kauffman bracket* type skein-relation [1, 59] or derived in other ways, but these are not reproduced here.

The skein relations of the HOMFLY polynomial $P(l, m)$ are [1, 36]

$$P(\bigcirc) = 1 , \tag{4.7}$$

$$lP(\times) + l^{-1}P(\times) + mP(\times) = 0 . \tag{4.8}$$

This relation is strictly more discriminating than both the Alexander and Jones polynomials, as $P(iq^{-1}, i(q^{-1/2} - q^{1/2})) = V(q)$ and $P(i, i(t^{1/2} - t^{-1/2})) = \Delta(t)$ [1, 36]. The latter case suggests that the Alexander polynomial may be described by a skein relation, an observation popularised by [6], though this is less numerically convenient than the matrix construction.

This discriminatory ability is their main advantage over the Alexander polynomial; among knots with the same Alexander polynomial, far fewer have the same Jones polynomial (though infinite families where this is the case do exist [60]). The HOMFLY polynomial is more discriminating again. Both polynomials may also be applied to links (unlike the Alexander polynomial), and to composite knots in which case the result is a multiple of the polynomials of the component prime knots. A final improvement is that both polynomials detect knot chirality, knots may be distinguished from their mirror images if the curves are different under ambient isotopy, though we do not make use of this property.

We make use of these invariants where possible to confirm many of our knot or link identifications, but they are numerically far less convenient than the Alexander polynomial, being calculable only in exponential time in the number of crossings rather than the polynomial time of the matrix construction [56, 61]. Using the KnotTheory package of Mathematica, the Jones polynomial may be calculated for diagrams with several hundred crossings, and the HOMFLY for several tens. This makes these invariants useful for distinguishing otherwise similar tangles when necessary, but much less efficient than the Alexander polynomial or other invariants we discuss, variations of which remain practical up to many thousands of crossings. It is possible to calculate derivatives of the Jones polynomial in polynomial time while retaining some discriminatory power over the Alexander polynomial [62], but we do not use this method.

4.2.4 *Hyperbolic Volume*

As discussed in Sect. 4.1, the hyperbolic volume refers to the volume of the complement of a knot or link under that complement's hyperbolic metric if it admits one, otherwise being defined as zero. If the hyperbolic volume is non-zero, the link itself may be called hyperbolic. The hyperbolic volume is invariant to ambient isotopy and so is a topological invariant [40]. It is simple to construct infinite families of knots with the same hyperbolic volume [63], so the invariant is not perfect, but in practice it distinguishes prime knots quite efficiently. We do not reproduce here the algorithms for decomposing the knot complement as they are relatively involved; explanations may be found in [40, 41], or source code of [64].

Not all knots are hyperbolic; in fact, following Thurston's famous conjecture (discussed in Sect. 4.1, they may fall into the three distinct categories of torus knots, hyperbolic knots, and satellite knots [1]. In the first and last cases, the complement does not have a hyperbolic structure and the hyperbolic volume is defined to be 0. In practice, this eliminates all composite knots (these are satellite knots), but within the prime knots almost all prime knots are hyperbolic knots; for instance, of the 802 knots with 11 or fewer crossings, only 3_1 , 5_1 , 7_1 , 8_{19} , 9_1 , 10_{124} and $K11a367$ are torus knots and therefore non-hyperbolic.

We make use of the SnapPy numerical tools to calculate hyperbolic volumes for our knots [64]. These implement optimisations able to quickly return an approximate result, though the error about the true hyperbolic volume makes this a less powerful invariant than the 'true' hyperbolic volume without error. This is particularly valuable in complementing the Alexander polynomial approximations of Sect. 4.2.2, allowing composite knots to be readily identified; this procedure is not perfect, as satellite knots may not necessarily be composite, but we will find that in fact such knots are rare and non-hyperbolic knots can generally be easily identified as simple torus or composite knots.

4.2.5 *Minimal Crossing Number*

As explained in Sect. 4.1, the minimal crossing number is simply the smallest number of crossings in any single possible diagram of a knot. It is hard to ascertain its value from only a single projection, so in practice we obtain it by a lookup in knot tables of other invariants of the knot, particularly the Alexander polynomial and the hyperbolic volume. This is not perfect, it is often not possible to identify a space curve with a single knot as even among knots with just 15 crossings many invariants values occur more than once, just as we have seen in the case of the determinant in which case many knots have the same value.

We use the knot tabulations of the online Knot Atlas [19] and KnotInfo [65] resources. These include many invariants of knots with up to 15 crossings, including in particular the Alexander polynomial in almost all cases, though in some cases only the HOMFLY polynomial is provided and we recover the Alexander polynomial through a substitution of variables. These resources also include the hyperbolic volume for knots with up to (and including) 11 crossings, which we use to confirm identifications when possible. The Jones and HOMFLY polynomials are not always available, so we make no direct use of their statistics, but they are used where possible to confirm identifications.

4.2.6 The Linking Number

An integral formula for the calculation of linking number Ln is given in (1.2) [66], but in fact it is much more convenient to use the fact that it simplifies to the sum of signed (by orientation) crossings in projection. Unlike the writhe, this does not require an average over all projection directions (all are equivalent) and so the linking number is recovered by simply performing the signed sum of orientations from the numerically recovered Gauss code of a single projection. This makes it extremely simple to calculate, and so it is our main tool in the identification of links. We can also define the linking number of a link with more than two components as the sum of linking numbers of each distinct pair [66],

$$Ln_{\text{tot}} = \sum_{\mathcal{L}_i, \mathcal{L}_j, i \neq j} |Ln_{ij}|, \quad (4.9)$$

where \mathcal{L}_i and \mathcal{L}_j are each pair of curves in the collection and Ln_{ij} is the linking number between them.

We note again that the linking number is a relatively crude invariant, it is easy to construct non-trivial links whose components wind about one another with both signs and so do not accumulate linking number. This must be kept in mind in our later discussion.

Having introduced the Reidemeister moves in Sect. 4.1, the invariance of linking number (in contrast to the writhe) may be now also understood as arising from the fact that Reidemeister move one does not exist between pairs of loops, it modifies only a single loop. However, it is the only move that can change the sum of crossing signs, Reidemeister moves two and three may involve strands from multiple link components simultaneously but either create one crossing of each orientation or do not modify the orientations at all, leaving the sum of crossing signs and hence the linking number unchanged.

4.2.7 The Fary-Milnor Theorem

The final knot invariant of which we make practical use is geometrical, relating not to the crossings of the curve with itself in projection, but to the integrated curvature along the length of the curve. We use the so-called Fary-Milnor theorem which states that if

$$\oint_K \kappa(s) ds \leq 4\pi, \quad (4.10)$$

where the integral is along a single closed curve K (parameterised by arclength s) and κ is the curvature as in Sect. 3.3, then the curve must be an unknot [31, 32]. No curve may have such a low curvature if it is more topologically complex. Phrased in this way, the Fary-Milnor theorem is not a knot invariant, but we may instead consider the *minimal integrated curvature* of a knot. This is 2π for the unknot, but following the theorem may not fall below 4π for any other, and so any closed loop with integrated curvature below 4π must unambiguously be unknotted. In our numerical analysis we work not with smooth representations but polygonal curves, but the theorem holds as a statement about the sum of angles between piecewise linear segments [33].

This theorem gives only a very weak constraint on what curves may be known to be knotted, and so is rarely useful for topological classification. We instead make use of it in Sect. 4.3 as a numerically simple way to identify curve segments that are definitely unknotted; unlike more complex invariants, the calculation scales only linearly with the number of segments of the curve, and the disadvantage of not detecting many unknots is mitigated by the context of checking only short curve segments that will often have low integrated curvature.

This result is distinct from all of the other invariants we make use of in that it cannot be calculated from the crossings of a projective diagram alone (such as expressed in a Gauss code), but relies on the full three-dimensional structure of the curve.

4.3 Simplifying Space Curves

As explained in Sect. 4.2, it is important that knot invariants are calculated from the simplest possible representations of the knot with as many crossings removed as possible, as all our invariants take at least polynomial time in the number of crossings (for the more powerful polynomial invariants, exponential time). Since we deal with complex tangles that may have tens of thousands of self-intersections in projection, or several thousand even after applying Reidemeister moves one and two, it is convenient to simplify not just the projection but the three-dimensional representation of the space curve. The more tangling that can be removed at this stage, the simpler the projective diagram, and the reduction is of value as long as the three-dimensional manipulation scales better than the calculation of invariants. A secondary advantage is that removing vertices from the curve improves the speed

of crossing detection in projection, which can also be a significant improvement for long vortex loops.

The core idea is that ambient isotopies of the curve are guaranteed to preserve the knot type, and so by performing such operations to remove local vertices we will tend to simplify the curve in all projections. This type of procedure is particularly simple in lattice models, in which algorithms such as the so-called BFACF procedure modify the local passage of the curve via lattice moves that are known to preserve topology [52]. Unfortunately this is difficult to apply to our piecewise linear vortex representations as they track the unconstrained three-dimensional structure of the original vortex; it is possible to fit the curve to a lattice before attempting to simplify it, but without having already simplified the local geometry the scale of the lattice must be extremely small to be sure of correctly capturing the local topology of the curve.

We thus simplify our curves without fitting to a lattice, effectively working instead with arbitrary random polygons (the piecewise linear steps of our numerical sampling). In this case it is more complex to verify that a given local deformation preserves topology, but for piecewise linear curves an analogous operation is to use local deformations that add or remove a triangle to the curve (the elementary deformations introduced in Sect. 4.1 [9, 67]), illustrated in Fig. 4.12. As long as no other line segment passes through the triangle either before or after the move is applied, it preserves topology. Applying this method alone to remove triangles tends to geometrically simplify the curve, and has been used in different fields to simplify tangled lines [68, 69]. However, under a naive implementation the use of elementary deformations would be quite inefficient for long curves, as every local triangle to remove must be checked against every segment of the curve to ensure that they do not intersect.

Some similar procedure is always necessary to guarantee topological continuity, but other methods have been used in the literature to simplify knots in different contexts; for instance, in the study of ideal knots different energy functionals and methods of self-intersection detection may be used to minimise the knot conformation [34]. A simple example is to propagate the knot in an electromagnetic field with charges located on its vertices (and ideally along each linear segment), which in




Fig. 4.12 Topology-preserving elementary deformation of a polygon; two edges of the polygon may be removed and replaced by a straight line joining the vertex before the first edge and after the second, as long as no other line in the diagram passes through the triangle formed by the original two edges and the replacement

practice tends to simplify the curve [70]. However, these methods add significant complexity to the simplification algorithm, and since we are not concerned with obtaining minimal representations (only significant reductions) we do not make use of them.

Our basic procedure extends the above idea of elementary deformations, but rather than considering local triangles one-by-one we first partition space efficiently into smaller cells via an *octree* data structure explained below. This procedure could already speed up a naive algorithm of elementary deformation application by allowing triangles to be removed after checking only that other polygon segments in the same local cell do not pass through them (rather than comparing with every other segment of the full curve), but our algorithm instead simplifies the curve slightly differently by either leaving the shape within the local volume unchanged (no application of elementary deformations) or removing the entire sub-arc within a given cell if it can be shown unambiguously to be unknotted. The advantage of basing geometrical simplification around the partitioning of space rather than the checking of individual triangles or polygons is that the curve can be simplified more efficiently with fewer individual calculations, though this is only true under certain circumstances discussed later and would not necessarily work well for piecewise linear curves arising from different processes.

We use the following core procedure, to which the input may be any number of piecewise-linear curves.

1. Locate the maximal extents of the space curve(s) in each of three orthonormal axes.
2. Define three cut planes, halfway through this containing volume along each axis.
3. Isolate the line segments in each of the eight octants defined by splitting the curve where it crosses these planes.
4. For each octant, perform one of the following three options as appropriate
 - If the octant volume is below a given cutoff size (we use 0.1λ), do nothing.
 - If more than one line arc passes through the octant, repeat this procedure recursively starting from (2), where the containing volume is the volume of the octant and the line segments are cut at its boundaries.
 - If just one line arc passes through the octant, test if it is unknotted. If so, replace the arc with a straight line (i.e. remove all vertices but the first and last where it intersects the walls of the containing volume).
5. Join together all the arcs from the smallest recursive octants, as they must still meet at the planes through which they were cut.

After step (5), the recombined arcs are guaranteed to have the same topology as the inputs, but may have some number of points removed.

Step (4) of this procedure requires a numerical test of whether a line segment is unknotted. As discussed in Sect. 4.1, doing this robustly is difficult, and the most powerful invariants are algorithmically inefficient. For this reason, we adopt a simple geometrical test based on the Fary-Milnor theorem of Sect. 4.2.7; if the sum of angles

between line segments of the arc is less than 4π , including the closure of the last vertex with the first, then the curve cannot possibly be knotted. This test is extremely efficient, requiring summing only enough angles to exceed 4π or to reach the end of the segment. It is far from a perfect unknot detector, many unknotted segments have curvature sum greater than 4π , but is sufficient for identifying locally trivial segments within roughly straight regions of the curve. It would also be possible to augment this detection with calculations of other invariants under appropriate heuristics, but our own implementation does not do so.

Figure 4.13a–c shows the result of applying this algorithm once to a noisy trefoil knot, including in Fig. 4.13b the octree structure and in Fig. 4.13c the result after removing points in arcs detected to be trivial. Figure 4.13d shows the same knot after repeating this procedure 20 times; it has been reduced to 39 points (initially 100) with much less local noise, and has also shrunk relative to the original knot. After a few applications the algorithm has hit a state it cannot simplify efficiently. Further simplification might be better achieved with elementary deformations.

Figure 4.14 shows every stage of the same procedure applied 20 times to a long vortex curve from a simulated energy 675 eigenfunction of the 3-torus, with the orthonormal axes chosen randomly each time. In this case the reduction in points is more sustained; the curve is reduced to just 31% of its original vertices in a single algorithmic pass, but continues to only 1% of the original number after 20 applications. At this point the curve is visibly much simpler, with local tangling largely removed.

Although the simplification is dramatic, this example demonstrates a major limitation of this basic procedure; the reduction to rectangular octant volumes does not well resolve certain features where lines approach closely, as under most axis choices they will not be separated into different octants, and so will rarely be simplified by the algorithm. For instance, the final curve of Fig. 4.14 still has local tangles whose complex local vortex structure and high local curvature resists identification as being unknotted under our simple geometrical test. Many of these local areas are trivial (though some could be knots that have been pulled tight by the algorithm), and could likely be removed efficiently with a further process, such as the application of elementary deformations as discussed above.

Since our core concern is reducing the number of vertices (or equivalently line segments) in the curve representation, this algorithm is sufficient even with these limitations; in Fig. 4.14d, the vortex curve has been reduced from over 130000 crossings in projection to only around 1400. This result partially depends on the particular properties of our recovered vortex curves; their persistence length (Sect. 3.4) is large on the scale of the vertex-vertex separation arising from our simulations, so the piecewise-linear representation is fairly straight on the vertex scale and local segments can be efficiently isolated by the octree. In addition, the statistical repulsion of nearby vortex lines tends to reduce the frequency of multiple arcs being found in the same octree cell.

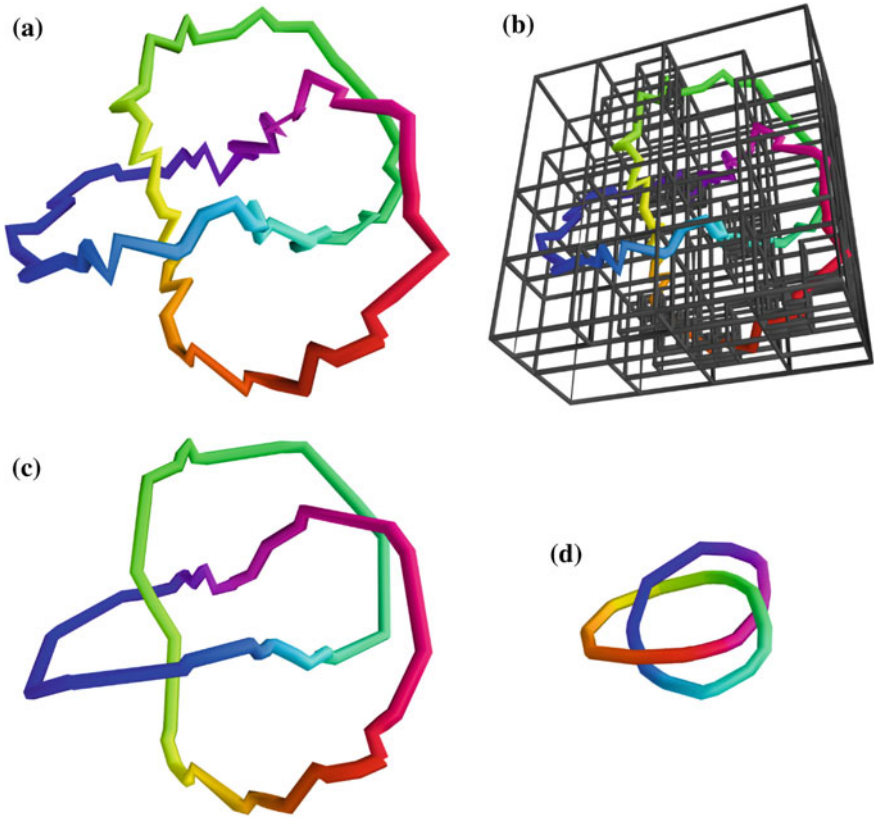


Fig. 4.13 Application of our octree simplification to a trefoil knot. **a** Shows the initial knot, while each cell drawn in **(b)** is an octant to which step (4) of the algorithm will be applied, and **c** is the result after doing so once. **d** Shows the result after applying the same algorithm a further 19 times, with the containing rectangular prism based on a random choice of orthonormal axes each time. The knot is initially made up of 100 points, with this reduced to 39 at the end of the computation

Figure 4.15a demonstrates the average performance of this algorithm in reducing the number of points in different curves, showing the average fraction of the curve's original vertex number remaining after the algorithm has been run multiple times, averaging over some of the longest and most complex tangles vortices in energy 675 simulations of 3-torus eigenfunctions. The figure also shows the same results applied to the random walks of Sect. 1.6. The results clearly demonstrate how the octree algorithm alone is highly suitable for our data, reducing the number of vertices very rapidly the first few times it is applied. In contrast, the random walk model does not respond well, with only a small reduction in the number of vertices even after many repeated applications of the algorithm. This reflects that the random walk includes much less redundant information; every new vertex may be a significant direction change and line segments may commonly approach much more closely. In




Fig. 4.14 Application of our octree simplification to a long vortex curve from our simulations. The curve is initially of length approximately 8700λ , composed of 132740 vertices. In **(b)**, the octree reduction has been applied once, reducing the curve to 45267 vertices. **c** Shows the curve after 6 simplifications with 8966 vertices remaining, and **d** after 20 simplifications with 1366 vertices remaining

some such cases, applying this method initially *increases* the number of vertices, as more are generated on the boundaries of the numerical octants than are removed by simplifying lines within them.

Finally, this algorithm can also be applied to multiple curves simultaneously, as it preserves not just the knotting of a given curve but the linking between any number of loops. In fact, the self-knotting is irrelevant for our primary goal of merely detecting links (it does not affect the linking number), and the calculation may be made drastically faster by removing the requirement that knotting of component loops be preserved. We instead assign a label each separate curve, and adjust step (4) in the original algorithm to

4. For each octant, perform one of the following four options as appropriate
 - If the octant volume is below a given cutoff size, do nothing.




Fig. 4.15 The effect of our octree reduction algorithm. Each *curve* shows the average fraction of the initial number of vertices remaining in a curve after it has been simplified some number of times. The average is taken over 10 different simplifications of each of 10 curves in each case, with different random axes chosen each time. Following the legend, *vortex lines* refers to single vortex loops of length around 1500λ , chosen to showcase a range of different knot types (this does not significantly affect the simplification). *Vortex links* refers to pairs of loops that are linked, with one curve of length over 1500λ and the other of length up to 300λ . Random walks refers to samples from the model of section A.2 with 10000 steps each

- If all line arcs in the octant have the same label, reduce every one to just its first and last vertices and the straight line between them, even if this rearrangement cannot be achieved without passing the line segments through one another.
- If line arcs with more than one different label pass through the octant repeat this procedure recursively starting from (2), where the containing volume is the volume of the octant.

The change means that now multiple arcs of the same line may potentially ‘pass through’ one another in the course of the simplification. This can change the type of the link by changing the knot type of this component, but cannot unlink loops that are topologically entangled, and so is highly convenient if performing measurements that only depend on separating unlinks from other types. In particular, our primary tool of the linking number is not affected at all. Figure 4.15 also shows the result of this new algorithm when applied to links in vortex curves from simulations on the 3-torus; these may involve some of the most complex tangles, in which long vortex loops link with adjacent copies of themselves through the periodic tiling, but are still simplified extremely rapidly.

The same algorithm may also be applied to simplify a collection of lines including passages through periodic boundaries, by choosing the containing volume of the simplification to be the same as the periodic cell volume. Lines passing through the periodic boundaries are made continuous by adding an extra straight line joining them (which ignores periodicity and passes straight across the volume of the periodic cell), but not including this line segment in the simplification. This means that no matter how the algorithm simplifies the line segments of the cell, they remain anchored to

the periodic walls and information about their periodicity is retained. Making the reduction efficient requires a slight modification to the octant selection; if octants are always chosen of uniform size, the planes at which the octants are cut are the same every time the algorithm is applied, since the containing volume is locked to the periodic cell rather than having new axes chosen every time. This may be resolved by giving the selection of octant cut planes a random distribution (we use a uniform distribution) across the length of the containing volume in each axis. This method is useful for the investigation of linking in the 3-torus, for which it is less costly to simplify the entire cell before checking for linking rather than simplifying each pair of vortex loops one by one including for every periodic copy if the extent of the loops exceeds the cell size.

4.4 Topological Complexity

A question that will be raised by the appearance of large and complex knots is how to judge topological complexity; what does it mean for a knot to be complex, and how should they be compared? The answer to this question is not obvious, as very different aspects of complexity may not necessarily be correlated.

Following the previous sections, the minimal crossing number of a knot may be considered an obvious type of topological complexity. The fewer minimal crossings in the knot, the more it is constrained to a small number of very simple topologies. In some cases, the minimal crossing number may directly relate to other physical characteristics; for instance, the rate of movement of DNA in gel electrophoresis is proportional to the minimal crossing number for chains of a given length [71], but such a direct relation is unusual, and other quantities such as the different knot invariants are not necessarily constrained by crossing number.

More general definitions have been suggested for what should be meant by a ‘good’ complexity measure; [72] propose functions whose value limits to 0 at the unknot, and whose value always grows on knot composition. Under this definition, many standard invariant properties are ‘good’ measures, including (for instance) the unknotting number, the number of prime factors on decomposition, the minimal crossing number, and the span of any non-trivial Laurent knot polynomial, though all of these may be generically difficult to compute in practice. These quantities are all also proven to diverge to infinity at least linearly with the length of a curve, but this still leaves great scope for the individual values to behave differently. It is also interesting to note that similar definitions may be met by geometrical measures of complexity such as the average crossing number of Sect. 3.7 [73].

Figure 4.16 shows how two particular invariants, the determinant of the knot and the hyperbolic volume of its complement (where defined), vary across all knots with 15 or fewer crossings (or 11 or fewer for hyperbolic volume). These invariants are *not* necessarily good in the sense of [72], in particular the hyperbolic volume is usually defined as 0 for non-hyperbolic knots and so clearly does not necessarily grow on composition. However, their distributions give some insight into how minimal

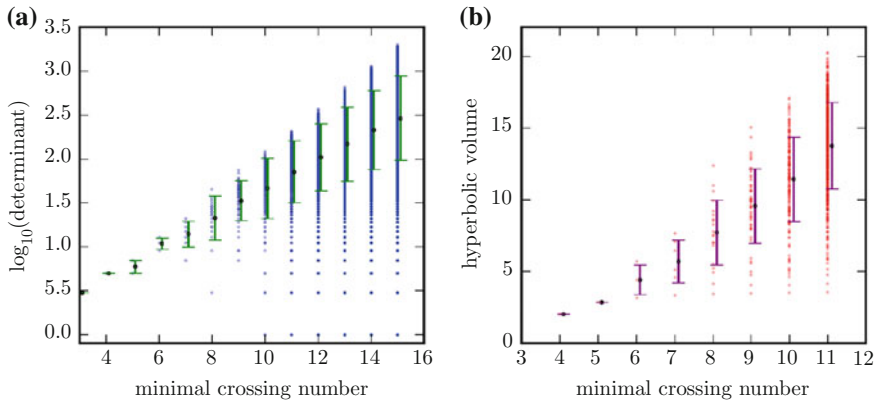


Fig. 4.16 Knots from the knot table, plotted by minimal crossing number against the determinant (in **a**) and hyperbolic volume (in **b**) invariants. Each knot is represented by a single coloured point, though these may not be distinguished where they overlap. In both cases, the plotted error bars show the average and standard deviation of the invariant values at each crossing number

crossing number may relate to other invariants; both quantities grow with this crossing number but at different rates (approximately exponentially and linearly respectively), and both exhibit a large variance.

For these reasons, we will confine our later discussion of complexity to the relation with minimal crossing number and specific knots where possible. Where knots exceed the known knot table, we will analyse only their determinant and hyperbolic volume, but these can only give a highly approximate relation to minimal crossing number as above.

References

1. C.C. Adams, *The Knot Book* (American Mathematical Society, 1994)
2. G. Burde, H. Zieschang, *Knots*. Volume 5 of *de Gruyter Studies in Mathematics* (Walter de Gruyter, 1985)
3. A. Sossinsky, *Knots* (Harvard University Press, 2002)
4. L.H. Kauffman, Virtual knot theory. *Eur. J. Combin.* **20**, 663–691 (1999)
5. C.H. Dowker, M.B. Thistlethwaite, Classification of knot projections. *Topol. Appl.* **16**, 19–31 (1983)
6. J.H. Conway, An enumeration of knots and links, and some of their algebraic properties, in *Computational Problems in Abstract Algebra*, ed. by J. Leech (Pergamon Press, 1970), pp. 329–358
7. K. Reidemeister, *Knotentheorie*, in *Ergebnisse der Mathematik und ihrer Grenzgebiete (Alte Folge 0, Band 1, Heft 1)*. (Springer, Berlin, 1932) (English transl., L. Boron, C. Christenson, B. Smith, BCS Associates, Moscow, Idaho, 1983)
8. C.F. Gauss, G.M. Rassias (eds.), *The Mathematical Heritage of C.F. Gauss* (World Scientific, 1991)

9. F. Comoglio, M. Rinaldi, A topological framework for the computation of the HOMFLY polynomial and its application to proteins. *PLoS One* **6**(4), e18693 (2011)
10. P.G. Tait, On knots I, II, III, in *Scientific Papers*, vol 1 (Cambridge University Press, 1900)
11. C.N. Little, On knots, with a census of order ten. *Trans. Connecticut. Acad. Sci.* **18**, 374–378 (1885)
12. J. Hoste, The enumeration and classification of knots and links, in *Handbook of Knot Theory* (Elsevier, 2005)
13. J.W. Alexander, G.B. Briggs, On types of knotted curves. *Ann. Math.* **28**(1), 562–586 (1926–1927)
14. J.H. Conway, An enumeration of knots and links, and some of their algebraic properties, in *Proc Conf Oxford*, 1967, pp. 329–358
15. D. Rolfsen (ed.), *Knots and Links* (AMS Chelsea Publication, 1976)
16. K.A. Perko Jr., On the classification of knots. *Proc. Am. Math. Soc.* **45**, 262–266 (1974)
17. Perko pair—from wolfram mathworld, <http://mathworld.wolfram.com/PerkoPair.html>. Accessed 01 Oct 2014
18. R. Elwes, *Mathematics 1001* (Firefly Books, 2010) (Corrected in online errata)
19. D.B. Natan, S. Morrison et al., The Knot Atlas, <http://katlas.org>. Accessed 15 Nov 2014
20. G. Burde, Knoten. *Jahrbuch Ueberblicke Mathematik*, 1978, pp. 131–147
21. J. Hoste, M. Thistlethwaite, J. Weeks, The first 1, 701, 936 knots. *Math. Intell.* **20**(4), 33–48 (1998)
22. S. Rankin, O. Flint, J. Schermann, Enumerating the prime alternating knots, part I. *J. Knot Theor. Ramif.* **13**(1), 57–100 (2004a)
23. S. Rankin, O. Flint, J. Schermann, Enumerating the prime alternating knots, part II. *J. Knot Theor. Ramif.* **13**(1), 101–150 (2004b)
24. C. Sundberg, M.B. Thistlethwaite, The rate of growth of the number of prime alternating links and tangles. *Pacific J. Math.* **182**(2), 329–358 (1998)
25. C. Ernst, D.W. Sumners, The growth of the number of prime knots. *Proc. Camb. Philos. Soc.* **102**, 303–315 (1987)
26. A. Coward, M. Lackenby, An upper bound on Reidemeister moves. *Am. J. Math.* (accepted)
27. M. Lackenby, A polynomial upper bound on Reidemeister moves, <http://arxiv.org/abs/1302.0180> (pre-print)
28. W. Haken, Theorie der normalflächen. *Acta Math.* **105**, 245–375 (1961)
29. W. Haken, Über das Homöomorphieproblem der 3-Mannigfaltigkeiten. *Math. Z.* **80**, 89–120 (1962)
30. L. Kauffman, New invariants in the theory of knots. *Am. Math. Mon.* **95**(3), 195–242 (1988)
31. J.W. Milnor, On the total curvature of knots. *Ann. Math.* **52**(2), 248–257 (1950)
32. I. Fary, Sur la courbure totale d’une courbe gauche faisant un nœud. *Bulletin de la Société Mathématique de France* **77**, 128–38 (1949)
33. J.M. Sullivan, Curves of finite total curvature, in A.I. Bobenko, J.M. Sullivan, P. Schröder, G. Ziegler (eds.), *Discrete Differential Geometry*, volume 38 of *Oberwolfach Seminars* (Birkhäuser, Basel, 2008), pp. 137–161
34. A. Stasiak, V. Katritch, L.H. Kauffman (eds.), *Ideal Knots*, volume 19 of *Knots and Everything*, chapter all chapters (World Scientific, 1998)
35. V.F.R. Jones, A polynomial invariant for knots and links via Von Neumann algebras. *Bull. Am. Math.* **12**, 103–111 (1985)
36. P. Freyd, S. Yetter, J. Hoste, W. Lickorish, K. Millett, A. Ocneanu, A new polynomial invariant of knots and links. *Bull. Am. Math. Soc.* **12**, 239–246 (1985)
37. O.T. Dasbach, Does the Jones polynomial detect unknottedness? *Exp. Math.* **6**, 51–56 (1997)
38. T. Kanenobu, Infinitely many knots with the same polynomial. *Am. Math. Soc.* **97**, 158–161 (1986)
39. H.R. Morton, P.R. Cromwell, Distinguishing mutants by knot polynomials. *J. Knot Theor. Ramif.* **05**, 225 (1996)
40. C. Adams, M. Hildebrand, J. Weeks, Hyperbolic invariants of knots and links. *Trans. Am. Math. Soc.* **326**, 1–56 (1991)

41. W. Thurston, Three-dimensional manifolds, Kleinian groups and hyperbolic geometry. *Bull. Am. Math. Soc.* **6**, 357–381 (1982)
42. C. McA Gordon, J. Luecke, Knots are determined by their complements. *J. Am. Math. Soc.* **2**, 371–415 (1989)
43. N. Reshetikhin, V.G. Turaev, Invariants of 3-manifolds via link polynomials and quantum groups. *Inventiones Mathematicae* **103**(1), 547–97 (1991)
44. V.A. Vassiliev, Cohomology of knot spaces, in *Theory of singularities and its applications*, volume 1 of *Adv Soviet Math* (American Mathematical Society, 1990), pp. 23–69
45. D. Bar-Natan, On the Vassiliev knot invariants. *Topology* **34**, 423–72 (1995)
46. P. Ozsváth, Z. Szabó, Holomorphic disks and knot invariants. *Adv. Math.* **186**(1), 58–116 (2004)
47. J. Rasmussen, *Floer homology and knot complements*. PhD thesis, Harvard, 2003
48. M. Khovanov, A categorification of the Jones polynomial. *Duke Math. J.* **101**, 359–426 (2000)
49. D. Bar-Natan, On Khovanov’s categorification of the Jones polynomial. *Algebr. Geom. Topol.* **2**, 337–370 (2002)
50. P.B. Kronheimer, T.S. Mrowka, Khovanov homology is an unknot-detector. *Publications mathématiques de l’IHÉS* **113**(1), 97–208 (2010)
51. S. Chmutov, S. Duzhin, J. Mostovoy, *Introduction to Vassiliev Knot Invariants* (Cambridge University Press, 2012)
52. E. Orlandini, S.G. Whittington, Statistical topology of closed curves: some applications in polymer physics. *Rev. Mod. Phys.* **79**, 611–642 (2007)
53. J.W. Alexander, Topological invariants of knots and links. *Trans. Am. Math. Soc.* **30**(2), 275–306 (1928)
54. A.V. Vologodskii, A.V. Lukashin, M.D. Frank-Kamenetskii, V.V. Anshelevich, The knot problem in statistical mechanics of polymer chains. *Zh Eksp Teor Fiz*, **66**, 2153–2163 (1974)
55. J. des Cloizeaux, M.L. Mehta, Topological constraints on polymer rings and critical indices. *J. Phys.* **40**, 665–670 (1979)
56. E.J. Janse van Rensburg, The knot probability in lattice polygons. *J. Phys. A* **23**, 3573–3590 (1990)
57. M.L. Mansfield, Knots in hamilton cycles. *Macromolecules* **27**, 5924–6 (1994)
58. N.T. Moore, R.C. Lua, A.Y. Grosberg, Topologically driven swelling of a polymer loop. *PNAS* **101**, 13431–35 (2004)
59. L.H. Kauffman, State models and the Jones polynomial. *Topology* **26**, 395–407 (1987)
60. S. Eliahou, L.H. Kauffman, M.B. Thistlethwaite, Infinite families of links with trivial Jones polynomial. *Topology* **42**, 155–169 (2003)
61. T. Deguchi, K. Tsurusaki, Numerical application of knot invariants and universality of random knotting. *Banach Center Publ.* **42**, 77–85 (1998)
62. T. Deguchi, K. Tsurusaki, A new algorithm for numerical calculation of link invariants. *Phys. Lett. A* **174**, 29–37 (1993)
63. D. Ruberman, Mutation and volumes of knots in S^3 . *Invent Math.* **90**, 198–215 (1987)
64. M. Culler, N.M. Dunfield, J.R. Weeks, SnapPy, a computer program for studying the topology of 3-manifolds, <http://snappy.computop.org>. Accessed 28 Jan 2014
65. J.C. Cha, C. Livingston, Knotinfo: table of knot invariants, <http://www.indiana.edu/~knotinfo>. Accessed 14 Dec 2014
66. C.F. Barendhi, R.L. Ricca, D.C. Samuels, How tangled is a tangle? *Physica D* **157**(3), 197–206 (2001)
67. J.W. Alexander, G.B. Briggs, On types of knotted curves. *Ann. Math.* **28**, 562–586 (1926)
68. W.R. Taylor, A deeply knotted protein structure and how it might fold. *Nature* **406**, 916–9 (2000)
69. K. Koniaris, M. Muthukumar, Knottedness in ring polymers. *Phys. Rev. Lett.* **66**(17), 2211–2214 (1991)
70. R.G. Scharein, *Interactive Topological Drawing*. PhD thesis, Department of Computer Science, The University of British Columbia, 1998

71. A. Stasiak, V. Katrich, J. Bednar, B. Michoud, J. Dubochet, Electrophoretic mobility of dna knots. *Nature* **384**, 122 (1996)
72. C.E. Soteris, D.W. Sumners, S.G. Whittington, Entanglement complexity of graphs in Z^3 . *Math. Proc. Camb.* **111**, 75–91 (1992)
73. E.J.J. van Rensburg, D.A.W. Sumners, E. Wasserman, S.G. Whittington, Entanglement complexity of self-avoiding walks. *J. Phys. A* **25**, 6557–6566 (1992)

Chapter 5

Knotting and Linking of Vortex Lines

In this Chapter, we apply the knot theoretic concepts of Chap. 4 to the vortex tangles in all three of our random wave systems; the 3-torus, 3-sphere and quantum harmonic oscillator. This will include direct comparisons between them as a measure of how each differs from the limit of the isotropic random wave model, as well as with random walks from the literature whose topological statistics are well studied, and with other physical systems where possible.

We first ask in Sect. 5.1 what knots exist at all in each type of Gaussian random eigenfunction, confirming that topological considerations make sense in the context of each of their constraints. Section 5.2 more quantitatively addresses how probable knots are as a function of arclength for a given vortex curve, and Sect. 5.3 compares the probabilities of specific knots in different systems. We then move to linking, Sect. 5.4 confirms which systems linking is well defined and checks that links exist, while Sect. 5.5 investigates precise probabilities of linking with length. The remaining three subsections investigate specific issues of knotting; in Sect. 5.6 we discuss how knotting may be properly defined for lines with non-trivial homology on the 3-torus, in Sect. 5.7 the way in which knotting constrains the radius of gyration of the loop, and in Sect. 5.8 the question of what is the *simplest* wave superposition that supports a knotted vortex. We conclude with brief discussion in Sect. 5.9.

5.1 Finding Knots

We begin our analysis of topological entanglement in vortices by investigating the most basic statistics of knotting; what knots occur in different systems, and what immediate patterns appear in their knot types and distributions?

5.1.1 The 3-torus

As previously discussed, much of the arclength in cells of the 3-torus is consumed by lines closing with non-trivial homology, wrapping around the torus some non-trivial number of times. Although such lines may intuitively be topologically entangled, it is difficult to quantify topology in these cases, and so we do not analyse them here. The precise problems and possibilities for such analysis are discussed in Sect. 5.6.

In contrast, loops that close with trivial homology may be analysed without restriction as normal knotted (or not) space curves, and freely form many different types of knot. We have already mentioned in Sect. 2.4.1 that a closed vortex loop is recovered by unwrapping through the periodic boundaries, but we must be careful here because this procedure can have topological implications. Figure 5.1 shows a projected version of the 3-torus where the vertical periodicity (i.e. out of the page) is irrelevant, with the area within the drawn cell corresponding to the simulated unit cell. Within this region is a loop that is somehow linked with itself, which has been unwrapped on both sides to form closed loops.

The topological issue is how their tangling should now be interpreted; if we consider multiple unwrapped copies of the same loop, it is linked with itself infinitely, with each pair of adjacent loops having linking number 2 between themselves. This is not the only way to resolve the topology of a loop in a space with periodic boundaries, it could instead be directly analysed as a periodic *virtual* knot with two crossings (a distinct object from the Hopf link) [1]. Virtual knots can be considered an extension of the tools of knot theory to projective knot diagrams drawn on surfaces with genus that may not be zero, such as in this case what is effectively a planar diagram on a 2-torus in Fig. 5.1 (ignoring the depthwise periodicity of the 3-torus that the example does not traverse), and can capture topological entanglement that depends on wrapping around the periodic boundaries. More of the mathematics of virtual knots is discussed in Sect. 5.6, but they are not investigated here for multiple reasons; their mathematics




Fig. 5.1 A loop in a periodic cell projected vertically, in which direction the cell would also be periodic in our simulations. The periodic cell contains the right half of the red loop and the left half of the green (dashed) loop. We unwrap the periodicity by considering these to represent components of periodic copies of a single loop, which is recovered by shifting the other components by multiples of the periodicity vector. Each loop is linked with its analogues after such a periodic shift, and if all space is tiled with periodic cells the links form an infinite chain

does not extend to periodicity in three directions as in the 3-torus, and it is also not clear that doing so would be desirable, as although it is necessary for a full description of periodic loops it does not seem to correspond to the notion of tangling in an infinite random wave model. In this case a loop linking with its periodic neighbour is a better representation of how vortices would behave were periodicity removed (and the periodic neighbour replaced by some other vortex loop). We thus only consider knotting and linking after unwrapping through the periodic boundaries.

After following this procedure, Fig. 5.2 shows examples of knotted vortex loops from 3-torus eigenfunctions with energy 675 with different lengths and complexities. (a) is an unknotted loop of length around 30λ for reference, while (b) is a trefoil knot of similar length; in fact, knotting of such short curves is extremely rare. Long vortex loops are far more commonly knotted, (c) shows a trefoil knot of length roughly 2700λ while (d) shows one of the longest vortex loops with length 9550λ , determinant 1708845643395 and hyperbolic volume approximately 144.2. Both of these latter two curves are visibly very tangled, and this is normal among long vortex loops.

Although demonstrating that knots of different types exist, these images do little to demonstrate the range of knots that appear. Figure 5.3 instead plots all the different knots appearing in our torus simulations of energy 243 and 675 against their determi-




Fig. 5.2 Example knotted vortices in simulated cells of the 3-torus. **a** and **b** both show vortices of length around 30λ , with **a** being an unknot and **b** a trefoil knot; extremely rare for such a short vortex curve. **c** shows another trefoil knot of length around 2700λ , approximately the maximum length of any vortex in cells with energy 273; such vortices have enough arclength to also form more complex knots. **d** A very complex knot with arclength 9550λ , and determinant 1708845643395. This knot is amongst the longest we find in 3-torus simulations with energy 675

nant invariants, additionally coloured approximately (there may be a small number of misidentifications) by the knot type as torus, satellite or hyperbolic; each point corresponds to a single knot. Unsurprisingly, longer knots commonly have larger determinants (and also larger hyperbolic volumes where applicable, not plotted), corresponding to an increased complexity afforded by the increased arclength, and correspondingly to larger minimal crossing numbers (though their direct geometric projections are far from minimal) and other complexity measures. For instance, the most complex entry for both invariants is the vortex loop pictured in Fig. 5.2d, whose full Alexander polynomial is

$$\begin{aligned}
& 12t^{24} - 354t^{23} + 5142t^{22} - 49086t^{21} + 346717t^{20} - 1933988t^{19} + 8877226t^{18} - \\
& 34496671t^{17} + 115873228t^{16} - 341839898t^{15} + 897061221t^{14} - 2116111767t^{13} + \\
& 4527306953t^{12} - 8852904216t^{11} + 15931003118t^{10} - 26543071054t^9 + 41166709732t^8 - \\
& 59711559562t^7 + 81319776742t^6 - 104311589246t^5 + 126330734951t^4 - 144697469530t^3 + \\
& 156913715807t^2 - 161200770953t - 144697469530t^{-1} + 156913715807 + 126330734951t^{-2} - \\
& 104311589246t^{-3} + 81319776742t^{-4} - 59711559562t^{-5} + 41166709732t^{-6} - 26543071054t^{-7} + \\
& 15931003118t^{-8} - 8852904216t^{-9} + 4527306953t^{-10} - 2116111767t^{-11} + 897061221t^{-12} - \\
& 341839898t^{-13} + 115873228t^{-14} - 34496671t^{-15} + 8877226t^{-16} - 1933988t^{-17} + \\
& 346717t^{-18} - 49086t^{-19} + 5142t^{-20} - 354t^{-21} + 12t^{-22}
\end{aligned} \tag{5.1}$$

This polynomial could be factored into multiple trefoil components $(1-t+t^2)$ after which the remainder would apparently represent a complex prime knot component, but in fact the non-zero hyperbolic volume means that the full knot must itself be prime, so its large determinant comes directly from its complexity as a prime knot. Following the patterns in invariants discussed in Sect. 4.4, this could correspond to approximately 70 crossings in minimal projection, though with a wide variance, and this guess is so far beyond the known knot table that such extrapolation may not be valid.

These figures show a total of 1255 knots at energy 243, and 1157 knots at energy 675, from a total of 1191708 and 2580009 loops respectively; clearly, knots are not in general common, although most vortex loops are relatively short and so would not be expected to be knotted. The statistics of this are discussed in the next section. The invariants we employ also allow a basic count of the number of knots of each different type, though we have only calculated the hyperbolic volume at energy 675. Of the 1157 knots at this energy, 494 are trefoil knots, 565 are hyperbolic, and 14 are the torus knot 5_1 . This means that at most 84 are composite, though these numbers are based on analysis only from roots of the Alexander polynomial and the hyperbolic volume and so it is possible that there are a small number of misidentifications. The detailed probabilities of different knots are compared with other systems in Sect. 5.3.

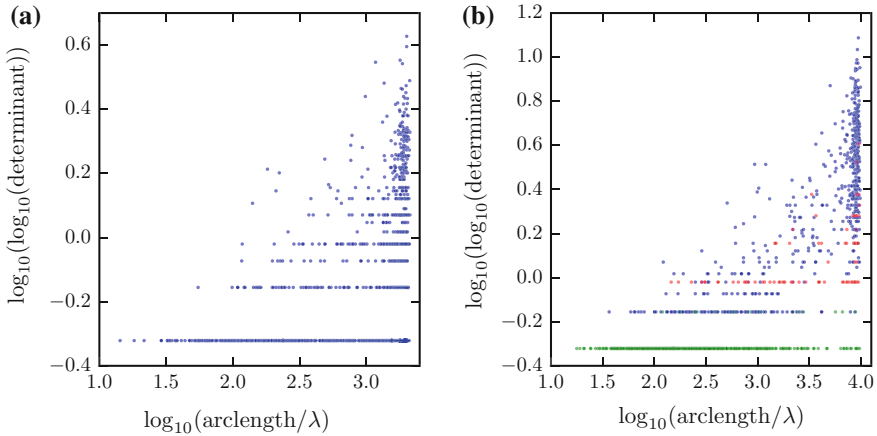


Fig. 5.3 Knots found in 3-torus simulations at different energies, plotted against their determinant invariants. **a** shows energy 243, plotting every knot with a blue point, as information beyond the knot determinant was not investigated. **b** shows energy 675, and in each plot every green dot represents a torus knot, every blue dot a hyperbolic knot, and every red dot the remainder (almost entirely composite knot). Other statistics of these datasets are enumerated in Sect. 2.5

5.1.2 The 3-sphere

Despite the additional complexity of its background metric, the 3-sphere is in fact the system in which topology is most simply defined of all those we examine. Every vortex loop is closed without appealing to periodicity or an external potential, and knotting can be analysed via any continuous map to \mathbb{R}^3 , for which we utilise stereographic projection.

We draw our topological results from simulations of 3-spherical harmonics with both energies 120 and 255. Some basic statistics of the recovered knotted loops given in Table 5.1, where the ‘max composite’ column indicates the number of knots that could not be directly identified as torus or hyperbolic types and so could include a small number of non-composite satellite knot. The basic trend is the same as that of the 3-torus, with trefoil knots being most individually common but most knots by type being hyperbolic.

Four example knots from eigenfunctions of energy 255 are shown in Fig. 5.4. (a) shows a particularly short trefoil knot of length around 30λ , occupying only a small region of the 3-sphere (for comparison, the equatorial circumference of the 3-sphere at this energy is around 16λ). (b) and (c) are longer at 130λ and 150λ , being another trefoil knot and a composite double-trefoil respectively. These loops are also large on the scale of the total 3-sphere volume. (d) shows one of the longest 3-sphere vortex knots, with length 1345λ , hyperbolic volume approximately 332.65 and determinant around 6.11×10^{19} (so high that the numerical floating point calculation loses precision on later digits, though these are not necessary in

Table 5.1 Basic statistical counts of knot occurrences in our 3-sphere simulations. These numbers are not normalised to any other quantity, and list the statistics directly from our simulations

Energy	Number	Num knots	Num trefoils	Num hyperbolic	Max composite
120	30910	2139	478	1341	298
255	107856	8768	1178	7010	474




Fig. 5.4 Example knotted vortices in simulated eigenfunctions of the 3-sphere with energy 255, plotted via stereographic projection. Each vortex is coloured by mapping the colour wheel along its length. **a** shows one of the shortest knotted vortex loops, of length around $30 \lambda_{\text{eff}}$, which forms a trefoil knot. **b** shows a loop of length roughly $130 \lambda_{\text{eff}}$, another trefoil knot, with this length comparable to the examples in other systems **c** is of length approximately 150λ and is a composite knot formed of two trefoils. **d** is one of the longest vortex loops at this energy, with length 1345λ , determinant approximately 6.11×10^{19} , and hyperbolic volume around 332.65. In **b**, **c** and **d** some of the octant boundaries from our numerical algorithm are included to demonstrate how different volumes in the projection compare, while in **a** the entire loop is contained within one of the central octants (not plotted)

our analysis). Following our minimal crossing number extrapolations this might most likely correspond to between 120 and 170 crossings in minimal projection (extrapolating from the determinant and hyperbolic volume respectively), though with the same large error as in our 3-torus example. This last loop occupies a large fraction of the total vortex length in its eigenfunction.

Figure 5.5 shows each knot recovered from these simulations as an individual point, separating torus, satellite and hyperbolic knots by colour. The distribution is distinctly different to that of the 3-torus; at both energies, knots cluster much more clearly about a straight line, such that at higher arclengths the knot determinant tends to always grow rather than long knots retaining a high chance of still being a simple trefoil knot. These long knots are also overwhelmingly hyperbolic in nature, despite that composite knots are (for random walks) expected to make up an arbitrarily large

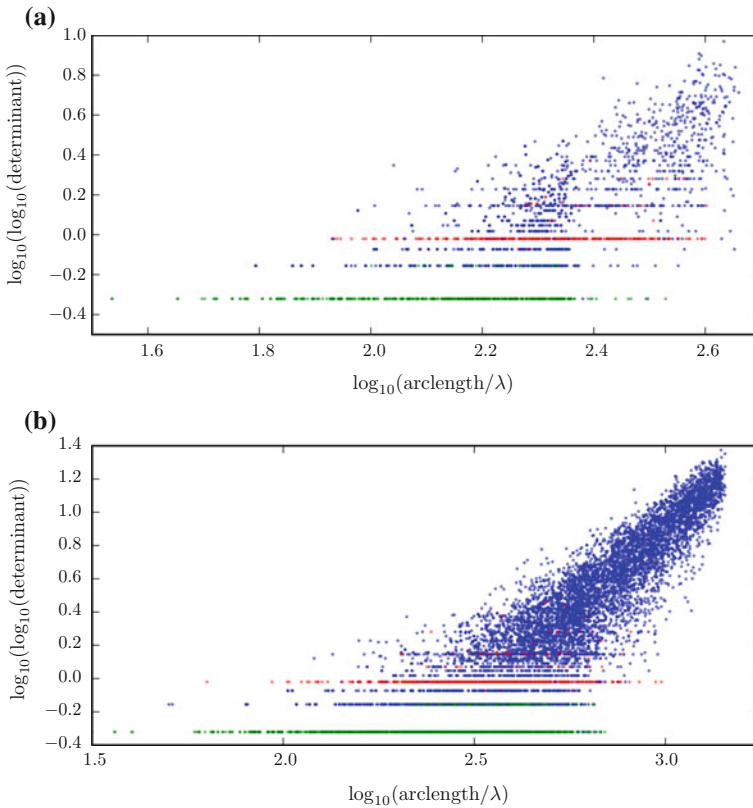


Fig. 5.5 Knots found in 3-sphere simulations at different energies, plotted against their determinant invariants. In each plot, every red dot represents a torus knot, every blue dot a hyperbolic knot, and every green dot a remaining knot, of which all specimens are likely composite. The data in **a** and **b** comes from 3-sphere datasets with energies 120 and 255 respectively, with other statistics of this data enumerated in Sect. 2.5.2

fraction of knots found as energy increases. We discuss this result in the following sections.


5.1.3 The Quantum Harmonic Oscillator

Unlike either of the previous systems, the background potential of the QHO forces vortex loops to either close locally within a ‘classical’ radius, or to leave this classical region and travel in straight lines to infinity. Figure 5.7 shows examples of tangles in both types of line. In (a) a vortex loop of length around $130 \lambda_{\text{eff}}$ forms a trefoil knot, being one of the shortest closed loops to do so.

Although such closed loops within the bulk of the tangle may clearly be considered as normal knots, it is less clear how to correctly define knotting in vortex curves that extend to infinity after leaving the classical radius. We observe that such lines can contain well defined topological tangling by treating them as closing along the boundary of an enclosing 3-ball, and for any given curve all such closures will give the same topological result as long as the vortex intersects the boundary of the ball in only two places, which is guaranteed if its radius is sufficiently large. This method of closure is the simplest possibility, but may be insufficient if there is a sense in which lines going to infinity may be associated with one another, being considered instead as segments of the same curve that join at infinity. Such a notion would relate to what is physically realised under perturbation of parameters, in which case such lines *could* join to become hairpins, in a fashion that could reveal more topology. Figure 5.6 shows an example of a pair of such lines where an appropriate association at infinity would detect a trefoil knot instead of two unknots. We do not use such an association in our analysis, instead only closing individual lines, but this means any knots we detect are only a minimal count of topological non-triviality.

Figure 5.7b–d shows three examples of vortex curves that are knotted under this closure, taken from energy $25 + \frac{3}{2}$ eigenfunctions of the QHO. (b) has length only around $30 \lambda_{\text{eff}}$ and forms a figure-eight knot; its entire length is near to the classical

Fig. 5.6 Examples of potential vortex lines in the QHO that are individually unknotted but could form a knot if joined at infinity. Each of the red and green lines has trivial conformations within the classical radius indicated by the grey circle, but the purple dashed closure forms a trefoil knot






Fig. 5.7 Example knotted vortices in simulated cells of the QHO with energy 25. **a** shows a closed vortex loop of length around $130 \lambda_{\text{eff}}$, which never leaves the bulk of the tangle and forms a figure-eight knot. **b** shows a vortex of length just $30 \lambda_{\text{eff}}$, closed at infinity, and forming a figure-eight knot despite its short length. **c** is another vortex closed at infinity, of length close to $130 \lambda_{\text{eff}}$ comparable to that of the closed loop in **a**, and forming a trefoil knot. The final example **d** also closes at infinity but is amongst the longest lines at $1130 \lambda_{\text{eff}}$, consuming the bulk of vortex arclength in its eigenfunction and with determinant 49539932485

boundary of the QHO, and it even includes a hairpin where the vortex curve leaves this region before quickly turning around and returning. (c) shows a longer trefoil knot, with length $130 \lambda_{\text{eff}}$ comparable to that of the closed loop in (a), though in this case the trefoil knot is quite localised. (d) shows one of the longest vortex loops in this system, with length $1130 \lambda_{\text{eff}}$, determinant 49539932485 and hyperbolic volume around 181; under the same minimal crossing number extrapolations as in the examples from other systems, this might correspond to a minimal crossing number of approximately 70 to 90, though once again with the caveats that this includes a large (but unknown) error, particularly since it is so far from the catalogued knots. This vortex curve is periodic, with the vortex curve being its own mirror image under inversion about the origin.

Under these definitions of knotting, both closed loops and lines joined at infinity form a wide range of knots. Figure 5.8 shows each knot plotted against its determinant

in QHO systems of four different energies, with lines coloured by their torus, satellite or hyperbolic nature in the same manner as the previous sections. As with both other systems, knots show an unsurprising increase in complexity with length, but other features of the distribution may be both sphere-like or torus-like; knots become significantly more common at higher arclengths, but simple knots do not become correspondingly unlikely as rapidly as in the case of the 3-sphere. We discuss the causes of these differences in Sect. 5.3.

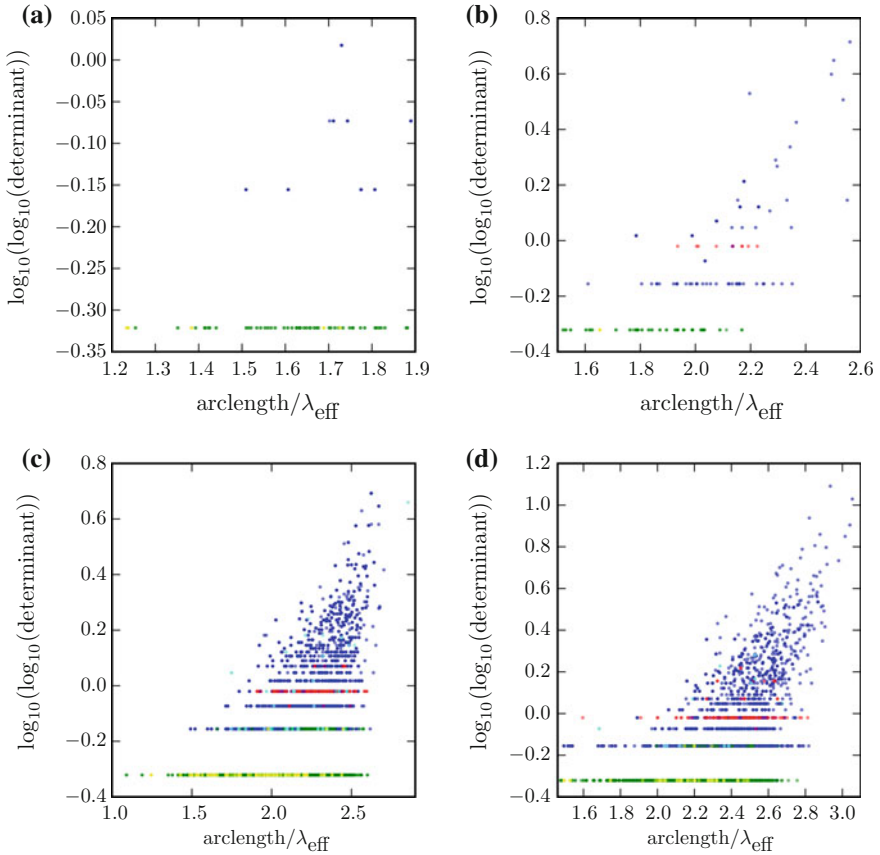


Fig. 5.8 Knots found in QHO simulations at different energies, plotted against their determinant invariants. In each plot, lines to infinity are plotted in green for torus knots, blue for hyperbolic knots and red for the remainder (primarily composite knots). Closed loops are plotted in yellow for torus knots, cyan for hyperbolic knots and purple for the remainder. Individual points are drawn semitransparent, so darker points denote multiple knots of similar or identical lengths, often the result of symmetry in the QHO. The data in plots **a–d** comes from QHO datasets with energies $10 + \frac{3}{2}$, $15 + \frac{3}{2}$, $20 + \frac{3}{2}$ and $25 + \frac{3}{2}$ respectively, with other statistics of this data enumerated in Sect. 2.5

Table 5.2 Basic statistical counts of knot occurrences in our QHO simulations with energy 20. These numbers are not normalised to any other quantity, and list the statistics directly from our simulations

Type	Number	Num knots	Num trefoils	Num hyperbolic	Max composite
Infinite	112384	5398	2987	2136	141
Loop	65318	189	122	59	2

Table 5.3 Basic statistical counts of knot occurrences in our QHO simulations with energy 25. These numbers are not normalised to any other quantity, and list the statistics directly from our simulations

Type	Number	Num knots	Num trefoils	Num hyperbolic	Max composite
Infinite	37964	2629	1226	1247	155
Loop	22062	78	50	23	1

These patterns both appear in both closed loops and lines closed at infinity, though the former are less numerous. Tables 5.2 and 5.3 list the precise statistics of these numbers for energies 20 and 25 respectively, those with the most numerous knots, along with a basic count of different knot types as with the other systems. The only major difference is that the number of trefoil knots exceeds that of the hyperbolic knots, which indicates at least some different structure to the 3-torus or 3-sphere. This could be simply due to the length distribution of vortex loops being different in the QHO; we make a more normalised comparison in Sect. 5.3.

5.2 Probabilities of Knotting

Before making deeper comparisons between the knot distributions of Sect. 5.1, we compare the probability of knotting as a whole with what is expected from normal random walks, providing a direct mechanism for relating our vortex curves with both manually constructed random tangles and other systems of filamentary tangle. A standard mechanism for such a comparison is the *unknotting probability*, the probability as a function of length that loops in the system will *not* be knotted. This has long been known to tend to zero in the infinite limit for classes of random curves [2, 3], and it has become standard to anticipate an exponential decay in random walks [4, 5]. Thus, the probability of unknotting is fit to the form

$$P(\text{unknot}) \propto \exp\left(\frac{-s}{\alpha_0}\right), \quad (5.2)$$

where s is the total arclength of a loop and α_0 is a scaling parameter whose value controls the lengthscale of knotting. Larger values correspond to systems that are less likely to be knotted assuming a consistent measure of s . We will measure α_0 in our vortex systems.

Beginning with the 3-torus, the unknot probability for energies 243 and 625 is shown in Fig. 5.9a–b, with corresponding values for α_0 of $(1820 \pm 30)\lambda$ and $(1990 \pm 20)\lambda$. These results are similar within their errors, and any difference cannot be detected here. Both scaling factors are reported in units of λ despite the discussion of persistence length in Sect. 3.4, for reasons explained further below. Although the large-length fit to the exponential model is good, it does not hold at the shortest length scales, with the first knots appearing only at lengths of around 150λ , illustrated in Fig. 5.10 for 3-torus eigenfunctions of energy 675, though our fit is to greater lengths so this does not present an issue. At the largest available lengthscales the lack of long vortex loops means that unknot probability cannot be tracked beyond around 36% unknots with energy 243, or 13% unknots with energy 625.

In the later discussion, we take these values of α_0 as most representative of the true random wave model in the limit of the periodicity tending to infinite lengths, for the same reasons that justified our geometrical comparisons in Chap. 3 (the topology could in principle follow the isotropic RWM less well than the geometrical statistics we have confirmed, but we see no reason to expect this). Although our long loops are much larger than a single periodic cell and so must feel the periodicity, it is unlikely that they would be *more* knotted as the periodicity is relaxed; if anything, α_0 might be expected to increase as loops are not constrained by the maximum arclength in a cell and may perhaps range more widely. Such a result is supported by the value of α_0 from our simulations, slightly higher in the case of energy 625 as opposed to 243, though the error on these values prohibits a stronger claim.




Fig. 5.9 Unknotting probability with wavelength in 3-torus eigenfunctions. **a** and **b** show the results at energy 243 and 675 respectively. The best fit line in **a** has gradient $(5.5 \pm 0.1) \times 10^{-4}$ based on 1255 knots from a total of 19845 simulations, and in **b** gradient $(5.03 \pm 0.04) \times 10^{-4}$ based a total of 2304 knots from 9525 simulation cells




Fig. 5.10 Unknotting probability in short 3-torus vortices. The dataset is the same as that of Fig. 5.9b, but covering only the region in which knotting first appears. This does not fit the ideal exponential distribution, knotting instead appears more gradually

Figure 5.11a, b shows the same result from simulations of the 3-sphere with energy 120 and 255 respectively. The latter dataset includes many more knots, allowing the full distribution of unknot probabilities from 1 to 0 to be well recovered, and the unknot probability to be distinguished even when only 0.003 % of vortex loops are unknotted. In both cases the unknot probability displays a distinct kink, so we include lines of best fit to both the initial straight line scaling and the later regime when unknot probability falls below around 10 %. All such best fit lines have a very different lengthscale to that of the 3-torus; the initial scaling factors are $\alpha_0 = (143 \pm 2)\lambda$ and $\alpha_0 = (171.2 \pm 0.6)$ at energies 120 and 255 respectively, and the later scaling factors $\alpha_0 = (71 \pm 5)\lambda$ and $\alpha_0 = (82.6 \pm 0.2)\lambda$ for these same energies.




Fig. 5.11 Unknotting probability with wavelenght in 3-sphere eigenfunctions. **a** and **b** show our numerical results at energy 120 and 255 respectively; in **a** we fit to an early exponential in purple with gradient $-(6.96 \pm 0.06) \times 10^{-3}$ and a later exponential in red with gradient $-(1.4 \pm 0.1) \times 10^{-2}$, while in **b** we fit to an early exponential in purple with gradient $-(5.8 \pm 0.1) \times 10^{-3}$ and a later exponential in red with gradient $-(1.11 \pm 0.02) \times 10^{-2}$

One possible interpretation for the shape of these results is that the ideal form of (5.2) is not reached until the later portion of the distribution, with non-random small scale behaviour of the vortex ensemble dominating up to around 500λ (a similar observation that has been made in other random walk investigations [6]). However, comparison with similar results for linking (described in Sect. 5.5) suggests that the reality is the opposite of this, and that later scaling is rather a finite-size effect representing scaling in a region where the initial random walk behaviour no longer applies. The decrease in α_0 is explained as being due to vortex lines becoming large enough that they span the entire 3-sphere; in a larger system they would simply travel further and interact (including potentially linking) with new vortex lines further from their origin, but in the 3-sphere once they pass a certain scale this linking is replaced by more tangling with themselves. This naturally increases the probability of knotting, similarly to how knotting in the 3-torus would increase if loops tangling with periodic copies of themselves were interpreted as knotting rather than linking. Thus, the earlier (larger) α_0 represents better the scaling of (5.2).

It is natural to ask if confinement produces such a visible kink in knotting probability in other systems. The probability of knotting with arclength has been studied for different types of random curves in different types of confinement [7–9], with the general result that knotting is significantly more likely at a given arclength than for unconfined random walks, consistent with our observation. These results from the literature do not seem to observe the same kinked shape in the probability of knotting with arclength, but this may be obscured even if it exists as they do not closely investigate loops that are much smaller than the confinement volume.

It is the nature of such finite-size effects that they should become less significant as the effective size of the system increases; in this case, as the energy is raised. In the limit of very large energies, the local vortex distribution should be described by the random wave model of (1.6). We leave a comparison with our results on the 3-torus to the discussion below, but observe here that the observed values of α_0 scale as we would expect with energy; the initial scaling in both cases is similar, but slightly higher for energy 255, consistent with a slow drift towards the much larger 3-torus results in the limit of high energies. Likewise, the point at which the scaling changes from this value to the finite-size regime of more rapid knotting should occur at longer lengthscales as the energy increases. The intersection of best fit lines is at approximately 360λ for energy 120 and 513λ for energy 255, consistent with this expectation although insufficient on its own to confirm this cause.

The same result for the QHO with energies $20 + \frac{3}{2}$ and $25 + \frac{3}{2}$ is shown in Fig. 5.12a, b respectively. These distributions fit the anticipated exponential relatively poorly, but we characterise two different regions this way. The first, before around $400 \lambda_{\text{eff}}$ is fit to $\alpha_0 = (164 \pm 3) \lambda_{\text{eff}}$ at energy $20 + \frac{3}{2}$ and $\alpha_0 = (310 \pm 4) \lambda_{\text{eff}}$ at energy $25 + \frac{3}{2}$. The error here may be larger than these errors indicate as it is not clear that the decay of unknot probability is truly exponential; we pick out this region instead because it is the only comparable area at the two different energies. In eigenfunctions of energy $25 + \frac{3}{2}$ there is a distinct kink similar to that of vortices in the 3-sphere after which knotting becomes common more rapidly, which we fit to $\alpha_0 = (137 \pm 4) \lambda_{\text{eff}}$. It is not clear why this kink is specific to the higher energy (although the final data point

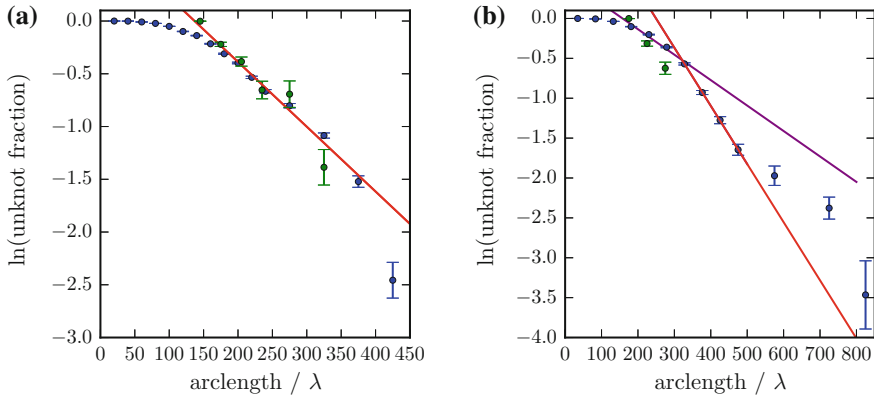


Fig. 5.12 Unknotting probability with wavelength in QHO eigenfunctions. **a** and **b** show our numerical results at energy $20 + \frac{3}{2}$ and $25 + \frac{3}{2}$ respectively. We fit **a** to an early exponential in purple with gradient $(6.1 \pm 0.1) \times 10^{-3}$, and **b** to an early exponential in purple with gradient $-(3.19 \pm 0.06) \times 10^{-3}$ followed by a later exponential in red with gradient $-(7.3 \pm 0.1) \times 10^{-3}$. Results are shown for both lines closing at infinity (blue) and loops closed within the tangle (green), but the best fits are calculated using only lines closing at infinity

at energy $20 + \frac{3}{2}$ hints at a potential similar effect), but we suggest that knotting in the QHO may simply be particularly sensitive to its classical radius. At the lower energy this may simply be so small that vortices are strongly affected in a way that encourages knotting even for relatively short vortices, leading to the much shorter α_0 fit to small lengthscales. There might still be a kink at higher lengthscales that our numerical results do not detect, but other results are also possible.

Although not such a good fit to the theoretical exponential, these results are consistent with the higher energy more closely approaching the isotropic random wave model, represented here by our results from 3-torus eigenfunctions. The early α_0 fit at energy $25 + \frac{3}{2}$ is much longer than that of energy $20 + \frac{3}{2}$, and even the shorter α_0 of the later fit sets a comparable lengthscales. The kinked shape is also consistent with our suggestion that vortices in the 3-sphere are characterised beyond some cutoff by occupying such a large fraction of its volume that they interact with themselves more significantly more than would a vortex in flat space; although the QHO does not have a directly comparable finite volume, its classical radius serves the same effect, in a sense imposing the even stronger constraint that vortices leaving the central region will quickly turn around. This could lead to higher knot probability at large vortex lengths for precisely the same reason, though this different mechanism may also be responsible for the unknot probability having a more complex form.

The QHO at energy 25 also fits the exponential poorly at the highest lengthscales lengthscales, where α_0 increases again; knotting returns to growing more slowly. It is unclear what should cause this regime, and the number of vortices at these lengths is too low to investigate deeply; it may be that the error in this region is underestimated because there are too few vortices in the results to properly sample low probabilities.

The QHO also includes the second class of knotted loops, closed within the bulk of the tangle rather than at infinity, whose statistics are not included in these values for α_0 . They are much less numerous so it is not possible to retrieve a separate exponential fit, but some of their unknot probabilities are shown in Fig. 5.12a, b. It seems that these are broadly similar to those of the lines closed at infinity, up to the given errors, with the exception that at short wavelengths below $130 \lambda_{\text{eff}}$ they are much less frequently knotted than the infinite lines.

We suggest that this may be understood by reference to the local considerations of the QHO; it is plausible that the smallest knotted infinite lines are effectively shorter than closed loops with an outwardly similar λ_{eff} , because the closed loop must spend arclength on closing within the tangle whereas the infinite line effectively terminates at two (potentially far apart) points on the sphere of classical radius, with our measure of λ_{eff} not accounting for the intervening space. This would also explain why the closed loops seem to quickly become knotted just as frequently as the infinite lines; once the length of the vortex curves is significantly longer than any potential closure on the surface of the sphere, both types of vortex curve experience the same pressures on their topology.

As mentioned above, we give all these results in units of wavelength, or effective wavelength on the QHO, despite the discussion in Sect. 3.6 regarding the persistence length. This is because it is unclear if the persistence length is truly the correct lengthscale by which to measure vortex curves; it describes a certain correlation length over which the direction of a vortex filament becomes unpredictable, but as described in our results on scaling relations this lengthscale is well below the cutoff for random walk behaviour. In addition, the persistence length is likely similar in all three of our systems since all are random eigenfunctions of similar practical sizes; maximum vortex lengths in each systems match within an order of magnitude. We also note that parameterising by persistence length would *increase* the value of α_0 in our systems, which would not modify our general observations.

Following these caveats, we are still easily able to compare the broad differences between the unknot scaling in each system. It is clear that knotting probability grows relatively slowly in the 3-torus. We cannot rule out that this is due to the periodic boundary conditions somehow discouraging knotting, but are unable to propose a reason to expect this. Instead, we suggest that this behaviour is actually closest to the knotting probability in the true random wave model without adulteration to ensure periodicity.

The higher knotting rates in both sphere and QHO eigenfunctions may be explained instead with reference to their particular additional constraints; in the 3-sphere, the background metric guarantees that vortices will interact with themselves more frequently than need be the case in the 3-torus, e.g. a roughly straight vortex line will circle around the 3-sphere and return to itself rather than simply travelling further away. Perhaps this curvature is sufficient to drastically encourage the formation of knots. We have already used a similar argument to justify the super-exponentially increasing probability of long vortex loops knotting, but in this case the effect would not require that the vortex fills a large volume of the 3-sphere, being instead a weaker encouragement for self-interaction in vortices of any length.

The QHO would experience a similar but less uniform effect from its classical radius encouraging vortices to quickly turn and re-enter the tangle bulk, while in the centre of the tangle $1/2(x^2 + y^2 + z^2) \ll E$ and the system is more similar to the 3-torus. This might explain why the QHO unknotting probability clearly decays at a rate between that of the 3-torus and 3-sphere, though it is unclear how to predict exactly what the effect of the boundaries should be.

None of these exponential fits is good at small arclengths, where the probability of unknotting remains 1 for at least several tens of wavelengths in every system. A certain minimal arclength is necessary for knotting to become common and fit the random expectation. Extrapolating backwards from our best fit lines, we can compare the intercept $P(\text{unknot}) = 1$ in each system and each energy. These results are shown in Table 5.4, given only approximately due to the inherent error in this measurement; the intercept inherits not only the error in the gradient, but an additional error due to the extrapolation to small lengthscales at which behaviour could vary in different systems. For these reasons, we note only that the broad trend matches our conclusions from the previous results; lines in the 3-torus must be relatively long (in wavelengths) before knots appear at all, while the intercept is much earlier in all other systems. The results also match the conclusion that the 3-torus best represents the true random wave model and that other models will tend towards it as the energy increases; in both the 3-sphere and QHO, the small scale knots become less likely with energy, and the intercept approaches the much larger value of the 3-torus.

We may also compare these values for α_0 to the many results for other random walks in the literature. Some of these are shown in Fig. 5.13, including our own results, those for walks made of edge permutations of collections of randomly oriented equilateral triangles [13], a rod-bead model with harmonic potentials [12], Gaussian random vector models biased to close [10, 11], the quaternionic random walk model we explain in A.2 (from which we have extracted the value of α_0 numerically), a rod-bead model of polymers (in the give case, without excluded volume) [5], and walks made from Hamiltonian walks on lattice grids [6] (for which a length-dependent scaling exponent is suggested, similar to our own observation [4, 6]). These results span a wide range of random walk constructions and measurements, and yet the results are largely well clustered around $200 < \alpha_0 < 450$, where the

Table 5.4 Vortex lengthscales at which knotting first appears. Each result is the approximate intercept of our exponential knotting fits with probability 1. These results are only approximate, as knotting at short lengthscales does not perfectly fit to an exponential

System	Energy	Approximate intercept/ λ
Torus	675	530
Sphere	120	150
Sphere	255	170
QHO	20	250
QHO	25	140

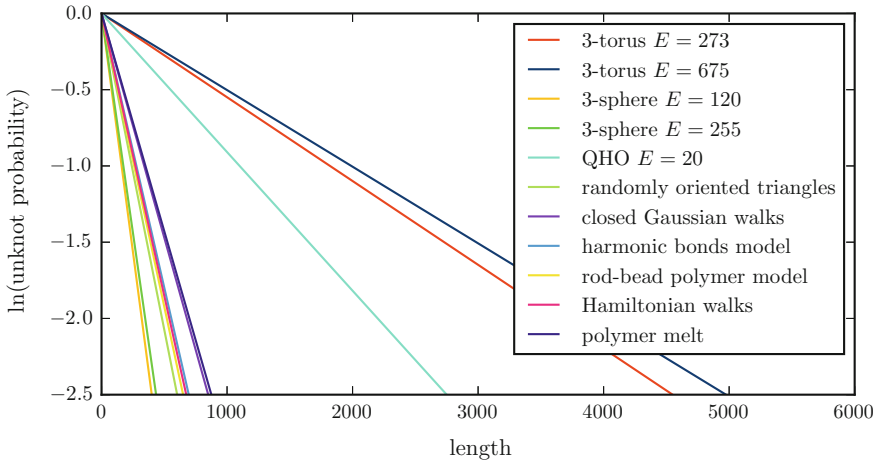


Fig. 5.13 Unknot probability with length for random walks with different constructions. The different walks use different length parameterisations; λ for our own results, and mixtures of the persistence length and step number for others. The sources include each of our 3-torus, 3-sphere and QHO systems, closed walks with Gaussian random step vectors [10, 11], walks with a harmonic bond potential between vertices [12], the random walks of Section A.2 from which we extracted the scaling factor directly, the randomly oriented triangle edges of [13], a rod-bead model of polymers [5], random Hamiltonian walks on lattices [4, 6], and the polymer simulations of [14] (discussed in the text)

length parameterisation is largely in units of the number of sides in the polygon of the curve (with the quaternionic walks an exception as described in Appendix A.2); a number that is likely close to the persistence length since these curves are primarily of almost-uncorrelated fixed-length segments.

In the context of these results, it is clear that knotting is not just rare in the 3-torus, but far more rare than in all the random walks we compare with. We suggest that this is likely a result of the fact that vortices in random waves are not individual random walks constrained only by any statistical repulsion with themselves, but rather are interacting with the full ensemble of other vortex loops. This would explain a regime of behaviour in between the persistence length scale and the true random walk scale, in which the vortices are decorrelated with themselves but still strongly constrained by the surrounding vortex tangle and not free to display the unconstrained conformations of a true random walk. This would make both the wavelength and persistence length naturally poor measures of the lengthscale that is relevant for random walk behaviour and large scale knotting, unlike the random walks from the literature which are constructed directly and so do not experience such a regime of semi-self-interaction. It might instead be the case that the crossover scale for the random walk regime is the right lengthscale to measure random walks; we observed in Sect. 3.5 that this scale is around 3.5λ in the 3-torus. If we use this lengthscale as our parameterisation on wavelength then $\alpha_0 \approx 550 \lambda$, closer to that of the other random walks we investigate but still not entirely consistent.

It is not clear how to choose a similarly appropriate scale in the case of the 3-sphere and QHO. Their line length distributions do not display the same characteristic shape as in the 3-torus (Sect. 3.6), and in addition their extra conditions of a background curvature or classical radius may affect the strength of the effect, though it is unclear how much; for instance, the rapid knotting of the 3-sphere may be the result of a more complex interaction of this random walk lengthscale with the changes in local behaviour enforced by its metric, which might make the persistence length relatively good for parameterising the unknot probability and lead to a relatively small change in α_0 with this taken into effect.

It is also unclear how to choose an appropriate scale for the QHO, but we observe that it may be appropriate to compare it to random polygons confined inside a volume (e.g. a cubic cell or sphere). Although the scaling will depend on the nature of this confinement, the broad effect is well known to be an increase in the probability of knotting with length [15, 16], consistent with the QHO scaling as compared to that of the 3-torus. We check a particular result of this idea in Sect. 5.3.

Our suggestion that vortices display an important intermediate regime between the persistence lengthscale and random scaling is not a unique idea, being reminiscent of models for polymer melts such as we discussed briefly in Sect. 3.8. We noted that it is common to consider the polymer strands in a tube model, in which each polymer of small diameter occupies an effective tube of a much larger diameter, marking the effect of surrounding chains [17, 18]. The polymer strand may be geometrically non-trivial within this tube in accordance with its physically determined persistence length, but with its radial motion restricted.

This model is clearly highly focused on the physics of polymers, in which entanglement is directly physical (the chains may not pass through one another). In contrast, our vortex lines are not directly confined; each eigenfunction is static but, under a change of parameters, the vortices may move and reconnect with one another rather than being individually confined to a small local region. However, the correlation functions of vortices do suggest a statistical repulsion [19], so we consider that a conceptually similar model may apply; in any given eigenfunction most vortices will be surrounded by vortex segments from different loops, or far away (by arclength) sections of the same loop. This would define an effective tube not dissimilar to that of the polymer, potentially wide enough not to inhibit the persistence length on the wavelength scale, but not so wide as to allow random walk behaviour on this scale. The apparently-longer (roughly 3.5λ) length scale for random walks would be understood as relating to the large scale geometry of the tube itself, rather than the interior vortex on the wavelength scale. The intermediate regime would correspond to the semi-random behaviour of the vortex within the tube radius.

Given this idea that a tube-like mode increases the lengthscale of knotting, it is natural to seek to compare our vortices with knotting in polymer melts. They cannot generally be directly compared because polymer melts are ensembles of open curves, whereas closed loops may conceivably behave quite differently and may be difficult to simulate appropriately. Qin and Milner [14] uses a closed loop polymer model designed to properly represent the melt when equilibrated, beginning with random walks and following a progressive relaxation method that allows ‘rebridging’ moves

to change the topology. The unknotting probability with number of segments in the relaxed loops is shown in Fig. 5.13, with $\alpha_0 = 350 \pm 10$ vertices. It is unclear what the persistence length of these polymer loops should be, but it is unlikely to be significantly lower than the average vertex-vertex length, and so it seems that knotting in this model is not significantly rarer than in the other random walks we have mentioned.

This result does not match that suggested above for the 3-torus, but a deeper investigation requires a more precise comparison of the results of the tube model and for our vortex systems. It may be that even if the basic model applies, it is sensitive to parameters such as the effective radius of the tube and the *entanglement length* of the polymer, a number representing the average distance between entanglements and related to the knotting probability in [14].

5.3 Comparison of Knot Type Statistics

The analysis of Sect. 5.2 compares the global knotting probability between our systems and with other types of random walk, but ignores the probability of seeing different specific knots. We address this here.

In principle, we could plot graphs such as those in the previous section not just for knotting probability, but for the probabilities of each different knot type. This has been investigated in a range of different models, and [20] suggest by observation a general relation for the probability of a given knot, $P_k(N)$

$$P_K(N) = C(K)(N/N(K))^{m(K)} \exp(-N/\alpha_k), \quad (5.3)$$

where K is some knot parameterised by length N , α_k is an exponential scaling factor specific to the knot just as for the unknot above, and $C(K)$ and $m(K)$ are fitting parameters with $C(K)$ suggested from observations to in fact be independent of the knot type, and similarly $m(K)$ independent of the model. Unfortunately, our results are do not include enough data to resolve their values in our random vortices, and we instead make some simpler comparisons.

Figure 5.14 shows histograms of the 10 most common knots in the random walks of Sect. 1.6, and in all the data from our systems at different energies. The random walk is included only for basic comparison; following (5.3) the exact results will vary in different random walk constructions. This comparison is completely unnormalised, and is based on the raw data alone, so the results may be considered as knot probabilities weighted by the number of vortex lines at each arlength (in (a), the random walk, it is an approximately even sampling up to $\approx 3\alpha_0$). This Figure is included only to demonstrate the basic statistics of knots found in our eigenfunctions beyond our initial count of those that are torus, hyperbolic or satellite. We see that in all cases, the trefoil knot is most common, with the next three by probability being the figure-eight (4_1), composite double-trefoil (3_7^2) and 5_2 (one of two 5-crossing knots) in some order. This result is as much down to the distribution of vortex arlengths

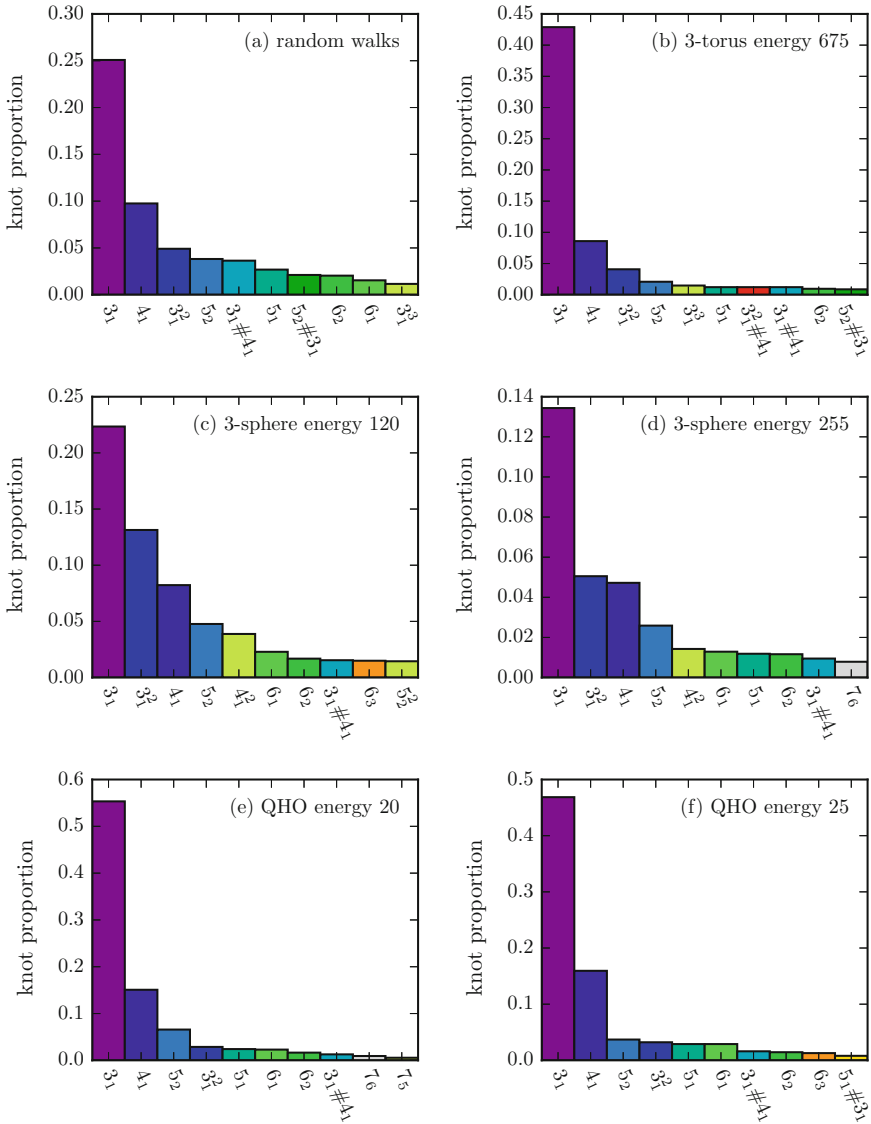


Fig. 5.14 Statistics of the most common knots found in different random wave systems. These results are not normalised, they count only the fractions of knots found throughout the datasets, without accounting for the different distributions of vortex lengths. **a** shows the result for the random walk model of Section A.2 sampled evenly up to $3\alpha_0$, while **b-f** are labelled with the system and energy that they are taken from. In each case, the histogram bar colour shows the knot type consistently

as the distribution of knotting probabilities; all our systems are biased towards low arclengths and so the knots that are common at these lengths are overrepresented. The other common knots somewhat reflect the individual systems, but with too much error to infer the influence of their different background conditions.

Figure 5.15 shows the same results under a very simple normalisation; each histogram is plotted only for knots with lengths between the intercept of the system's best fit line with probability 1 (i.e. approximately the length of the shortest possible knotted vortex), and this point plus $2\alpha_0$ in the system. The probabilities are crudely normalised by grouping this region into 6 bins, calculating the fraction of each knot type in each bin, and returning the average across all bins. This weights the results more evenly by arclength, though we must assume a large error in the resulting values, and will only comment on general trends.

Under this comparison, the different systems display the same similarity in that the trefoil, double-trefoil and figure-eight knots are most common, but there are more differences beyond this point. First, none match the distribution of knots in our example random walk model beyond the three most common configurations. The 3-torus actually only exhibits 8 different knots within the measured range, of which the last four or five are rare enough to likely only be statistical artefacts. We do note that the double-trefoil (3_1^2) and figure-eight (4_1) are switched in order from the random walk; while this could also be down to statistical fluctuation, it could perhaps also be a natural result of the correlation function for vortices discouraging their approaching closely. This would make the vortex curve in a sense self-avoiding, even if this averages out statistically when considering fractality, effectively reducing the arclength available for knotting (the vortex might behave more like a 'thick' rope than the infinitely thin one) and encouraging prime knots above composite ones.

Our results for the 3-sphere contain a larger number of more complex knots even at these low lengthscales, which is consistent with the previous conclusion that the manifold of the geometry encourages more complex knotting. The 3-sphere also has an extra condition in its construction; as discussed in Sect. 1.7.2, we simulate eigenfunctions over its entire volume and so they are periodic with every vortex pattern repeating identically on the 'other side' of the sphere. This might be expected to increase the number of 'squared' composite knots where a single rotation-symmetric vortex loop forms both prime components, which would be consistent with the observation that 3_1^2 is again the second most common knot, and also with the appearance of 4_1^2 in 3-sphere eigenfunctions of both energies even though it is not common in the random walk model or in any of our other eigenfunction examples. It might also explain why other composite knots are not common (only $3_1\#4_1$ appears in these plots) despite occurring relatively frequently in our random walk; having multiple knots in a single vortex loop consumes a lot of arclength, and if the prime knots do not form a mirrored pair then the entire composite knot must occur twice, which would naturally be relatively unlikely. If these explanations are correct then this would likely be visible in the chirality of the component prime knots, the symmetric pairs would be of opposite chirality in system-symmetric vortex loops, but we do not collate this information and investigate it here.

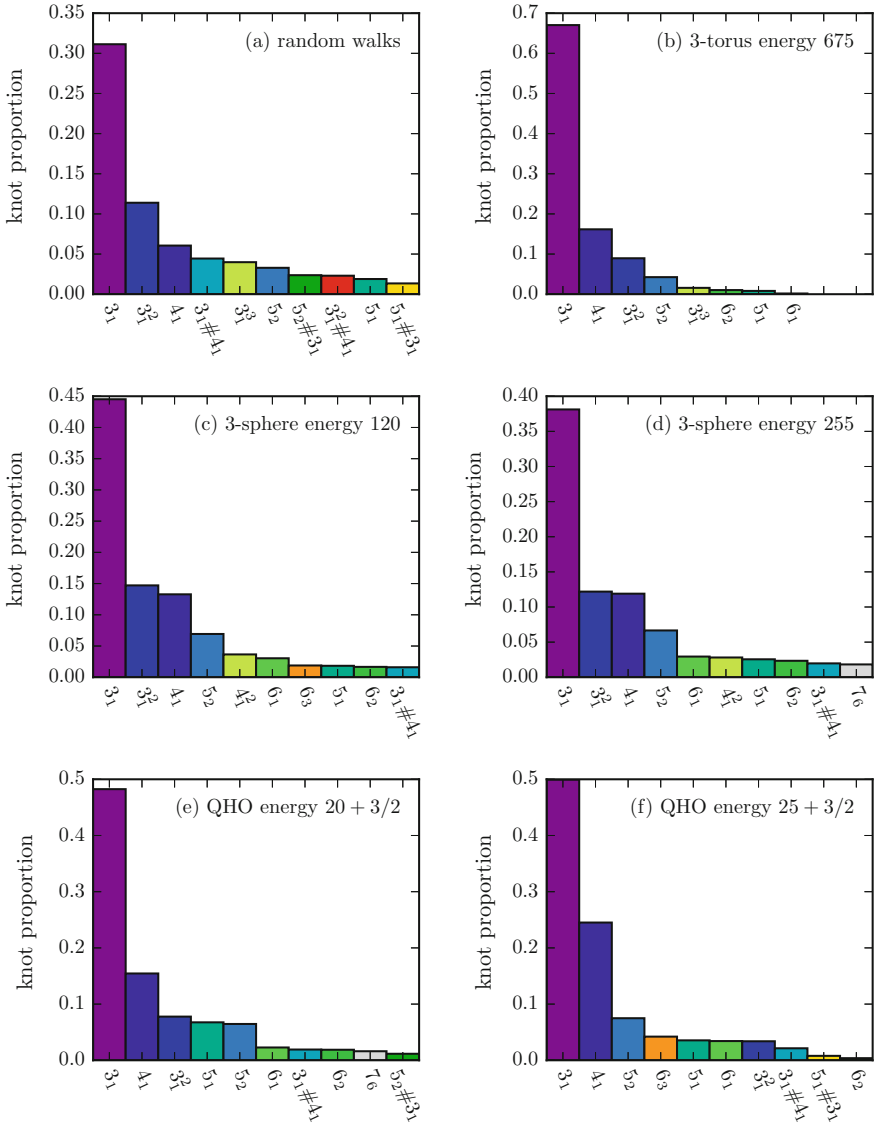


Fig. 5.15 Normalised statistics of most common knots found below $2\alpha_0$ in different random wave systems. These results are normalised by binning all knots between the probability 1 best fit lines of Sect. 5.2 and this value plus $2\alpha_0$ into 6 bins, calculating the probability of knots in each bin, and taking the average of these results

Our results for the QHO are not entirely consistent with these suggestions. Both energies include more complex knots than the 3-torus, again consistent with the extra conditions of the QHO encouraging more complex knotting, but although the QHO

also has a symmetry condition requiring vortex symmetry under inversion about the origin we do not see this expressed as a higher probability of squared knots. 4_1^2 does not appear at all here, and 3_1^2 is not especially common (at energy $25 + \frac{3}{2}$ it is relatively *uncommon*, though this plot is the least statistically reliable). It is not clear why the QHO should be so different to the 3-sphere in this respect; perhaps it is just the accumulation of errors (there are more data points available for the 3-sphere), or perhaps the potential of the QHO simply expresses itself quite differently to the metric of the 3-sphere. A stronger conclusion would require more data.

In the previous section, we suggested that the much higher probability of knotting in the 3-sphere and QHO as opposed to the 3-torus was due to the vortex being effectively confined, which might be analogous to other methods of confinement. Micheletti et al. [21] examine a model of confined random chains and find distinct signatures of this effect, in particular a favouring of chiral knots over those that are achiral—in the simplest case, a suppression of 4_1 in favour of (amongst others) 5_1 . We observe no such trend in our own results, 4_1 is actually more common in the QHO than in our other systems, and the only hint of an increased probability for 5_1 is its being relatively common in the QHO at energy $20 + \frac{3}{2}$, though the errors on these results would permit this to be only a statistical fluctuation. We thus cannot conclude that chirality is encouraged by the type of confinement of the QHO or 3-sphere, though our results are under quite different conditions to this previous work, and it is possible such an effect could be recovered under different parameters or at different lengthscales.

5.4 Finding Links

Measuring linking requires defining topology not for individual curves, but for multiple curves simultaneously. This poses different problems to knotting, and we address here the question of whether vortex curves in each of our systems can be said to be linked together.

Unlike with knotting, we are not able to define linking satisfactorily in every system, and will instead focus on the 3-sphere and 3-torus. We also note that the only invariant used is the linking number, as described in Sect. 4.2.6. This is a much less discriminating invariant than the Alexander polynomial or hyperbolic volume, with relatively simple links being indistinguishable from the unlink, and so we stress that the following results are strictly lower bounds on the frequency of linking.

5.4.1 The 3-torus

Linking is well defined between vortex loops on the 3-torus once they have been unwrapped through the periodic boundaries, with the extra condition that we must check linking also with all periodic copies after translation by the sidelength of the

3-torus some integer number of times in each axis. This poses a numerical problem in that the number of loop comparisons is increased by a non-trivial factor, but this does not limit our analysis at our chosen energy. We have also already shown an example of how a loop may form an infinite link with periodic copies of itself, in Fig. 5.1. It is clearly useless to try and sum the linking number over the entire infinite link, but we will deal only with the linking number of individual loops and so this does not pose a problem.

Figure 5.16 shows two example links from eigenfunctions of the 3-torus with energy 243. The first is a simple Hopf link between two relatively short loops, and the second part of a periodic Hopf link between two much longer vortices with length 1185λ , which would repeat infinitely if we included the other periodic copies. In fact such long loops are very commonly linked not just once but with multiple periodic copies as well as other different vortex loops; we do not recover all of this information in general, but Fig. 5.17 shows in directed graph form all the linking of vortices in the eigenfunction in which Fig. 5.16b appears, with every node denoting a loop of the given length and every edge denoting linking, labelled by the number of unit cell translations for the second loop along each Cartesian axis needed before the loops link in this way. Linking is quite common, with the largest loop linked 13 times with its own periodic copies as well as with two shorter loops. Such behaviour appears generic amongst 3-torus eigenfunctions; almost all long loops link with themselves as standard, as well as with others. This graph shows only what can be characteristic for eigenfunctions of energy 243; even at energy 675 there are many more links.

We do not investigate here linking in 3-torus eigenfunctions with energy 675, as linking is extremely common amongst vortex lines that are several hundreds or thousands of wavelengths in length, which also have a large extent and so must be compared with several periodic copies of every vortex curve. This makes numerical analysis impractical. For the same reason, we do not record individual linking





Fig. 5.16 Example linked vortex loops in the 3-torus. In each case, one vortex is plotted in green, and the other in red. **a** shows two vortices of lengths 81λ and 276λ in green and red respectively, forming a Hopf link. **b** shows a longer vortex loop of length 1185λ which is Hopf-linked with a periodic copy of itself; although only two copies are shown here, the same pattern repeats in an infinite chain if we tile space with 3-torus cells

Fig. 5.17 Graph of all linking in an example 3-torus eigenfunction. The eigenfunction has energy 243 and is the one in which the link of Fig. 5.16 appears. Every node marks a vortex loop of the given length, while every directed edge marks a link between the two loops after a periodicity translation the given number of times along the Cartesian axes labelled x , y , z . The specific numbers are arbitrary, depending on how the loops are numerically unwound



numbers of vortex loops in the 3-torus (instead only whether they are non-trivially linked with any other loop or not), and so do not plot their distribution here. We also investigate linking only amongst closed loops, by postselecting eigenfunctions that include no lines of non-trivial homology (roughly 1/20 of those with energy 243), in order to simplify the numerical detection of linking. These are numerous enough to determine the lengthscale of unlinking in our later analysis, though a full analysis would also differentiate between linking with NTH lines and closed loops.

With these caveats, our results on linking come from 1332 eigenfunctions with energy 243, which include 4412 loops with non-zero linking number. It is immediately clear that linking is more common than knotting, with an average of several links per eigenfunction. This also must be an underestimate, since there undoubtedly exist linked loops with linking number zero, though it is unclear what fraction of links these should make up and so our results are only a lower bound on linking.

5.4.2 The 3-sphere

As previously discussed, linking in the 3-sphere is defined simply and without ambiguity because all vortex loops occupy a simply connected volume that maps continuously to \mathbb{R}^3 , and we therefore investigate its statistics just as for knots. Figure 5.18 shows two examples of links, though in fact linking is extremely common and almost all long loops (including those demonstrated to be knotted in earlier sections) are linked with at least one other. Of these two, Fig. 5.18a is a Hopf link between two very short vortices and (b) a link with linking number 4 between the two symmetric copies of a single vortex loop under inversion of hyperspherical coordinates.

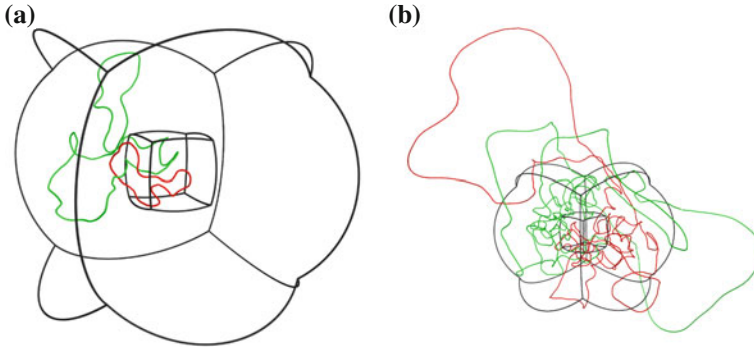


Fig. 5.18 Example linked vortex loops in the 3-sphere, in stereographic projection. In each case, one vortex is plotted in green, and the other in red. *a* shows vortices of length 11λ and 18λ (though the projection distorts their relative proportions), forming a Hopf link. *b* shows vortices both of length 80λ with linking number 4 between them; these are actually a linked pair of mirror copy vortices, where each point on one vortex is separated by π in every hyperspherical angle from a point on the other. In both cases, some of the octant boundaries from our hypercube of simulation cells are plotted to show the scale of the projection

Figure 5.19 shows the full distribution of linking numbers detected in 3-spherical eigenfunctions of energy 120 and 255. Linking is extremely common in both cases, with 13066 detected links amongst 30716 lines at energy 120 and 18213 detected links amongst 105990 lines at energy 255. Rather than showing each link with an individual point as with our sparser knotting data, we show the density of links of a given length against their sum of linking numbers, the total linking number of the given curve with every other vortex loop in the system. This number regularly rises to several tens at both energies, indicating that loops wind around one another many times.

5.4.3 The Quantum Harmonic Oscillator

The nature of infinite lines on the QHO makes it difficult to properly capture linking. Figure 5.20 demonstrates this with an invented example of two infinite lines closed via two different pairs of paths on an enclosing shell. Unlike with the knotting of an individual curve, different paths can give different results, in this case a Hopf link or unlink. These are not the only possible closures; the lines could spiral around the enclosing ball multiple times, resulting in other kinds of link when closed.

The linking of closed loops in the QHO is well defined, including when they are threaded by an infinite line, in which case closing the infinite line at infinity does give a unique result. Such links do exist, but due to the difficulty of defining linking in the system as a whole we do not examine them further.

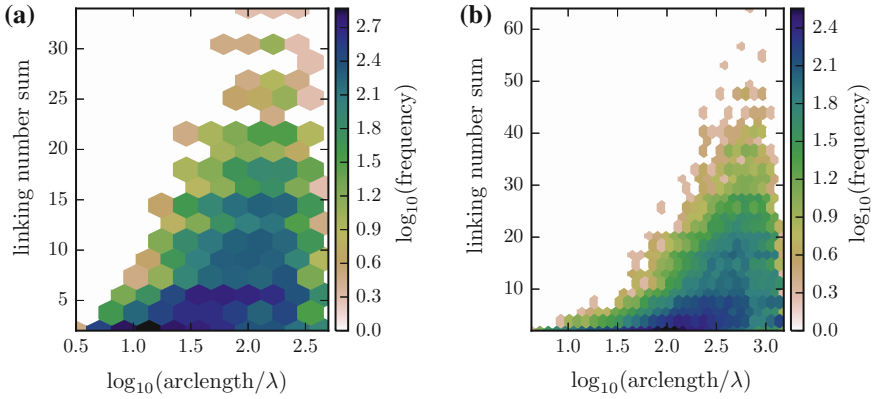


Fig. 5.19 Total linking number against frequency for links between vortices in eigenfunctions of the 3-sphere. **a** shows linking numbers at energy 120 and **b** at energy 255. The total linking number is the sum of linking numbers of the given loop with any other loops




Fig. 5.20 Two topologically different ways to close infinite lines on the QHO. The red and green lines each represent individually unknotted vortex lines, while the purple closures in each case show paths that could join the curves on an enclosing shell. The left and right closures yield a Hopf link and unlink respectively

5.5 Probabilities of Linking

Just as it is normal to expect the probability of unknotting with length to decay exponentially, we expect a similar relation for the probability of linking. Thus, the unlink probability would have the form

$$P(\text{unlink}) \propto \exp\left(\frac{-s}{\beta_0}\right), \quad (5.4)$$

where β_0 is a scaling factor quantifying the linking length scale, analogous to α_0 in the case of knots. Such a quantity is less discussed in the literature than the knotting


analogue, since it is inherently a property of multiple loops rather than applicable even to individual random walks but an exponential relation has still been observed [22].

Figure 5.21 shows probabilities of linking in the torus fit to this distribution. Linking at scales larger than in the plot is also common, but the number of such long loops is too low for reliable data analysis. The recovered β_0 is $370 \pm 10 \lambda$, only around 20% of the knotting lengthscale, suggesting that linking is far more probable and in keeping with our initial observation that the number of links per eigenfunction is quite high.

The shape of the plot also suggests a more complex behaviour at lower length scales; linking does not occur until around 130λ (the same lengthscale as for knotting), but then increases rapidly before settling to the slower exponential fit. It seems to be that 130λ marks a loose lengthscale at which the vortex suddenly has enough arclength to leave large enough holes that others can pass through it. As soon as these holes appear, other vortices wind through them, and linking jumps rapidly, then much beyond this lengthscale these small scale details average out and the exponential fit applies.

We can compare this result directly to that of [22] for linking in a (trily-periodic) optical speckle propagation model. They recover a lengthscale of approximately 30Λ for linking of loops with NTH lines or 130Λ for linking of loops with other loops, where Λ is an adjusted wavelength to take account for their Gaussian spectrum and anisotropic propagation model. Since we have postselected for eigenfunctions with no NTH lines, it is the former number that is most relevant. This is outwardly much lower than the β_0 of our system. The persistence length of the optical model is much smaller than a wavelength, in contrast to our own result and potentially making the wavelength the wrong lengthscale of comparison, but even adjusting by this factor gives around 100Λ ; still much lower than for vortices in our eigenfunctions. This may be an expression of how more fundamental properties of the field control knotting, and understanding how this ratio changes would require a more detailed investigation of vortex tangles with different parameters. O’Holleran et al. [22] also note that for

Fig. 5.21 Unlink probability with length for vortex loops in eigenfunctions of the 3-torus with energy 243. The best fit is to gradient $-(2.7 \pm 0.1) \times 10^{-3}$, and the data comes from 109341 closed loops amongs 1332 eigenfunctions



them to have found no knots in simulations of optical speckle, the lengthscale of knotting must be very large, on the order of $10^5 \Lambda$ or several thousand times their β_0 , another contrast to our monochromatic eigenfunctions in which for the 3-torus the difference is only a factor of approximately five. It is unclear why knotting should be so much rarer here, as in other measures the optical system does not appear dissimilar to our random waves [23].

Figure 5.22a, b show the unlinking probability with length for links in 3-spherical eigenfunctions with energy 120 and 255 respectively. At energy 120, linking is *extremely* probable, with an initial fit to $\beta_0 = (11.1 \pm 0.6)\lambda$, far less than the $\alpha_0 \approx 140$ in the same system. However, we suggest below that this comparison may not be between the same regimes.

The same results for energy 255 are more consistent with those for knots; the initial link scaling goes with $\beta_0 = (152 \pm 1)\lambda$, a slightly lower lengthscale than for knotting. This is to be expected in general, since forming a knot imposes strong statistical restrictions on local geometry unless the arclength available is far greater than necessary for a tight knot, while forming a link is much easier even for much shorter loops as they need only encircle an area large enough for another vortex line to pass through. As discussed above the true linking lengthscale will be lower, accounting for the many links whose linking number is zero, but it is unclear how large this effect will be.

The linking distribution in the 3-sphere displays yet another kink, this time directly related to that in the unknotting probability distribution of Fig. 5.11. In the case of knotting at energy 255, the value of α_0 *decreases* at a length of around 510λ , which we interpreted as loops beginning to occupy the full 3-sphere and so interacting more with themselves and gaining a higher probability of self entanglement. This change in β_0 is at a similar lengthscale but instead shows a significant *increase* to





Fig. 5.22 Unlink probability with length for vortex loops in eigenfunctions of the 3-sphere. **a** shows results at energy 120, with an initial best fit in red to gradient $-(9.03 \pm 0.06) \times 10^{-2}$ and later fit in purple to gradient $-(2.3 \pm 0.1) \times 10^{-2}$. **b** shows results at energy 255, with an initial best fit in red to gradient $-(6.58 \pm 0.06) \times 10^{-3}$ and later fit in purple to gradient $-(8.9 \pm 0.1) \times 10^{-4}$

Fig. 5.23 Average square linking numbers for vortex loops of different lengths in 3-sphere eigenfunctions of energy 255. The main scaling region is fit to a straight line with gradient 0.55 ± 0.1



$\beta_0 = (1123 \pm 3) \lambda$. This demonstrates the opposite effect to the increase in knotting probability; as the vortex line begins to tangle more with itself, its probability of tangling with other vortices grows more slowly. In an isotropic tangle of infinite volume, the line would have a greater chance of winding around new vortices after travelling for a while, and the rate of linking would remain more constant.

The kink in linking probability at energy 120 behaves similarly, this time with the fitted β_0 increasing from $(11.1 \pm 0.6) \lambda$ to $(43.5 \pm 0.1) \lambda$. These are both very significantly shorter lengthscales than those of knotting, but we suggest that in fact they are not the same regimes as observed at energy 255 as the observed vortex lengths are well below 200λ ; the cutoff is probably not between the main exponential scaling and the large-size effects that increase knotting, but instead between the initial regime where loops grow large enough to link (and linking suddenly jumps very quickly, as seen in the 3-torus) and then the normal exponential fit regime. We would expect that there is a later kink corresponding to that seen at energy 255, arising because vortices begin to occupy the whole 3-sphere, but following its location in the knotting results of Fig. 5.11 would not be seen until lengths of around 300λ . These lengths are not sampled here because linking is already so common that the difference cannot be detected.

A different measure of linking probability is to look at the distribution of linking number with length of the vortex loop. Following Fig. 5.19, there is a clear increase in the possible linking numbers with length, but it is not clear how the average behaves.

Figure 5.23 shows the mean square linking sum $\langle Ln^2 \rangle$ recovered from vortices in 3-sphere eigenfunctions with energy 255; this is the total linking number of each given curve with all others in the same eigenfunction. Like the unlinking probability, it displays at least three regimes; at low arclengths, linking is uncommon and the linking number grows only slowly. At higher lengths (in this case, $200 \lambda < s < 450 \lambda$), linking grows more rapidly, corresponding to the region identified as exponential in our discussion of unlinking probability. Finally, just as the unlink probability grows more slowly beyond a certain lengthscales, the linking number sum falls below this linear fit at this same lengthscales. Under our interpretation of this region as a finite




Fig. 5.24 Example vortices with non-trivial homology in the 3-torus. Both vortices are from the same energy 675 eigenfunction; **a** is a relatively short vortex loop with periodic length around 90λ , travelling just once along a Cartesian axis before closing, while **b** is a much longer loop with periodic length 3560λ that travels one cell in one axis and two in both of the others before closing. Each vortex is coloured along its length periodically, such that the open red ends close the curve when taking periodicity into account

scaling effect, this is a natural consequence of a vortex loop growing large enough that it begins to interact with itself more, rather than interacting with other new loops that it meets as its extent grows; just as this reduces the rate of linking, it is unsurprising that it should reduce the total linking number for loops that are linked.


The central regime, corresponding to the ideal exponential scaling of unlinking, is fit to a straight line of gradient 0.55 ± 0.1 . It is not clear what determines this gradient; it has been proposed that when considering pairs of just two random walks then $\langle \text{Ln}^2 \rangle \propto N^{1/2}$ for random walks with N segments [24], but the argument does not seem to directly extend to ensembles of many loops. In addition, these results come from linking in the 3-sphere which, following our discussion of unknotting and unlinking probabilities, may behave quite differently to an unbiased random walk.

5.6 Knotting in Lines with Non-trivial Homology

In previous sections, we explicitly ignored vortices of non-trivial homology in the 3-torus; those that wrap around the periodic boundaries some non-zero total number of times before closing. We will demonstrate here how, although they may clearly be topologically entangled in an intuitive sense, the standard tools of knot theory fail to quantify this topology. This is important because, following Sect. 3 (particularly 3.6), these lines often make up the majority of the arclength in a given eigenfunction, and so almost certainly include many knots.

Figure 5.24 shows two example NTH vortices in eigenfunctions of the 3-torus with energy 675. (a) is a particularly short example, wrapping around just one periodic side one time before closing with total length around 90λ , and (b) is a much longer NTH

Fig. 5.25 Examples of lines with non-trivial homology on the 3-torus constructed to depend on projection direction. In each case, the edges of the plotted cell are identified following their arrows, but the plotted region is also depthwise-periodic. Points where the line passes through this depthwise periodic interface are indicated with a dot. Each line is oriented such that the orientation points away from any depthwise crossing ‘down’ into the cell



loop with length 3560, wrapping once around one periodic axis and twice around both of the others before closing; these are only examples at some of the shortest and longest lengthscales, other NTH lines occupy the full spectrum of homology vectors in between, or may be even longer. Both are geometrically very similar to the long vortex loops that we have shown to be very commonly knotted in previous Sections. We will demonstrate that topology is not so easily characterised via the diagrams of Fig. 5.25.

Beginning with Fig. 5.25a, this simple tangle takes advantage of the periodic boundaries only in a very simple way; its ends close only on the boundary with no depthwise traversals, but the tangle in the middle looks very much like a trefoil knot. If we write down a Gauss code for this projection, we obtain 1+, 2-, 3+, 1-, 2+, 3- (ignoring crossing sign), exactly as would be expected from this understanding. However, unlike for knots in \mathbb{R}^3 , this projection direction is privileged, and we could instead look at the tangle along the axis of its periodicity, shown in Fig. 5.25b. The line is now oriented by the arrow such that the dot is an entry into the ‘top’ of the represented 3-torus cell, and when the line reaches the dot again it has exited through the ‘bottom’. This curve represents the same tangle as in (a), but the projection now has just a single crossing! Although we can recover the ‘obvious’ trefoil by cutting the periodic tails and closing along some bounding shell, the existence of a privileged projection is quite different to curves in \mathbb{R}^3 for which all projections are equivalent, and it is suddenly much less clear that the list of crossings necessarily captures the topology in all situations.

Figure 5.25c shows a more complex example, in which the periodic segment still only crosses the torus once in total, but now interacts with its periodic neighbours if

we imagine space tiled with periodic cells in this projection. Again, this represents a true kind of entanglement in that if we could manufacture such a system of string the lines would be woven together and could not be separated, but there is no longer an ‘obvious’ knot. It also does not seem that we can simply declare this a type of linking, as a single periodic segment of the unwrapped curve winds about multiple others simultaneously. We can also write down the direct Gauss code of the line counting the crossings only in a single periodic cell of the drawn projection to obtain $1 + 2 + 3 - 4 + 2 - 3 + 4 - 1 -$, but it is no longer clear if this method of counting crossings is even still meaningful; this Gauss code would describe the trefoil knot if considered as a representation of a curve in \mathbb{R}^3 , but this result would not be the same if a different projection were chosen.

These tangles *can* be analysed topologically (to a limited extent) using the concept of virtual knots [1]. These can be understood in multiple different ways, of which a convenient choice is the idea of extending the methods of knot theory to Gauss codes (or equivalently other knot representations) that could not be obtained via the planar projection of a normal knot. For instance, the Gauss code $1+, 2-, 1-, 2+$ could not be obtained in such a projection, and it would not even be possible to draw a normal knot diagram with this code without requiring at least one extra crossing to close the curve. However, this represents a valid virtual knot, and virtual knot theory includes analogues of many standard knot invariants that distinguish these objects. The connection to our three-dimensional periodicity is that virtual knots could be considered as representing planar diagrams drawn not on an infinite plane but on a surface with potentially non-zero genus; for instance a diagram with the above Gauss code can be drawn without difficulty on a 2-torus by closing the diagram through the hole of the torus without adding an additional self-intersection. Since virtual knots can describe certain kinds of knotting in periodic boundary conditions it is then natural to apply these ideas to our curves in the 3-torus, but unfortunately they do not resolve our general problem, and only classify the topology of our simple cases if we privilege projection direction to form a virtual knot by projection to a 2-torus. This is not even possible if the vortex loop wraps around the 3-torus in all three directions, as in Fig. 5.25d; this tangle wraps once around all three periodic axes of the torus, so there is no axis in which the projection maps simply to a ‘normal’ knot or even to one that could be analysed virtually.

We are not able to suggest a satisfactory extension of knotting or virtual knotting to capture this extra dimension of periodicity. However, it is possible to construct a consistent topological quantifier by unwrapping the NTH loop and cutting at its periodic boundaries (the representation used in Fig. 5.24) to leave an *open* curve. Such a curve cannot be classically knotted, but statistical methods have been suggested for performing closures and assigning different knot types with probabilities related to the frequency at which the closures (weighted in some appropriate fashion) form the given knot type [25]. These techniques have the advantage of (amongst other properties) largely agreeing with topological intuition where available, with ‘obvious’ knots such as the trefoil in Fig. 5.25a being identified as such except for some small probability of other types depending on the closure if the method allows closing lines to pass through the bulk of the tangle. There is no longer an obvious

knot type if the ends of the open curve are within the tangle bulk rather than well outside it, as may be the case for many NTH line, but the closure average would assign fractional knot types appropriately. Fully parameterising an NTH line in this way also requires averaging periodic cuts along its entire length, as the knot type returned by statistical closure depends on the line geometry and so will vary if the periodic cutoff in different places. For instance, if the period was cut in the centre of the trefoil knot of Fig. 5.25a, the average would likely return a high unknot fraction, in contrast to if cut at the periodic boundaries of the diagram in which case it would have a very high trefoil fraction. This would be an unusual extension of such open knotting algorithms, and it is not clear how well it would behave in general.

We suggest this method only as an example of how some tangling characteristics of NTH lines could be captured. In practice, such an algorithm is highly impractical; a statistical average over closures requires many knotting calculations even without having to additionally average over choices of periodic cut, and performing these calculations for real examples of long NTH lines consumes too much time to be useful. It is also unsatisfactory that performing such an algorithm gives only a statistical map to normal topology in \mathbb{R}^3 , when the curves seem to be tangled in a real way that we expect should permit some rigorous characterisation just as is provided for knots in $T^2 \times \mathbb{R}$ by the concept of virtual knotting.

5.7 Postselection Relations

It has been observed in many different random walks that the statistics of curves postselected by knot type do not always match those of the class of all loops. In particular fractal measures of the radius of gyration are different for specific knot types than for the ensemble of all curves [13, 26], and the geometrical tangling quantities of the writhe and average crossing number also scale differently [27]. We examine the first of these observations in our vortex system, in which the radius of gyration of the full loop ensemble was verified as consistent with a random walk in Sect. 3.5.

It was apparently first argued in [28] that topological constraints on an otherwise random walk would lead to an increased effective excluded volume about the curve, effectively leading to self avoiding scaling in fractal measures. This effect has been confirmed [29] and verified in several different random walk models [13, 26]. The core observation we will use is that

$$\langle r_{g,k}(s) \rangle \propto s^{-\nu} , \quad (5.5)$$

where $r_{g,k}$ is the radius of gyration of the knot type k (which may be any knot type) with arclength s , and the exponent $\nu \approx 0.588 \approx 3/5$ is standard for a self-avoiding walk (or walk with excluded volume) [29]. This relation coexists with the $\nu = 0.5$ scaling when considering all loops without postselecting by knot type; as s increases and r_g for simpler knots grows far larger than this average, many more complex knots

appear with lower r_g [26]. Although not postselected by knot type, r_g can be seen to vary significantly in Fig. 3.7, even at shorter lengths where knotting is uncommon.

We investigate this relation in our simulations of vortices in eigenfunctions of the 3-torus. Knots are more common in our other systems, but it is difficult to define a comparable fractal measurement accounting for their more complex geometry, and as discussed in previous sections it may be the case that knotting is so significantly affected by the curvature of the 3-sphere or potential of the QHO that (5.5) no longer applies.

Figure 5.26 shows the radius of gyration scaling both for all loops as discussed in Sect. 3.5, and for only the trefoil knots from this dataset. The smaller error bars on trefoil-only points at high arclength are because although the trefoils account for a small number of the total loops, their radius of gyration is more consistent than the full family.

The gradient of radii of gyration for trefoil knots is 0.55 ± 0.01 . This does not directly match the 0.588 expected for self-avoiding curves. However, we have performed the same analysis for generated random curves as described in Section A.2, and for random walks based on edge permutations of randomly oriented equilateral triangles as described and used in [13], recovering gradients of 0.557 ± 0.003 in the first case and 0.561 ± 0.002 in the second. These results are consistent with our scaling in random wave vortices. Although these results are more consistent, it is not clear why they do not match the expected 0.588, though it may simply be that (5.5) is an improper fit; it is really the first term of a power series which may include corrections at small arclengths, which may better fit our own curve. Equation (5.5) would only definitely be recovered at arclengths much larger than α_0 [29]. In contrast, the maximum length in Fig. 5.26 is 1000λ , well below the 3-torus α_0 that is about 1800λ . A stronger claim about the scaling would require considering significantly longer knots, though it remains unclear if this would significantly change our result; our scaling values recovered from other random walks seem constant even well above their α_0 , despite being well below the anticipated 0.588.

These results also make accessible the concept of *equilibrium length*, as described in [26]. This is the lengthscale at which a given knot type has ROG equal to that of the full loop ensemble, the value for arclength at which its line of best fit intercepts the line with gradient 0.5. Below this value, the knots of the given type are generally smaller than the average loop, while above it they are generally larger. The equilibrium length provides another measure of lengthscale for the random walk, and may be expected to reflect the differences in α_0 in different systems. In [26] it is identified as 170 for random walks as described in [30].

Comparing our trefoil scaling with the value for all random walks as described in Sect. 3.5 gives an intercept of 4300λ , around 25 times the result in [26], and similarly large even if considering distances in persistence length or some similar scale. The ratio is also significantly larger than that of α_0 , which is only around 10 times smaller in random walk models as discussed in Sect. 5.2.

Plotting radii of gyration in this way also reveals that the scaling gradient of 0.52 ± 0.01 determined previously may not well describe the longest knots that are insufficiently numerous to otherwise affect the measurement. In Fig. 5.26, r_g for these




Fig. 5.26 Radius of gyration against loop length, with loops postselected by topology. The original fit line refers to that obtained in Sect. 3.5, which is not strongly affected by the small number of high-arclength vortices. The large length fit is to only vortices of length greater than 100λ , and has gradient 0.48 ± 0.01 . The trefoil fit is to the radii of gyration of trefoil knots alone, and has gradient 0.55 ± 0.01

longer curves appears to consistently deviate from such a fit (although within errors), and we also fit the radius of gyration of these longer loops to the gradient 0.48 ± 0.06 . The intersection with this line would give an equilibrium length of approximately 1000λ , but may not represent the true large scale behaviour of the random wave model if the deviation is due somehow to the periodic boundary conditions. These two equilibrium lengths clearly bound a large range, and it is not possible to quantify the behaviour of our random wave vortices beyond the observation that their large equilibrium length is consistent with the lengthscale of knotting also being large compared to other random walks.

5.8 The Smallest Knot

Our previous analysis has focused on large scale questions about the statistical nature of knotting in random waves with no more bias than is enforced by the nature of each system. There are also natural questions arising from the opposite direction; what are the *simplest* wave superpositions that can produce a knotted vortex, creating the knot from the smallest possible number of interfering waves? This directly addresses the nature of wave chaos, being directly dependent on how each new wave in superposition increases the available space of local geometries, and eventually on how these combine to allow topological self-entanglement.

This has previously been investigated in [31] for superpositions of plane waves, beginning from the simplest system of a single plane wave and increasing the complexity through two (phase singularities are surfaces), three (the singularities become

filamentary vortices) and four superposing waves (vortex loops may form loops, or lines, or disappear depending on the wavevectors and amplitudes). Beyond four waves, the method of phasor analysis used becomes intractable. However, no four-wave superposition contains knotted vortex loops and so four waves is a minimum bound for superpositions to contain knotted vortices.

We approach the question instead from the direction of far more complex waves, with the systems investigated in previous sections being superpositions of hundreds of components. By reducing the energy, we reduce the number of waves in a way specific to each system, as explained in Sect. 1.7. We will investigate how the distribution of knots changes under such a reduction, in order to find a new (smaller) bound on the number of waves necessary.

Considering first the 3-torus, this system is highly unsuitable for such a search. There is no reason to expect that the high α_0 of Sect. 5.2 should significantly reduce as the scale is lowered, and so knotting is simply unlikely. In addition, since most of the arclength of an average cell is in lines of non-trivial homology most topological entanglement is difficult to detect. For example, Fig. 5.27 shows a relatively long example NTH line from a 3-torus eigenfunction with energy 75, a degeneracy of 56 different plane waves; although long enough to support a knot, its local geometry is unlikely to produce one, and it is roughly 90% unknotted following the statistical measure of Sect. 5.6. This seems typical, and we do not further investigate the 3-torus here.

Our other systems are both further removed from the direct random wave model. This has the disadvantage that the question of how many waves are needed to knot is not quite the same, depending on quite different sets of basis eigenfunctions, but the general nature of the problem remains the same and has the advantage that we have found knotting to be automatically more common in both the 3-sphere and the harmonic oscillator.

Continuing with the 3-sphere, we move to considering much lower energies than in previous discussion, with $N = 3, 4, 5$ such that energies are $E = 15, 24, 35$ and degeneracies 16, 25, 36 respectively. Figure 5.28 shows vortices in example 3-sphere eigenfunctions at each of these energies. All three clearly include much less vortex arclength than the higher energies in which statistical knotting was considered.

In 2710 eigenfunctions of energy 35, we find 32 knots, an average of one per approximately 85 eigenfunctions. In 2592 eigenfunctions of energy 25 we find 28 knots, an average of approximately one per 92 eigenfunctions, only slightly rarer, and with some of the difference undoubtedly due only to the lower vortex arclength in the eigenfunction. Finally, in 11656 eigenfunctions of energy 15 we find no knots; knots must be extremely rare, or not exist at all, and certainly are much rarer than at energies 25 and 35. This places a lower bound of 25 component waves to create a knot, though in this case these components are the hyperspherical eigenfunctions, not plane waves.

The rarity of knots is likely at least partially due to the periodicity of the vortex pattern under inversion of all the hyperspherical angles (explained in Sect. 1.7.2). This symmetry seems to forbid the existence of a single symmetrical trefoil or figure-eight knot, and instead each knotted eigenfunction contains either (at least) two mirrored




Fig. 5.27 An example NTH line in a torus eigenfunction of energy 75. The vortex curve has length around 250λ (of its periodic segment), and is roughly 90 % unknotted via the probabilistic closure of Sect. 5.6




Fig. 5.28 Vortices traced in low-energy eigenfunctions of the 3-sphere. **a–c** have energies 15, 24 and 35 respectively. Each vortex is plotted with a different random hue

knots, or a single composite knot including both components. Thus, to form any knot requires at least twice the arclength of vortex than might naively be expected. It may be that the reduction in energy to 15 lowers the average available arclength below some critical value to make knotting far less likely, if not impossible.

To further lower the bound on the necessary number of waves, we could track vortices in eigenfunctions that are not random, instead biasing the vortex distribution by setting the amplitude of one or more eigenfunction components to zero. It is not clear how the removal of each individual eigenfunction would affect knotting. We do not perform these measurements here.

Another alternative is to try to more directly construct a function containing knotted singularities of the phase. A method has been suggested for torus knots, but appears to require $N = 8$ or higher [32] (links do not require such a high value but we do not investigate this). This method is related to the mathematical construction used to synthesise knots in optical vortices in [33], in which an initial phase pattern in a two-dimensional plane fixes both the forwards and backwards propagation of its

singularities. We can also ask if this construction would map directly to the 3-sphere by placing the initial two-dimensional phase slice (which under optical propagation is a slice through a vortex knot) as a plane through the origin in stereographic projection of the 3-sphere; any results of the mathematical construction could then in principle describe knots in the 3-sphere if this plane obeys the correct hyperspherical inversion symmetry. In fact, the first knotted field configuration of this type with the correct symmetry is the (3, 5) torus knot, where these numbers indicate the standard indices of winding; in this case three times about the torus' axis of rotational symmetry, and five times through its hole. This once again would require $N = 8$, higher than the $N = 4$ at which we have found knotting, so neither of these constructions gives an improved bound. These observations are also consistent with our not seeing single simple knots such as the trefoil or figure-eight appear just once in an eigenfunction by including the hyperspherical inversion symmetry in their geometry. However, this is only a preliminary investigation; the method of [34] produces only one very characteristic braided representation of the knot and does not rule out the existence of representations with different symmetries that might in principle match those of the 3-sphere.

Considering instead the QHO makes available a similar strategy that may increase knotting; rather than using an unbiased random eigenfunction we can manipulate one or more components in order to demand that no vortex loops leave to infinity. If the main effect is to simply replace these ends with hairpins near to or outside the classical radius, the probability of knotting will be increased, as what might otherwise appear as linking becomes knotting of a single strand.

Since we are concerned with adjusting only the radial behaviour of the eigenfunction, it is natural to decompose the harmonic oscillator in spherical coordinates rather than the Cartesian coordinates that give Hermite modes as used throughout the rest of this work. Biasing the Hermite modes would in turn bias the vortex lines along the Cartesian axes, but in spherical coordinates we must modify only the single radial mode with the greatest radial extent; if its magnitude exceeds those of all the other modes, it will dominate as $r \rightarrow \infty$.

Under decomposition in spherical coordinates similar to that in Sect. 1.7.3, degenerate random eigenfunctions of the QHO (1.32) are

$$\psi_{klm}(r, \theta, \phi) = \sqrt{\frac{2}{\pi}} \frac{2^{k+2l+3} k!}{(2k+2l+1)!!} r^l e^{-r^2} L_k^{(l+\frac{1}{2})}(2r^2) Y_{lm}(\theta, \phi), \quad (5.6)$$

where $L_k^{(l+\frac{1}{2})}$ are the generalised Laguerre polynomials and Y_{lm} are the spherical harmonics as also used in the 3-sphere. The energy is given by $E = (2k + l + \frac{3}{2})$. Thus, we can write random degenerate eigenfunctions as

$$\Psi_N(r, \theta, \phi) = \sum_{k=0}^{N-2k} N/2 \sum_{m=-N+2k}^{N-2k} a_{k,m} \psi_{klm}, \quad (5.7)$$

where N is equivalent to $2k + l$ and therefore $l = N - 2k$, and the degeneracy is $(N + 1)(N + 2)/2$. This ensemble of eigenfunctions gives access to exactly the same space of solutions as the Cartesian Hermite modes used earlier, but by choosing N even we can prevent vortices travelling to infinity by manipulating only the mode with highest k (for which there is no m degeneracy). By setting its coefficient to exceed those of all other waves in the sum, all vortices must close.

Figure 5.29 shows examples of this procedure carried out at different energies, $N = 4, 6, 10$. In each case, the magnitude of the maximal k component was set to approximately exceed those of all other components, and vortices are successfully forced to close without significantly leaving the classical radius. However, in several hundred simulations we have found no knots or even loops with geometry that is visually complex. It seems that the procedure to prevent infinite lines has a strong secondary effect of encouraging vortices to lie along the zeros of the component eigenfunction with maximal k , which are a series of shells about the origin. This significantly affects the radial behaviour of vortices, and seems to strongly discourage knotting, even at energy 11 for which several knots were found in our unbiased simulations.

We do not further investigate the QHO here, although it would be possible to carry out a similar analysis to that of the 3-sphere. It is unlikely that the results would be significantly different, since the harmonic oscillator potential appears to exert a similar effect in practice to the spherical metric in our analysis of knotting probability. The QHO would not otherwise be expected to show significantly more knotting, unless it is possible to bias it less naively to increase the number of hairpins without trivialising the geometry as is the case in our method.




Fig. 5.29 Vortices traced in eigenfunctions of the harmonic oscillator that have been biased such that no vortices travel to infinity. In **a**, the energy is $4 + \frac{3}{2}$ and the Laguerre mode with $k = 4$ has its amplitude set to $6.5 + 6.5i$, producing 4 distinct vortex loops. In **b**, the energy is $6 + \frac{3}{2}$ and the Laguerre mode with $k = 6$ has its amplitude set to $6.5 + 6.5i$, producing 12 distinct vortex loops. In **c**, the energy is $10 + \frac{3}{2}$ and the Laguerre mode with $k = 10$ has its amplitude set to $12.5 + 12.5i$, producing 26 distinct vortex loops

5.9 Discussion

We have confirmed that knotting not only occurs in Gaussian random waves, but that at large lengthscales it is very common. In high energy eigenfunctions of the 3-torus that we consider likely closely approximate the isotropic random wave model, any line longer than a few thousand wavelengths will almost certainly be knotted. We have seen also that such vortex lines will far more commonly be linked, particularly after a certain minimal lengthscale.

Knots have also been shown to exist in random degenerate eigenfunctions of the 3-sphere and QHO as models of random wave systems that will approximate the behaviour of the isotropic wave model in the limit of large size, but whose strong constraints at lower energies give rise to some quite different features of the tangle. In both of these systems knotting was far more likely than in eigenfunctions of the 3-torus, with vortices needing only to be a few hundred wavelengths long to be almost certainly knotted. We have suggested that this should be understood as a result of the finite volume in each case; the 3-sphere has metric such that vortices travelling ‘away’ from a point will eventually come back towards it, while the QHO has a progressively larger potential with radius which encourages vortices beyond it classical boundary to turn back to the centre of the tangle in highly curved hairpins. We have seen that the effect of increasing knotting is not surprising, as confinement is recognised to have this effect on random walks.

What this does not make clear is exactly how the 3-sphere and QHO systems differ other than in the general sense. They seem potentially quite different in terms of what knots are most common, and it is unclear to what extent this is a small-sampling error or whether it results instead from some fundamental difference between the geometry imposed by each system’s metric. An interesting question that we have only briefly addressed is what are the more specific differences in these systems; are the results we see down to very specific small-scale effects, are these generic or unique to each system, and is it possible to predict at what energy the eigenfunctions of these manifolds would converge on the behaviour of isotropic random waves? Our investigation of eigenfunctions with only a small number of different energies is insufficient to ask these questions. These questions can also be framed in a more general context, systems such as the hydrogen atom or Airy function multiplied by sine curves in the orthogonal axis [35] also admit degenerate eigenfunctions with some of their own particular structure, and it is unclear how our results would be adapted to fit them.

A related question is what would happen if considering Gaussian random waves with different power spectra. All our systems are degenerate eigenfunctions and so inherently monochromatic, but the general random wave model does not demand this, and previous investigation of some vortex statistics has quantified significant geometrical differences between spectra [36]. Such a superposition could no longer describe an eigenfunction, but could describe for instance a non-monochromatic light field under the Helmholtz equation. Vortices in such a field may be quite geometrically different to our monochromatic examples, having different persistence lengthscales

with respect to the wavelength analogue and potentially exhibiting quite different topological statistics; in preliminary investigations we have found that knotting initially appears more common in Gaussian spectra (with random multiplier modulated by a Gaussian based on the wavevector magnitude), but it is unclear what determines this. Loosening the monochromatic requirement opens up continuous ranges of spectrum parameters that may allow this question to be investigated in detail.

Such investigations could still be carried out on the 3-torus for numerical convenience, by another restriction of wavevectors to particular points of a lattice, and this also suggests relations of these notions of random tangle to the generic tangle of [37] that we discussed in Sect. 3.6. The tangle in these general models has accessible its own spectra of behaviour, and it is possible that modifying our power spectrum might make more directly clear the relation between random wave vortex tangle and the statistical model. In particular, this might help to make clear the question of whether knotting displays universality in the same way as the geometrical scaling relations we discussed.

The relation with the generic tangle of Sect. 3.6 also returns to the problem of vortex loops with non-trivial homology in the 3-torus, which we have seen make up a natural part of such a periodic system, but whose topology we were not able to satisfactorily quantify. Understanding this may be important to understanding what aspects of topology are truly important to the system, not all of which are necessarily expressed in the occasional random wave superpositions in which all loops are closed.

We mentioned briefly in Sect. 3.8 that models of the mechanical tangle of polymer melts in a tube model might be applied to our vortex tangle, though we did not investigate such a relation in this thesis. An extra aspect of this relation is that [14] are able to relate the characteristic entanglement length of the polymer melt (which can be obtained via a geometric calculation) to the α_0 scaling of the knotting probability. Understanding this relation in our own system would give a direct relation between geometrical and topological constraints, and also be a measure of whether our quantised vortex tangle is similar to that of the polymer melt as this would demand a particular relationship between knotting and geometrical tangling lengthscales.

A core idea amongst all of these suggestions is the quantification of what vortex tangle characteristics are down to the specifics of a given system, and to what extent they must be shared by any tangle in an appropriate statistical limit; to return to our very first example of classical turbulence as observed by Leonardo da Vinci, is there some universality of hairiness? The answer seems less clear for topological measures than for geometrical quantities; we have seen that some results of random walks may naturally be applied, but questions remain about the lengthscales we obtain for knotting and linking, and how small scale geometrical considerations such as correlations beyond the persistence length may express themselves. Similarly, there seems to have not yet been any investigation of topology in the statistical models of [38], though these are expected to describe all of our different vortex systems (and many more) in an appropriate limit. Tying together these different ideas may be a natural direction for further investigation.

References

1. L.H. Kauffman, Virtual knot theory. *Eur. J. Comb.* **20**, 91–663 (1999)
2. D.W. Sumners, S.G. Whittington, Knots in self-avoiding walks. *J. Phys. A* **21**, 94–1689 (1988)
3. N. Pippenger, Knots in random walks. *Discret. Appl. Math.* **392**, 273 (1989)
4. M.L. Mansfield, Knots in hamilton cycles. *Macromolecules* **27**, 6–5924 (1994)
5. K. Koniaris, M. Muthukumar, Knottedness in ring polymers. *Phys. Rev. Lett.* **66**(17), 4–2211 (1991)
6. R. Lua, A. Borovinskiy, A.Y. Grosberg, Fractal and statistical properties of large compact polymers: a computational study. *Polymer* **45**, 717 (2004)
7. Y. Diao, C. Ernst, U. Ziegler, Random walks and polygons in tight confinement. *J. Phys. Conf. Ser.* **544**, 012017 (2014)
8. J. Arsuaga, M. Vázquez, S. Triguers, D.W. Sumners, J. Roca, Knotting probability of dna molecules confined in restricted volumes: dna knotting in phage capsids. *PNAS* **99**(8), 77–5373 (2002)
9. L. Dai, J.R.C. van der Maarel, P.S. Doyle, Effect of nanoslit confinement on the knotting probability of circular dna. *ACS Maroc. Lett.* **1**, 6–732 (2012)
10. T. Deguchi, K. Tsurusaki, Topology of closed random polygons. *J. Phys. Soc. Jpn.* **62**, 14–1411 (1993)
11. J. des Cloizeaux, M.L. Mehta, Topological constraints on polymer rings and critical indices. *J. Phys.* **40**, 70–665 (1979)
12. J.P.J. Michels, F.W. Wiegel, On the topology of a polymer ring. *Proc. R. Soc. A* **403**, 84–269 (1986)
13. N.T. Moore, R.C. Lua, A.Y. Grosberg, Topologically driven swelling of a polymer loop. *PNAS* **101**, 13431–35 (2004)
14. J. Qin, S.T. Milner, Counting polymer knots to find the entanglement length. *Soft Matter* **7**(22), 10676–93 (2011)
15. C. Micheletti, D. Marenduzzo, E. Orlandini, D.W. Sumners, Knotting of random ring polymers in confined spaces. *J. Chem. Phys.* **124**, 064903 (2006)
16. J. Arsuaga, M. Vázquez, S. Trigueros, D.W. Sumners, J. Roca, Knotting probability of DNA molecules confined in restricted volumes: DNA knotting in phage capsids. *PNAS* **99**(8), 5373–7 (2014)
17. T.C.B. McLeish, Tube theory of entangled polymer dynamics. *Adv. Phys.* **51**(6), 527–1379 (2002)
18. S.F. Edwards, D.A. Vilgis, The tube model theory of rubber elasticity. *Rep. Prog. Phys.* **51**, 243 (1988)
19. M.R. Dennis, Nodal densities of planar Gaussian random waves. *Eur. Phys. J. Spec. Top.* **145**, 191–210, June 2007. Conference on Nodal Patterns in Physics and Mathematics, Wittenbrg, Germany, 24–28 July 2006
20. T. Deguchi, K. Tsurusaki, Numerical application of knot invariants and universality of random knotting. *Banach Cent. Publ.* **42**, 77–85 (1998)
21. C. Micheletti, D. Marenduzzo, E. Orlandini, D.W. Sumners, Simulations of knotting in confined circular DNA. *Biophys. J.* **95**, 3591–3599 (2008)
22. K. O’Holleran, M.R. Dennis, M.J. Padgett, Topology of light’s darkness. *Phys. Rev. Lett.* **102**, 143902 (2009)
23. K. O’Holleran, M.R. Dennis, F. Flossmann, M.J. Padgett, Fractality of light’s darkness. *Phys. Rev. Lett.* **100**, 053902 (2008)
24. J.F. Marko, Linking topology of tethered polymer rings with applications to chromosome segregation and estimation of the knotting length. *Phys. Rev. E* **79**, 051905 (2009)
25. K.C. Millett, Tying down open knots: a statistical method for identifying open knots with applications to proteins. *Ser. Knots 17–203* (2005)
26. A. Dobay, J. Dubochet, K. Millett, P.E. Sottas, A. Stasiak, Scaling behavior of random knots. *PNAS* **100**, 5611–5615 (2003)

27. J. Portillo, Y. Diao, R. Scharein, J. Arsuaga, M. Vazquez, On the mean and variance of the writhe of random polygons. *J. Phys. A* **44**, 275004 (2011)
28. J. des Cloizeaux, Ring polymers in solution: topological effects. *J. Phys. Lett.* **42**, L433 (1981)
29. A.Y. Grosberg, Critical exponents for random knots. *Phys. Rev. Lett.* **85**(18), 61–3858 (2000)
30. K.V. Klenin, A.V. Vologodskii, B.B. Anshelevich, A.M. Dykhne, M.D. Frank-Kamenetskii, Effect of excluded volume on topological properties of circular dna. *J. Biomol. Struct. Dyn.* **5**(6), 85–1173 (1988)
31. K. O'Holleran, M.J. Padgett, M.R. Dennis, Topology of optical vortex lines formed by the interference of three, four and five plane waves. *Opt. Express* **14**, 44–3039 (2006)
32. T. Aissiou, S. Dyda, D. Jakobson, *Priv. Commun.* (2008)
33. M.R. Dennis, R.P. King, J. Barry, K. O'Holleran, M.J. Padgett, Isolated optical vortex knots. *Nat. Phys.* **6**, 118–121 (2010)
34. M.R. Dennis, Gaussian random wavefields and the ergodic mode hypothesis, in *New Directions in Linear Acoustics and Vibration*, ed. by M. Wright, R. Weaver (Cambridge University Press, 2010), pp. 59–76
35. W.E. Bies, E.J. Heller, Nodal structure of chaotic eigenfunctions. *J. Phys. A* **35**, 85–5673 (2002)
36. M.V. Berry, M.R. Dennis, Phase singularities in isotropic random waves. *Proc. Roy. Soc. A*, **456**, 79–2059 (2000). Including erratum
37. A. Nahum, J.T. Chalker, P. Serna, M. Ortuno, A.M. Somoza, Length distributions in loop soups. *Phys. Rev. Lett.* **111**, 100601 (2013)
38. A. Nahum, J.T. Chalker, Universal statistics of vortex lines. *Phys. Rev. E* **85**, 031141 (2012)

Chapter 6

Conclusions

This thesis has been an investigation of the large-scale tangle formed by vortices of the complex phase in wave chaos, as expressed in random degenerate eigenfunctions of the Laplacian in the 3-torus, 3-sphere and harmonic oscillator. The conformations of these lines appear morphologically similar to filamentary tangle in a diverse range of systems, and we have quantified many of these relationships through numerically tracking vortex lines with both high geometrical precision and large scale topological accuracy. With these methods we have investigated the characteristic properties of vortex curves both in the context of wave chaos and by comparison to other physical systems. A large part of our analysis covered their geometrical character, verifying known analytic results on the isotropic random wave model as well as obtaining numerically extensions that are not analytically accessible, and investigating the large scale statistical scaling properties of the vortex ensemble. We also performed a topological analysis of vortex loops, verifying the existence of knots and links and using their statistics to measure properties of the curve.

Obtaining these results from numerical simulation required new numerical methods optimised to the tracking of vortices in complex fields of known form. Our recursively resampled Cartesian grid algorithm, described extensively in Chap. 2, approaches this problem by taking advantage of vortices forming closed loops. It first creates an approximate map of vortex locations then re-resolves regions where vortex paths are ambiguous in order to guarantee local topological accuracy. This avoids calculating many unnecessary field points (that are not near vortices) while still allowing high precision in the recovery of local vortex geometry by re-resolving these regions at a high resolution once the vortex paths have been approximately located. The method is flexible enough to recover vortices with tens or hundreds of tracked points per wavelength in smaller systems, or to track vortices with across large lengthscales without losing local topological accuracy by sampling at a relatively low resolution and resampling only enough to unambiguously follow vortex paths where they approach closely. It also generalises well to our different systems of wave chaos; the 3-torus is realised by the addition of periodic boundary conditions, the QHO by tracking vortices only around the classical region where energy exceeds the potential, and in the 3-sphere by covering its volume with eight locally-Cartesian

cells with appropriate faces identified with one another in order to track vortices across the full volume of the manifold. In this way, the advantages of the RRCG method are preserved in each case.

With these numerical methods, we analysed in Chap. 3, first aspects of the statistical geometry and scaling of vortex curves in the 3-torus, as the closest numerically-accessible approximation to the isotropic random wave model. On small scales their conformations are described locally by a curvature and torsion; we confirmed analytic results regarding the curvature, and recovered numerically the torsion (which appears analytically inaccessible), as well as analysing the correlations of different local geometric quantities. These results suggest that vortices have a random nature beyond a short lengthscale of approximately a wavelength, and through analysis of the fractality of individual vortex curves we resolved that vortices indeed have the highly characteristic scaling of a random walk, consistent with previous results on different random wave systems and a property in common with filamentary tangle with more diverse sources including polymer melts and models for cosmic strings. Analysing a different aspect of the tangling via the average crossing number and writhe of vortex curves was consistent with this random walk result.

These results apply to individual vortex curves, but we also investigated properties of the full set of vortex curves in a given eigenfunction. Analysing the fractality of this tangle revealed a box counting fractal dimension close to 3, indicating a level of uniformity and lack of voids at all scales, a result that may potentially vary in other systems of filamentary tangle even if each filament is a random walk. We compared also the statistical distribution of vortex loop lengths, whose scaling in a truly random, scale-invariant ‘loop soup’ is expected to follow a characteristic power law, but which in practice has been observed to include classes of vortex line that do not fit this pattern. These differences in our systems were demonstrated to match perfectly with more advanced predictions in recent literature that account for the effects of periodic boundary conditions or finite volume.

The second part of our analysis consisted of topological investigation. The relevant knot theoretic background was given in Chap. 4, with particular emphasis on the choices of algorithmic tools necessary to perform topological analysis on the very long and geometrically complex conformations of vortex curves. By choosing topological invariants based on the Alexander polynomial and hyperbolic volume we were able to adequately classify even these long curves, obtaining a loose notion of their complexity and information about whether the knot was composite. As an aid to this procedure it was also important to perform three-dimensional simplification of some vortex curves (in a topology-preserving way), for which we described an efficient octree-based algorithm that performs particularly well for numerically recovered vortex curves.

In Chap. 5 we applied these topological tools to the vortex loops of random wave tangle, establishing that knots can and do appear generically in all of our systems of wave chaos. However, comparison of their statistics with those anticipated for random walks in unbiased models revealed significantly different scalings of knotting as a function of vortex arclength; knots are comparatively extremely rare in the 3-torus as a function of geometrical decorrelation lengthscales, and by extension are

likely also rare in the isotropic random wave model. It was suggested that this arises from an intermediate vortex regime between the lengthscale of geometrical decorrelation and the lengthscale of random walk behaviour, in which the conformations of vortex lines are strongly influenced by their neighbours, in contrast to the behaviour of a truly unbiased random walk. Investigation of the 3-sphere and QHO revealed the opposite result, knots are extremely common even for short (on the scale of correlation distance) vortices, particularly beyond a characteristic cutoff lengthscale after which knotting is even more likely. This was explained as being due to an effective confinement, not via hard boundaries but the finite volume of the 3-sphere and finite classical volume of the QHO; once vortices become large enough, they 'fill' this region and tangle with themselves more than would otherwise be the case, an effect consistent with investigation of confinement in the literature. Analysis of topological linking supported these results and demonstrated the same effect in reverse; linking is still extremely common but is discouraged at larger lengthscales, as loops knot with themselves instead of leaving their local volume and linking with neighbours.

Our remaining analysis followed several individual strands of topological investigation. Comparing the specific knot types that appear in each different system revealed significant differences between the eigenfunctions and the behaviour of a selected model for an unbiased random walk, some of which was understood in terms of the physical constraints dominating behaviour as discussed above, although the causes of some patterns remain unclear. Separating curves by knot type also made accessible the relationship between topology and fractal scaling, observed in the literature for other random walk models and consistent in our own results, and giving another lengthscale for random knotting. Finally, we made use of the numerical tools of these investigations to investigate the quite different question of what is the smallest number of interfering waves whose vortex pattern may be knotted, which is known to be impossible with just five plane waves but has not been demonstrated with fewer than at least several tens of waves. Though brute force simulation we established that apparently 25 waves are necessary in the 3-sphere (and knots are at least very rare with fewer), though an attempt to optimise this procedure on the QHO was not successful.

A natural direction for future work is to further some of this later analysis; our results indicate certain strong patterns in the statistics of different knot types in different systems, but it is unclear how exactly they depend on the physical constraints of the field beyond the most extreme effects like that of confinement. This has been found in the literature to have more subtle effects in different random walk scenarios, but it was not possible to distinguish these in the investigation described. Such effects may not even hold, or may be modified, since the vortex tangles of wave chaos involve many vortices in ways that we have already demonstrated to affect the statistics of an unbiased random walk. Our systems of wave chaos also have other properties that may affect their bulk statistics such as the symmetries of the 3-sphere and QHO under transportation to the antipodal or inversion points (through the origin) respectively. These did not appear to play a strong role in our analysis, but could affect the investigation of more subtle quantities.

These questions exist in the wider context of querying vortex tangle itself, not in the terms of any individual system but the general statistics of collections of random walks. Our analysis has suggested that various geometrical and topological statistics do not map directly between such tangles, including that the fractality of the ensembles may vary (even though each individual curve is a random walk) and that large scale topological scaling may depend on small scale interactions. However, it is not clear what governs these differences, whether there is some further spectrum of parameters that limit the possible results, and if any quantities of the full tangle may display universalities in the same way as the randomness of individual curves. Recent work on analytic investigation of scaling laws in random tangle, discussed in Chap. 3 in the context of loop length distributions, may provide a route for such investigation, but current results do not look at topological quantities and it is unknown if they admit similar universalities.

Appendices

A.1 Hyperspherical Harmonics

The following reproduces notes of Professor Mark Dennis on the derivation of hyperspherical harmonics in hyperspherical coordinates via Gegenbauer polynomials. It is this analytic form that we use in our numerical investigation of the 3-sphere, introduced in Sect. 1.7.2 and explained in Chap. 2, though other representations would be equivalent under an appropriate weighting.

We here describe some of the coordinate systems for the 3-sphere based on straightforward generalization of the 2-sphere, and the spherical harmonics, defined as the natural orthonormal eigenfunctions of the Laplace-Beltrami operator on the 3-sphere.

Most simply, the (unit) 3-sphere in 4-dimensional euclidean space with cartesian position (x, y, z, w) is defined as the locus

$$x^2 + y^2 + z^2 + w^2 = 1.$$

Positions on the 3-sphere are specified by three angles, ϕ , θ and ψ , with $0 \leq \theta, \psi \leq \pi$ and $0 \leq \phi < 2\pi$, in a way generalizing the coordinates of the 2-sphere:

$$\begin{aligned}w &= \cos \psi, \\z &= \sin \psi \cos \theta, \\x &= \sin \psi \sin \theta \cos \phi, \\y &= \sin \psi \sin \theta \sin \phi,\end{aligned}$$

which clearly satisfies the 3-sphere condition. Clearly, the ‘equator’ with $w = 0$, $\psi = \pi/2$ corresponds to a regular 2-sphere with coordinates θ, ϕ .

We also frequently use ‘complex coordinates’ u, v for the 3-sphere, with

$$u \equiv w + iz, \quad v \equiv x + iy,$$

clearly therefore satisfying $|u|^2 + |v|^2 = 1$. In particular, the argument of v is ϕ , which is useful in identifying $e^{i\ell\phi}$ behaviour in 3-sphere harmonics.

The most frequent representation of the 3-sphere (or the group of 3D rotations/SU(2), which has the 3-sphere or a related space as its group manifold) is by ‘stereographic projection’ into regular 3D euclidean space with position $\mathbf{R} = (X, Y, Z)$, most conveniently with the assignment $\mathbf{R} \equiv \tan \psi/2(\cos \phi \sin \theta, \sin \phi \sin \theta, \cos \theta)$, i.e. like usual spherical coordinates with $R = \tan \psi/2$. This can be thought of as ‘south-pole projection’ of the w -axis of the 3-sphere, as may be readily checked from the complex coordinates assignment

$$u = \frac{1 - R^2 + 2iZ}{1 + R^2}, \quad v = \frac{2(X + iY)}{1 + R^2}.$$

We note it has often been convenient in previous work to use north-pole projection ($\text{Re } u = (R^2 - 1)/(R^2 + 1)$), which would lead to $\cot \psi/2$ in the stereographic projection of our chosen angles; we adopt south-pole projection here.

These different coordinates give rise to a metric

$$\begin{aligned} ds^2 &= g_{ij} dx^i dx^j \\ &= d\psi^2 + \sin^2 \psi (d\theta^2 + \sin^2 \theta d\phi^2) \\ &= \frac{4}{(1 + R^2)^2} (dR^2 + R^2 (d\theta^2 + \sin^2 \theta d\phi^2)), \end{aligned}$$

which has volume element $\sin^2 \psi \sin \theta d\psi d\theta d\phi$ (or, stereographically projected, $8(1 + R^2)^{-3} R^2 \sin^2 \theta dR d\theta d\phi$), yielding a volume of $2\pi^2$.

This metric allows the Laplace-Beltrami operator Δ to be defined (we use Δ rather than ∇^2 as we are on a curved manifold), where as usual $g = \det g_{ij}$,

$$\begin{aligned} \Delta &= \frac{1}{\sqrt{g}} \partial_i (\sqrt{g} g^{ij} \partial_j) \\ &= \frac{1}{\sin^2 \psi} \partial_\psi (\sin^2 \psi \partial_\psi) + \frac{1}{\sin^2 \psi} \left(\frac{1}{\sin \theta} \partial_\theta (\sin \theta \partial_\theta) + \frac{1}{\sin^2 \theta} \partial_\phi^2 \right), \end{aligned}$$

where in the last line, the term in large brackets is the usual laplacian on the 2-sphere.

We now deduce the form of a natural set of eigenfunctions of Δ in $\psi\theta\phi$ coordinates. Since the regular 2-sphere laplacian occurs in our expression for Δ , we assume the desired functions have the form

$$\mathcal{Y}_{n,\ell,m}(\psi, \theta, \phi) = \Psi_n^\ell(\psi) \sin^\ell \psi Y_\ell^m(\theta, \phi),$$

for the usual spherical harmonics $Y_\ell^m(\theta, \phi)$, which are orthonormal on the 2-sphere and have eigenvalue $-\ell(\ell + 1)$, with $2\ell + 1$ -fold degeneracy. The $\sin^\ell \psi$ factor has been included following the angular forms of x, y, z above: each Y_ℓ^m can be

expressed as an order- ℓ complex homogeneous polynomial of cartesian coordinates on the 2-sphere.

Writing $c \equiv \cos \psi$, $\Psi(\psi) = X(c)$, we can rewrite

$$\begin{aligned} \Delta \mathcal{Y} &= \frac{Y_\ell^m}{1-c^2} \left(-\sqrt{1-c^2} \partial_c \{ -(1-c^2)^{3/2} \partial_c \} [(1-c^2)^{\ell/2} X(c)] - \ell(\ell+1)(1-c^2)^{\ell/2} X(c) \right) \\ &= (1-c^2)^{\ell/2} Y_\ell^m \left[(1-c^2) X'' - (2\ell+3)cX' - \ell(\ell+2)X \right], \end{aligned}$$

and recognising the similarity of the last line with the Gegenbauer differential equation

$$(1-x^2)y'' - (2a+1)xy' + n(n+2a)y = 0$$

whose solutions are the Gegenbauer polynomials $C_n^{(a)}(x)$, we conclude that

$$\mathcal{Y}_{n,\ell,m}(\psi, \theta, \phi) = \sqrt{\frac{2^{2\ell+1} n! (1+\ell+n)}{\pi (1+2\ell+n)}} \ell! \sin^\ell \psi C_n^{\ell+1}(\cos \psi) Y_\ell^m(\theta, \phi),$$

where the appropriate Gegenbauer normalization factor has been included. We thus have

$$\Delta \mathcal{Y}_{n,\ell,m}(\psi, \theta, \phi) = -(n+\ell)(n+\ell+2) \mathcal{Y}_{n,\ell,m}(\psi, \theta, \phi),$$

so the eigenvalue is (unsurprisingly) labelled by $n + \ell \equiv N$, and we rewrite

$$\mathcal{Y}_{N\ell m}(\psi, \theta, \phi) = \sqrt{\frac{2^{2\ell+1} (N-\ell)! (1+N)}{\pi (1+\ell+N)}} \ell! \sin^\ell \psi C_{N-\ell}^{\ell+1}(\cos \psi) Y_\ell^m(\theta, \phi),$$

with Laplace-Beltrami eigenvalue $N(N+2)$. Since each ℓ is $2\ell+1$ -fold degenerate, the total degeneracy for the eigenvalue labelled by N is

$$\sum_{\ell=0}^N (2\ell+1) = (1+N)^2.$$

A.2 A Random Walk Construction

We give here the algorithmic details of generating closed random walks via the method of [1]. A full explanation of the mathematical details and relevant proofs is given in this original source.

The algorithm in the general case returns an unbiased random walk of N linear segments (which therefore generically fails to close by some random distance), from the space of all random walks with a beta-distribution on edge lengths which is mapped to the $(4N+1)$ -sphere. The algorithm is as follows:

- Select at random a point on the $(4N + 1)$ -sphere, i.e. an array of $4N$ real entries with modulus 1.
- Partition the vector as N sets of 4 numbers. Each of these is identified as the real, i, j and k components of a quaternion; we label them p_1, p_2, p_3 and p_4 .
- Convert each quaternion to a point on the 2-sphere via the Hopf map. In practical terms, the x, y and z Cartesian positions are given by

$$x = p_1^2 + p_2^2 - p_3^2 - p_4^2, \quad (\text{A.1})$$

$$y = 2(p_2p_3 - p_1p_4), \quad (\text{A.2})$$

$$z = 2(p_1p_3 - p_2p_4). \quad (\text{A.3})$$

- Each of these positions defines one edge of the random walk, of length $\sqrt{x^2 + y^2 + z^2}$. Recover the random curve by joining edges as sequential vectors.

To create a closed curve, as we need for an investigation of knotting and comparison with our vortices, [1] isolate positions on the $(4N - 1)$ -sphere corresponding to random walks that end with distance 0 from their origin. The above algorithm is modified to select only such points via a construction of complex numbers, as follows:

- Generate two orthogonal Gaussian random vectors with scalar product zero, each of N entries. By first generating two independent Gaussian random vectors labelled \mathbf{u}_i and \mathbf{v}_i , we orthogonalise these as

$$\mathbf{u}_f = \mathbf{u}_i / \sqrt{|\mathbf{u}_i|^2}, \quad (\text{A.4})$$

$$\mathbf{v}_1 = \mathbf{v}_i - \mathbf{u}_f^* \cdot \mathbf{u}_f, \quad (\text{A.5})$$

$$\mathbf{v}_f = \mathbf{v}_1 / \sqrt{|\mathbf{v}_1|^2}, \quad (\text{A.6})$$

such that \mathbf{u}_f and \mathbf{v}_f are normalised and orthogonal.

- Perform the Hopf map on these components to retrieve a set of edges. This now has the form

$$x = \text{Re}(\mathbf{u}_f \cdot \mathbf{u}_f^* - \mathbf{v}_f \cdot \mathbf{v}_f^*), \quad (\text{A.7})$$

$$y = \text{Re}(i(\mathbf{u}_f \cdot \mathbf{v}_f^* - \mathbf{v}_f \cdot \mathbf{u}_f^*)), \quad (\text{A.8})$$

$$z = \text{Re}(\mathbf{u}_f \cdot \mathbf{v}_f^* + \mathbf{v}_f \cdot \mathbf{u}_f^*). \quad (\text{A.9})$$

- Join each of these edges as before, though now they will form a closed polygon.

The result is a closed polygon of N segments, still with a beta distribution on segment lengths.

Reference

1. J. Cantarella, T. Deguchi, C. Shonkwiler, Probability theory of random polygons from the quaternionic viewpoint. *Commun. Pur. Appl. Anal.* **67**, 1658–1699 (2014)