



Būtaitė, Unė Gabrielė (2020) Enhanced optical tweezing: from hydrodynamic micro-manipulation to optimised optical trapping. PhD thesis.

<https://theses.gla.ac.uk/81662/>

Copyright and moral rights for this work are retained by the author

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge

This work cannot be reproduced or quoted extensively from without first obtaining permission in writing from the author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Enlighten: Theses

<https://theses.gla.ac.uk/>
research-enlighten@glasgow.ac.uk

UNIVERSITY OF GLASGOW

DOCTORAL THESIS

Enhanced optical tweezing: from
hydrodynamic micro-manipulation
to optimised optical trapping

Author:

Unė Gabrielė BŪTAITĖ

Supervisors:

Dr. Jonathan M. TAYLOR

Dr. David B. PHILLIPS

Dr. Johannes COURTIAL

*A thesis submitted in fulfilment of the requirements
for the degree of Doctor of Philosophy*

in the

School of Physics and Astronomy
College of Science and Engineering

2nd September 2020

UNIVERSITY OF GLASGOW

ABSTRACT

Enhanced optical tweezing: from hydrodynamic micro-manipulation to optimised optical trapping

by Unė Gabrielė BŪTAITĖ

Optical tweezers, with their ability to trap particles at the focus of a laser beam and control their motion, have provided unparalleled insight into the inner workings of the micro-world of colloids, cells and biomolecules. However, not all materials yield to optical trapping, and living organisms can be damaged by direct light exposure.

Many optical tweezing experiments are performed in aqueous environments. This offers a route to indirect particle manipulation via the surrounding fluid. We develop, study, and experimentally demonstrate an approach which, by employing optically controlled micro-rotors to induce flow currents in the surrounding fluid, exerts near field hydrodynamic control over freely diffusing particles, irrespective of their material.

With our optically actuated hydrodynamic manipulation we were able to suppress the thermal motion of single sedimented micro-sized target particles of various materials in both translational and rotational degrees-of-freedom, translate individual particles over complex local trajectories or transport them over long distances across the holding sample cell, and exert control over multiple particles simultaneously.

The biggest challenge in our hydrodynamic manipulation technique is the accuracy with which the optical tweezers can control the actuator motion. This boils down to optical trapping stiffness. We employ the generalised Wigner-Smith operators alongside an optimisation scheme to upgrade the optical trapping field from the standard Gaussian beam to a three-dimensionally stiffness-enhanced trap.

Within simulations we demonstrate light fields with an order of magnitude stiffness enhancement in all three dimensions simultaneously. Such fields, as well as the techniques used to develop them, can find applications throughout the wide community of optical trapping and manipulation.

Publications

Unė G. Būtaite, Graham M. Gibson, Ying-Lung D. Ho, Mike Taverne, Jonathan M. Taylor and David B. Phillips. [Indirect optical trapping using light driven micro-rotors for reconfigurable hydrodynamic manipulation](#), Nature Communications **10**, 1215 (2019)

Conference proceedings

Une G. Butaite, Graham Gibson, Jonathan Taylor and David B. Phillips. [Hydrodynamic micro-manipulation using optically actuated flow control](#), SPIE Proceedings, Optical Manipulation Conference, (24 April 2018)

Une G. Butaite, Graham M. Gibson, Ying-Lung D. Ho, Mike Taverne, Jonathan M. Taylor and David B. Phillips. [Enhanced optical trapping](#), SPIE Proceedings, Optical Manipulation and Structured Materials Conference, (15 June 2020)

Declaration of Authorship

I, Unè Gabrielè BŪTAITĒ, declare that this thesis, titled ‘Enhanced optical tweezing: from hydrodynamic micro-manipulation to optimised optical trapping’, and the work presented in it are my own, unless explicitly stated otherwise in the text or in the Contributions section. I confirm that:

- This work was done wholly while in candidature for a research degree at this University.
- This thesis has not previously been submitted for a degree or any other qualification at this University or any other institution.
- Where I have consulted the published work of others, this is always clearly attributed.
- I have acknowledged all main sources of help.

Signed:

Date: 25 June 2020

Acknowledgements

First and foremost I would like to express my gratitude towards my supervisors. Dave Phillips, without your brilliant mind this project would not exist, and without your untiring enthusiasm it would certainly not have progressed as far as it did. Jonny Taylor, thank you for agreeing to take on this project as the main supervisor, and for your invaluable guidance during it, especially through the mathematical landscapes we encountered. I learned a lot from both of you and I look forward to learning even more in the future.

I would like to thank Miles Padgett whose lab I frequented for my experiments and which, of course, cannot be imagined without Graham Gibson - without his expertise and help my project would have tripped over quite a few hurdles. I would also like to thank Richard Bowman, whom I only met very briefly, but whose work proved to be of great inspiration and value to me. Further thanks go to Johannes Courtial, who was somewhat unexpectedly thrown into the official position of my second supervisor in the middle of the project, and who was happy to chat whenever needed.

And last, but certainly not least, I would like to thank FK for greeting me with a big smile every night I came back late after a day of working on this dissertation, and who, during the time it took me to write up, managed to create an entire novel more than double in volume of my thesis.

Contents

Abstract	i
Publications	ii
Declaration of Authorship	iii
Acknowledgements	iv
1 Introduction	1
Part I	4
2 Background theory	5
2.1 Where it all begins - the Langevin equation	5
2.2 Brownian forces	7
2.3 Hydrodynamic forces	7
2.3.1 Fluctuation-dissipation theorem	9
2.3.2 The hydrodynamic mobility tensor	10
2.4 Optical forces	13
2.4.1 The operating principle of optical tweezers	13
2.4.2 The tweezing force	15
2.5 Revisiting the Langevin equation	17

2.5.1	Behaviour of an optically trapped particle	18
2.5.2	Near a surface	19
2.6	Computational methods	21
2.7	The practical part of optical tweezers	23
2.7.1	Basic optical tweezers	23
2.7.2	Holographic optical tweezers	24
2.7.3	Data acquisition and applications	27
2.8	Concluding remarks	29
3	Optically actuated hydrodynamic manipulation	30
3.1	Designing feedback for hydrodynamic manipulation	31
3.1.1	Designing the actuators	31
3.1.2	Developing the maths	34
3.1.3	Summary	39
3.2	Simulating hydrodynamic manipulation	40
3.2.1	Simulating a system of an arbitrary number of spherical particles	40
3.2.2	Implementing hydrodynamic feedback	44
3.3	Tailoring hydrodynamic feedback	47
3.3.1	Different forms of μ_f	47
3.3.2	The need to limit actuator rotation rates	48
3.3.3	Influence of optical trapping stiffness	50
3.3.4	Diffusion effects	52
3.3.5	6D hydrodynamic manipulation	53
3.4	Our tweezers setup	55
3.4.1	Sample preparation	58
3.5	Experimental results	59
3.5.1	Hydrodynamic clamping	59
3.5.2	Hydrodynamic micro-manipulation	62
3.5.3	Simultaneous control of multiple particles	66
3.5.4	Limitations	67
3.6	Discussion	69
3.6.1	Wider context	69
3.6.2	Summary	70
Part II		72
4	More background theory	73
4.1	The generalised Wigner-Smith operator	74
4.2	Generalised Lorenz-Mie theory	78

4.2.1	Properties of VSWFs	78
4.2.2	Field decomposition into VSWFs	79
4.2.3	Beam shape coefficients and beam representation	81
4.2.4	Optical force	84
4.2.5	Full definition of VSWFs and BSCs	86
5	Optimising optical trapping stiffness	90
5.1	Methods	91
5.1.1	The Bessel basis	92
5.1.2	The optimiser	95
5.1.3	GWS	97
5.1.4	Evaluating new traps, and additional notes	98
5.2	Optimising 1D stiffness in the transverse plane	100
5.2.1	Using the GWS	100
5.2.2	Using the optimiser	106
5.3	Optimising 1D stiffness in the axial direction	108
5.4	Multi-dimensional stiffness enhancement	111
5.4.1	In the transverse plane	111
5.4.2	Full 3D trap	115
5.4.3	Optimal field for our hydrodynamic actuators	118
5.5	Phase-only optimisation	122
5.6	Discussion	124
5.6.1	Irregularities in our GWS operators	124
5.6.2	Summary of results	128
5.6.3	Applications	130
5.6.4	Comparison to the work of others	131
5.6.5	Improvements and future work	132
6	Conclusion	135
	Contributions	139
	Bibliography	141

List of Figures

2.1	Random walk	7
2.2	Schematic of notations used for defining μ_{ij}	12
2.3	Light ray interaction with a spherical bead	13
2.4	Light beam interaction with a spherical bead	14
2.5	Optical trapping with a focused beam	15
2.6	Optical gradient force	16
2.7	Trace of an optically trapped particle	18
2.8	Diffusion gradient	20
2.9	A single-beam inverted microscope optical tweezers setup	23
2.10	The working principle of HOT	26
2.11	Image processing in optical tweezers	27
2.12	PSD of an optically trapped 5 μm radius silica bead	29
3.1	Actuator influence on a target	32
3.2	Different types of hydrodynamic actuators	33
3.3	Bead constellation rotor geometry	36
3.4	Symmetry implications for a hydrodynamic manipulation setup	38
3.5	Flowchart of the multiple particle simulation	41
3.6	Fluid flow maps for different actuator types	42
3.7	Validation: correlations between optically trapped beads	43
3.8	Validation: diffusivities near a surface	44

3.9	Flowchart of the hydrodynamic feedback implemented into the simulation code	44
3.10	Simulation results for 2D clamping	46
3.11	Limiting actuator rotation rates	49
3.12	Effect of optical trapping stiffness on hydrodynamic clamping	51
3.13	Hydrodynamic clamping of different size targets	52
3.14	Influence of a surface on hydrodynamic clamping	53
3.15	6D hydrodynamic manipulation setup	54
3.16	Our HOT setup	55
3.17	Ghost orders in optical tweezers	56
3.18	Feedback loop within an experimental setup	57
3.19	Experimental hydrodynamic clamping using differently sized micro-rotors	59
3.20	Comparison of micro- and constellation rotors	60
3.21	Hydrodynamic clamping of asymmetric objects	62
3.22	Hydrodynamic translation techniques: locally moving the target	63
3.23	Hydrodynamic translation techniques: moving the sample stage	65
3.24	Multiple target hydrodynamic manipulation	67
3.25	Main limitations of hydrodynamic micro-manipulation	68
4.1	Properties of VSWFs	79
4.2	Effect of n_{\max}	81
4.3	Formation of a Bessel beam	82
4.4	Description of a plane wave travelling at an angle to the z -axis	83
5.1	Relation between the pupil and focal planes via focusing by a high NA lens	91
5.2	Geometry of Bessel beams	92
5.3	Conservation of power in Bessel beams	95
5.4	Constraining the objective function of the optimiser	96
5.5	Output basis for the GWS method	97
5.6	The main colour-maps used within Chapter 5	99
5.7	Polarisation ellipse representation	100
5.8	Optimising optical force with the GWS operator	101
5.9	Optimising optical trapping stiffness with K_x	102
5.10	Force profiles of the two-lobe and Gaussian beam traps	103
5.11	Make-up of the two-lobe trap	103
5.12	Effect of polarisation on optical trapping stiffness	104
5.13	Optimising optical trapping stiffness with K	105

5.14	Dependency of the optimiser on the starting point: κ_x	106
5.15	Dependency of the optimiser on the starting point: κ_y	107
5.16	Enhancing κ_z with K_z	108
5.17	Enhancing κ_z indirectly with Q_z	109
5.18	Enhancing κ_z with the optimiser	110
5.19	Enhancing 2D stiffness in the transverse plane with GWS	112
5.20	Optimising the four-lobe field	112
5.21	Optimising circularly polarised four-lobe field	113
5.22	Enhancing $\kappa_x + \kappa_y$ with the optimiser	114
5.23	Combining fields for stable 3D trapping: A	116
5.24	Intensity distribution in field A	117
5.25	Combining fields for stable 3D trapping: B	118
5.26	Combining fields for stable 3D trapping: C	119
5.27	Field C in motion	120
5.28	Dependence of hydrodynamic clamping on κ	121
5.29	Phase-only stiffness enhancement in the xy -plane	123
5.30	Phase-only stiffness enhancement along z	124
5.31	Non-perfect linearity: x	125
5.32	Non-perfect linearity: z	126

List of Tables

3.1	Effect of the form of the mobility tensor used in the feedback loop . . .	47
3.2	Experimental hydrodynamic clamping results	61
4.1	Functions used in defining the VSWFs	86
5.1	Summary of optimised 2D trapping in the transverse plane	115
5.2	Summary of the main traps in Chapter 5	128

List of Abbreviations

au	arbitrary unit
BSC	Beam Shape Coefficient
GLMT	Generalised Lorentz Mie Theory
GPU	Graphics Processing Unit
GWS	Generalised Wigner Smith
HOT	Holographic Optical Tweezers
MSD	Mean Squared Displacement
NA	Numerical Aperture
OAM	Orbital Angular Momentum
ODE	Ordinary Differential Equation
OE	Optical Eigenmode
PSD	Power Spectral Density
SLM	Spatial Light Modulator
SVD	Singular Value Decomposition
VSWF	Vector Spherical Wave Function

Physical Constants and Material Properties

Boltzmann constant	$k_B = 1.380\,648\,52 \times 10^{-23} \text{ m}^2 \text{ kg s}^{-2} \text{ K}^{-1}$ (exact)
Density of silica	$\rho_{\text{Si}} = 2.2 \times 10^{-3} \text{ kg m}^{-3}$
Dynamic viscosity of water ¹	$\eta_w = 1.002 \times 10^{-3} \text{ kg s}^{-1} \text{ m}^{-1}$
Electric permittivity of vacuum	$\epsilon_0 = 8.85 \times 10^{-12} \text{ F m}^{-1}$
Refractive index of silica ²	$n_{\text{Si}} = 1.5$
Refractive index of water ³	$n_w = 1.33$

¹ at room temperature

² at 800 nm wavelength

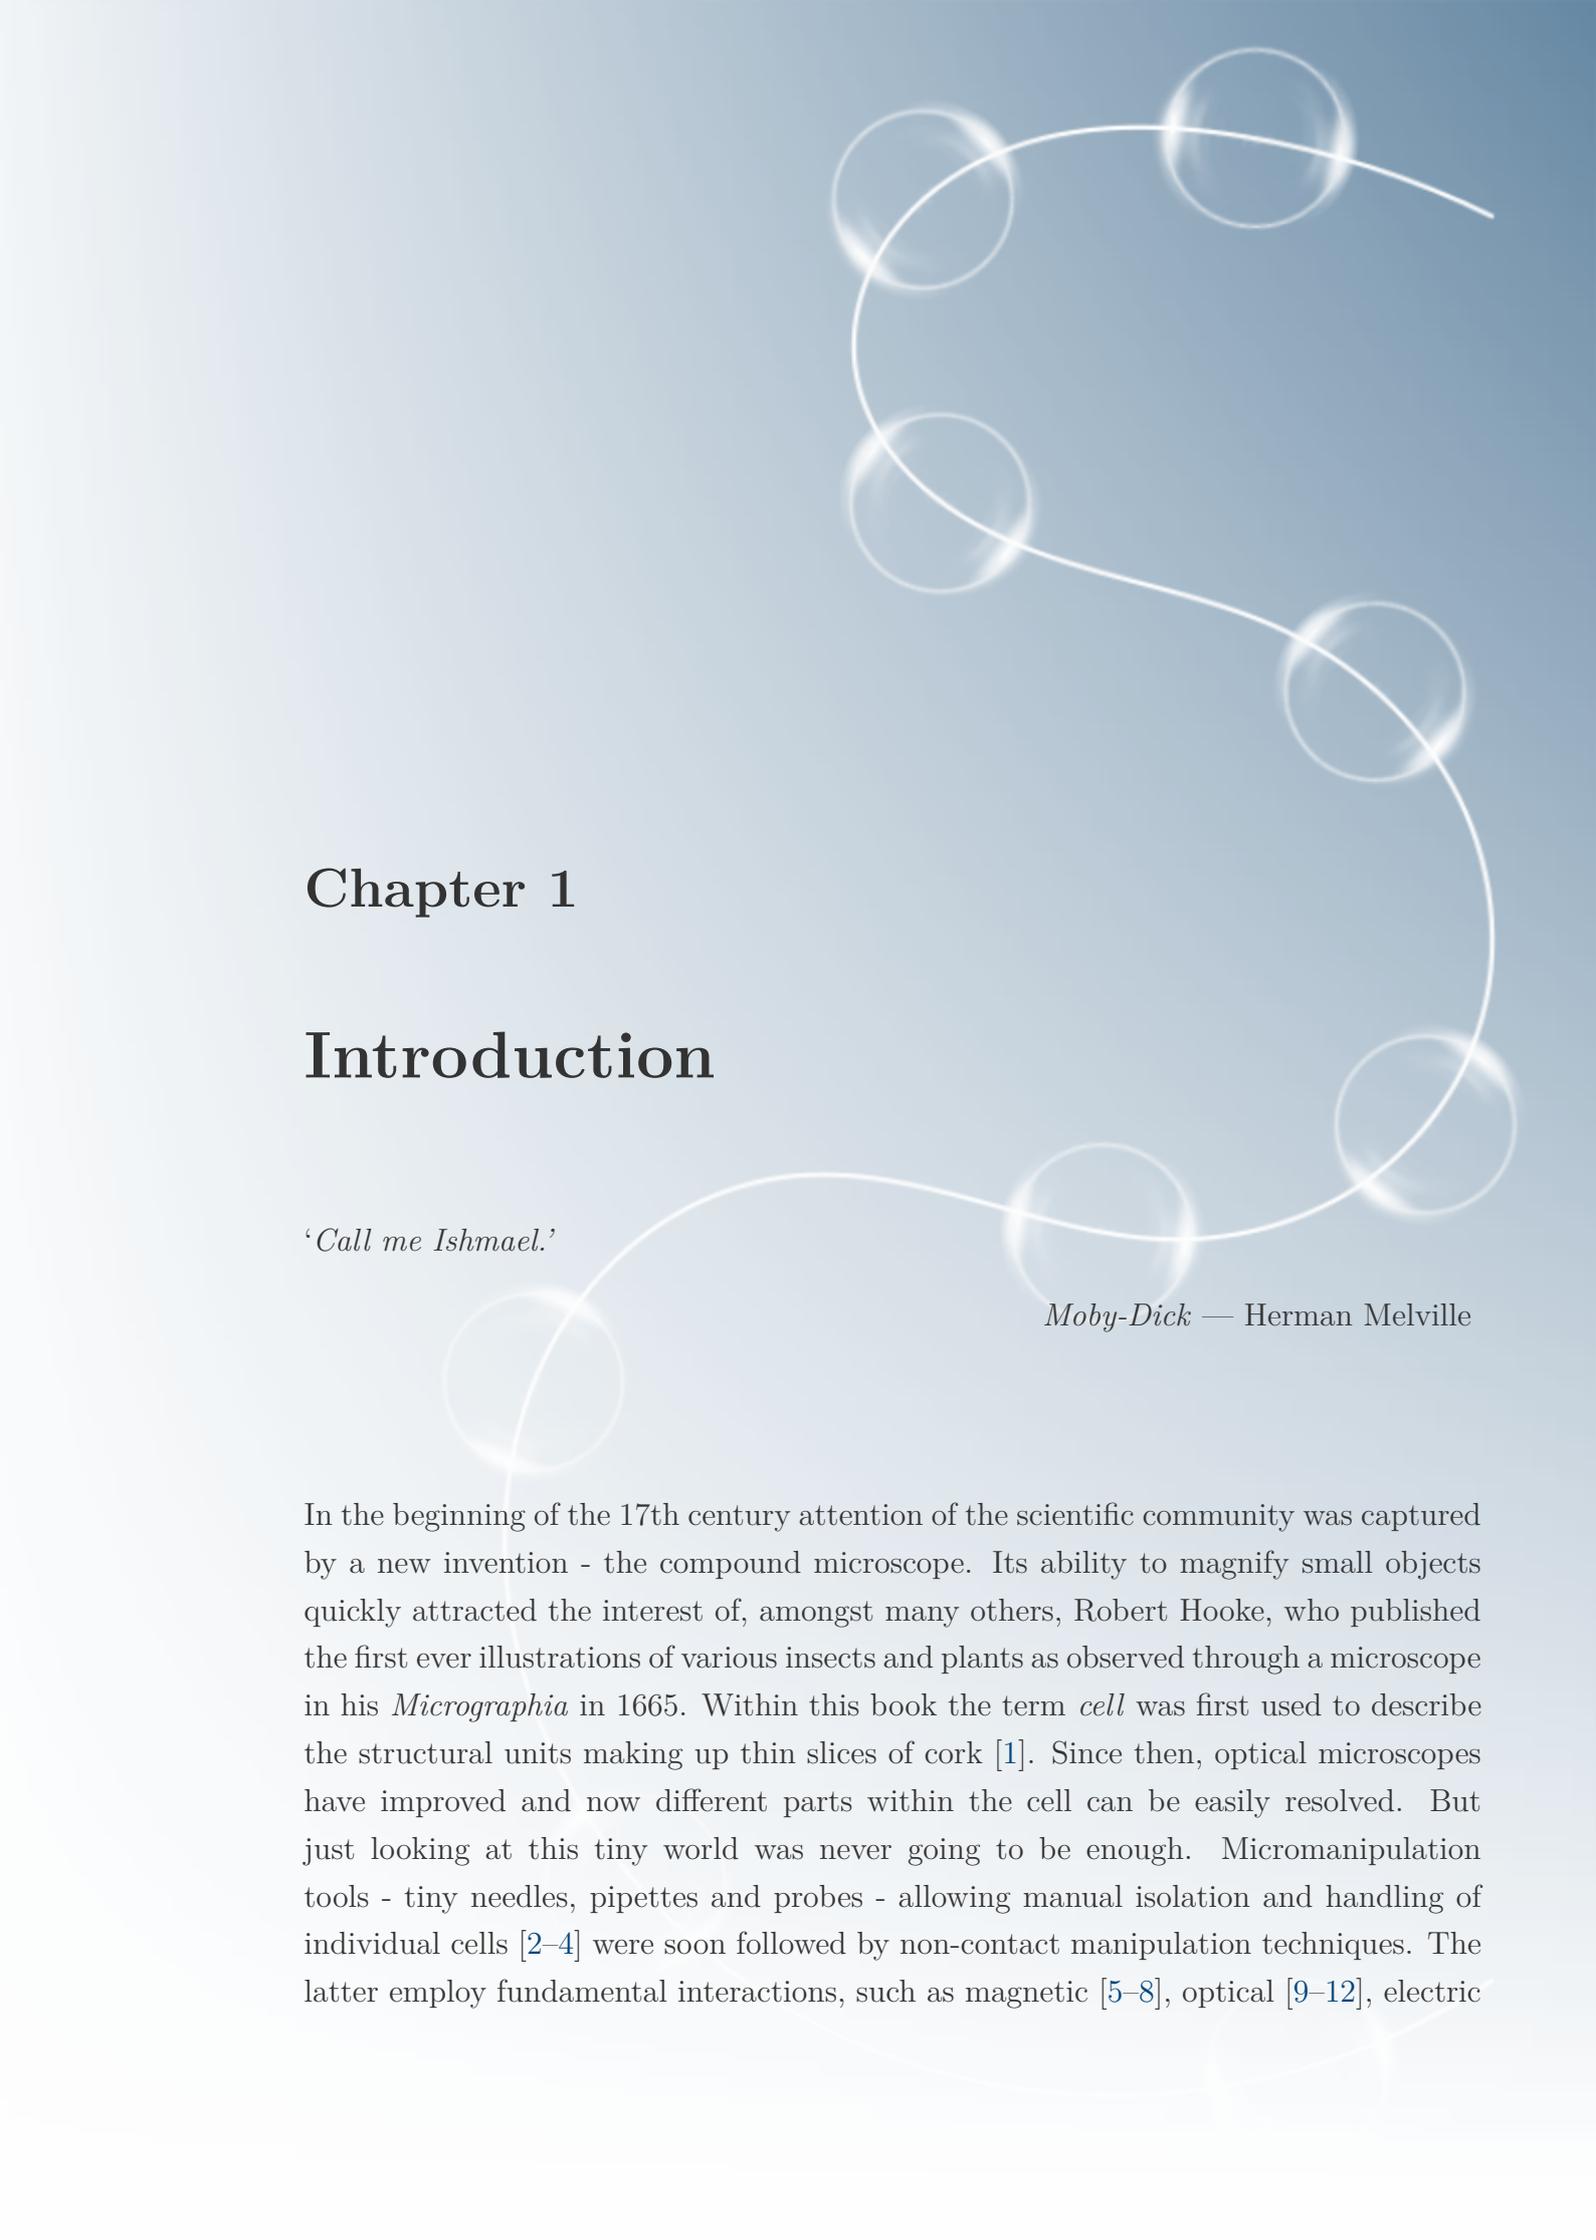
³ at 1064 nm wavelength

List of Symbols

a	bead radius	m
D	diffusion coefficient	$\text{m}^2 \text{s}^{-1}$
E	electric field amplitude	V m^{-1}
F, f	force	N
K	GWS stiffness operator	
k	wave-number	rad m^{-1}
L	orbital angular momentum number	
m	mass	kg
n	refractive index	
Q	GWS force operator	
q	generalised coordinate	
\mathcal{R}	stiffness enhancement factor	
S	scattering matrix	
T	absolute temperature	K
W	white noise	
β	hydrodynamic friction coefficient	kg s^{-1}
ϵ	Levi-Civita symbol	
ϵ_r	relative electric permittivity	
η	dynamic viscosity	$\text{kg s}^{-1} \text{m}^{-1}$
κ	optical trap stiffness	N m^{-1}
λ	eigenvalue of a GWS operator	
$\boldsymbol{\mu}$	hydrodynamic mobility tensor	mixed units
$\boldsymbol{\xi}$	hydrodynamic friction tensor	mixed units
σ	standard deviation	
ω	angular velocity	rad s^{-1}

‘Experience is merely the name men gave to their mistakes.’

The Picture of Dorian Gray — Oscar Wilde



Chapter 1

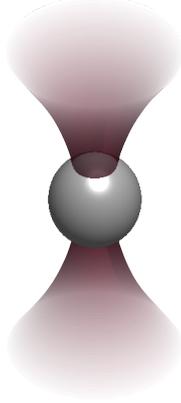
Introduction

‘Call me Ishmael.’

Moby-Dick — Herman Melville

In the beginning of the 17th century attention of the scientific community was captured by a new invention - the compound microscope. Its ability to magnify small objects quickly attracted the interest of, amongst many others, Robert Hooke, who published the first ever illustrations of various insects and plants as observed through a microscope in his *Micrographia* in 1665. Within this book the term *cell* was first used to describe the structural units making up thin slices of cork [1]. Since then, optical microscopes have improved and now different parts within the cell can be easily resolved. But just looking at this tiny world was never going to be enough. Micromanipulation tools - tiny needles, pipettes and probes - allowing manual isolation and handling of individual cells [2–4] were soon followed by non-contact manipulation techniques. The latter employ fundamental interactions, such as magnetic [5–8], optical [9–12], electric

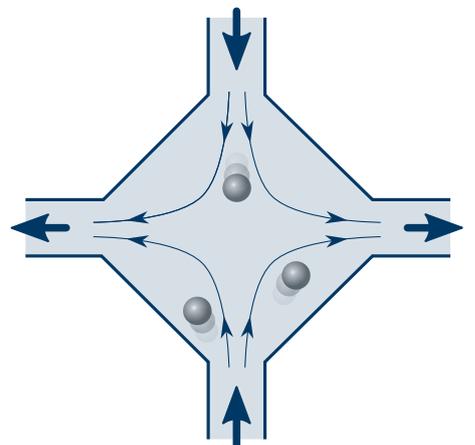
[13–16], hydrodynamic [17–19], and acoustic [20–23], to move and trap, not just cells, but colloids, synthetic micro-tools, and individual molecules.



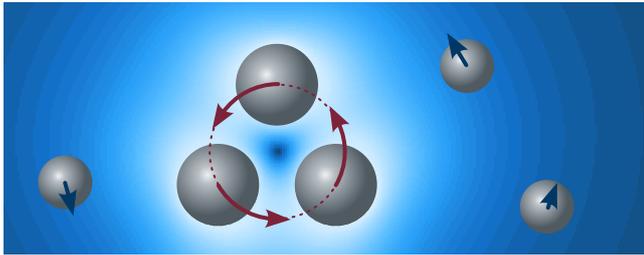
One of such non-contact instruments is optical tweezers, first demonstrated by Arthur Ashkin in 1970 [24]. The main part of conventional optical tweezers is a focused laser beam, which ‘attracts’ certain nano- or micro-sized particles towards its focus. Typically optical tweezers are used to control particles submerged in water, although air and vacuum trapping have also been demonstrated [25–27]. We will cover the details of how optical tweezers operate later; for now, in a nutshell - as the laser beam passes through a particle, the beam gets deflected, which means that its momentum has changed. The particle is then required, by conservation laws, to change momentum too. If the refractive index of the particle is higher than that of the surroundings, this gained momentum will direct it towards the beam focus, trapping it there. An optical trap can pick up one or several objects and move them around in 3D. Conventionally, the trapped object is a spherical bead ranging in size from a few tens of nano-metres to tens of micro-metres [12], but living cells [28–32] and specially designed micro-tools [33–36] are also routinely trapped. However, not all materials yield to optical trapping and direct light exposure can be damaging to living cells [37–40], which limits the range of applicability of optical tweezers.

Hydrodynamic tweezers, on the other hand, are micro-manipulation tools which are not limited by choice of material. As the name implies, they use water flow to trap and control an object. Conventional microfluidic hydrodynamic tweezers take the form of a fixed arrangement of several converging channels, where the pressure can be adjusted in each

channel individually, making the water flow in or out of it [17, 41–44]. Complex fluid flow patterns can be thus created in the channel convergence area, capable of controlling the motion of one or several target particles. This technique is non-invasive, it works for all materials indiscriminately and would not damage living cells. But, while the optical tweezers are very localised in their action, hydrodynamic ones inevitably affect the entirety of the sample. This means that they cannot act on just a single particle and leave the rest undisturbed.



Could the optical and fluidic approaches be combined to eliminate their respective disadvantages? By merging the configurational adaptability of optical tweezers with the material independent versatility of hydrodynamic control, we propose optically actuated hydrodynamic manipulation. With this technique the fluid is controlled locally by optically trapped particles. For example, rotating a circular arrangement of several optically trapped spherical beads will create a vortex flow, entraining nearby free particles. With appropriate mathematical description of the hydrodynamics and optical forces involved, this principle can be extended to a closed-loop control system. Development



and demonstration of near field hydrodynamic manipulation will constitute Part I of this thesis.

We will find that hydrodynamic manipulation is limited by how tightly one can optically trap the fluid actuators. This will lead us down the

path towards enhancing the stiffness with which spherical particles can be optically trapped and controlled.

Ever since Ashkin demonstrated the first optical tweezers, there has been a lot of research aiming to improve and adapt them to better suit some specific needs. Often this took the form of engineering the properties of the trapped object in some way. Particles have been intentionally shaped to create desired optical force fields [45] and maximise the axial optical force [46]. Optical force has been enhanced by applying an anti-reflection coating to particles [47]. Special shape designs have achieved rotation [48] and lift [49] in optical tweezers.

Alternatively, instead of manipulating the properties of the object to be trapped, one can engineer the light field itself. For example, light field wave-fronts have been corrected to minimise effects of propagating through highly distorting media [50]. 3D refractive index maps have been measured in complex assemblies of colloidal particles and then used to create a 3D light field distribution to stably trap and control the position and orientation of arbitrarily shaped particles [51]. Iterative optimisation approaches have been employed to enhance optical force and torque [52] that a laser beam can exert on a particle. These were followed by computationally faster analytical approaches [53]. In Part II of this thesis, we too will take on light field shaping - with the intention to enhance three-dimensional optical trapping stiffness.

PART I



Chapter 2

Background theory

‘One always has exaggerated ideas about what one doesn’t know.’

The Outsider — Albert Camus

2.1 Where it all begins - the Langevin equation

To describe our problem in rather loose terms - we are interested in a system of micro-particles submerged in water and exposed to laser light - we wish to know how they will interact. As every undergraduate knows, in such cases we begin by writing down Newton’s 2nd law of motion - the vector sum of all the forces acting on a body is

equal to its mass multiplied by its acceleration. So what are the forces involved in our system?

First of all, any object moving through a fluid will experience a hydrodynamic force. It is perhaps easiest understood as friction - a race car will be slowed down by the air as it is running down the track, and a small particle moving through a mass of water will experience resistance of the same nature. Next, we have optical force; now a commonplace concept - dating back to predictions of Johannes Kepler in 1600s and James Clerk Maxwell in the 19th century - that can intuitively be understood in the particle picture of light. A photon in a laser beam will impart momentum, albeit a tiny amount, onto an object it is incident on, much like a cue ball directing a red towards the pocket. The particle description does not paint a full picture though, and we will see later on in Section 2.4 that momentum can also be transferred in such a manner that compels the object to move *towards* the light beam, not away from it. And finally, we have a thermal force. As vividly explained by Richard Feynman, wherever temperature is non-zero, all the atoms and molecules will be jiggling. An observation that puzzled Robert Brown [54] and many before him, was that inanimate pollen grains and ash in water were behaving as if though they were alive, constantly moving. In fact, it is the jiggling of the water molecules, stemming from their thermal energy, that induces the seemingly random behaviour in aqueous solutions, referred to as Brownian motion.

Knowing the above, we can write down:

$$\mathbf{m} \odot \frac{d^2 \mathbf{q}}{dt^2} = \mathbf{f}^{\text{hydro}} + \mathbf{f}^{\text{opt}} + \mathbf{f}^{\text{Brn}}, \quad (2.1)$$

which is the law that will govern the behaviour of our system. Here the vectors contain information about all the degrees-of-freedom of all the particles, e.g. generalised coordinate $\mathbf{q} = [\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N, \boldsymbol{\alpha}_1, \dots, \boldsymbol{\alpha}_N]$, where for the i -th particle $\mathbf{r}_i = [x_i, y_i, z_i]$ is its position vector and $\boldsymbol{\alpha}_i = [\alpha_{xi}, \alpha_{yi}, \alpha_{zi}]$ is its orientation about the x, y and z axes respectively. Similarly, $\mathbf{m} = [\mathbf{m}_1, \dots, \mathbf{m}_N, \mathbf{I}_1, \dots, \mathbf{I}_N]$, where $\mathbf{m}_i = [m_i, m_i, m_i]$ lists the mass of the i -th particle and $\mathbf{I}_i = [I_{xi}, I_{yi}, I_{zi}]$ is its moment of inertia about each of the axes. We also define the generalised force $\mathbf{f} = [\mathbf{f}_1, \dots, \mathbf{f}_N, \boldsymbol{\tau}_1, \dots, \boldsymbol{\tau}_N]$ where $\mathbf{f}_i = [f_{xi}, f_{yi}, f_{zi}]$ is the force acting on particle i and $\boldsymbol{\tau}_i = [\tau_{xi}, \tau_{yi}, \tau_{zi}]$ is the torque acting on this particle. The symbol \odot indicates a Hadamard product (i.e. element-wise multiplication).

The above equation is known as the Langevin equation. We will now have a closer look at each of the forces involved, so that we can write down this equation in a solvable form.

2.2 Brownian forces

What sets the Langevin equation apart from other forms of Newton's 2nd law is the last term - the Brownian, stochastic force. f^{Brn} describes the random thermal 'kicks' that a particle in a fluid receives from the molecules that are constantly bombarding it, and thus causing it to undergo a random walk. This involves not only translation but also re-orientation of the particle, as illustrated in Figure 2.1.

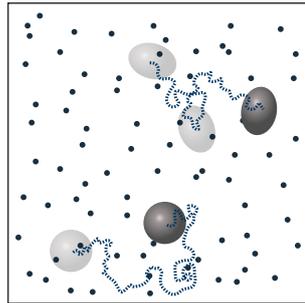


Figure 2.1: Random walk. Two particles (opaque grey) immersed in a fluid continuously collide with the oscillating fluid molecules (dark blue dots), which causes the particles to undergo a random walk - both in translation (trace in blue) and in orientation (translucent grey).

The Brownian force can, in principle, be predicted, but this is practically impossible as it is determined by an immense number of variables, e.g. a cup of water contains roughly 10^{24} continuously jiggling molecules affecting each other's motion and generating the random walk. Instead of trying to keep track of all the variables, one can take a statistical approach to model the stochastic force, using what is known as white noise W . Considering one degree-of-freedom we can write [55]:

$$f^{\text{Brn}}(t) = \sqrt{\frac{2(k_{\text{B}}T)^2}{D}} W(t), \quad (2.2)$$

where $\frac{2(k_{\text{B}}T)^2}{D}$ is the intensity of the noise, D is the diffusion coefficient, T is the absolute temperature and k_{B} is the Boltzmann constant. White noise has the following properties [56]: it has a zero mean, $\langle W(t) \rangle = 0$, unitary power, $\langle W^2(t) \rangle = 1$, and $W(t_1)$ and $W(t_2)$ are independent of each other for $t_1 \neq t_2$, where t denotes time. White noise can be viewed as the continuous form of a discrete sequence of independent random numbers. We will see how it can be computationally modelled in Section 2.6.

2.3 Hydrodynamic forces

When we refer to hydrodynamic forces, we mean the drag that a particle experiences as it is traversing a fluid. Unlike dry friction, which is almost independent of velocity, drag

is directly proportional to it. The hydrodynamic force acting on a spherical particle travelling at velocity v relative to the surrounding fluid is given by Stokes' law:

$$f^{\text{hydro}} = 6\pi\eta a v, \quad (2.3)$$

where a is the particle's radius and η is the dynamic viscosity of the surrounding fluid. We can also write down the relation for the hydrodynamic torque on a spherical particle:

$$\tau^{\text{hydro}} = 8\pi\eta a^3 \omega, \quad (2.4)$$

where ω is the angular velocity of the particle with respect to the fluid.

The above relations are valid only in environments with low Reynolds number Re , where inertial effects are negligible compared to viscous forces. Reynolds numbers for translation and rotation, Re^{T} and Re^{R} respectively, are defined as ratios of inertia over viscosity:

$$Re^{\text{T}} = \frac{\rho v a}{\eta}, \quad Re^{\text{R}} = \frac{\rho \omega a^2}{\eta}, \quad (2.5)$$

where ρ is the fluid density. Low Reynolds number systems ($Re \ll 1$) exhibit smooth, laminar flows, as opposed to turbulent ones.

Coefficients $6\pi\eta a$ and $8\pi\eta a^3$ are known as friction coefficients and are inverse to the mobility $\boldsymbol{\mu}$ of the particle. Note that Equations 2.3, 2.4 are one-dimensional - we can write them as such because for a spherical particle different degrees-of-freedom behave independently. In general though, a particle will be characterised by a 6×6 mobility tensor $\boldsymbol{\mu}$:

$$\boldsymbol{\mu} = \begin{bmatrix} \boldsymbol{\mu}^{\text{TT}} & \boldsymbol{\mu}^{\text{TR}} \\ \boldsymbol{\mu}^{\text{RT}} & \boldsymbol{\mu}^{\text{RR}} \end{bmatrix}, \quad (2.6)$$

where the four 3×3 sub-tensors describe the coupling between different degrees-of-freedom: translational $\boldsymbol{\mu}^{\text{TT}}$, rotational $\boldsymbol{\mu}^{\text{RR}}$, and translational with rotational $\boldsymbol{\mu}^{\text{TR}}$ (and we have $\boldsymbol{\mu}^{\text{TR}} = [\boldsymbol{\mu}^{\text{RT}}]^{\text{T}}$, where superscript T indicates a matrix transpose). An isolated spherical particle, for example, will be described by:

$$\boldsymbol{\mu} = \begin{bmatrix} \frac{1}{6\pi\eta a} \mathbf{I}_3 & \mathbf{0} \\ \mathbf{0} & \frac{1}{8\pi\eta a^3} \mathbf{I}_3 \end{bmatrix}, \quad (2.7)$$

where \mathbf{I}_3 is the 3×3 identity matrix. We see that translation of a sphere does not affect its rotation ($\boldsymbol{\mu}^{\text{TR}} = \mathbf{0}$) and its translation in x is unaffected by its translation along y (off-diagonal terms in $\boldsymbol{\mu}^{\text{TT}}$ are zero). For more complex non-spherical shapes this is not generally true.

Matters become even more complicated as we enter systems with multiple particles. For N particles $\boldsymbol{\mu}$ increases in dimensionality to $6N \times 6N$. For example, in the case of two particles i, j :

$$\boldsymbol{\mu}_{ij} = \begin{bmatrix} \boldsymbol{\mu}_{ii}^{\text{TT}} & \boldsymbol{\mu}_{ij}^{\text{TT}} & \boldsymbol{\mu}_{ii}^{\text{TR}} & \boldsymbol{\mu}_{ij}^{\text{TR}} \\ \boldsymbol{\mu}_{ji}^{\text{TT}} & \boldsymbol{\mu}_{jj}^{\text{TT}} & \boldsymbol{\mu}_{ji}^{\text{TR}} & \boldsymbol{\mu}_{jj}^{\text{TR}} \\ \boldsymbol{\mu}_{ii}^{\text{RT}} & \boldsymbol{\mu}_{ij}^{\text{RT}} & \boldsymbol{\mu}_{ii}^{\text{RR}} & \boldsymbol{\mu}_{ij}^{\text{RR}} \\ \boldsymbol{\mu}_{ji}^{\text{RT}} & \boldsymbol{\mu}_{jj}^{\text{RT}} & \boldsymbol{\mu}_{ji}^{\text{RR}} & \boldsymbol{\mu}_{jj}^{\text{RR}} \end{bmatrix}, \quad (2.8)$$

where $\boldsymbol{\mu}_{ij}^{\text{TT}}$ is the mobility sub-tensor that describes the coupling between the translational degrees-of-freedom of the two particles, etc. It can be very difficult to write down a mobility tensor for a large number of complex-shaped objects because of how involved the mathematics would be. For spheres, however, symmetry significantly simplifies things.

By analogy to Stokes' law we saw in Equation 2.3, $f^{\text{hydro}} = v/\mu$, we see that to obtain all the $6N$ hydrodynamic forces and torques acting on all the particles we need to multiply the inverse of $\boldsymbol{\mu}$ by the $6N \times 1$ particle velocity vector $\frac{d\mathbf{q}}{dt}$:

$$\mathbf{f}^{\text{hydro}} = \boldsymbol{\mu}^{-1} \frac{d\mathbf{q}}{dt}. \quad (2.9)$$

We must note here that this equation assumes an instantaneous effect - if a particle is exposed to an external force and moves, the fluid responds everywhere immediately, neglecting the time it takes for the flow to propagate through the fluid. As we will mostly be working in the near field where the particles are close to one another, we do not foresee this as a significant issue.

2.3.1 Fluctuation-dissipation theorem

Hydrodynamic forces are very closely related to Brownian forces, via the fluctuation dissipation theorem [57], which states that for a particle at thermodynamic equilibrium:

$$\mathbf{D} = k_{\text{B}}T\boldsymbol{\mu}. \quad (2.10)$$

In other words, this connects the intensity of random fluctuations D which give energy to a particle, with the energy dissipation μ due to particle's motion.

We can use the above relationship in Equation 2.2 to obtain a new expression for the stochastic force, which we now extend to multiple particles:

$$\mathbf{f}^{\text{Bm}} = \sqrt{2k_{\text{B}}T} \boldsymbol{\mu}^{-\frac{1}{2}} \mathbf{W}, \quad (2.11)$$

where $\boldsymbol{\mu}^{1/2}$ is the square root of the mobility tensor, defined by $\boldsymbol{\mu} = \boldsymbol{\mu}^{1/2} \left[\boldsymbol{\mu}^{1/2} \right]^T$.

We will see in the following section that for multiple particles the mobility tensor is dependent on their relative positions. This makes the Brownian force term multivariate [58], i.e. the random ‘kicks’ that a particle receives from nearby fluid molecules are no longer entirely independent of each other, they will be correlated due to the presence of other particles in the fluid.

2.3.2 The hydrodynamic mobility tensor

Deriving an expression for the mobility tensor $\boldsymbol{\mu}$ of a system of N spherical particles is far from trivial. The usual method for investigating the behaviour of multiple particles in a fluid is the method of reflections. The idea behind it is to first consider the motion of each object as if no other particles were present, and then calculate the corrections in the local fluid velocity field around each sphere due to the motion of the other spheres [59–61]. As a drawback, this method requires knowledge of the fluid velocity field. In contrast, Mazur and Van Sarloos use a method in which results are obtained directly from the boundary conditions for the fluid velocity field at the surfaces of the spheres [62, 63]. It is their derived expression for the mobility tensor that we will be using. This derivation is too involved to be repeated here, so we will simply describe it conceptually and quote the final results.

They begin by considering the equations of motion of the fluid and the spheres. These equations have to satisfy the boundary conditions at the surface of the spheres - in particular the no-slip conditions, which state that at a solid boundary the fluid has zero *relative* velocity to that boundary. This set of equations is then reformulated in such a way that eventually allows, without explicit knowledge of the fluid velocity field, calculation of the hydrodynamic forces acting on the spheres in terms of their velocities (or vice versa). The two are linked via the mobility tensor, which is estimated using methods developed in [64, 65].

The first terms in their derived expansion for $\boldsymbol{\mu}$ between two spherical particles i and j in an unbounded fluid with multiple other particles are the following:

$$\begin{aligned} 6\pi\eta a_i \boldsymbol{\mu}_{ij}^{\text{TT}} &= \mathbf{I}_3 \delta_{ij} + \left[\frac{3a_i}{4r_{ij}} (\mathbf{I}_3 + \hat{\mathbf{r}}_{ij} \hat{\mathbf{r}}_{ij}) - \frac{3a_i}{4r_{ij}^3} (a_i^2 + a_j^2) (\hat{\mathbf{r}}_{ij} \hat{\mathbf{r}}_{ij} - \frac{\mathbf{I}_3}{3}) \right] (1 - \delta_{ij}) \\ 12\pi\eta a_i^2 \boldsymbol{\mu}_{ij}^{\text{RT}} &= -\frac{3a_i^2}{2r_{ij}^2} \boldsymbol{\epsilon} \cdot \hat{\mathbf{r}}_{ij} (1 - \delta_{ij}) \\ 8\pi\eta a_i^2 a_j \boldsymbol{\mu}_{ij}^{\text{RR}} &= \mathbf{I}_3 \delta_{ij} + \frac{3a_i^2 a_j}{2r_{ij}^3} (\hat{\mathbf{r}}_{ij} \hat{\mathbf{r}}_{ij} - \frac{\mathbf{I}_3}{3}) (1 - \delta_{ij}), \end{aligned} \quad (2.12)$$

where $\hat{\mathbf{r}}_{ij}$ is a unit vector in the direction of the bead separation vector $\mathbf{r}_{ij} = \mathbf{r}_j - \mathbf{r}_i$, with magnitude $r_{ij} = |\mathbf{r}_{ij}|$, \mathbf{r}_i is the position vector of particle i , δ_{ij} is the Kronecker delta, and $\boldsymbol{\epsilon}$ is the Levi-Civita tensor defined as:

$$\epsilon_{k,l,m} = \begin{cases} +1 & \text{if } (k, l, m) \text{ is } (1, 2, 3), (2, 3, 1) \text{ or } (3, 1, 2) \\ -1 & \text{if } (k, l, m) \text{ is } (3, 2, 1), (2, 1, 3) \text{ or } (1, 3, 2) \\ 0 & \text{if } k = l, \text{ or } k = m, \text{ or } l = m, \end{cases} \quad (2.13)$$

and the convention for the tensor product notation $\hat{\mathbf{r}}_{ij} \otimes \hat{\mathbf{r}}_{ij} \equiv \hat{\mathbf{r}}_{ij} \hat{\mathbf{r}}_{ij}$ was adopted.

The first term in the square brackets of the expression for $\boldsymbol{\mu}_{ij}^{\text{TT}}$ is known as the Oseen tensor, which is independent of particle's size, and the second is known as the Rotne-Prager tensor. We note that $\boldsymbol{\mu}_{ij}^{\text{RT}}$ is independent of particle's size up to the first order, since a_i only appears in the expansion as a third order term. $\boldsymbol{\mu}_{ij}^{\text{RT}}$ is non-zero only if $i \neq j$, in line with Equation 2.7 which says that translation of a sphere in a fluid does not couple to its orientation.

The equations quoted here only include pairwise interactions between the spheres, but they can be extended to include three-sphere contributions. These first appear as follows: in $\boldsymbol{\mu}_{ij}^{\text{TT}}$ of order r^{-4} , in $\boldsymbol{\mu}_{ij}^{\text{RT}}$ of order r^{-5} , and in $\boldsymbol{\mu}_{ij}^{\text{RR}}$ of order r^{-6} . We also note that the mobility tensor is independent of the particle's orientation as this quite clearly does not affect the physical behaviour of a sphere.

In addition, the proximity of a planar surface can also be taken into account [63], which results in the following terms being added to Equation 2.12:

$$\begin{aligned}
\boldsymbol{\mu}_{ij}^{\text{TT}} : & -\frac{3a_i}{4r_{ijs}^3} \left[\mathbf{I}_3 + \hat{\mathbf{r}}_{ijs} \hat{\mathbf{r}}_{ijs} - \frac{2l_i}{r_{ijs}} \hat{\mathbf{r}}_{ijs} \hat{\boldsymbol{w}} + \frac{2l_j}{r_{ijs}} \hat{\boldsymbol{w}} \hat{\mathbf{r}}_{ijs} + \frac{2l_i l_j}{r_{ijs}^2} (\mathbf{I}_3 - 2\hat{\boldsymbol{w}} \hat{\boldsymbol{w}} - 3\hat{\mathbf{r}}_{ijs} \hat{\mathbf{r}}_{ijs}) \right] \\
& + \frac{3a_i}{4r_{ijs}^3} (a_i^2 + a_j^2) (\hat{\mathbf{r}}_{ijs} \hat{\mathbf{r}}_{ijs} - \frac{1}{3} \mathbf{I}_3) \\
& - \frac{3a_i}{2r_{ijs}^4} (a_i^2 - a_j^2) (l_i \hat{\mathbf{r}}_{ijs} \hat{\boldsymbol{w}} + l_j \hat{\boldsymbol{w}} \hat{\mathbf{r}}_{ijs}) \\
& + \frac{3a_i}{2r_{ijs}^5} (a_i^2 l_j + a_j^2 l_i) (l_i + l_j) (\mathbf{I}_3 - 2\hat{\boldsymbol{w}} \hat{\boldsymbol{w}} - 5\hat{\mathbf{r}}_{ijs} \hat{\mathbf{r}}_{ijs}) \\
\boldsymbol{\mu}_{ij}^{\text{RT}} : & + \frac{3a_i^2}{2r_{ijs}^2} \left[\boldsymbol{\epsilon} \cdot \hat{\mathbf{r}}_{ijs} + \frac{2l_j}{r_{ijs}} (\boldsymbol{\epsilon} \cdot \hat{\boldsymbol{w}} - 3(\hat{\mathbf{r}}_{ijs} \hat{\boldsymbol{w}}) \hat{\mathbf{r}}_{ijs}) \right] \\
\boldsymbol{\mu}_{ij}^{\text{RR}} : & - \frac{3a_i^2 a_j}{r_{ijs}^3} \left[\hat{\mathbf{r}}_{ijs} \hat{\mathbf{r}}_{ijs} - \frac{\mathbf{I}_3}{3} + \frac{2(l_i + l_j)^2}{r_{ijs}^2} \mathbf{I}_3 - 2(\hat{\mathbf{r}}_{ijs} \hat{\boldsymbol{w}})(\hat{\mathbf{r}}_{ijs} \hat{\boldsymbol{w}}) \right], \tag{2.14}
\end{aligned}$$

where $\hat{\boldsymbol{w}}$ is a unit normal to the wall, $l_i \equiv \hat{\boldsymbol{w}} \cdot \mathbf{r}_i$ is the distance from particle i to the wall, $r_{ijs} \equiv |\mathbf{S} \cdot \mathbf{r}_j - \mathbf{r}_i| = (r_{ij}^2 + 4l_i l_j)^{1/2}$ is the magnitude of a vector pointing from particle i to the mirror image of particle j with respect to the wall, as illustrated in Figure 2.2, $\mathbf{S} = \mathbf{I}_3 - 2\hat{\boldsymbol{w}} \hat{\boldsymbol{w}}$ is called the mirror tensor, and $\hat{\mathbf{r}}_{ijs} = \hat{\mathbf{r}}_{ijs} + 2\hat{\boldsymbol{w}}(l_i + l_j)/r_{ijs}$ is the unit vector pointing from the mirror image of particle i to particle j . The mobility tensor obeys the following symmetry relations: $\boldsymbol{\mu}_{ij}^{\text{TT}} = [\boldsymbol{\mu}_{ji}^{\text{TT}}]^{\text{T}}$, $\boldsymbol{\mu}_{ij}^{\text{RR}} = [\boldsymbol{\mu}_{ji}^{\text{RR}}]^{\text{T}}$, $\boldsymbol{\mu}_{ij}^{\text{RT}} = [\boldsymbol{\mu}_{ji}^{\text{TR}}]^{\text{T}}$.

We now have all the information necessary to describe $\mathbf{f}^{\text{hydro}}$ and \mathbf{f}^{Brn} as dictated by Equations 2.9 and 2.11.

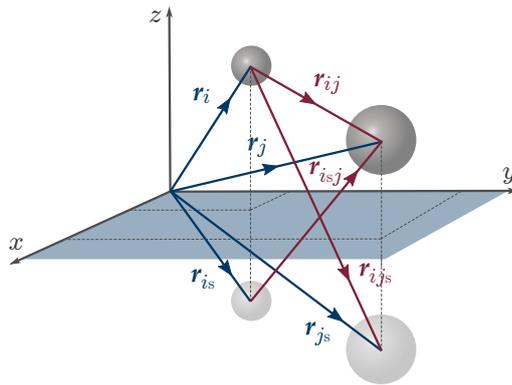


Figure 2.2: Schematic of notations used for defining $\boldsymbol{\mu}_{ij}$. Position vectors (blue) of the particles (upper two beads) and their mirror images with respect to the surface (lower two beads) and their separation vectors (red) are indicated as used in Equations 2.12, 2.14.

2.4 Optical forces

2.4.1 The operating principle of optical tweezers

Normally, we think of optical force as a repulsive force, like in the idea of solar sails [66]. A spacecraft equipped with ‘sails’ built from mirrors would reflect the radiation incident from the sun and would thus be propelled across space - a phenomenon referred to as radiation pressure. This is due to conservation of momentum. Take, for example, a reflective spherical bead in Figure 2.3(a). If we expose it to a uniform beam of light and consider one ray in that beam with momentum \mathbf{p}_{in} , we see that upon reflection the ray will have momentum $\mathbf{p}_{\text{out}} = -\mathbf{p}_{\text{in}}$. Now, the total amount of momentum must be the same before and after reflection, or, in other words, the bead must gain momentum which is equal and opposite to the change in momentum of the ray: $\mathbf{p}_{\text{Bead}} = -(\mathbf{p}_{\text{out}} - \mathbf{p}_{\text{in}})$. Hence, by Newton’s 2nd law, the light ray exerts a force on the bead.

But what happens if the bead is perfectly transparent? We now briefly revisit the law of refraction. A light ray travelling in between two media of different refractive indices will change its direction. If the refractive index of the first medium n_1 is smaller than that of the second medium n_2 , the ray will bend towards the normal to the surface, otherwise, it will bend away from the normal. Consider now a bead that has a larger refractive index than the surrounding environment. In Figure 2.3(b) we see that a light ray entering such a bead will be deflected after exiting it. Vector addition shows that the bead will gain momentum pulling it towards the ray and the bead will stabilise once the ray is passing right through the middle of it. The higher the contrast between the refractive indices of the bead and its environment is, the more deflection the rays will experience and the more momentum the bead will gain.

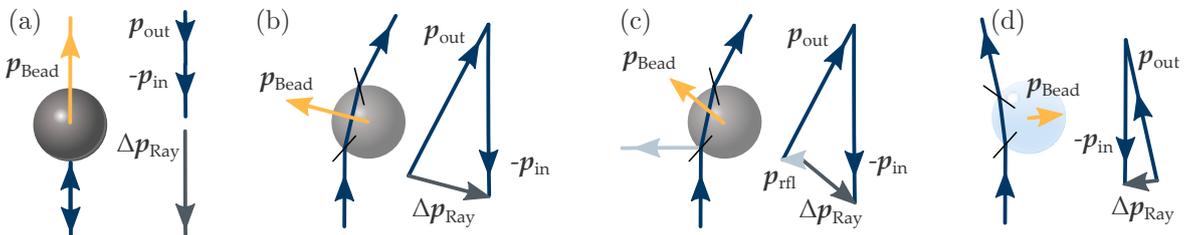


Figure 2.3: Light ray interaction with a spherical bead. In (a) a particle fully reflects the incident light ray, resulting in momentum transfer pushing the bead in the direction of the incident ray. In (b) a transparent particle with a refractive index larger than that of the surroundings deflects an incident light ray and thus gains momentum pulling it towards the axis of the ray. (c) Is the same as (b), but with partial reflection taken into account. (d) Illustrates the situation when the bead is of a lower refractive index than the surroundings.

But we must not forget that a small portion of the light will still be reflected. The effect of reflection is to reduce the magnitude of the momentum of the outgoing ray, though not the angle by which it is deflected. The overall effect, illustrated in Figure 2.3(c), is that the momentum of the bead now points more in the direction of travel of the ray. Once the bead aligns itself so that the ray is passing through its centre, the bead will be pushed forwards along the ray. Note that we have ignored the reflections happening inside the bead. Internally reflected rays would also be deflected upon exiting the bead, thus affecting the direction and magnitude of the momentum imparted on the bead. We also note briefly that if the refractive index of the bead is smaller than that of the surroundings, the bead will be pushed away from the incident light ray, as seen in Figure 2.3(d).

Let us now consider the full beam made up of multiple rays. If the beam has a uniform intensity profile, a bead within it will only experience the force pushing it along the beam; it will not gain any sideways momentum because the rays on the opposite sides of the bead will cancel each other out, as seen in Figure 2.4(a). If, however, the beam has an intensity gradient, e.g. a collimated Gaussian beam, the rays passing through the middle of the beam will have greater momentum, and so, will act to pull the bead towards the highest intensity point, as in Figure 2.4(b). This force acting to push the bead up the intensity gradient is referred to as the gradient force, and the force that pushes the bead downstream in the beam is called the scattering force.

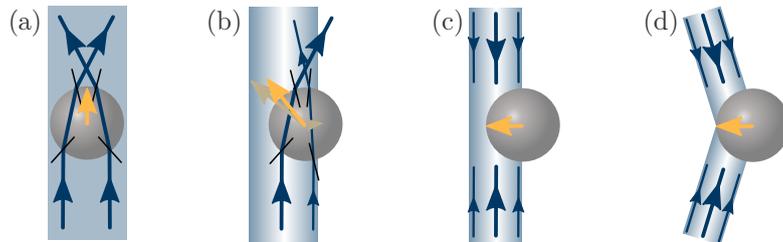


Figure 2.4: Light beam interaction with a spherical bead. In (a) the bead is within a light beam of uniform intensity - the overall effect of individual rays is to push the bead in the direction of the beam. The beam in (b) has a Gaussian intensity distribution, resulting in a larger momentum transfer from the rays in the middle of the beam, thus creating 2D trapping. Stable trapping in 3D can be created with two collimated beams incident on the bead from opposite directions (c), thus cancelling out the scattering force. Stability of 3D trapping is still maintained if the two beams are at a high angle at each other (d).

We thus have stable optical trapping in 2D provided by the gradient force, and instability along the third dimension because of the scattering force. The most natural instinct for stabilising the bead's motion in the axial direction is to introduce a second

beam, travelling in the opposite direction to the first, as in Figure 2.4(c). The scattering forces of the two beams will cancel out, stabilising the bead in 3D. In fact, this still works if the two beams are at a high angle to each other, as in Figure 2.4(d).

The great breakthrough of Arthur Ashkin [67] was to realise that two laser sources are not necessary - the effect can be reproduced with a single highly focused Gaussian beam, shown in Figure 2.5(a). This was first termed ‘the single beam gradient force trap’, and later became known as optical tweezers. For successful tweezing one needs to have a lens of a very high numerical aperture (NA)¹ capable of tightly focusing the light beam. A bead in the near vicinity of the focus of such a beam will be pulled towards the focus regardless of whether it is in front of, Figure 2.5(a), or behind it, Figure 2.5(b). If the bead starts too far away behind the focus, Figure 2.5(c), it will only be pushed downstream, because it will no longer sample enough high angled rays. We also note that the stable trapping position will be slightly behind the focus, not at it, because of the scattering force. If the focusing lens does not have a high enough NA, the situation will look more like that of a collimated beam and there will be no stable trapping in the axial direction, as shown in Figure 2.5(d).

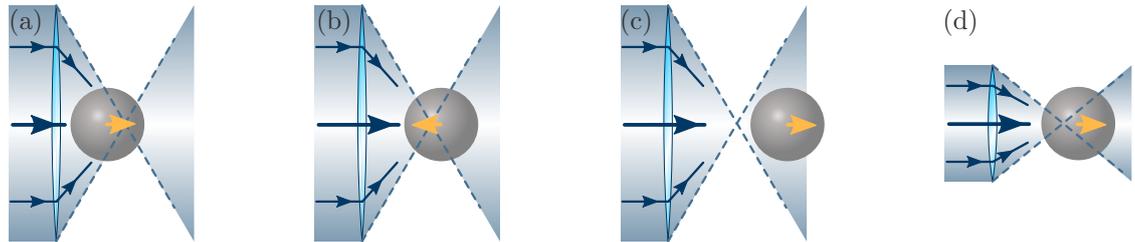


Figure 2.5: Optical trapping with a focused beam. 3D optical tweezing can be achieved with a single tightly focused Gaussian beam (a-b). If the particle is too far ahead of the focus to sample enough high angled rays it will not be pulled towards the focus of the beam (c). Likewise, if the focusing lens does not have a high enough NA, the particle will move downstream the beam (d).

2.4.2 The tweezing force

We now understand the working principle of the optical tweezers, but what we need for our Langevin equation is the optical force. While it is possible to get an approximate expression for the total force acting on the bead by considering the individual forces that each ray in the beam exerts during every reflection and refraction it undergoes within the bead [68], we will reach the same result faster via the dipole approximation.

¹ Numerical aperture is defined as $NA = n \sin \theta$, where n is the refractive index of the medium in which the lens is working, and θ is the maximum half-angle subtended by the lens.

A note before we begin - the ray optics picture that we painted above is technically only valid for particles which are significantly larger than the wavelength of the trapping light beam. At the other end of the spectrum, where the particles are significantly smaller than the wavelength of light, we have the dipole approximation - we will use results from this framework to arrive at the approximate expression for the optical force, which holds true in the ray picture as well.

We already introduced the concept of the gradient force and we will now see that it is indeed proportional to the gradient of the electric field intensity I . The gradient force at location r , within the context of dipole approximation, is given by [55]:

$$f_{\text{DA,grad}}(r) = \frac{1}{2} \frac{\alpha'}{c\epsilon_0} \nabla I(r), \quad (2.15)$$

where c is the speed of light, ϵ_0 is the electric permittivity of free space, and α' is the atomic polarisability of the bead, which is dependent on the particle's volume, and hence its radius. A Gaussian beam, most commonly used for optical trapping, has the following transverse intensity profile, to the first approximation:

$$I(r) = I_0 e^{-2r^2/w_0^2}, \quad (2.16)$$

where r is the radial coordinate in the transverse plane with the origin placed at the centre of the beam, I_0 is the maximum intensity, and w_0 is the beam waist. Substituting this into Equation 2.15 and differentiating we obtain:

$$f_{\text{DA,grad}}(r) = -\kappa_r r e^{-2r^2/w_0^2}, \quad (2.17)$$

where we have defined the optical trap stiffness along the direction of r as $\kappa_r = 2 \frac{\alpha'}{c\epsilon_0} \frac{I_0}{w_0^2}$. We plot this force as a function of position in Figure 2.6 for two different stiffness values.

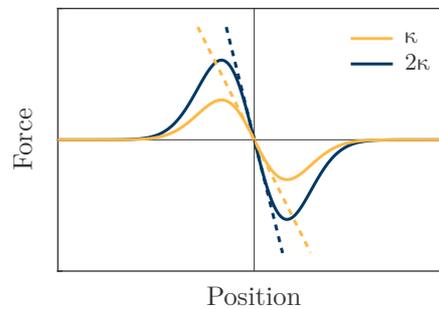


Figure 2.6: Optical gradient force as a function of position of a bead with respect to the trapping beam focus, for two different optical stiffness κ values. The dashed lines indicate the linear approximation of the force which is valid for small bead-trap displacements.

We see that the particle will experience a negative restoring force pulling it towards the highest intensity of the beam. The force increases with displacement from the beam axis up to a certain point, but then decreases exponentially, reaching zero when the bead is too far away from the beam to interact with it. We observe that for small displacements the force is linear, just like in Hooke's law for a mass on a spring, which can also be shown by doing a Taylor series expansion of the exponential term in Equation 2.17 about $r = 0$. This gives us the final expression for the optical force:

$$\mathbf{f}^{\text{opt}} = -\kappa \delta \mathbf{r}, \quad (2.18)$$

where $\delta \mathbf{r}$ is the displacement of the bead from its stable trapping location and we assume that the stiffness is symmetric in all directions². This is a very convenient simple approximation and we will use it for the rest of Part I, but will revisit the optical force in Part II in much more detail. So far we have not mentioned optical torques at all - they will be briefly discussed in Section 3.1.1.

2.5 Revisiting the Langevin equation

Now that we understand all the forces acting in our system, we can revisit the Langevin equation (Equation 2.1) in more detail. First, we consider the left-hand-side term - inertia $\mathbf{m} \odot \frac{d^2 \mathbf{q}}{dt^2}$. We mentioned already that we are working in a low Reynolds number environment, where viscous forces dominate over inertial forces. In other words, inertial effects decay over a very short time scale and can therefore be neglected when considering the behaviour of particles at low Reynolds number [69]. It is therefore common practice to set the inertial term to zero in the Langevin equation. Rearranging for particle velocities we are left with

$$\frac{d\mathbf{q}}{dt} = \boldsymbol{\mu} [\mathbf{f}^{\text{opt}} + \mathbf{f}^{\text{Brn}}]. \quad (2.19)$$

This is a first order differential equation that describes the time evolution (time dependence is implicit here) of our system, known as the over-damped Langevin equation.

² Generally this is not the case, and stiffness in the direction parallel to the beam axis is smaller than in the plane transverse to the beam. More on this in Section 2.5.1

For clarity and future reference we explicitly write down how one would calculate the velocities of a single particle residing at \mathbf{q}_i in a multiple particle system:

$$\frac{d\mathbf{q}_i}{dt} = \sum_{j=1}^N \boldsymbol{\mu}_{ij} \mathbf{f}_j^{\text{opt}} + \mathbf{f}_i^{\text{Brn}}. \quad (2.20)$$

We also present the one dimensional case of a *single particle* system with one degree-of-freedom, say x . Writing out the terms in Equation 2.19 we get:

$$\frac{dx(t)}{dt} = -\mu_{xx} \kappa_x x(t) + \sqrt{2k_B T \mu_{xx}} W_x(t), \quad (2.21)$$

where we have, for now, taken the location of the optical trap to be the origin so that we can write the trap-bead displacement δx as x .

2.5.1 Behaviour of an optically trapped particle

A single particle in an optical trap will still experience diffusion, which acts to take it further away from the trap. But the gradient force will act to prevent this, thus confining the particle to a limited volume that it can explore. Something we have not mentioned so far is the fact that the trapping stiffness along the beam axis, i.e. in the z -direction, is several times lower than the stiffness in the transverse plane. This is because the focus of a Gaussian beam is tighter in the transverse than in the longitudinal plane, resulting in a weaker restoring gradient force along the beam axis. Because of this, the motion of a trapped particle is bounded to an ellipsoidal volume as can be seen in Figure 2.7. Here the trajectory of the particle was simulated using methods that will be covered in Section 2.6.

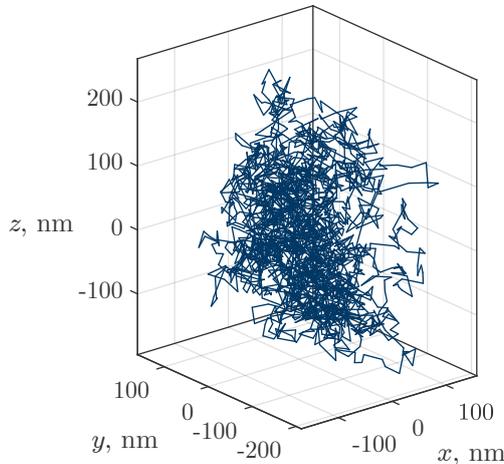


Figure 2.7: Trace of an optically trapped particle simulated for a $1\ \mu\text{m}$ radius silica bead in an optical trap of stiffness $[1, 1, 0.5] \times 10^{-6} \text{ N m}^{-1}$ in the x, y, z directions respectively. The particle covered the shown trajectory in 1 s.

2.5.2 Near a surface

We have already seen in Section 2.3 that the mobility tensor is altered if the particles are near a surface, but this is not the only adjustment needed in the presence of a boundary. Up until now we treated diffusion D as a constant, however, near boundaries we encounter diffusion gradients - where particles diffuse less as they approach a surface [70, 71]. The mobility tensor is now state-dependent, i.e. $\mu \rightarrow \mu(x(t))$. This means that the noise term is, in turn, dependent on the particle locations. The end result is that the over-damped Langevin equation will require some modification - an additional term, let's call it $d(x)$ - which will depend on the conventions used when evaluating the noise.

Let us write the intensity of the noise in Equation 2.21 as $g^2(x(t)) = 2k_{\text{B}}T\mu(x(t))$. To solve the over-damped Langevin equation (for one particle with one degree-of-freedom) one needs to integrate:

$$x(t) = - \int_t^{t+\Delta t} \mu(x(t))\kappa x(t)dt + \int_t^{t+\Delta t} g(x(t))W(t)dt, \quad (2.22)$$

where we have dropped the x subscript. We call the second integral on the right-hand-side $I(t, \Delta t)$. The noise term $W(t)$ is not continuous in the interval $[t, t + \Delta t]$, but we can deal with $g(x(t))$ using the first integral mean-value theorem [72] to write:

$$I(t, \Delta t) = g(\alpha x(t + \Delta t) + (1 - \alpha)x(t)) \int_t^{t+\Delta t} W(t)dt, \quad (2.23)$$

where the choice of parameter $\alpha \in [0, 1]$ will determine the value of the integral I . There are three common conventions. The Ito convention, with $\alpha = 0$, means that $g(x(t))$ is evaluated *before* any noise in the interval $[t, t + \Delta t]$ occurs. The Stratonovich convention, with $\alpha = \frac{1}{2}$, evaluates $g(x(t))$ at the midpoint of the interval $[t, t + \Delta t]$. And the isothermal convention, with $\alpha = 1$, when $g(x(t))$ is estimated *after* the noise occurs in the integration interval.

It can be shown that, by requiring that the system evolves to a Boltzmann distribution at long times, i.e. is consistent with thermodynamics, the additional term in the over-damped Langevin equation is dependent on the spatial gradient of the mobility [72]:

$$d(x) = 2(1 - \alpha)k_{\text{B}}T \frac{\partial \mu(x)}{\partial x}. \quad (2.24)$$

We see that in the isothermal convention, where $\alpha = 1$, this additional term conveniently vanishes and is therefore a tempting choice [73]. The Stratonovich convention is

often favoured in the physics community [74–76]. However, the Ito convention is particularly convenient for numerical simulations using the finite difference method [55], since in this scheme the noise estimation is required *before* the new particle locations can be computed [77] - we will therefore adopt the Ito convention. With the required addition of the correction term $d(x)$, the over-damped Langevin equation for a particle near a surface now reads:

$$\frac{dx}{dt} = -\mu(x)\kappa x(t) + \sqrt{2k_{\text{B}}T\mu(x)}W(t) + 2k_{\text{B}}T\frac{\partial\mu(x)}{\partial x}. \quad (2.25)$$

The last term in the above is referred to as the spurious drift - somewhat misleadingly as it suggests a nonphysical nature, while it is observable in experiments [78]. This term can be understood as follows [79]. If we consider a Brownian particle at some initial location x_0 in a constant diffusion D environment, the probability of it moving to the left or to the right is the same, i.e. the probability distribution of the particle's final location x_{fin} is symmetric about x_0 . If the diffusion has a gradient, for example if there is a wall perpendicular to the x -axis as in Figure 2.8, D will have different values at x_0 and x_{fin} and evaluation of particle's displacement is not unambiguous. Choosing $D = D(x_0)$, the Ito convention, will not change the probability distribution - it will still be symmetric. If, on the other hand, we choose $D = D(x_{\text{fin}})$, the isothermal convention, the probability will become skewed with a longer tail in the direction of the diffusion gradient, because the particle will experience greater displacements in the regions of higher diffusion. In the Stratonovich convention, the probability distribution is also skewed, but not by as much as in the isothermal. In this context, the isothermal convention describes the physical behaviour the best - the particle will have a slight preference of moving towards the region with a higher diffusion gradient, i.e. away from the surface, which leads to a skewed position probability distribution. Therefore, the isothermal convention does not require the spurious drift term, while it is needed for the Ito and Stratonovich conventions to account for the asymmetry of the probability distribution [79].

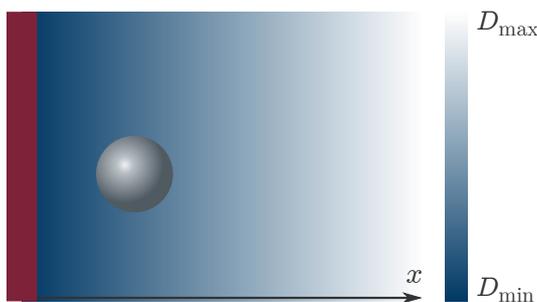


Figure 2.8: Diffusion gradient. Presence of a surface (dark red) creates a gradient in the diffusion D (as indicated by the intensity of blue colour) along the direction perpendicular to the wall, which affects the probability distribution of a freely diffusing particle.

In a multidimensional multi-particle case Equation 2.25 becomes:

$$\frac{d\mathbf{q}}{dt} = \boldsymbol{\mu} \mathbf{f}^{\text{opt}} + \sqrt{2k_{\text{B}}T} \boldsymbol{\mu}^{\frac{1}{2}} \mathbf{W}(t) + 2k_{\text{B}}T \nabla_{\mathbf{q}} \cdot \boldsymbol{\mu}, \quad (2.26)$$

where dependence of $\boldsymbol{\mu}$ and \mathbf{f}^{opt} on \mathbf{q} is implicit and the gradient of mobility becomes the divergence of $\boldsymbol{\mu}$ [75, 76].

We also note that mobility of multiple particles, even in the absence of a boundary, is state-dependent. However, the divergence of our $\boldsymbol{\mu}$ for multiple particles in an unbounded fluid is several orders of magnitude smaller than divergence of $\boldsymbol{\mu}$ in the presence of the surface (which we have verified with explicit computations); and in both cases the spurious drift term is significantly smaller than the other terms in the over-damped Langevin equation. We therefore neglect the drift term if there is no surface in proximity.

2.6 Computational methods

Equation 2.19 is a 1st order differential equation and can be solved using a finite difference method. However, unlike ordinary differential equations (ODEs), this one contains a stochastic term in the form of white noise W , making its solution process considerably more complicated. Outlined below is an implementation of the Euler's method for solving Equation 2.19. Note that higher order methods, such as ordinary Runge-Kutta extended for stochastic equations [80–82], can also be employed, if higher accuracy is desired [83]. These methods, however, are computationally more expensive and considerably more difficult to implement, making the simpler first order method, developed by Ermak and McCammon [74], the method of choice when simulating particles in optical tweezers [84–86].

Finite difference methods are implemented by approximating a continuous solution with a discrete one. If we index the current time step with a subscript l , the non-stochastic variables in Equation 2.21 become: $x(t) \rightarrow x_l$, $\frac{dx(t)}{dt} \rightarrow \frac{x_{l+1} - x_l}{\Delta t}$, $\mu_{xx}(t) \rightarrow \mu_{xx,l}$, where Δt is the length of the time step.

We deal with the white noise term by replacing it with an appropriate sequence of random numbers $W_{x,l}$ that mimics the properties of $W_x(t)$ detailed in Section 2.2. Since $W_x(t)$ has a zero mean $W_{x,l}$ will also have a zero mean; to satisfy the unitary power condition we require that $\langle (W_{x,l} \Delta t)^2 \rangle = 1$, so that $W_{x,l}$ has a variance of $1/\Delta t$; and because white noise is uncorrelated in time, $W_{x,l}$ and $W_{x,m}$ must be independent of each other for $l \neq m$. In practice, we can obtain $W_{x,l}$ by generating a sequence $w_{x,l}$ of

random Gaussian distributed numbers with zero mean and unitary variance (functions that do this are readily available in many programming languages) and then rescaling it to have a sequence with variance $1/\Delta t$. So that $W_x(t) \rightarrow \sqrt{\Delta t} w_{x,l}$ and our finite difference equation is:

$$x_{l+1} = x_l - \kappa_x \mu_{xx} x_l \Delta t + \sqrt{2k_B T \Delta t \mu_{xx}} w_{x,l}. \quad (2.27)$$

Time-evolution behaviour of the particle is obtained by iteratively solving this equation at discrete time steps, $t_l = l\Delta t$.

Numerical instability can be a problem for the stochastic Euler method, unless an appropriately sized time step Δt is chosen [87]. In Equation 2.27 Δt is limited by two factors. The lower limit is imposed when choosing to ignore inertial effects. Inertia is characterised by momentum relaxation time of a bead, $t_m = m\mu$, which is typically very short (e.g., of the order 10^{-7} s for microscopic silica beads in water at room temperature). The upper limit is imposed by the optical trap relaxation time, $t_{ot} = 1/(\kappa\mu)$, which describes the time scale on which the restoring force acts. Choosing a time step larger than t_{ot} would result in a numerically unstable solution which either oscillates or diverges. As long as the time step satisfies $t_m \ll \Delta t \ll t_{ot}$, the simulation will produce physically accurate behaviour.

Next, we extend the above description to the multidimensional multi-particle case:

$$\mathbf{q}_{l+1} = \mathbf{q}_l + \boldsymbol{\mu}_l \mathbf{f}_l^{\text{opt}} \Delta t + \sqrt{2k_B T \Delta t} \boldsymbol{\mu}_l^{\frac{1}{2}} \boldsymbol{w}_l, \quad (2.28)$$

where \boldsymbol{w}_l is a vector containing the random Gaussian numbers for each degree-of-freedom in the system making it a $6N \times 1$ column vector. $\boldsymbol{\mu}_l^{\frac{1}{2}}$ is a lower triangular matrix [88] that we defined in Section 2.3.1, which can be computed using Cholesky decomposition. Algorithms for this are readily available in most programming languages.

We must also discuss the computation of the spurious drift term emerging in the presence of a nearby surface, which contains divergence of the mobility tensor $\nabla \mathbf{q} \cdot \boldsymbol{\mu}$ (Equation 2.26). We employ the *random finite difference* scheme [75, 76], which states that:

$$\nabla \mathbf{q} \cdot \boldsymbol{\mu} = \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon} \langle \boldsymbol{\mu}(\mathbf{q} + \epsilon \Delta \mathbf{q}) \Delta \mathbf{q} - \boldsymbol{\mu}(\mathbf{q}) \Delta \mathbf{q} \rangle, \quad (2.29)$$

where ϵ is a very small number, e.g. to the order of 10^{-12} (but not too small, to avoid numerical round-off problems), $\Delta \mathbf{q}$ is a standard Gaussian random variable with zero mean and unitary variance, and the average (indicated by angle brackets) is over the

number of additional evaluations of $\mu(\mathbf{q} + \epsilon\Delta\mathbf{q})$. The main advantage of the random finite difference method, compared to explicit evaluation of the divergence, is that only a small number of additional evaluations of μ are needed per iteration.

2.7 The practical part of optical tweezers

2.7.1 Basic optical tweezers

An optical tweezers setup can take different forms but it will most likely be organised around a high NA lens in an inverted microscope construction. Illustrated in Figure 2.9, the inverted microscope part consists of an illumination source throwing light on a sample cell, which is usually filled with an aqueous suspension of Brownian particles, and is held on a translation stage. Placed below the sample is a high NA objective lens that collects the illumination light after it passes through the sample.

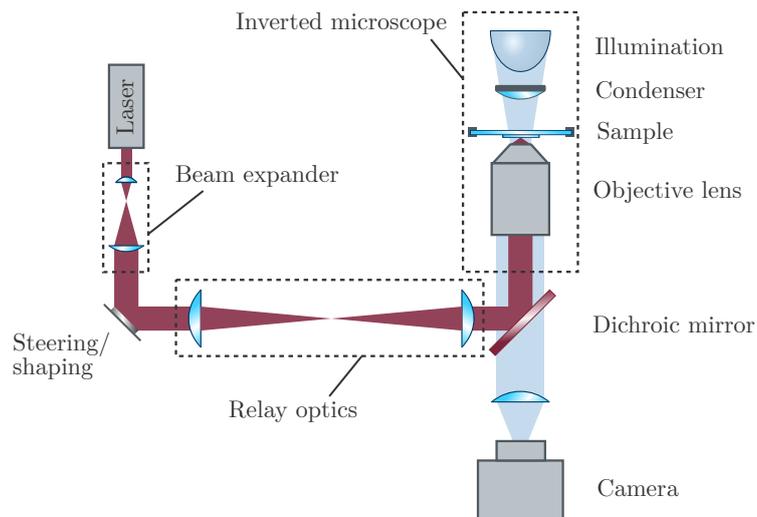


Figure 2.9: A single-beam inverted microscope optical tweezers setup.

To convert such an inverted microscope into optical tweezers, a laser beam has to be coupled to it. In their simplest form, the tweezers will generate a single stationary optical trap. The laser beam usually needs to be expanded and re-imaged onto the sample, which is achieved using relay optics, consisting of two lenses separated by a distance equal to the sum of their focal lengths, as illustrated in Figure 2.9. It will then be focused onto the sample by the same objective lens to create an optical trap - it is for this purpose that the objective needs to have a high NA, as explained in Section 2.4.1. We can now ‘capture’ optically trappable objects - conventional choices being silica or polyester micro-beads ($\sim \mu\text{m}$ in diameter). Usually laser sources with

wavelengths of 1064 nm and about 820 nm are used. These wavelengths coincide with local minima of the optical absorption for water, thus minimising the heating of the sample [89, 90].

If the optical trap is to be moved across the sample without translating the sample itself, some method for directing the beam is needed. A steerable mirror placed ahead of the re-imaging optics in the plane conjugate to the back aperture of the objective lens (as shown in Figure 2.9) can be used to move the beam in the xy -plane, while location of the focus along the z -direction can be changed by displacing one of the relay lenses along the beam axis. Combining two steerable mirrors with a couple of polarising beam splitters can result in creation of two movable optical traps [91]. Alternatively, multiple traps can be generated by time-sharing the single laser beam with the use of a scanning mirror or an acousto-optic deflector [92, 93]. Non-Gaussian beam shapes can also be created by using specialised optics. Bessel beams, for example, can be generated with the use of an axicon [94, 95], and Laguerre-Gaussian beams can be produced using specially designed holographic masks [96]. All of the above, however, are single purpose optical setups. An alternative method exists that can simultaneously steer, split and shape the beam using a computer-controlled diffractive optical element (DOE) - known as holographic optical tweezers (HOT).

2.7.2 Holographic optical tweezers

Holographic optical tweezers work on the principle that the phase of a light beam can be locally altered, causing the different parts of the beam to interfere with each other (after propagating some distance away from the plane where the phase modulation was applied), which results in some particular intensity pattern [12, 97, 98]. A familiar example of this is a diffraction grating with its interference pattern, which ‘splits’ the incident light beam into several diffraction orders. In HOT a diffraction grating can be created by using, for example, a spatial light modulator (SLM) [99]. An SLM has a liquid crystal display, subdivided into a rectangular grid of pixels, each a few or tens of microns across. A voltage can be applied across each pixel individually to change the orientation of the liquid crystals on that pixel, which in turn changes the refractive index (since liquid crystals are birefringent). So the light incident on different pixels will experience different phase shifts. Choosing an appropriate distribution of refractive indices across the SLM one can shape the incident light in a variety of ways.

One of the mechanisms for the user to address an SLM is by displaying a greyscale hologram on its screen (although not all SLMs have this feature). Usually each pixel

in the hologram image will have a value in the range from 0 to 255, which corresponds to the voltage that the user requests to be applied across each pixel on the SLM. Or - looking at things from another angle - each pixel in the hologram corresponds to the phase shift that the light incident on the corresponding SLM pixel will experience; this is in the range from 0 to 2π . Figure 2.10 shows some examples of such holograms along with how they change the trap appearance in the focal plane of the objective.

Naturally, a flat hologram (Figure 2.10(a)) will have no effect and we will see the original zero-order Gaussian beam trap. No matter what hologram is displayed on an SLM, it will not change the position of the zero-order beam, but it can direct some light away from it. Figure 2.10(b-d) shows a trap being deflected along x , y , and combined xy directions using phase ramp patterns. A phase ramp is wrapped at intervals of 2π , so it appears as multiple lines where the colour is changing from white to black across the width of each line. The steeper the phase ramp, the narrower the spacing between the consecutive phase wraps, the further away the trap will be shifted. Changing the sign of the gradient of a ramp will move the trap in the opposite direction.

It is also possible to move the trap in the z -dimension, i.e. change the location of the trapping beam focus with respect to the image plane of the objective lens. This is done with a hologram where phase ramps are circularly symmetric as in Figure 2.10(e), imitating the phase shift created by a Fresnel lens. Again, this can be combined with xy translation as seen in Figure 2.10(f).

These holograms for a single trap can be generated by calculating the phase pattern ϕ^S (where superscript S stands for ‘single’) using the following equation:

$$\phi^S(u, v) = \left(\frac{ku}{f}\right)x + \left(\frac{kv}{f}\right)y + \left(\frac{k(u^2 + v^2)}{2f^2}\right)z, \quad (2.30)$$

where k is the wavenumber of the laser beam and f is the focal length of the objective lens, the resulting trap will be located at (x, y, z) in the focal plane, and u, v denote the pixel coordinates on the SLM. The first two terms on the right-hand-side correspond to a diffraction grating shifting the trap in x and y , and the third applies a defocus by imitating a Fresnel lens.

Examples of holograms that create multiple traps with varying intensity and at different locations are shown in Figure 2.10(g,h). Such trap arrangements can be created using the gratings and lenses algorithm [100], which superposes the individual holograms of each trap:

$$\phi^M = \arg\left(\sum_i A_i \phi_i^S\right), \quad (2.31)$$

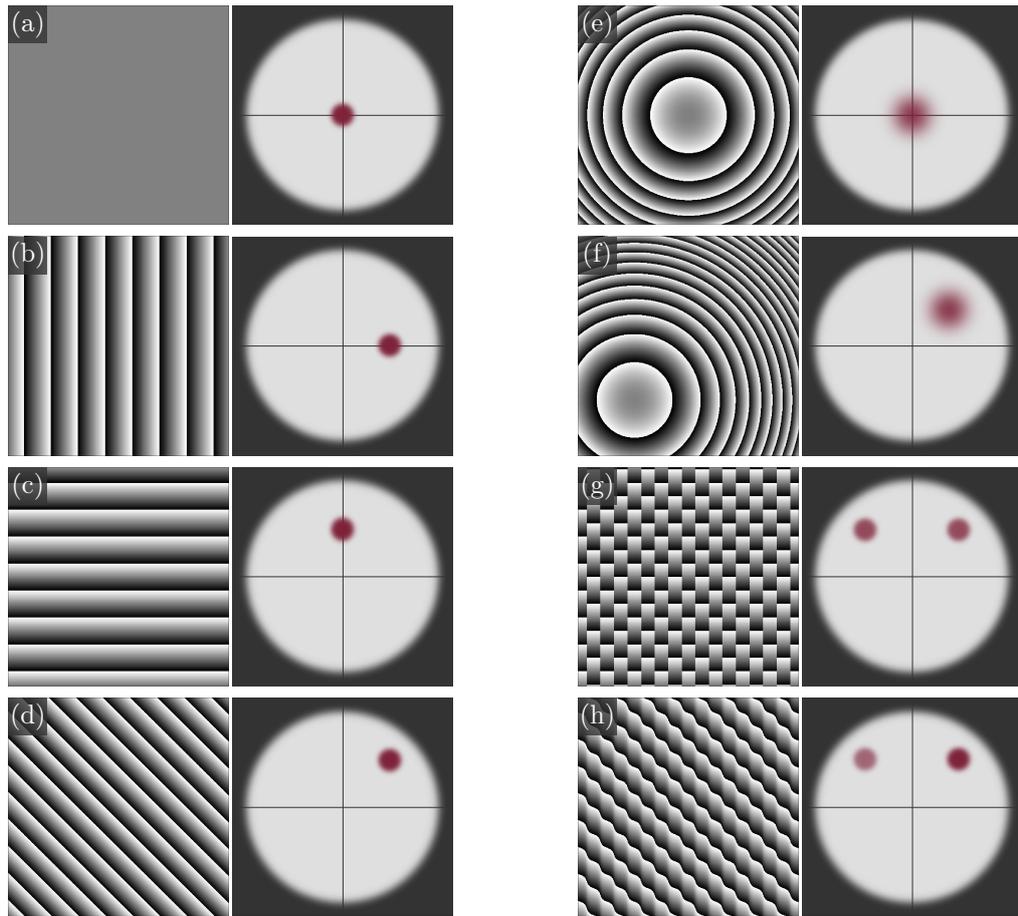


Figure 2.10: The working principle of HOT: holograms that can be placed on an SLM and schematics showing how the traps that they create would appear in the inverted microscope image plane. (a) Shows the original undisturbed Gaussian beam trap. The trap can be shifted along x (b), y (c) or combined xy (d) directions. The focus of the trapping beam can be moved out of the focal plane (e) and simultaneously translated in xy (f). Generation of two traps of equal (g) and different intensities (h). Zero-order beam and ghost traps are not shown.

where A_i is the relative amplitude of each trap and superscript M stands for ‘multiple’. This method is very computationally efficient, but it produces ghost traps - unwanted bright spots in the focal plane, which were never meant to be there. There are two reasons behind the emergence of ghost traps. When multiple traps are created by placing a diffraction grating pattern on an SLM, higher order interference maxima will also be generated. The second reason is that an SLM can only control the phase but not the amplitude of the incident beam, therefore some parts of the hologram will have more light reflected from them leading to imperfect interference in the focal plane. Methods for optimising the hologram pattern to perfect the intensity distribution in the focal plane exist, such as Gerchberg-Saxton [101] or the adaptive-additive [102], but they are computationally slower.

2.7.3 Data acquisition and applications

We have so far only discussed how optical tweezers can trap and move an object - more advanced experiments are possible, most of which require knowledge of the motion of the trapped objects. The position of one or several particles can be tracked by using digital video microscopy. As shown in Figure 2.9 the illumination light leaving the objective lens can be focused onto a digital camera to observe and record what is happening in the sample. These video images can be analysed either after recording or in real time during the experiment. High speed data acquisition is often desired in optical tweezers experiments (to observe events happening faster than typical optical trap relaxation times) requiring high video frame rates. Several kHz rates can be achieved even when using relatively inexpensive CMOS cameras and the 2D position of a particle can be tracked with an accuracy as good as several nm [103, 104], corresponding to around one hundredth of a pixel [105].

Image processing itself is a fairly simple procedure. When a spherical particle is slightly above the focal plane of the objective lens and is imaged from underneath, it appears as a bright spot surrounded by a dark ring (since it will act as a lens and focus the illuminating light onto the focal plane), as seen in Figure 2.11(a). Background noise can be thresholded out and the mean position of the intensity in the remaining image gives the location of the particle. This is done by weighting each pixel coordinate by its intensity before calculating the mean position and is computationally fast to perform. The particle trace can then be recreated from this data, see Figure 2.11(b). Tracking can be extended to 3D, with simple image analysis using stereoscopic imaging [106]. This technique uses two cameras to image the sample from two slightly different angles.

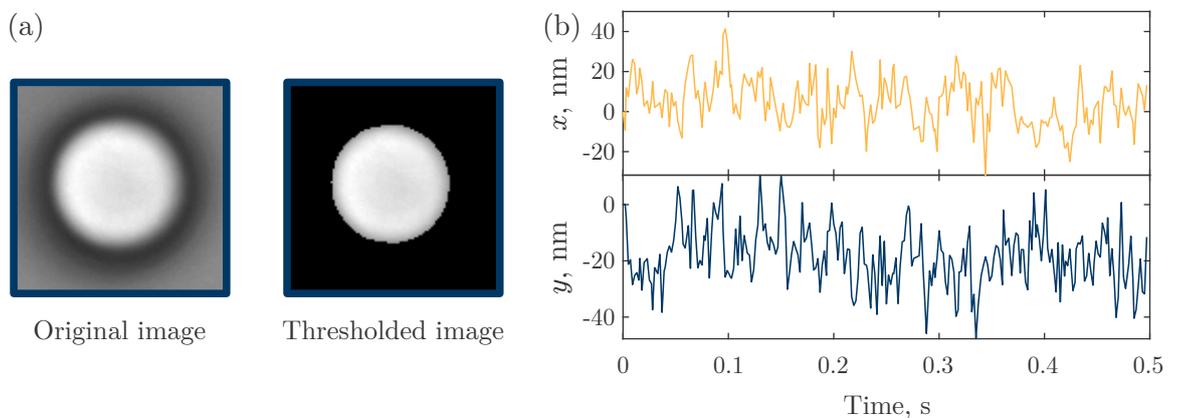


Figure 2.11: Image processing in optical tweezers. A camera image (a) of a $5\ \mu\text{m}$ radius silica bead held in an optical trap before (left) and after (right) it has been thresholded. The position trace of the bead (b) has been extracted from a video recorded at a frame rate of 600 Hz.

The lateral displacement of a particle in each image is linearly proportional to its depth (axial location). By finding the difference in the lateral location of the particle in the two images, information about its axial position can thus be determined.

Different kinds of information can be extracted from a recorded particle trace. Some are particularly useful for calibrating the tweezers [107–109]. Others provide information about the surrounding environment, for example, fluid viscosity [110–113] or its velocity field [114–116], and external forces (other than optical and hydrodynamic) acting on a trapped object [117–119]. Here we discuss two methods of analysing the particle trace data to determine optical trap stiffness: equipartition and power spectral density.

Equipartition method

We saw in Section 2.4.2 that in the limit of small bead-trap displacement the trap exerts a linear restoring force on the bead (as long as we assume a conservative system). This corresponds to quadratic energy stored in the trap $U(x)$:

$$U(x) = \frac{1}{2}\kappa(\delta x)^2. \quad (2.32)$$

Equipartition theorem states that a system at thermodynamic equilibrium has, on average, $\frac{1}{2}k_{\text{B}}T$ of energy in each degree-of-freedom. Equating the two energies gives an expression for stiffness:

$$\kappa = \frac{k_{\text{B}}T}{\langle(\delta x)^2\rangle}, \quad (2.33)$$

where $\langle(\delta x)^2\rangle$ is a time average of squared particle-trap separation, obtained from an experimental data set. Averaging over several data sets would further improve precision of this method.

Power spectrum analysis

Analytically, power spectral density (PSD) $P(\omega)$ of the over-damped Langevin equation is obtained by taking the squared modulus of its Fourier transform [120], giving:

$$P(\omega) = \frac{2k_{\text{B}}T\mu}{\omega_{\text{c}}^2 + \omega^2}, \quad (2.34)$$

where the corner frequency ω_{c} is defined as $\omega_{\text{c}} = \kappa\mu$. The PSD has a characteristic Lorentzian shape that can be seen in Figure 2.12. We can interpret this shape by considering the motions that a trapped particle undergoes. Just like a free object,

the trapped particle will undergo diffusion, taking it further away from the trap; this happens at short time scales (high frequencies). After some period of time, longer than diffusion time scales, the trap will exert enough force to pull the particle back - diffusion will be hindered on long time scales (low frequencies).

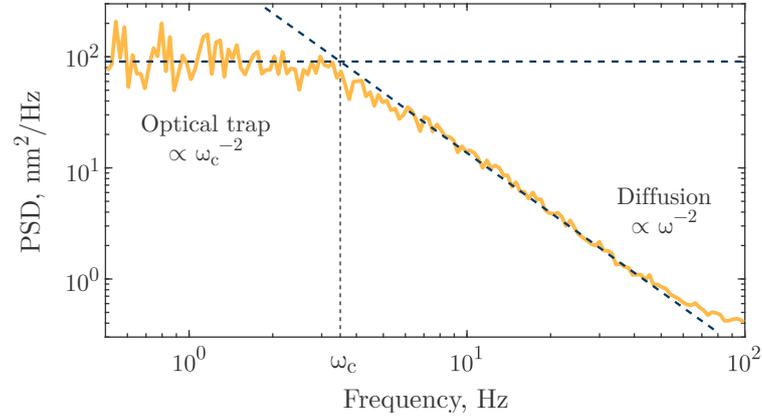


Figure 2.12: PSD of an optically trapped 5 μm radius silica bead. The plateau at low frequency is caused by the optical trap impeding the diffusion. At frequencies higher than the response time of the trap the diffusion signal dominates. Deviation from this trend at the high frequency end is caused by noise. Corner frequency ω_c is approximately 3.7 Hz here.

If needed, a least square fitting can be performed on an experimentally obtained PSD to determine its corner frequency and, hence, optical trap stiffness. This is considered to be the most accurate method for calibrating optical tweezers. Note that the expression given above is only valid for a particle trapped far away from any boundaries. Close to a surface hydrodynamic corrections need to be taken into account [55, 120].

2.8 Concluding remarks

In this chapter we saw how three different theories - Brownian motion, hydrodynamics and light-matter interaction - come together in an optical tweezers system. Equipped with the mathematical models of the three, and with the practical knowledge of how optical tweezers operate, we can now proceed to the main task of Part I of this thesis - designing and demonstrating optically driven hydrodynamic micro manipulation.

Chapter 3

Optically actuated hydrodynamic manipulation

‘ROSENCRANTZ: Why don’t you go and have a look?’

GUILDENSTERN: Pragmatism?! - is that all you have to offer?’

Rosencrantz and Guildenstern Are Dead — Tom Stoppard

The core idea of hydrodynamic manipulation is very simple - local currents within a fluid are used to move selected freely diffusing objects (this encapsulates both translation and rotation). The fluid itself can be set into motion by a number of means. Imagine, for example, a rubber duck that you can keep at the far end of the bathtub by pushing the water with your hand, whenever you feel like the duck is approaching too close. Except, in our case, instead of the bathtub we will have a water droplet on a microscope slide; instead of the rubber duck - a micro-sized particle; instead of the hand - some optically trapped object that can be moved. And instead of actively

making the decision to move it ourselves, we will have a feedback loop that will do that for us.

Let us break down the contents of this chapter. In Section 3.1 we will introduce the optically trappable actuators, which will be used to control the fluid flow, and develop the mathematical model of the feedback loop. In Section 3.2 we will computationally simulate an aqueous system of multiple particles, some of which are optically trapped, and then implement the hydrodynamic feedback within this simulation. We will further explore various aspects of our feedback performance in Section 3.3 within simulations, before proceeding to implement and test it experimentally in Sections 3.4 and 3.5.

3.1 Designing feedback for hydrodynamic manipulation

Our hydrodynamic feedback loop should have the following three functionalities. In order to direct the particle of interest, let's call it the target, towards some desired location, we need to know where it currently is. The first feature we need is, therefore, real-time video tracking. Second, we need to be able to calculate what actuator motion would drive the target towards this desired location - some sort of an equation will be needed for that. And third, the actuators have to be set into motion, by updating the optical tweezers.

The first and third steps - particle tracking and optical trapping are well researched and ready to be implemented (in fact, we had the benefit of an already existing tweezers setup with accompanying software [121] to perform all the basic tasks). The remaining thing to figure out then, is the details of step two - the maths that will tell us how to convert actuator motion into the desired target behaviour, and what exact shape the actuators will take.

3.1.1 Designing the actuators

Before we begin to think about a mathematical description of the hydrodynamics involved, we need to decide what kind of actuators we are going to use. Preliminary work investigating different options has already been done by my supervisor Dave Phillips in [122]. Suppose we choose to exert hydrodynamic control on a single degree-of-freedom of the target by translating a nearby optically trapped bead, as in Figure 3.1(a). Say, we find the target on the right of the desired location O, then the action that would

bring it back to O would be to drag the actuator to the left. Intuitively, we can tell that the distance the actuator has to cover would be larger than the separation between the target and point O . We could continue applying this principle for a while, as the target is continuously diffusing away from O , but eventually the actuator would either end up so far away that its motion would no longer produce a significant hydrodynamic force on the target, or the two beads would crash into each other. We therefore choose to have stationary (translation wise) actuators, which would control the fluid by rotating.

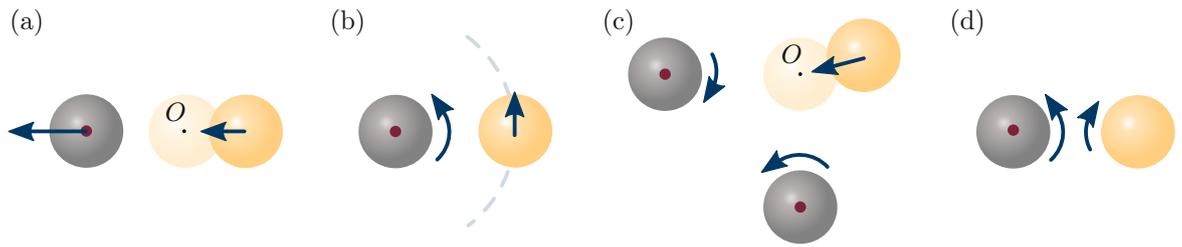


Figure 3.1: Actuator influence on a target. Translating an optically trapped actuator (grey) to the left (a) will drag the fluid and the free target (yellow) towards the desired location O . A spinning actuator (b) will affect two degrees-of-freedom of the target, pushing it along a circular trajectory, but introducing a second spinning actuator (c) will enable simultaneous control of the target’s translation in 2D. Orientation of the target can also be affected by a spinning bead (d). Red dots indicate optical traps.

For example, if we place a single spherical bead on the left of the target, Figure 3.1(b), the y -motion of the target can be addressed just by spinning the actuator - clockwise (about the z -axis) to move the target down in y , and anti-clockwise to move it up. This would, however, affect the x -motion as well. Addition of a second actuator displaced in y from the target, Figure 3.1(c), equips us with 2D translational control. This can be extended to include other degrees-of-freedom as well. The spinning actuators would also affect the orientation of the target, see Figure 3.1(d), by making it spin about the z -axis in the opposite direction to that of the actuator, due to the difference in flow velocity on opposite sides of the target. A bead spinning about the x - or y -axis would change the z -location of the target. With enough actuators strategically placed in the vicinity of the target, 6D¹ hydrodynamic manipulation could in principle be performed.

Causing spherical beads to spin is no trivial matter, though. It has been demonstrated that absorbing [123, 124] or birefringent [110, 125] particles can be rotated while in optical tweezers. Most notably, some of these methods make use of elliptically polarised light, allowing control of direction and speed of the rotation of the trapped

¹ By 6D here we mean six degrees-of-freedom (three translational and three rotational), not six dimensions.

bead. Reported rotation rates are, however, lower than what we anticipate to be necessary for our purposes. Another significant limitation of this technique is that an elliptically polarised beam can only induce rotation about its direction of propagation. It has been recently shown, though, that light beams with purely transverse angular momentum can be generated [126], which are capable of inducing rotation in trapped beads around axes transverse to that of the incident light. Together with recent developments in polarisation control [127] this would allow creation of multiple optical traps, each of different polarisation. While experimentally feasible, these methods would have required substantial modifications to our tweezers setup, and without a strong promise of success, we therefore decided to look for alternatives.

Drawing inspiration from the world of micro-tools, my supervisor Dave Phillips has designed wheel-shaped micro-rotors, shown in Figure 3.2(a). These micro-rotors consist of a smooth ring with three inner spokes, each with a prolate handle - optical trapping can be achieved by directing a laser beam through each of the handles, see Figure 3.2(b). Moving the three optical traps in a circle will spin the wheel. The smoothness of the outer ring ensures that the fluid flow around the rotor is mostly independent of its orientation², echoing the effect of a spinning sphere. These micro-rotors were fabricated at the University of Bristol by our collaborators, using direct laser writing [128]; for more information on this technique see Section 3.4.1.

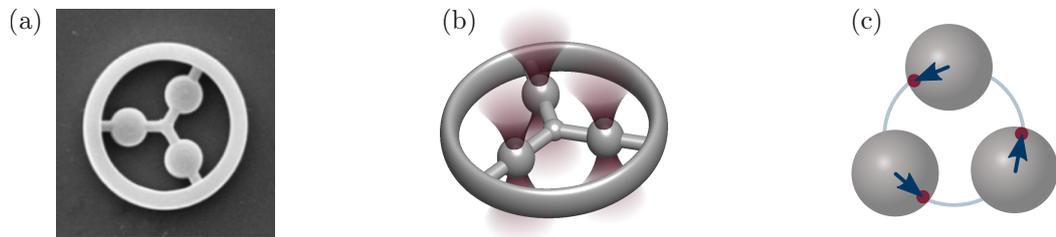


Figure 3.2: Different types of hydrodynamic actuators. A scanning electron microscope image (a) of a laser-printed micro-rotor (diameter of the handle beads is $4\ \mu\text{m}$) which can be optically manipulated with three laser beam traps (b), and a schematic of a constellation rotor (c) formed with three optically trapped spherical beads, confined to move along a circular trajectory (red dots indicate optical traps).

As we will see later on in this chapter, the laser-printed micro-rotors proved to be very effective at hydrodynamic manipulation, but they are not readily available for purchase and require specialised equipment to be fabricated. We therefore also suggest a different type of rotor - a constellation rotor - consisting of several spherical beads (mostly we will use three) held in a circular arrangement with optical traps, as shown

² There might be some minute effects due to the non-isotropic structure within the outer ring.

in Figure 3.2(c). We confine the beads to move along a ring, mimicking the behaviour of the micro-rotors, but the fluid flow they create is not isotropic and depends on rotor orientation - because of this we anticipate the need of a thorough mathematical description of the hydrodynamics involved. Note that here we use the term ‘micro-rotor’ to refer exclusively to the laser-printed rotors, while the term ‘rotor’ is more general and includes both micro- and constellation rotors.

3.1.2 Developing the maths

Now that we know what our hydrodynamic control setup will consist of, we can proceed to the maths. An important note before the work that follows - while in principle our wheel-shaped micro-rotors could be modelled mathematically using the bead-shell model [129, 130], this would be computationally very time consuming; we therefore approximate them as spheres, with confidence that any qualitative miss-estimations would be corrected for by the nature of closed-loop feedback. We saw in Chapter 2 in Equation 2.20 that velocity of a single particle, our target bead in this context, is related, via the mobility tensor, to the forces and torques acting on the other beads, the actuators. But it would be more convenient if we could relate the *rotation rates* of the actuators to the way in which we want the target to move. We would like to do this in a conceptually simple way, so we propose the following equation:

$$\mathbf{v}_t = \mathbf{C}_{tr}\boldsymbol{\omega}_r, \quad (3.1)$$

as it can be easily solved by inversion. Here \mathbf{v}_t is a column vector describing the velocity of the target (both translational and rotational), with one element for each degree-of-freedom (df) that we wish to control, making it a column vector with N_{df} elements. This vector incorporates both speed and direction that we want the target to achieve. $\boldsymbol{\omega}_r$ is a column vector containing the rotation rates of each actuator, with one entry per actuator totalling N_{act} elements. \mathbf{C}_{tr} is an $N_{df} \times N_{act}$ matrix expressing hydrodynamic coupling between the target and the rotors. As an example, say we wish to perform 2D hydrodynamic manipulation on the x and y motion of a single target using two rotors r_1 and r_2 . The above equation will take the following form:

$$\begin{bmatrix} v_{t_x} \\ v_{t_y} \end{bmatrix} = \begin{bmatrix} C_{t_x r_1} & C_{t_x r_2} \\ C_{t_y r_1} & C_{t_y r_2} \end{bmatrix} \begin{bmatrix} \omega_{r_1} \\ \omega_{r_2} \end{bmatrix}, \quad (3.2)$$

which we can solve by inversion to determine the required rotation rates $\boldsymbol{\omega}_r$, given the desired velocities v_{t_x}, v_{t_y} . To derive an expression for the matrix elements C in

Equation 3.2 we need to consider two things: which elements of the mobility tensor $\boldsymbol{\mu}$ are relevant, and what the constraints are on the rotor motion. Both of these will depend on the type of rotors used.

Let's first deal with micro-rotors; consider the 1D case of a target being manipulated along the x -dimension with one spinning micro-rotor. The micro-rotor r_1 is set into motion by applying a torque \mathbf{T}_{r_1} via translation of the handle spheres, so we need to consider the $\boldsymbol{\mu}^{\text{TR}}$ part of the mobility tensor. If we model each particle as a sphere, then the target's velocity is given by:

$$v_{t_x} = -\boldsymbol{\mu}_{t_x r_1}^{\text{TR}} \mathbf{T}_{r_1}, \quad (3.3)$$

where $\boldsymbol{\mu}_{t_x r_1}^{\text{TR}}$ is a 1×3 segment from the mobility tensor. We rewrite this equation by using the known relation for the torque magnitude on a single sphere, $T_{r_1} = 8\pi\eta a_{r_1}^3 \omega_{r_1}$, and writing the torque vector as $\mathbf{T}_{r_1} = \hat{\mathbf{b}} T_{r_1}$, where $\hat{\mathbf{b}}$ is a unit vector in the direction of the axis around which the micro-rotor is spinning. We have:

$$\begin{aligned} v_{t_x} &= -\boldsymbol{\mu}_{t_x r_1}^{\text{TR}} \hat{\mathbf{b}}_{r_1} 8\pi\eta a_{r_1}^3 \omega_{r_1} \\ &= C_{t_x r_1} \omega_{r_1}, \end{aligned} \quad (3.4)$$

Thus, we have derived $C_{t_x r_1}$ and can likewise obtain the rest of the \mathbf{C} -matrix in Equation 3.2, for the case of micro-rotors.

We note that there is an implicit assumption in the above derivation. While the solution for $\omega_{r_{1,2}}$ does take into account the *combined* effect of the two rotors on the target, it does not take into account the effect the rotors have on each other. This is because the \mathbf{C} -matrix does not contain terms from the mobility tensor that couple one rotor to the other - in this sense, from the perspective of rotor r_1 , rotor r_2 does not exist. Which is equivalent to saying that, from the point of view of rotor r_1 , rotor r_2 is a freely floating particle, i.e. the flow field generated by rotor r_1 is not affected by rotor r_2 . At first it seems that this assumption makes our approach invalid because, quite clearly, rotor r_2 is anchored in place by optical tweezers and will interfere with the fluid flow generated by rotor r_1 . However, the optical relaxation time of the micro-rotors is on the order of 100 ms (estimated from the corner frequency of experimentally obtained PSD), which is considerably longer than the control-loop feedback time period of ~ 5 ms (as will be explained in Section 3.2). Therefore, over the timescale of each feedback iteration, it is reasonable to assume that the rotors are able to freely drift in each others flow. If the optical relaxation time of the micro-rotors was significantly shorter than the 5 ms feedback time, the optical forces acting on the micro-rotors would have to be accounted for in the feedback model, which would require exact knowledge of the position of each

micro-rotor relative to the location of its trap.

The non-isotropic flow fields created by rotation of the constellation rotors can also be captured by Equation 3.1. However, in this case the derivation of the matrix elements of \mathbf{C}_{tr} is a little more involved than above. For simplicity, we consider a setup where one rotor, consisting of three beads b_1 , b_2 and b_3 , is used to control one degree-of-freedom of the target, v_{t_x} . Here, we are only concerned with how x translation of the target couples to translation of each rotor bead, so we only need to consider the $\boldsymbol{\mu}^{\text{TT}}$ part of the mobility tensor. To begin with, each of the three optically trapped beads has three translational degrees-of-freedom, so Equation 2.20 can be expanded as (ignoring the stochastic term):

$$v_{t_x} = - \left[\boldsymbol{\mu}_{t_x b_1}^{\text{TT}} \mathbf{f}_{b_1} + \boldsymbol{\mu}_{t_x b_2}^{\text{TT}} \mathbf{f}_{b_2} + \boldsymbol{\mu}_{t_x b_3}^{\text{TT}} \mathbf{f}_{b_3} \right]. \quad (3.5)$$

Since we are trying to determine the rotor motion required to translate the target at v_{t_x} , this is a single equation with nine unknown forces \mathbf{f} . However, in a rotor formation, the beads are confined to a circular trajectory. This means that all beads in the same spinning rotor will experience the same *magnitude* of the optical force, $F = \kappa\delta \equiv |\mathbf{f}_{b_{1,2,3}}|$, where κ is the trap stiffness (assuming isotropic traps) and δ is the distance between the trap and its bead. We will now see how, by considering geometry, the optical forces on each bead can be related to a single free parameter, the rotation rate of the rotor.

In 2D the direction of these vector forces can be written down relatively easily, but in 3D the derivation becomes more complicated. We have taken a general approach that is compatible with future extensions to arbitrary planes of rotation. We begin by observing that between successive updates the trap constellation is rotated by an

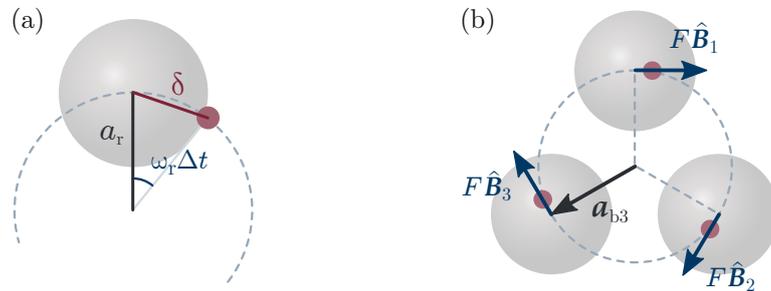


Figure 3.3: Bead constellation rotor geometry. In (a) an optical trap is translated by δ from the bead's equilibrium location. Assuming the displacement is small we can write $\delta = a_r \omega_r \Delta t$, where a_r is the rotor radius, ω_r is its rotation rate and Δt is the trap update time. Since all traps in the rotor will be moved by the same amount, all the beads will be subjected to forces of the same magnitude F (b).

angle of $\omega_r \Delta t$, where Δt is the update rate and ω_r is the chosen angular velocity of the optical traps (and hence of the rotor). From geometry, see Figure 3.3(a), we find that the magnitude of the force, F , on all three beads immediately after the SLM update is $F = \omega_r \kappa a_r \Delta t$, where a_r is the rotor radius. In order to determine the direction of the force, we express the optical force acting on bead b_j as $\mathbf{f}_{b_j} = \hat{\mathbf{B}}_j F$, where $\hat{\mathbf{B}}_j = [B_{b_{jx}}, B_{b_{jy}}, B_{b_{jz}}]$ is a unit vector pointing in the direction of the optical force acting on bead b_j , as illustrated in Figure 3.3(b), and the B 's are coefficients dependent on bead position. The B 's can be determined from the following three conditions. First, $\hat{\mathbf{B}}_j$ has magnitude one:

$$1 = B_{b_{jx}}^2 + B_{b_{jy}}^2 + B_{b_{jz}}^2 \quad (3.6)$$

Second, $\hat{\mathbf{B}}_j$ is tangential to the ring along which the bead is moving (assuming small displacements between trap and bead), i.e. it is perpendicular to the vector \mathbf{a}_{b_j} pointing from the centre of the rotor to the centre of the bead b_j , as shown in Figure 3.3(b). Hence, the dot product of $\hat{\mathbf{B}}_j$ and \mathbf{a}_{b_j} is zero:

$$B_{b_{jx}} a_{b_{jx}} + B_{b_{jy}} a_{b_{jy}} + B_{b_{jz}} a_{b_{jz}} = 0 \quad (3.7)$$

And third, $\hat{\mathbf{B}}_j$ lies in a plane perpendicular to the axis of rotation of the ring, defined by a unit vector $\hat{\mathbf{u}}_j$:

$$B_{b_{jx}} \hat{u}_{jx} + B_{b_{jy}} \hat{u}_{jy} + B_{b_{jz}} \hat{u}_{jz} = 0 \quad (3.8)$$

We know the location of the beads in the rotor (assuming they never lag behind their traps too much, so that we can use the trap locations; see Section 3.3.3 for more details) and our traps impose the rotor's axis of rotation, therefore we know \mathbf{a}_{b_j} and $\hat{\mathbf{u}}_j$, and can solve Equations 3.6-3.8 to determine the B -coefficients for each bead in a rotor. Consequently, the number of unknowns in Equation 3.5 is reduced from nine to one, and we can rewrite it to fit the form of Equation 3.2:

$$\begin{aligned} v_{t_x} &= - \left[\mu_{t_x b_1} \hat{\mathbf{B}}_{b_1} + \mu_{t_x b_2} \hat{\mathbf{B}}_{b_2} + \mu_{t_x b_3} \hat{\mathbf{B}}_{b_3} \right] F \\ &= - \sum_{j=1}^3 \mu_{t_x b_j} \hat{\mathbf{B}}_j \kappa a_r \Delta t \omega_r \\ &= C_{t_x r} \omega_r. \end{aligned} \quad (3.9)$$

The derivation demonstrated in this section can be extended to any number of degrees-of-freedom we wish to control, any type and number of actuators, and for any number of targets.

Different forms of the hydrodynamic feedback equation

Now that we know what \mathbf{C}_{tr} looks like, we can consider some implications of Equation 3.1. This system of linear equations can take three forms: overdetermined, underdetermined or critical.

In the case of an overdetermined system the number of rotors is smaller than the number of degrees-of-freedom we are trying to control, for example, if we try to control the motion of the target in both x and y with a single rotor. In this case there is most likely no solution for ω_r , except for some special cases.

A critical system will have a rotor for each degree-of-freedom, making \mathbf{C}_{tr} a square matrix, like in the previously discussed case presented in Figure 3.1(c). Provided that \mathbf{C}_{tr} is invertible, this case will have a solution. However, for some target-rotor configurations there will be no solution. For example, if we have four rotors manipulating x and y motion of two targets, and find them in the configuration illustrated in Figure 3.4(a), where the two targets have to be pulled apart. This example tells us that designing the hydrodynamic manipulation setup is not as straightforward as assigning one actuator for one degree-of-freedom. The two target problem could be solved in two ways. We can either move the rotors to break the symmetry whenever we find ourselves in a no-solution configuration; this would require an additional clause in the feedback and the targets would inevitably be disturbed while the rotors are re-configuring. Or, we can break the symmetry by introducing a fifth rotor, like in Figure 3.4(b).

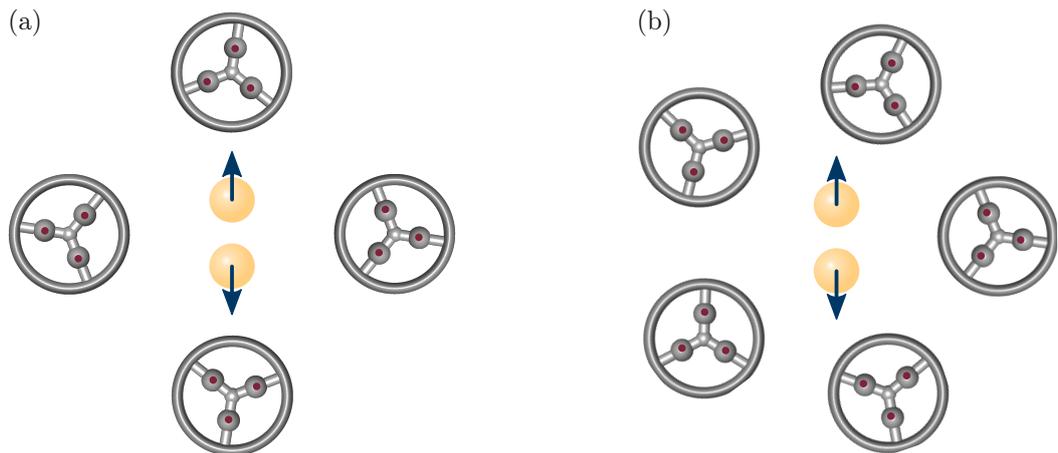


Figure 3.4: Symmetry implications for a hydrodynamic manipulation setup. In a setup where we have four actuators to control four degrees-of-freedom, we can end up in a situation which is impossible to solve (a). One way out is to break the symmetry by introducing a fifth actuator (b).

The presence of an extra rotor takes us to the final case - that of an underdetermined system - where we have more rotors than target degrees-of-freedom. As long as an underdetermined system is not inconsistent, i.e. having no solution, it will have infinitely many. This poses the question - are all solutions equally good or do we need to be clever about which solution we choose? We found that the conventional algorithms for solving systems of linear equations (such as MATLAB solvers ‘mldivide’ and ‘linsolve’, and LabVIEW solver ‘Solve Linear Equations’) almost always return a solution where one of the rotors is stationary. This is not a good use of the available laser power, as some of it would just be thrown away into trapping a rotor which is not even contributing to hydrodynamic manipulation. Since experimentally our system will be limited by the maximum possible rotation rate of the rotors (see Section 3.3.2 for more detail), we define the optimum solution to be the one in which the rotation rate of the fastest rotor is minimised. This way the highest contributor (the fastest rotor) will not be allowed to do too much work, thus encouraging the other rotors to be active as well.

We identify this solution by a simple trial-and-error method in which we input a range of different fixed rotation rates for one of the rotors. Once one of the rotation rates is fixed, we are back to solving a critically determined system. We obtain the solution of this system for each fixed input to acquire trial solutions for the rest of the rotors. Our estimate is refined with further rounds of trial solutions in the vicinity of the best solution of the previous iteration. The algorithm terminates if the solution is no longer changing, or after 11 iterations. We found that this method works well and is not computationally expensive.

3.1.3 Summary

We now have a complete description of our hydrodynamic manipulation setup. A freely diffusing target particle will be subjected to fluidic control in order to be moved along a desired trajectory or to be kept stationary. The fluid will be actuated using one of the two different types of rotors - wheel-shaped micro-rotors (Figure 3.2(a)) capable of generating a smooth fluid flow, or more readily available constellation rotors made up of several beads confined to move along a circle (Figure 3.2(c)). The needed actuator rotation rates and the desired motion of the target are linked via a matrix \mathbf{C}_{tr} , which encompasses information about the hydrodynamic behaviour of the system and the constraints on rotor motion. Computation of this matrix requires knowledge of the relative locations of the objects in the system. Since these positions are continually changing, the matrix will have to be recalculated in every iteration of the feedback loop. Once the necessary rotation rates are known, the actuators will be set into motion using

optical traps. The target will be pushed towards the desired location and the feedback loop will be repeated.

3.2 Simulating hydrodynamic manipulation

Before we jump straight to implementing hydrodynamic manipulation in an experimental tweezers setup, we wish to have a better idea of what to expect - a computational simulation will serve this purpose. Using mathematical models and numerical methods introduced in Chapter 2 we wrote a set of MATLAB scripts and functions that can simulate the behaviour of an arbitrary number of spherical particles - freely diffusing or optically trapped - in an unbounded fluid or near a planar surface. Within this simulation we then implemented the hydrodynamic feedback described in the previous section, to test how well it will work.

3.2.1 Simulating a system of an arbitrary number of spherical particles

To briefly recap the theory: the finite difference version of the over-damped Langevin equation (Equation 2.28) allows us to simulate the time evolution of N spherical particles immersed in a fluid. All particles will be subject to Brownian motion, modelled as white noise. Some of these particles will be optically trapped, by using the Hookean spring model to describe the restoring optical force. Unless otherwise stated, these particles will be spherical silica beads. And, as is common practice in low Reynolds number environments, we ignore the inertial effects.

Within our simulation all the particles have a corresponding optical trap, which we can ‘turn off’ for the freely diffusing ones simply by setting the trap stiffness to zero. Each bead is only responsive to its own trap, they are not affected by any other traps in the vicinity. Because the Hookean spring model allows the optical force to increase indefinitely with bead-trap separation, care must be taken to ensure that the simulation does not diverge from physical reality. This is done by throwing an error if the bead-trap distance exceeds one and a half of the bead radius. Then either the stiffness can be increased or the motion of the particle slowed down.

As illustrated in Figure 3.5, the simulation loop starts in a specified arrangement of optical traps and particles, described with \mathbf{x}_{trap} and \mathbf{q}_1 respectively. Note that \mathbf{q}_i includes orientation of every bead as well as its location; here subscript i denotes

the iteration number. Each iteration makes note of the current particle location to evaluate the $6N \times 6N$ mobility tensor $\boldsymbol{\mu}$. The displacement of each bead from its optical trap $\boldsymbol{\delta}$ is calculated as well, from which the optical forces $\boldsymbol{f}^{\text{opt}}$ can be inferred (using Equation 2.18). Brownian forces $\boldsymbol{f}^{\text{Brn}}$ are simulated as explained in Section 2.6. Positions of optical traps can be updated if desired; this is easiest to achieve by assigning trap locations for every iteration in advance. All elements are now in place to solve the Langevin equation for the location and orientation of the particles in the next iteration. Time-wise, the iterations are spaced apart by the time step Δt - care must be taken to ensure that this time step is significantly smaller than the optical relaxation time of the trapped beads, otherwise the solution will not converge, resulting in unrealistic oscillatory or diverging behaviour. The time step Δt must also be significantly larger than the momentum relaxation time t_m of the beads, because at this time scale the motion of the water molecules becomes ballistic and is no longer described by the uncorrelated white noise that we defined in Sections 2.2 and 2.6. The loop repeats for a set number of times N_{it} , or until some other specified condition is met.

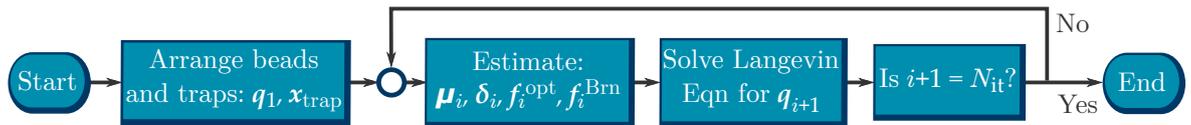


Figure 3.5: Flowchart of the multiple particle simulation.

Fluid flow maps

Our hydrodynamic manipulation technique does not explicitly require the knowledge of the fluid flow, nevertheless, visualising it is very useful for developing an intuition about the behaviour of our system. Flow fields created by translating and/or rotating particles can be estimated using Equation 2.20 in the limit $a_i \rightarrow 0$. In practice we do this by defining a 2D grid on the area of interest and placing a zero-size tracer particle at each point of the grid. The relevant mobility tensor terms are then calculated and the velocity of the tracer particles determined. We reiterate that the laser-printed micro-rotors are modelled as spheres, and, since we cannot place tracker particles inside the spheres, we assume a $1/r^2$ fluid velocity fall-off from the inside edge of the rotor disk, to plot the flow-fields inside the micro-rotors.

Figure 3.6 presents a comparison between flow fields generated by the two actuator types. A smooth isotropic flow is generated by a spinning micro-rotor, while we see

three-fold symmetry in the flow around a constellation rotor. The figures include flow streamlines which are everywhere parallel to the instantaneous fluid velocity.

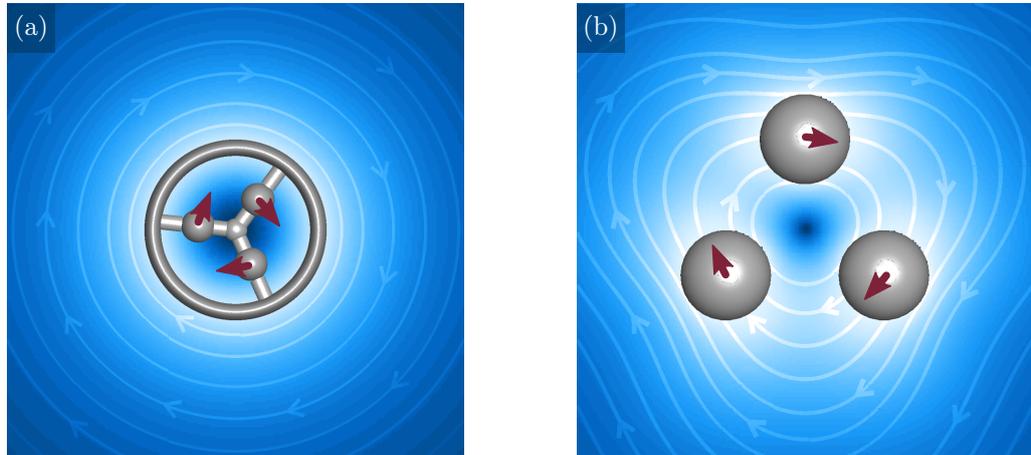


Figure 3.6: Fluid flow maps for different actuator types. The flow generated around a spinning micro-rotor (a) and a three bead constellation rotor (b). Background and streamline colour is proportional to fluid speed.

Validation

A limited number of experimental results characterising hydrodynamic interactions of spherical particles are available in the literature. In one of them, in order to obtain quantitative information about hydrodynamic interactions between colloidal particles, Meiners and Quake [131] observed two optically trapped latex beads. From measured position fluctuations of the particles they calculated the auto- and cross-correlation functions and compared them with analytical models. The theoretical functions showed very good agreement with experimental results and so we used the former to validate our simulation. These analytical functions were derived from the Langevin equation, using the Oseen approximation to the hydrodynamic mobility tensor and can be found in the cited paper as Equations (4-5).

Figure 3.7(a) shows the autocorrelation of the simulated behaviour of a single bead in an optical trap to be in excellent agreement with the theoretical prediction. Figure 3.7(b) shows the cross-correlation functions for beads separated by $8\ \mu\text{m}$ - also in excellent agreement with the theory. Thus we are confident that our simulation predicts correct behaviour in an unbounded fluid.

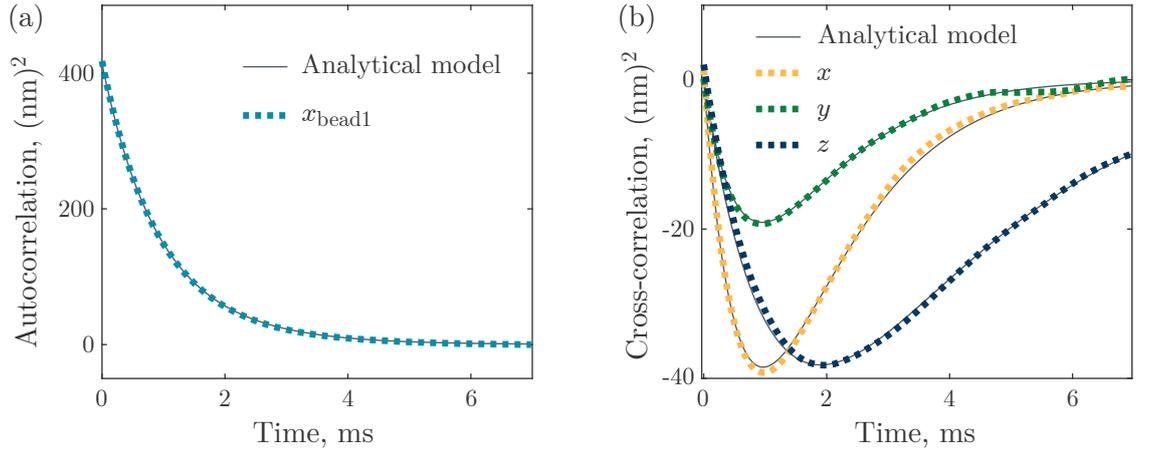


Figure 3.7: Validation: correlations between optically trapped beads. Autocorrelation of a single optically trapped silica bead (a) and cross-correlation between two optically trapped silica beads separated by $8\ \mu\text{m}$ along the x -direction (b). Correlations were estimated from simulation data of 600 s for $1\ \mu\text{m}$ diameter beads in optical traps of stiffness 10^{-5} , 10^{-5} and $0.5 \times 10^{-5}\ \text{N m}^{-1}$ in x , y and z respectively, with a time step $\Delta t = 2 \times 10^{-5}\ \text{s}$, at the temperature of 300 K, with other properties as specified in the front matter. The simulations are compared to analytical correlation predictions for the same system.

Diffusion of two spherical beads close to a planar surface was experimentally and theoretically investigated by Dufresne et al. in [132]. This paper considered diffusivity of two modes of motion of the spheres: relative D^R (to describe how the particles move relative to each other) and collective D^C (to describe the motion of the centre of mass). For each mode two directions were defined: parallel D_{\parallel} and perpendicular D_{\perp} to the bead separation vector; both of these are parallel to the surface, and diffusion in the direction perpendicular to the surface is not considered. Thus there are a total of four modes M . Theoretically derived diffusivities were shown to agree with experimentally obtained values.

Diffusivity of any mode $D^M(d, h)$ at a given bead separation d and height above the surface h can be obtained from the mean squared displacement (MSD) of the coordinate of that mode r^M :

$$\langle (r_{t_0+\tau}^M - r_{t_0}^M)^2 \rangle = 2D^M\tau, \quad (3.10)$$

where the angle brackets denote an ensemble average, t_0 is the start time and τ is the duration of the observation. To recreate Dufresne's experiments and calculate the diffusion coefficients $D(d, h)$ we simulated 8000 sets of two freely diffusing beads for different starting combinations of d and h , and of $\tau = 20\ \text{ms}$ duration. Calculating the MSD of each set and taking an ensemble average then gives the diffusivities. These results are compared with the theoretical model in Figure 3.8, and show very good agreement. We begin to see some divergence from theory in Figure 3.8(c) at smaller

bead separations and closer to the surface, as this is approaching the limit of the accuracy of the mobility tensor equations that are used in the simulation.

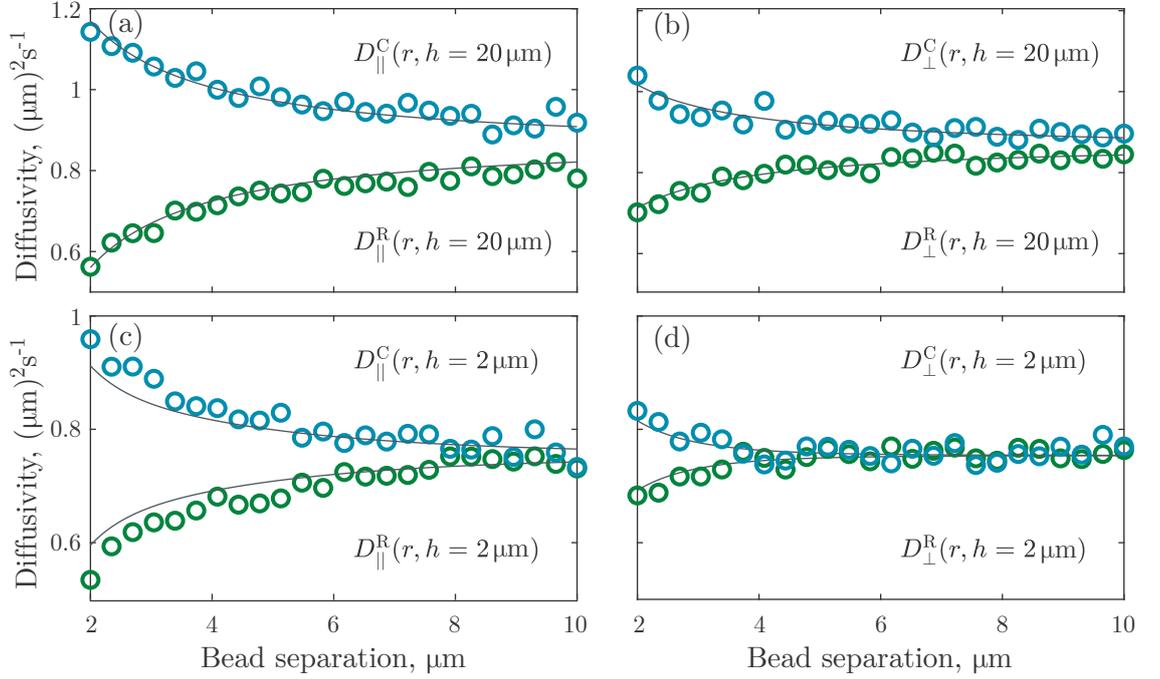


Figure 3.8: Validation: diffusivities near a surface. Comparison of diffusivities as defined by Dufresne (black lines), and estimated from 8000 sets of 20 ms duration simulations (coloured circles) of two 1 μm diameter beads diffusing near a planar surface at different bead separations and heights h . Blue circles refer to collective and green - to relative motions. The time step used was $\Delta t = 4 \times 10^{-5} \text{ s}$, with other parameters being identical to the ones in Figure 3.7

3.2.2 Implementing hydrodynamic feedback

Having validated the simulation code we proceeded to implement the hydrodynamic feedback equations from Section 3.1.2. The relevant flowchart is presented in Figure 3.9.

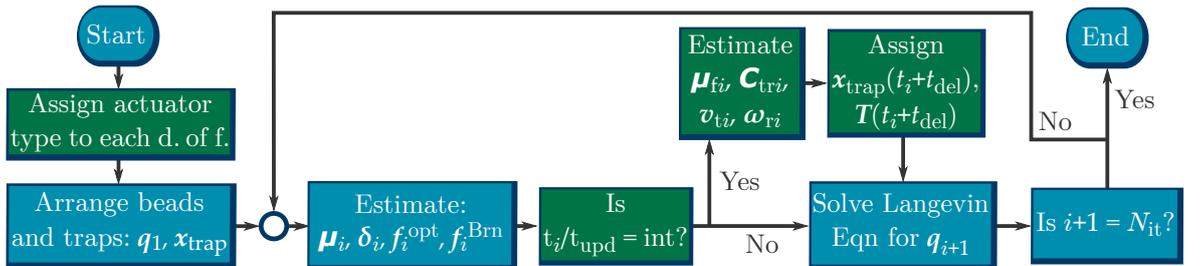


Figure 3.9: Flowchart of the hydrodynamic feedback implemented into the simulation code.

Before the simulation starts we need to choose what kind of actuators are going to be used - bead constellation rotors or micro-rotors, which, as discussed earlier, will be modelled as spinning spheres - and where we are going to place them. Constellation rotors can be rotated by translating the optical traps along a circular trajectory, but to spin the spheres we have to assign an ‘optical’ torque. Since we are not planning to use spinning beads in an actual experiment, this torque will simply be a number estimated from the required rotation rate using $T = 8\pi\eta a_r^3\omega_r$ as discussed in Section 3.1.2, without attempting to simulate optical origins of such a torque.

To mimic an experiment, we do not solve the hydrodynamic feedback equation, Equation 3.1, in every iteration of the simulation, but account for the presence of possible delays. First, we have the update rate of the SLM which determines how frequently the optical trap locations can be updated. Our SLM can be updated at ~ 200 Hz setting the trap update time t_{upd} to 5 ms. The feedback code is therefore executed every t_{upd} seconds. We must also account for the computational time required for image processing to track the target, performing hydrodynamic calculations and generating the hologram, which together amount to a delay time t_{del} of 10-15ms. This delay is accounted for in the simulation by updating the trap locations and torque, only t_{del} after the iteration during which they were estimated. During this delay time the target is able to diffuse away from its registered position, so the shorter the delay time is, the more accurate the feedback will be.

The trap update time also serves in estimating the desired target velocity, needed to move the target from its current location and/or orientation to the desired one: $\mathbf{v}_t = \frac{\mathbf{q}_{t,\text{desired}} - \mathbf{q}_t}{t_{\text{upd}}}$. This might, as t_{upd} is very short, result in unreasonably high requirements for actuator rotation rates ω_r , which will be discussed more in Section 3.3.2.

It is important to note that we need a new mobility tensor $\boldsymbol{\mu}_f$ for the feedback calculations, different from $\boldsymbol{\mu}$ which is used in solving the Langevin equation. The main reason is that in an experiment we will not know the actual locations of the actuators, unless we video track them along with the target. We will, however, have access to the trap locations - it is using these that $\boldsymbol{\mu}_f$ is estimated; Section 3.3.3 discusses the implications of this in more detail. We might also want to investigate the effect of using different equations for $\boldsymbol{\mu}_f$ (e.g. a different number of expansion terms) than those for $\boldsymbol{\mu}$, which will be based on the most accurate possible expansion to make the simulation physics as realistic as possible. We do not, however, need to calculate the entire $\boldsymbol{\mu}_f$, saving some computational time, only the parts which are needed for estimating \mathbf{C}_{tr} . The effects of choosing different methods for calculating $\boldsymbol{\mu}_f$ are discussed in Section 3.3.1.

Comparing actuator types

Before we discuss the influences of different parameters of the system in more detail, we present sample results to illustrate the performance of our hydrodynamic manipulation method. We assigned the feedback with the task of suppressing the Brownian motion of a target bead by keeping it at a constant location - we call this hydrodynamic clamping. For now, we only deal with 2D translation clamping in x and y directions. Figure 3.10 shows the position scatter plots together with 1D histograms of a hydrodynamically clamped target for both types of actuators. To quantify the effectiveness of clamping we consider the standard deviation σ of the target's position over the time of the simulation, defined as the average of σ_x and σ_y . We see that micro-rotors perform nearly 50% better than constellation rotors, which we attribute to two causes. As we will see in Section 3.3.3, constellation rotors do not actually behave exactly as the feedback prescribes, while the micro-rotor response to the applied torque is immediate (within a simulation). We also expect that the inherent orientation dependence associated with the constellations impairs their capabilities. However, we expect this difference between the two actuator types to become smaller in the experiments, with a less fruitful outcome for the micro-rotors, as they will also suffer from a non-perfect response to the optical traps.

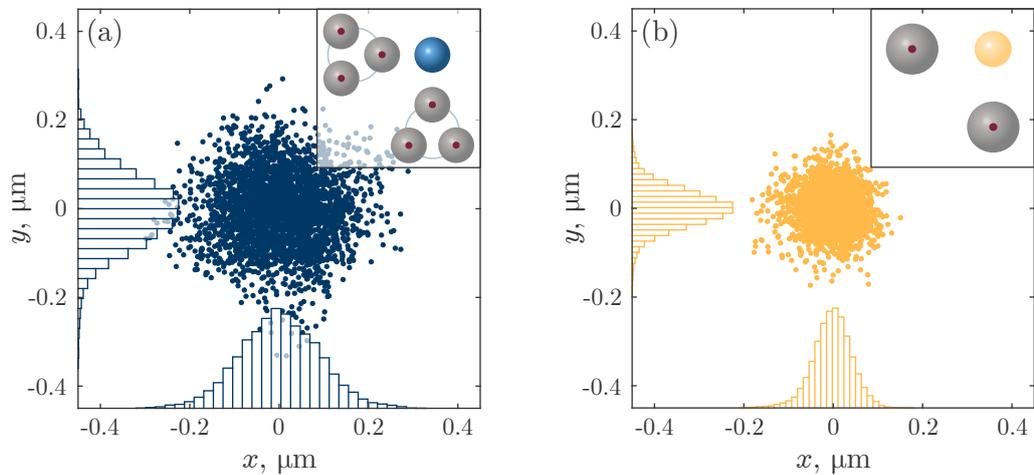


Figure 3.10: Simulation results for 2D clamping: using constellations (a) and micro-rotors, modelled as spheres, (b). The mean standard deviation of the $5\ \mu\text{m}$ radius target's motion in the two simulations was $\sigma = 83.3\ \text{nm}$ and $\sigma = 42.1\ \text{nm}$ for constellations and micro-rotors respectively.

We now proceed to examine a variety of variables and sources that can impact the results of our hydrodynamic manipulation technique.

3.3 Tailoring hydrodynamic feedback

In the previous section we have demonstrated the viability of our hydrodynamic manipulation scheme. We now wish to investigate various aspects of it in more detail to see if, and how, it can be improved, before being applied in an experiment. For example, we would like to know how hydrodynamic clamping in an unbounded fluid is different from one near the bottom of the sample cell; how many degrees-of-freedom we can expect to be able to control; what magnitudes of actuator rotation rates are necessary for achieving successful hydrodynamic manipulation. These and more questions will be investigated in this section.

3.3.1 Different forms of μ_f

We now return to the question of the exact form of the mobility tensor used in the feedback equations, μ_f . Since our experiments will be done with sedimented beads, the mobility tensor μ that governs the physical behaviour of the system will be based on Equations 2.14, which include the surface terms. But we may not necessarily want to use the same complex equations in the feedback. We have several different options for μ_f . The simplest form that the equations can take, the Oseen expansion³, accounts only for the first order terms. Including an additional term in μ^{TT} of the order a^2/r^3 gives us the Rotne-Prager tensor (a term of this order does not appear in μ^{TR} and is therefore not applicable to micro-rotors). And finally, we have a mobility tensor which includes the surface terms in the expansion. We also investigated what happens if the orientation dependence of the constellation rotors is ignored, i.e. if the feedback sees one constellation as a single spinning sphere.

μ_f	σ , nm	
	Constellations	Micro-rotors
Oseen	83.82	42.10
Rotne-Prager	85.55	—
Surface terms	85.00	43.06
Ignoring orientation dependence	106.6	—

Table 3.1: Effect of the form of the mobility tensor used in the feedback loop. Performance of hydrodynamic clamping, expressed as the standard deviation σ of the target's motion, with different equations for the mobility tensor μ_f in the feedback.

³ Note that technically ‘Oseen’ refers only to the μ^{TT} part of the tensor, and therefore applies only to the constellation rotors, but here we will use this term to also refer to the first order expansion of the μ^{TR} part.

The simulated results of hydrodynamic clamping set out in Table 3.1 show no significant difference between the alternative forms of $\boldsymbol{\mu}_f$. We conclude that, while the feedback is slightly more accurate with longer expansions for the mobility tensor, this improvement is not substantial and is lost in the noise of other limitations of the system. What stands out, however, is the effect of ignoring orientation dependence of the constellation rotors - a 25 % reduction in clamping performance. This highlights the importance of taking into account the hydrodynamic couplings between the target and all the actuator beads. The main take-away message here is that the lowest order expansion of $\boldsymbol{\mu}_f$ will be sufficient for experimental implementation.

3.3.2 The need to limit actuator rotation rates

There are two reasons why we might need to scale down the estimated actuator rotation rates before they are applied to the optical traps holding the actuators. The more obvious one stems from experimental limitations - if the optical traps are moving too fast, the actuator beads will not be able to keep up because of increasing hydrodynamic drag. This way we will either have the rotors spinning much slower than the traps (as the traps will jump from dragging one bead to dragging the next), or will simply lose the actuators, as they will be pushed upwards in the direction of the laser beam by an unfavourable light field. While the simulation is not actually limited by the above reasons, since each bead is aware of only one optical trap and the force response is linear, we wish to keep things as realistic as possible so as not to have far-fetched expectations in regard to experiments.

Another, more subtle, motive originates from the presence of delay time. Suppose that at some point the target was found to be some distance away from the desired location, point O, and the rotors are already spinning to push it back. When the next feedback iteration is entered, the target's location will be registered, based on which, new actuator rotation rates will be estimated. But these rotation rates will only be applied some 10 ms later, by which time the target will have been pushed further along towards O. So the new rotation rates will be applied when the target is closer to O than the feedback thought - therefore overestimating the necessary rotation rates. This leads to less efficient clamping because the target ends up continuously overshooting O.

Figure 3.11(a) shows the behaviour of a target clamped with two micro-rotors when the maximum available rotation rate of the actuators ω_{lim} is limited to three different values. At $\omega_{\text{lim}} = 710^\circ \text{s}^{-1}$ the target oscillates about the desired location ($x, y = 0$). Reducing ω_{lim} to 230°s^{-1} significantly suppresses the amplitude of target's motion.

But bringing ω_{lim} too low, e.g. 30°s^{-1} is no longer very effective as now it takes a lot of time for the rotors to push the target back to 0. It is not immediately obvious by how much the rotation rates need to be scaled down for best clamping results, requiring a bit of trial and error. The yellow data points in Figure 3.11(b) show the standard deviation of the clamped target, indicating the best results at $\omega_{\text{lim}} = 230^\circ \text{s}^{-1}$ (for the case of constellation rotors formed from three $5 \mu\text{m}$ radius beads located on a ring of radius $8 \mu\text{m}$ and placed $22 \mu\text{m}$ away from the target, the optimum rotation rate was found to be 150°s^{-1}).

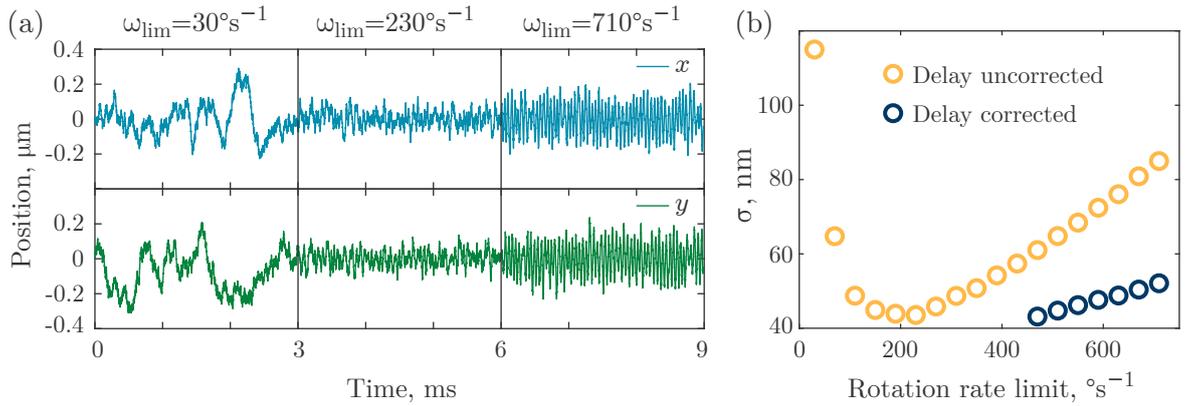


Figure 3.11: Limiting actuator rotation rates. Trace of a target particle (a) when rotation rates of two spinning actuators are limited to $\omega_{\text{lim}} = 30, 230$ and 710°s^{-1} . Standard deviation σ of the target (b) for different limits on actuator rotation rates without (yellow) and with delay-correction (blue). Each data point was estimated from a 60 s simulation.

Another, more rigorous, option for combating delay is attempting to predict the actual position of the target at the time when the rotation rates will be actually applied. We can do this by solving another finite difference equation to simulate the trajectory the target would take during t_{del} . Once the end point \mathbf{q}_{pred} of this trajectory is computed, $\boldsymbol{\omega}_{\text{r}}$ is updated based on this predicted position, rather than the one registered by tracking. The equation that needs to be solved here will look very much like the Langevin equation, but it will no longer have the stochastic term:

$$\mathbf{q}_{l+1} = \mathbf{q}_l - \boldsymbol{\mu}_{\text{fl}} \mathbf{f}_l^{\text{opt}} \Delta t, \quad (3.11)$$

following the same conventions as in Section 2.6. Note, that this method can only correct for the overestimation of $\boldsymbol{\omega}_{\text{r}}$, and not the fact that the target also diffuses away from its registered position during t_{del} , as we do not have the means to predict Brownian motion.

We do not want this additional computation to contribute to the already existing delay. For this reason, the time step Δt in Equation 3.11 is set to 10 ms - same as t_{del} - allowing

the equation to be solved in one step. Duration of 10 ms is still smaller than the optical relaxation time, so it can be used as a value for the time step when solving the finite difference equation. We estimated additional ~ 0.2 ms of computation time needed to solve Equation 3.11 for a system of three particles, which is negligible compared to t_{del} . This time increases up to ~ 1 ms if three-bead constellation rotors are used instead.

The blue data points in Figure 3.11(b) illustrate how effective this technique is - the standard deviation after including Equation 3.11 in the feedback is reduced by up to 40%, and shows the most improvement at the highest rotation rates. We still see, however, that even with the delay correction we do not obtain the same σ for different rotation rates. Another method worth exploring in the future for combating over-actuation is correction of actuator rotation rates based on how close the target is to its desired location. The shorter this distance, the lower the rotation rate, thus minimising the chances for the target to over-shoot the desired position. We expect that in conjunction with delay correction, this method would improve clamping even more. The rotation rates experimentally achievable with our optical tweezers system do not, however, reach the magnitudes that can significantly over-actuate the target and, therefore, we will not be implementing the delay correction technique in the experiments.

3.3.3 Influence of optical trapping stiffness

At first it might seem that to estimate μ_f we need the exact knowledge of the actuator positions (which would imply video tracking of all the particles in the experimental system, not just the target). We could, though, use the readily available trap locations instead, as an approximation.

The beads in rotating constellation rotors will always lag behind their traps - the lower the optical trapping stiffness, the higher the lag (Figure 3.12(b)), and the higher the inaccuracy of the feedback if trap locations are used to represent bead locations. Figure 3.12(a) compares (for different optical trapping stiffness values) the hydrodynamic clamping performance when the feedback equations use the actual actuator bead locations (grey) and when their optical trap locations are used instead (red). There is, surprisingly, hardly any difference between the two, and knowledge of precise bead locations does not necessarily result in tighter clamping (e.g., the first point on the graph, where the bead-trap lag becomes comparable to the bead size, shows substantially poorer performance when bead locations are used). We explain this by noting that the important quantity for estimating μ_f - separation between the beads - is an order of

magnitude larger than the miss-estimation caused by not knowing the exact actuator locations. It therefore does not matter which locations we use in the feedback.

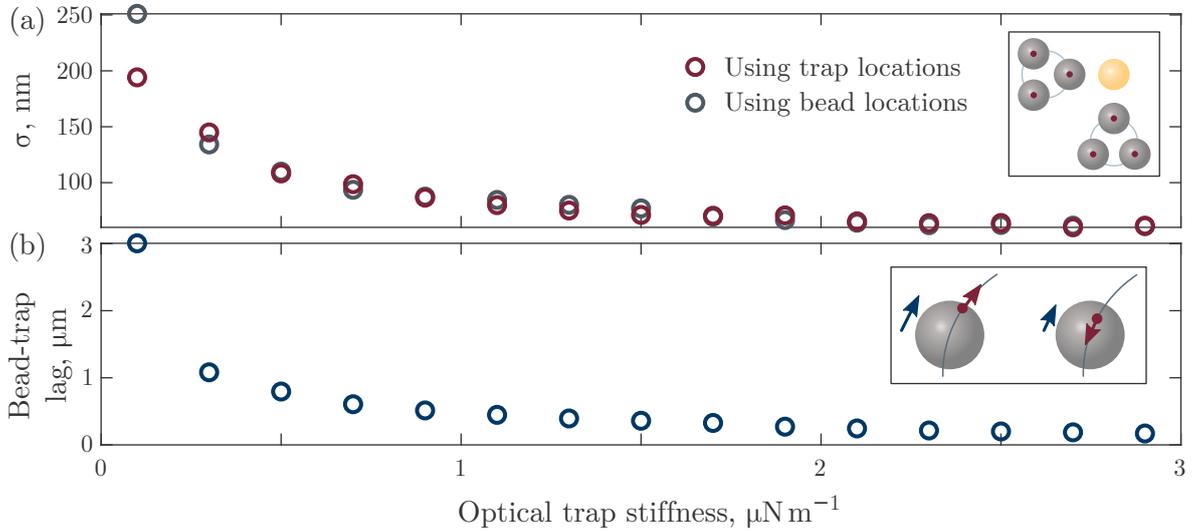


Figure 3.12: Effect of optical trapping stiffness on hydrodynamic clamping. Dependence of the standard deviation of a clamped bead on the optical trapping stiffness (a) in a setup with bead constellation rotors, as illustrated in the inset. Average lag between actuator beads and their optical traps, for different trap stiffness (b); inset: illustration of how the assigned rotation rate is not followed by a bead in a constellation rotor; blue arrows indicate the motion of the bead, red arrows show the direction in which the trap was translated. Each data point was estimated from a 60 s simulation.

We further notice, in Figure 3.12(a), a significant drop in performance going from high to low optical stiffness, regardless of whether the trap or actuator locations are used in the feedback. This can be explained by the fact that the beads which are loosely trapped are not repeating the motion of the traps with high accuracy, and the feedback is not aware of that. The equations we use assume that the beads will have always caught up with their traps at the end of each feedback iteration - the looser the trap the less true this assumption is. In some cases the rotor might even end up spinning in the opposite direction to the assigned one, as illustrated in the inset of Figure 3.12(b). Suppose a bead in a constellation rotor is being moved clockwise by a trap at its edge. In the next feedback iteration, the rotor is assigned a lower rotation rate in the opposite direction - the trap will make a step in the correct, anti-clockwise, direction, but it will still be positioned to move the bead clockwise. We could potentially account for the bead-trap lag caused by low stiffness, if we could estimate it in every feedback iteration, which would again require knowledge of actuator locations.

3.3.4 Diffusion effects

We expect the rate of diffusion of the target to be one of the main factors determining how well hydrodynamic clamping works. Here we investigate two sources that have an impact on diffusion - particle size and the presence of a boundary.

As dictated by the fluctuation-dissipation theorem ($D \propto a^{-1}$), smaller particles diffuse further in any given amount of time, so we expect the performance of hydrodynamic clamping to worsen with decreasing target particle size. This is indeed the case; Figure 3.13 shows that clamping is more effective for larger target particles, following the expected a_t^{-1} trend.

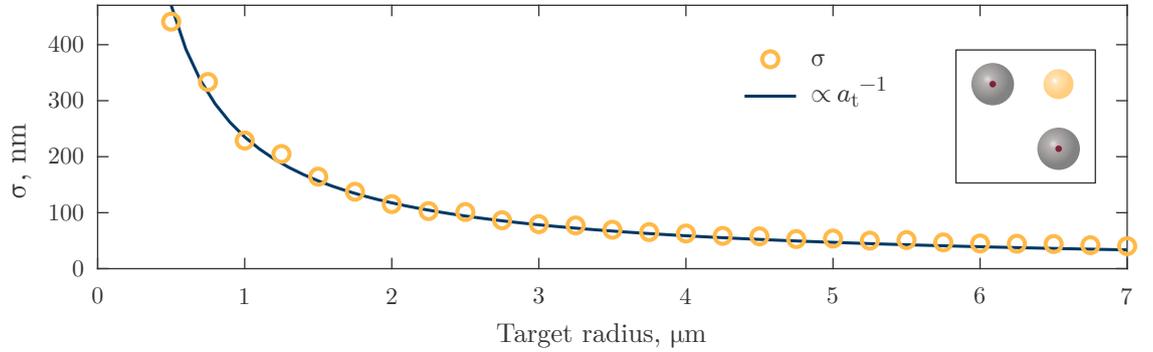


Figure 3.13: Hydrodynamic clamping of different size targets. Stiffness of hydrodynamic clamping expressed as the standard deviation σ of the target’s motion, in a setup with two spinning-sphere actuators as shown in the inset. The blue line is a fit inversely proportional to the target’s radius a_t , with a fitting coefficient of 235.2.

Diffusion gets hindered in the presence of surfaces, so we expect the fact that our experiments will be conducted with sedimented beads to work in our favour. In particular, diffusion in the direction parallel to a nearby planar boundary is given by the Faxen formula [133]:

$$D_{\parallel}(h) \propto 1 - \frac{9}{16} \left(\frac{a}{h}\right) + \frac{1}{8} \left(\frac{a}{h}\right)^3 - \mathcal{O}\left(\left(\frac{a}{h}\right)^4\right), \quad (3.12)$$

which is equivalent to the mobility tensor expansion in Equations 2.14.

As illustrated in Figure 3.14, which compares the standard deviation of a target clamped at different heights above the surface, simulations predict a $\sim 30\%$ clamping improvement in the near vicinity of the surface. We also see that clamping follows a similar trend to diffusion, as dictated by Equation 3.12. However, near the surface the standard deviation of the target’s motion does not quite reach the values we would expect based on diffusion theory. This can be explained by the fact that the maximum allowed actuator rotation rate ω_{lim} is the same for different h values, but the speed

of the target next to an actuator spinning at a constant rotation rate will be smaller closer to the surface (the relationship showing this (Equation 3.13) will be derived in Section 3.5.2). Therefore, hydrodynamic clamping near the surface, for the particular value of $\omega_{\text{lim}} = 90^\circ$ that we used here, is slightly worse than we would expect based on diffusion effects alone. We also note that close to the surface the maximum rotation rate of the actuators is decreased, as dictated by Equations 2.14. In addition to this, while in an unbounded fluid a particle spinning about the z -axis only induces fluid flow in the xy -planes, next to a surface it will also create some flow in the z -direction.

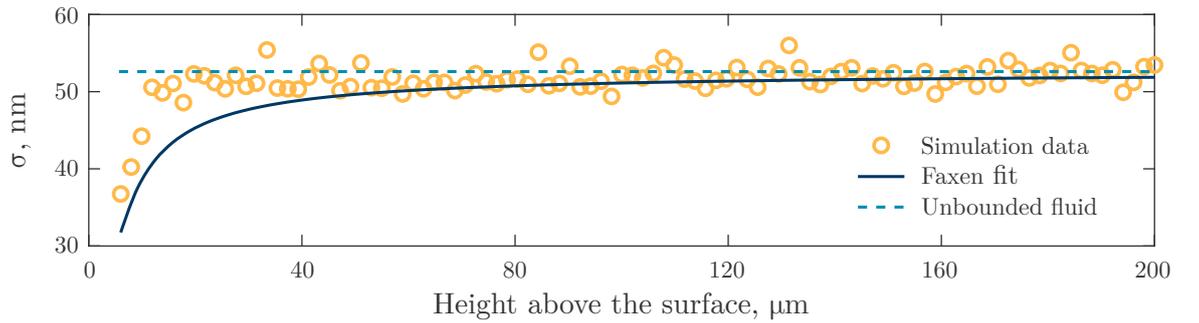


Figure 3.14: Influence of a surface on hydrodynamic clamping. Standard deviation σ of a $5\ \mu\text{m}$ radius target clamped at different heights h above a surface (yellow). The data is fitted with the Faxen formula with a fitting coefficient of 52.6, which is the value of σ in an unbounded fluid (estimated by averaging over a hundred 30s simulation sets, with standard deviation of 1.1 nm).

3.3.5 6D hydrodynamic manipulation

We are well aware that experimentally we will only be able to attempt 2D translational - and perhaps 1D rotational (if we have a target whose orientation about the z -axis can be easily tracked) - hydrodynamic manipulation. But lead by the spirit of curiosity, we explore the possibility of extending our framework to 6D control within simulations.

We propose using constellation rotors for orientation control in the manner illustrated in Figure 3.15(a). To control rotation of the target about the y -axis α_y , we orient a constellation rotor so that it spins around the y -axis - this arrangement only minimally affects the translation of the target. Placing another two rotors - one displaced along x and one directly below the target - completes orientation clamp setup. Before combining this with translation control we are able to achieve orientation clamping with an average standard deviation of $\sigma_\alpha = 8.4^\circ$ when $8\ \mu\text{m}$ radii constellation rotors, made up of three $5\ \mu\text{m}$ radii beads each, are placed $22\ \mu\text{m}$ away from a $2\ \mu\text{m}$ radius target.

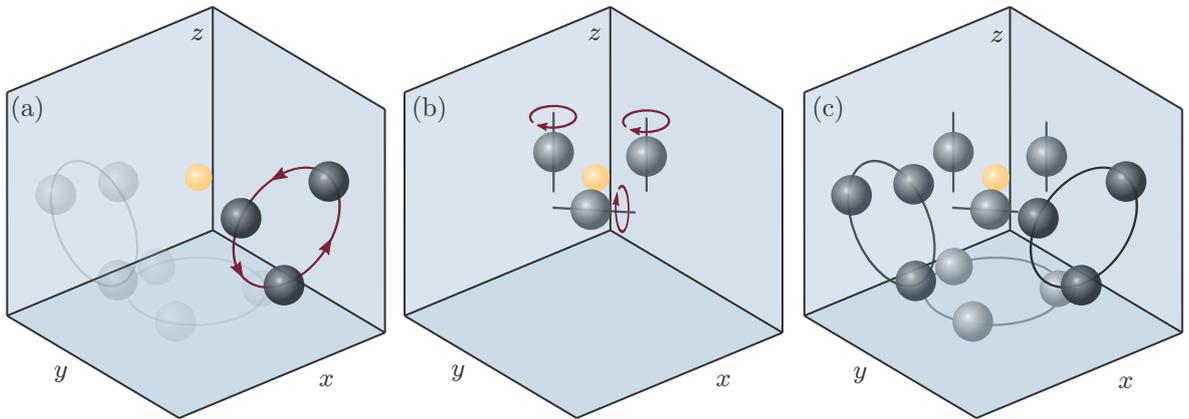


Figure 3.15: 6D hydrodynamic manipulation setup. Target's rotation about an axis can be controlled by a constellation rotor rotating about that axis (a). 3D translation control with spherical actuators (b). Complete 6D hydrodynamic manipulation setup (c) with three constellation rotors aimed at controlling orientation and three beads positioned for translation control.

For translation control, we complement the usual 2D arrangement by placing a third actuator that will spin about the $y = -x$ axis to address the z -motion of the target, as can be seen in Figure 3.15(b). Such a crowded environment, Figure 3.15(c), imposes several complications. The constellation rotors will have to be placed further away to make room for the spinning beads - this will reduce the impact they have on the target. To compensate for this we would need to increase the available rotation rates, which would, in turn, require higher stiffness optical traps. Furthermore, the fluid set in motion by the constellation rotors will inevitably affect the orientation of the spinning bead actuators, potentially acting against the torque applied onto them, thus interfering with translation control. This might be accounted for by modifying the feedback equations - currently the feedback takes into account the hydrodynamic interactions between the target and each actuator bead, but not between the actuators themselves.

Even in very optimistic conditions where we are able to spin the rotors at a rate of several revolutions per second, and ignore gravity, the results are underwhelming - either one type of clamping has to be sacrificed in favour of the other, or both can be balanced, but at the expense of increasing the standard deviation several times. Simulations show that an order of magnitude higher rotation rates are needed to achieve results similar to those of independent orientation and translation clamping.

In addition to the challenges mentioned above, an experiment would also face the difficulties of multidimensional tracking of the target's location and orientation, overcoming gravity, protecting the target from illumination in a complex scattered light landscape, and finally, somehow managing to ensure that scattered light does not interfere with

the optical trapping of the actuators. Considering all of this, we only attempt two-dimensional hydrodynamic manipulation in experiments, leaving the third dimension to an intrepid future researcher.

3.4 Our tweezers setup

Once we were confident in the viability and our understanding of the nuances of the proposed hydrodynamic manipulation technique, we went ahead with the experimental implementation of it. We had access to a functioning HOT setup (maintained by Graham Gibson of the Optics Group at the University of Glasgow) with accompanying Red Tweezers software [121] written in LabVIEW. Red Tweezers make use of a Graphics Processing Unit (GPU) to generate the holograms displayed on the SLM at very high speeds (down to 1 ms for a single trap hologram, while a similar algorithm running on a CPU would take from tens to hundreds of milliseconds [135, 136]). Each optical trap, in addition to having independent 3D locations, is characterised by a number of other customisable parameters, such as relative intensity, phase, vortex charge (for generating Laguerre-Gaussian beams), etc.

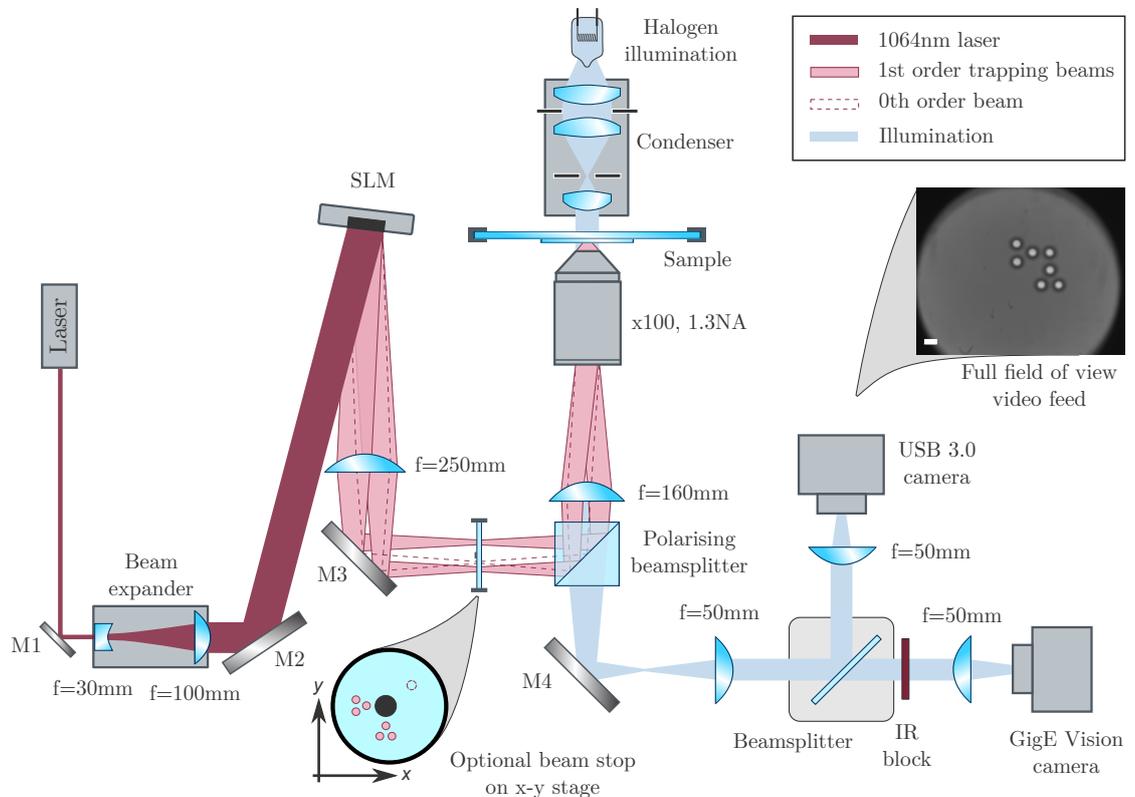


Figure 3.16: Our HOT setup. Adapted from Graham Gibson’s diagram in [134].

We presented a generic setup for optical tweezers back in Section 2.4.1. Now we show a detailed diagram of the particular HOT setup that we used in Figure 3.16. A 1064 nm laser beam is expanded and directed to slightly overfill a reflective SLM (Boulder Nonlinear Systems: XY-series, 512×512 resolution), where it is shaped into multiple traps. The trapping beams are then re-imaged with $4f$ relay optics onto the back aperture of a high NA oil-immersion objective lens (Nikon Plan Fluor: 100×, 1.3NA) which focuses the beams onto the sample, creating the traps.

The sample was imaged using a halogen illumination module (Zeiss Axiovert: 100W) fitted with a 0.55NA condenser. The illuminating light was collected using the same objective lens and then directed, via an IR filter to remove any residual laser light, onto camera 1 (GigE Vision, Teledyne DALSA Genie: HM1024) which was used for real-time high frame rate (several hundred frames per second) particle tracking.

The way in which the algorithms within Red Tweezers work required two minor modifications to the existing setup, to adapt it to our needs. First, the optical traps are generated using the ‘gratings and lenses’ [135] algorithm - it offers speedy hologram calculation, which is very appealing to our real-time feedback, but it suffers from ghost traps. These are additional unwanted traps created by the phase only modulation nature of the SLM. And even though Red Tweezers include methods for minimising light intensity of ghost spots [137], we found that in some trap configurations there is still too much light incident on the target particle, as illustrated by Figure 3.17(a). We wish to avoid this, as unwanted light might interfere with our hydrodynamic manipulation. To combat the issue we simply placed a beam-block in the laser path, which

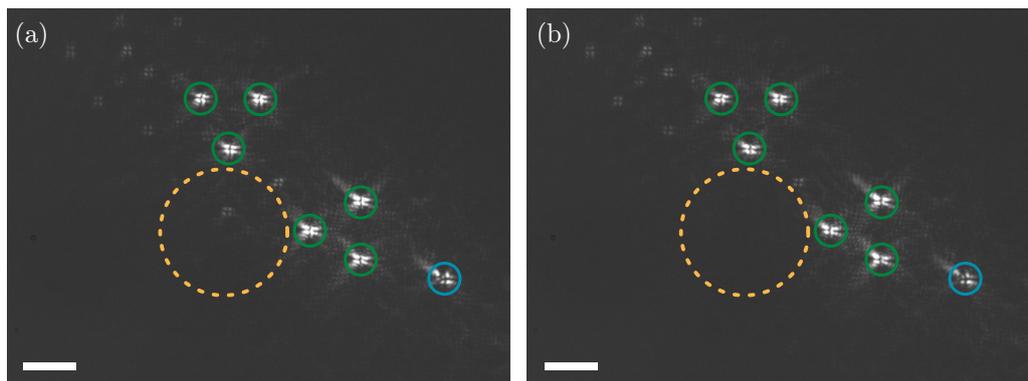


Figure 3.17: Ghost orders in optical tweezers. Optical trap setup for 2D hydrodynamic manipulation without (a) and with (b) the beam block present. The images show laser light reflected from a glass slide. Green circles indicate the traps which are meant to hold the rotors, and the yellow perimeter indicates the region where the target particle is to be present and where the mask blocks out the unwanted light. The blue circle indicates the original zero order beam. The scale bars are 10 μm in length.

obstructs the unwanted light from reaching the sample (see Figures 3.16 and 3.17(b)). The same approach could be used to mask the light outside of the functional actuator area, if the rest of the sample also needs to be protected from illumination. Second, to enable very fast particle tracking, the version of Red Tweezers that we used interrupts the live video feed during the tracking. We therefore direct some of the illuminating light from the main camera to a second USB camera for the purposes of observing the hydrodynamic manipulation in action and recording of videos.

With the minor issues cleared up, we next modified the Red Tweezers software to include our hydrodynamic feedback, as illustrated in Figure 3.18. Much like in the simulation, we use the latest available target's location (which can be tracked at several hundred frames per second) and optical trap locations to calculate the rotation rate of the actuators that would push the target towards a desired location. The actuators themselves are set into motion by assigning their optical traps with new locations. The new trap locations are then passed to the GPU engine, which calculates the hologram to be displayed on the SLM - trap locations are updated at a rate of 200 Hz. The rotors respond to the translation of optical traps and, via the movement of the surrounding fluid, propel the target. The feedback loop is then repeated for a given amount of time or until the experiment is manually stopped.

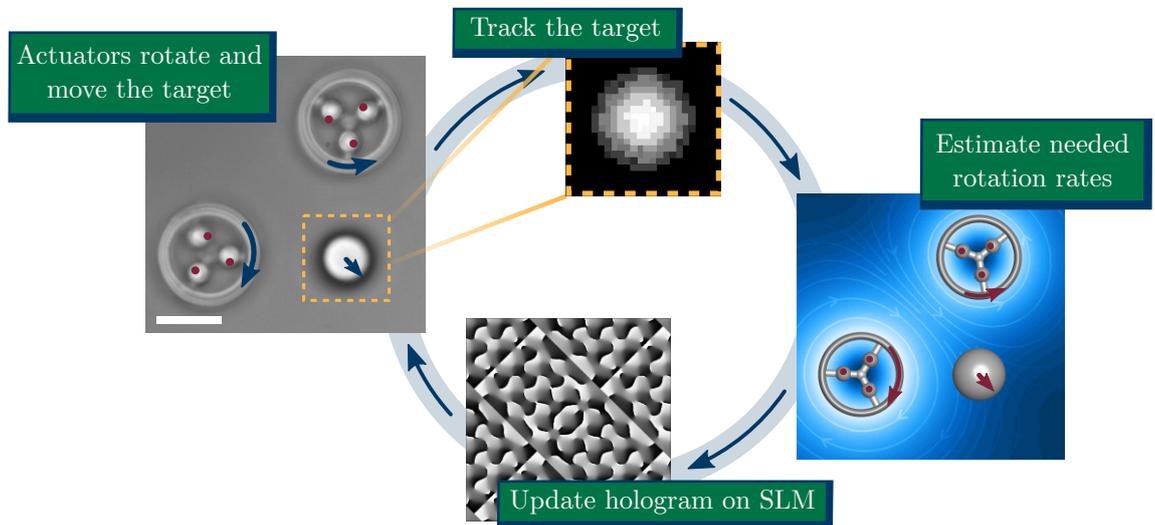


Figure 3.18: Feedback loop within an experimental setup. The left-most image shows a setup for 2D hydrodynamic manipulation of a silica bead using two micro-rotors (scale bar is 10 μm). Next, in the clockwise direction, is an illustration of a thresholded image of a target, used for tracking. The right-most image is a visualisation of the hydrodynamic interactions between the target and the rotors, which are used to determine the actuator rotation rates. Last is an example of a hologram displayed on the SLM.

3.4.1 Sample preparation

Our micro-rotors were manufactured at the University of Bristol using a 3D laser lithography system (Nanoscribe Photonic Professional). Direct laser writing relies on polymerisation - a photoresist material (Nanoscribe IP-L) interacts with an incident laser beam (wavelength of 780 nm was used) which causes it to solidify locally. This enables the ‘drawing’ of desired shapes with resolution of up to several hundred nanometres. Once an array of shapes is created, the remaining unpolymerised photoresist is dissolved away, leaving our micro-rotors stuck to a glass slip. The process of manufacturing 100 micro-rotors takes approximately 8 hours. We had the rotors made in three different sizes - 16, 22, and 30 μm in diameter. The height of the outer ring was 5 μm for all micro-rotors, with the width of the ring and the radii of the handle beads varying from 1.5 to 2 μm .

The micro-rotors then have to be transferred into a clean sample cell. We first place a droplet of 1 % water-Tween20 (SIGMA) solution on top of the array. Then, viewing the micro-rotors under a 5x magnification optical microscope we would manually scrape them off the surface using a copper wire, manipulated with an electronically controlled XYZ translation stage; we found that doing this by hand without the translation stage would push the rotors together too violently and would cause them to stick to each other. Using a pipette we would transfer the micro-rotor suspension into a prepared sample cell, of roughly 150 μm thickness. The rest of the sample would then be filled with an aqueous solution of target particles. We would seal the sample either with Norland Optical Adhesive 68 and 81 (Norland Products) or transparent nail varnish. Before being used in an experiment the sample cell would be stored on its edge for some time to allow the beads and micro-rotors to sediment in one location, where they could be easily accessed. We would then optically trap the particles we needed and drag them further away from the rest into a clear area to conduct the experiments.

For targets we would mostly use silica beads (Bangs Laboratories), which would also form the constellation rotors. The benefit of using silica beads as targets is that they are optically trappable. This made the setting up for the experiments fairly straightforward. We would trap the target silica bead and drag it to the starting location of an experiment, before deleting the trap and exposing the bead to the influence of hydrodynamic manipulation.

3.5 Experimental results

We have performed a variety of experiments testing and pushing the limits of our hydrodynamic manipulation technique. Most of the results presented here have been published in [134] and have accompanying videos, which can be found at the following [link](#)⁴. We believe these videos significantly ease the visualisation of our experiments, and will be referring to them throughout the text.

3.5.1 Hydrodynamic clamping

First, we demonstrate clamping, as discussed in Section 3.2, of a $5\ \mu\text{m}$ radius silica bead. We tried out three different micro-rotor sizes - $16\ \mu\text{m}$, $22\ \mu\text{m}$ and $30\ \mu\text{m}$ in diameter. Figure 3.19 presents the standard deviation σ of the clamped target for the different rotor sizes and for different rotation rate limits; as explained in Section 3.3.2 rotation rates need to be limited for two reasons: to ensure that the actuators do not fall out of the optical traps, and to not over-drive the target. We only reached rotation rates too high for the target with the $22\ \mu\text{m}$ actuators, though; with minimum σ at $100\ \text{° s}^{-1}$.

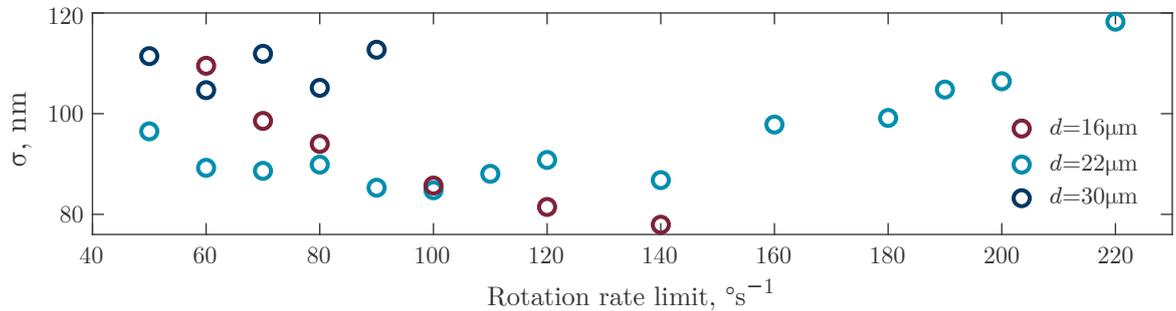


Figure 3.19: Experimental hydrodynamic clamping using differently sized micro-rotors. The graph shows the standard deviation σ of a clamped $5\ \mu\text{m}$ radius silica bead using micro-rotors of different diameters d and for different assigned actuator rotation rate limits.

Note, that both $16\ \mu\text{m}$ and $22\ \mu\text{m}$ rotors were placed $22\ \mu\text{m}$ away from the target bead, meaning that they had different edge-to-edge rotor-target separations. We believe that this distance is important because the fluid moves at the highest velocity right next to the edge of the spinning actuator, and then falls off with the square of the distance. Hence, the smaller rotors were generating lower fluid velocities at the target's location -

⁴ The figures in this text which are associated with a video will have this icon , and a direct link in the caption.

this explains why, at low rotation rates, the 22 μm rotors outperformed the 16 μm ones. Nevertheless, it was the smallest rotors that were, overall, able to clamp the best. This can be explained by the fact that they have a shorter optical relaxation time, i.e. they are quicker to respond to the motion of their optical traps - hence their actual rotation rates follow the prescribed rotation rates the closest. The very same reason explains why the large 30 μm rotors (placed 26 μm away from the target) showed a comparatively poor performance. Their motion was effectively low-pass filtered in comparison to the trap motion, rendering them less adequate in responding to the target's movement.

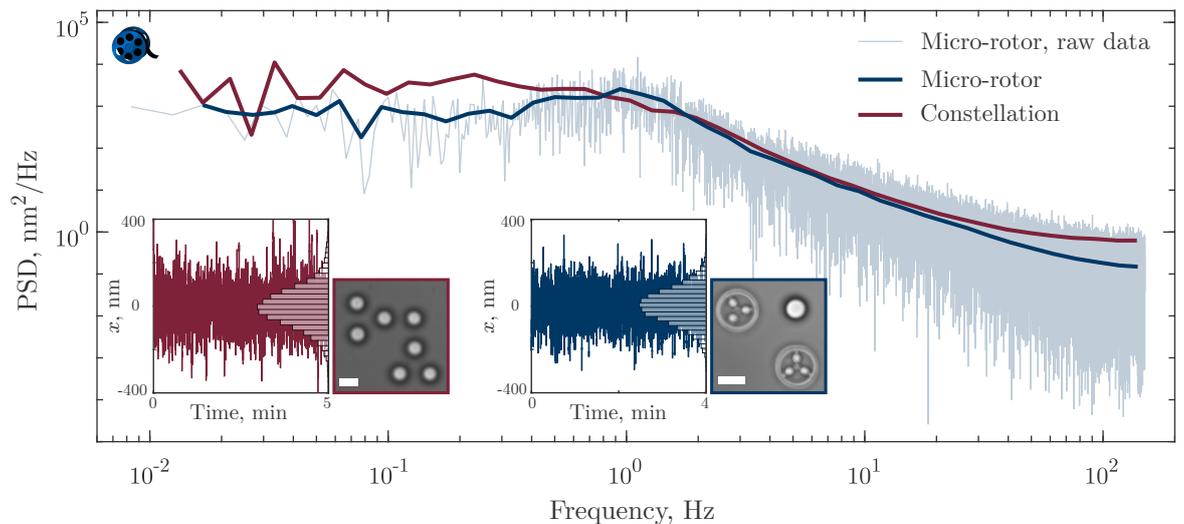


Figure 3.20: Comparison of micro- and constellation rotors. The main figure shows the PSD of a hydrodynamically clamped 5 μm radius silica bead for different actuator types - laser-printed micro-rotors (raw data in light blue, exponential average in blue) and three-bead constellations (exponential average in red, raw data not shown for visual clarity). The left inset shows the trace and histogram of the x -coordinate of the silica bead clamped with constellations for 5 min, in the setup shown on the right of the inset; standard deviation achieved in this experiment was $\sigma = 89$ nm (see [Video 5](#)). The right inset shows the same information but for micro-rotors, with $\sigma = 79$ nm (see [Video 1](#)). Both scale bars are 10 μm .

We now make our way to a more detailed analysis. Figure 3.20 presents the power spectral density (PSD) of a hydrodynamically clamped bead. In blue we see the data for the best clamping result using the 16 μm diameter micro-rotors and in red - when using constellation rotors formed with three 5 μm radius silica beads each, positioned on a ring of 8 μm radius (this is the distance from the rotor centre to the centre of each bead). We analyse the trace and PSD as one would for an optically trapped bead. The corner frequency of the PSD is proportional to the inverse of the clamping response time - this is the time it takes for the target to respond to the motion of the actuators. The corner frequency for the constellation rotors is not as well-defined, but it is slightly shifted to the left of the micro-rotor corner frequency - suggesting less successful clamping. Indeed, the standard deviation of the silica bead clamped with

constellations is $\sigma = 89$ nm compared to the micro-rotor result of $\sigma = 79$ nm (these results are also quoted in Table 3.2 for easy comparison with other targets and setups). We attribute this to the orientation dependent performance of the constellation rotors and the fact that the silica beads demonstrated longer optical relaxation times. We use the Equipartition Theorem, described in Section 2.4.1, to estimate the ‘hydrodynamic clamping stiffness’ to be $6.6 \times 10^{-7} \text{ N m}^{-1}$ and $4.7 \times 10^{-7} \text{ N m}^{-1}$ for the micro-rotors and constellations respectively. Both are comparable to a weak optical trap.

Hydrodynamic clamping gets progressively worse with decreasing target size - compare lines 2-4 in Table 3.2 - although the expected trend of $\sigma \propto a_t^{-1}$ (see Section 3.3.4) is not exactly followed. If we take the result of the $5 \mu\text{m}$ target as a reference, the standard deviation of the 2.5 and $1 \mu\text{m}$ beads should have been around 178 and 445 nm respectively - instead, particularly for the $1 \mu\text{m}$ target, we observe larger values. This can be explained by noting that smaller particles diffuse more out of the image plane of the camera, making video tracking less accurate, whereas in the simulations the exact location was always available.

Target	Radius, μm	Actuator	σ , nm
Si bead	5	2× 	79
Si bead	5	2× 	89
Si bead	2.5	2× 	186
Si bead	1	2× 	588
Chromium piece	1.5	2× 	482
Yeast cell	2	2× 	452
Two Si beads	5	5× 	120
Three Si beads	5	7× 	427

Table 3.2: Experimental hydrodynamic clamping results for different targets and actuators.

We confirm our system’s independence of target material by clamping an optically untrappable piece of chromium (GoodFellow) for 30 min with $\sigma = 482$ nm, see Figure 3.21(a,b). We believe this is worse than clamping of a spherical bead of a similar size because of the irregular shape of the piece of chromium, which reduces tracking accuracy. We also note that the feedback equations assume all particles are spheres, which might disadvantage asymmetric targets as their translational and rotational motions are coupled - we expect, though, this to be a smaller source of inaccuracy compared to tracking, computation delays and inability to apply the exact rotation rates on the actuators.

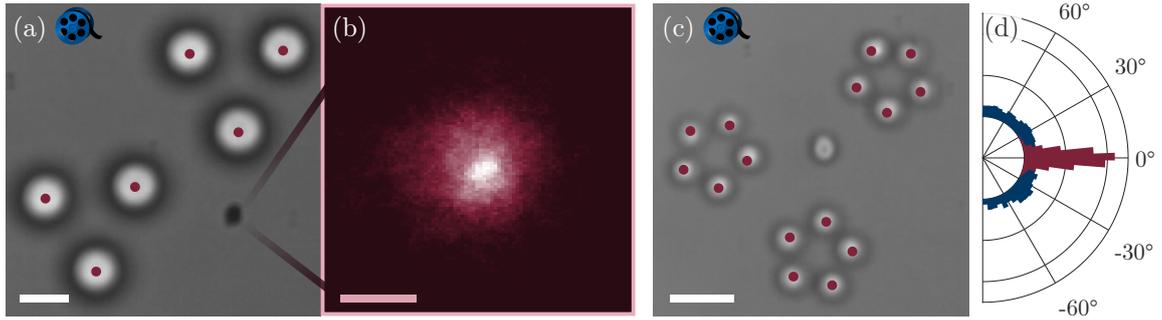


Figure 3.21: Hydrodynamic clamping of asymmetric objects. Experimental setup (a) and 2D occupancy histogram (b) of a clamped chromium piece, roughly $3\ \mu\text{m}$ across (see [Video 6](#)). Setup for clamping translation and in-plane orientation of a yeast cell (c) using constellation rotors formed with yeast (see [Video 8](#)). (d) Compares clamped yeast orientation (red) with orientation when only translation clamping is performed. White scale bars are 10 and pink is $1\ \mu\text{m}$.

To show compatibility of our platform with biological systems we clamp another asymmetric target - a yeast cell (we used fast action dried bakers yeast from Sainsbury's). This time each constellation rotor is formed with 5 'sacrificial' optically trapped yeast cells, removing the need to introduce foreign specimens into the sample. With a roughly $2\ \mu\text{m}$ radius yeast cell we achieved standard deviation of $452\ \text{nm}$, comparable to that for chromium. The prolate shape of the yeast cell inspired us to attempt orientation clamping. We do this simultaneously with position clamping by introducing another five-yeast rotor, shown in Figure 3.21(c,d). We achieve a 7.5° standard deviation of the in-plane orientation of the yeast over the period of 10 min. And we do this without compromising translation clamping with $\sigma = 385\ \text{nm}$. Orientation tracking was performed using the inbuilt LabVIEW function 'IMAQ Particle Analysis'.

3.5.2 Hydrodynamic micro-manipulation

So far we have only demonstrated suppression of the Brownian motion of an otherwise freely diffusing object. But to truly perform hydrodynamic manipulation we need to be able to translate the target particle along a prescribed trajectory. We first demonstrate this by translating a yeast cell along a circular path, shown in Figure 3.22(a,b). To achieve translation we simply update the desired location coordinates once the target has approached the current desired location to within some specified distance. We identify two strategies for prescribing trajectories. To impose small-scale motions of high precision onto the target, like the circle example above, we define the path with many closely spaced points and only switch to the next point if the target is close, e.g. a $1/4$ of its radius, to the current one. If we simply wish to transport the target from

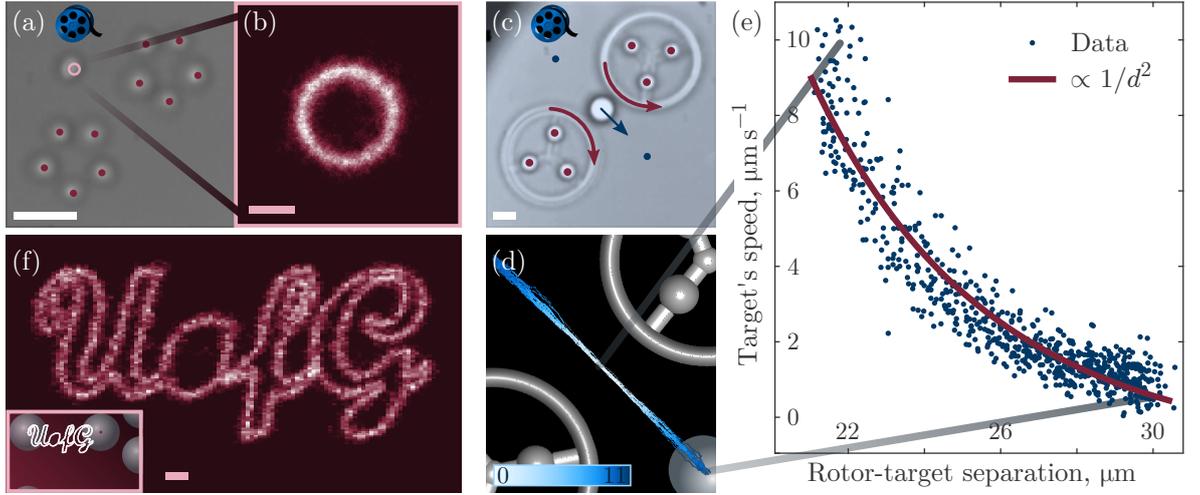


Figure 3.22: Hydrodynamic translation techniques: locally moving the target. A yeast cell (a) moved along a circle (indicated in pink) using two five-yeast constellation rotors with the corresponding occupancy histogram (b) (see Video 7). A 5 μm silica bead (c) transported in between two points (blue dots) by two micro-rotors (see Video 2). (d) Shows the speed of the target (colour) along its trace; the colour-bar indicates the speed in $\mu\text{m s}^{-1}$. (e) Shows the speed data plotted against rotor-target separation d , along with a theoretical fit. 2D position histogram of a complex trajectory of the University of Glasgow logo traced out by a bead (f); the inset shows the trajectory in comparison to the bead sizes. White scale bars are 10 μm and pink - 1 μm .

one location to another as quickly as possible, and are not particularly interested in the exact shape of the trajectory it takes, it is sufficient to define only a few points along the way. For example, to investigate the speed that a target can achieve, we assign the feedback to repeatedly move it in between two points, as illustrated in Figure 3.22(c). We extract the target's speed from the data and plot it in Figure 3.22(d). Unsurprisingly, the target is at its fastest when it is directly in between the two rotors - and closest to them. We analyse the speed data in more detail and note that for the case of a single target bead located a distance d away from a spinning micro-rotor, exposed to torque T_r , target's velocity will be given by:

$$v_t = -\mu_{tr}^{\text{TR}} T_r.$$

To obtain an accurate fit we use the expansion for μ_{tr}^{TR} in Equation 2.14, which takes into account not only d but also the the distance away from the bottom of the sample h . Substituting this into the above and simplifying suggests the following trend:

$$v_t \propto \frac{c_1}{d^2} - \frac{c_2}{d^2 + 4h^2}, \quad (3.13)$$

where we have introduced scaling factors c_1 and c_2 . We fit this equation to the experimental data in Figure 3.22(e) with $c_1 = 37\,440\ \mu\text{m}^2\ \text{s}^{-1}$ and $c_2 = 41\,030\ \mu\text{m}^2\ \text{s}^{-1}$. To obtain the fit we assume the target to always be equidistant from both actuators and the distance from the bottom of the sample h to be equal to the radius of the target particle, $5\ \mu\text{m}$. We also note that the equations used dictate that the velocity of a target particle being translated by a spinning actuator is nominally independent of the target's size.

Unsurprisingly, such 'local' translations, where the target is still close to the actuators, work very well. But they are restricted to a very small working region as we can only position the actuators so far away before they lose their influence. What if we wish to move the target over longer distances? From the above analysis we conclude that to maximise the actuator influence, it is beneficial for the actuators to follow the target as it is being translated. Or, rather, we make the actuators follow the current desired location, so as not to jiggle them continuously with the Brownian motion of the target. As an example of a long trajectory which still requires high accuracy, we task the feedback to push a $5\ \mu\text{m}$ silica bead along a $17\times 8\ \mu\text{m}$ University of Glasgow logo outline, as seen in Figure 3.22(f). This trajectory was traced out in about an hour, during which time the constellation rotors followed the target by keeping a fixed $22\ \mu\text{m}$ distance.

By fixing the target-rotor separation throughout an experiment we are free to translate particles along any trajectories and to any location within the camera's field-of-view. But sometimes one might wish to translate an object across the entire sample. We do this by returning to the idea of clamping. We fix the target's desired location with respect to the camera's frame of reference and translate the stage holding the sample cell. Unlike in the clamping strategy with two rotors demonstrated earlier, to keep the target stationary we now use a set of three reconfigurable micro-rotors. Since the aim now is to transport the target from one part of the sample to another, and we are not particularly concerned about how accurately we control the target's location, we use the large $30\ \mu\text{m}$ diameter micro-rotors because they will create the strongest current within the fluid. To maximise the target's speed, we use two 'driving' micro-rotors positioned on opposite sides of the target, shown in Figure 3.23(a), spinning in opposite directions to push the target against the bulk of fluid which is moving together with the translating stage. To make sure that the target doesn't drift into either of these actuators, we place a third 'regulating' micro-rotor to be orthogonal to the other two.

Since the direction of travel might change, we program the actuators to reposition

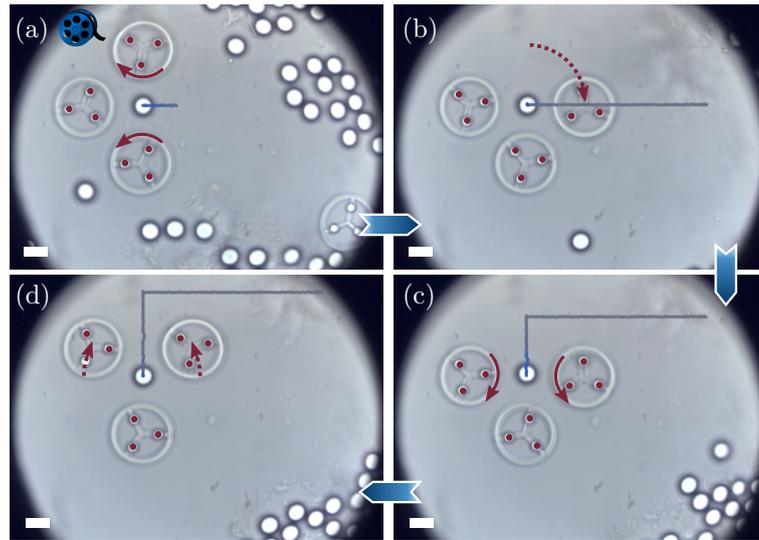


Figure 3.23: Hydrodynamic translation techniques: moving the sample stage. The target can be transported across a sample cell using a reconfigurable three rotor setup (see Video 4). When the stage is in motion (a,c) two of the rotors are driving the target to move it against the bulk of fluid, while the third rotor regulates the position of the target to prevent it from crashing into the actuators. If direction of translation needs to be changed (b) the stage is stopped and the rotors are rearranged depending on the new direction. When translation is complete (d) the rotors are distributed around the target for hydrodynamic clamping. Trajectory of the target through the sample is shown as a trailing blue-grey line. Scale bars are $10\ \mu\text{m}$.

themselves, as in Figure 3.23(b), so that the line connecting the centres of the driving rotors is perpendicular to the line of travel, and the regulating rotor is either behind or in front of the target on the line of travel. While the rotors are moving along a circle to reposition themselves, the sample is kept stationary. Once reconfiguration is complete, the stage is set into motion in the new direction, Figure 3.23(c). We are able to translate the sample at the speed of $1\ \mu\text{m s}^{-1}$ without losing the target bead for an extended period of time. If we wish to stop the translation and simply perform clamping, the rotors rearrange themselves to surround the target as in Figure 3.23(d). Note that in this setup we have more actuators than the target degrees-of-freedom and we use techniques described in Section 3.1.2 to deal with the mathematics involved. Unlike other forms of flow control which operate only within predefined regions of the sample [41], this approach gives us access to any areas of the sample that can be reached by translating the stage in 2D.

We draw the reader's attention to the fact that with this technique we were not able to achieve target translation speeds as high as the $10\ \mu\text{m s}^{-1}$ suggested by the previous stationary sample experiment in Figure 3.22(c-e). First of all, here we used smaller micro-rotors ($26\ \mu\text{m}$ diameter instead of $30\ \mu\text{m}$) which were positioned further away

from the target ($25\ \mu\text{m}$ instead of $21.2\ \mu\text{m}$) in order to accommodate three actuators. We note that the minimal distance at which we could position the actuators was also limited by the use of the beam block - the target had to be shielded from the influence of ghost traps which created an exclusion region where micro-rotors simply could not be optically trapped. On the other hand, we were able to spin these smaller rotors at higher rotation rates ($150^\circ\ \text{s}^{-1}$ compared to $90^\circ\ \text{s}^{-1}$). Considering the listed factors and using the data in Figure 3.22(e) as a reference, we anticipated achievable translation speed to be $\sim 7\ \mu\text{m}\ \text{s}^{-1}$. However, as already stated, the speed of the stage had to be limited to $1\ \mu\text{m}\ \text{s}^{-1}$ in order to successfully transport the target over large distances. This was because, even with the use of the beam block, the target was surrounded by an unfavourable light field environment, where even a small deviation towards the edge of the area shielded by the block was enough to destabilise the target in the z -direction, posing a risk of losing it altogether.

3.5.3 Simultaneous control of multiple particles

Encouraged by the success of our experiments, we next face the challenge of simultaneously hydrodynamically manipulating multiple particles. As explained in Section 3.1.2 we now need more rotors than the target degrees-of-freedom and therefore we use five micro-rotors to control 2D translation of two targets, and seven micro-rotors for three targets, in setups shown in Figure 3.24(a,g). First, we simultaneously clamp freely diffusing $5\ \mu\text{m}$ radius silica beads. Compared to the clamping of a single bead (see Table 3.2), the outcome is unsurprisingly poorer, with $\sigma = 120\ \text{nm}$ and $\sigma = 427\ \text{nm}$ for two and three target beads respectively - for several reasons. To distribute more actuators around the targets we had to increase the target-rotor separation from $22\ \mu\text{m}$ to $27\ \mu\text{m}$. For comparison, a single bead clamped $26\ \mu\text{m}$ away from the micro-rotors achieved $\sigma = 105\ \text{nm}$, indicating that, at least for the two target case, the performance of hydrodynamic clamping is not substantially reduced with the increase in degrees-of-freedom that need to be addressed (although we note that this comes at the expense of increasing the laser power, in order to trap more rotors). Furthermore, we have used smaller, $16\ \mu\text{m}$ instead of $22\ \mu\text{m}$ diameter, micro-rotors for the three targets, thus reducing the strength of the hydrodynamic interaction between the silica beads and the actuators.

Next, we were able to orchestrate different kinds of motions of the two targets, illustrated in Figure 3.24: the two beads moving anti-clockwise on the same ring (a), being pushed together (b) and pulled apart (c), following concentric trajectories in opposite

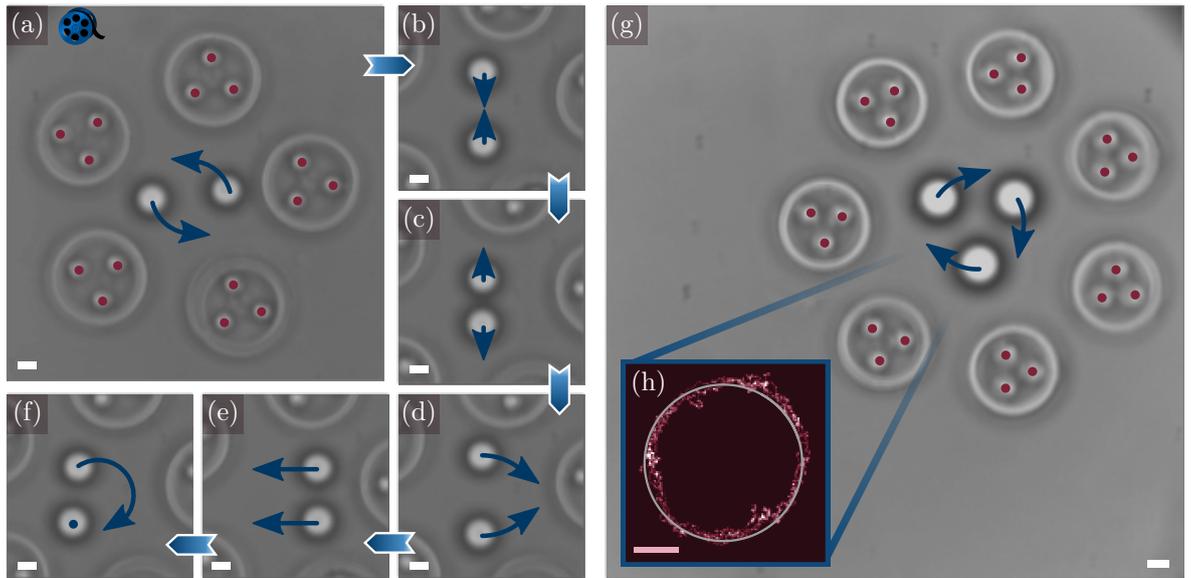


Figure 3.24: Multiple target hydrodynamic manipulation. Different motions (a-f) in which we can translate two target beads using five micro-rotors (see [Video 9](#)). Three target beads (g) set to follow a circular trajectory, indicated in grey in (h) and their occupancy histogram (data for all beads shown together). All scale bars are 10 μm .

directions (d), moving the centre of mass with fixed relative particle separation (e), and finally, holding one target stationary while the other one is tracing out an arc (f).

Moving up from two to three target manipulation has proven to be a much larger step in complexity than going from one to two. We observed that the three particles tended to huddle together, and it was nearly impossible to pull them apart if they found themselves clustered in the centre of the rotor ring. This made the preparation of the setup, before hydrodynamic manipulation could even be activated, rather tedious and requiring some dexterity. We were only able to move the three beads in a circular trajectory, with fairly low precision, as shown in [Figure 3.24\(h\)](#).

3.5.4 Limitations

We have presented a wealth of successful proof-of-concept experiments that demonstrate our hydrodynamic manipulation technique. So what are the limitations of our method?

The most obvious limitation of the effectiveness of hydrodynamic clamping comes from the feedback delay time - the time between when the target's location is registered and when the optical trap locations are updated. During this time the target diffuses away from the registered position, so the actuator rotation rates determined by the feedback

are not quite accurate. This delay includes the computational time required to estimate the \mathbf{C} -matrix, solve Equation 3.1 for new trap locations, generate a hologram based on these trap locations and send it to the SLM, as well as the time for the liquid crystals on the SLM to change their orientation. We estimated the delay time to be 10-15ms.

Optical trapping stiffness is a major factor in determining the tightness of hydrodynamic clamping. As discussed in Section 3.3.3 low optical stiffness results in a slow actuator response to the trap motion - rotor dynamics is effectively suppressed in comparison to the trap dynamics, as shown in Figure 3.25(a). This results in the decrease of the corner frequency of the clamped target's motion. Our laser-printed micro-rotors were designed to have small prolate-shaped handles instead of spherical ones, which, by maximising the overlap between the handle and the elongated shape of the Gaussian beam focus along the axial direction, improves their optical trapping stiffness. Further optimisation of actuator interaction with the trapping light could be achieved by re-designing the shape of the trapping light beam - Part II of this thesis will concentrate on that. Potentially it would also be possible to account for the imperfect rotor response if it could be predicted in advance, and thus corrected for.

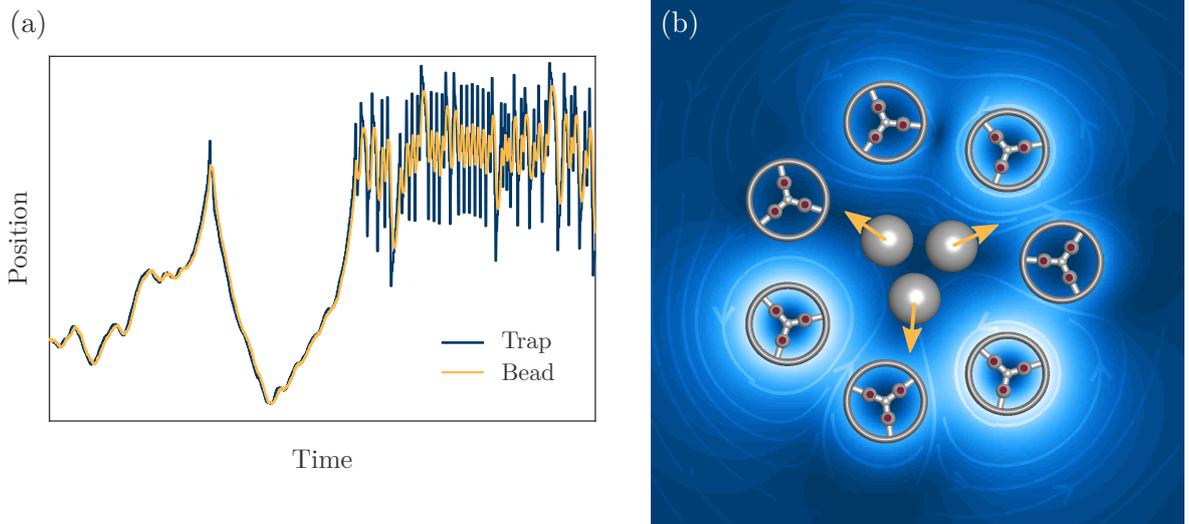


Figure 3.25: Main limitations of hydrodynamic micro-manipulation. In (a) we see that an optically trapped bead follows the trajectory of a slow-moving trap very closely, but cannot keep up with high frequency motions (the trajectories shown were simulated). In (b) we see a flow field with a stagnation point close to the target beads, which limits the speed that the fluid can achieve at the target locations to pull them apart.

We saw that simultaneously manipulating several targets leads to reduced performance. First of all, more targets require more actuators, which pushes the limits of the space available for them - to accommodate more actuators they will have to be positioned further away from the targets, reducing their hydrodynamic influence. Furthermore,

increasing the number of actuators means that each gets less laser power, cutting the achievable rotation rates. Less evident but more fundamental is another reason - certain motions of multiple target particles require the flow to contain stagnation points, i.e. points where the fluid is stationary. For example, to pull three beads apart there has to be a stagnation point in between them, as seen in Figure 3.25(b). This places the beads close to a zero-velocity point which inevitably limits the fluid velocity that can be achieved at the bead locations.

3.6 Discussion

3.6.1 Wider context

As stated in Section 3.5.1 our hydrodynamic clamping, by limiting the motion of a $5\ \mu\text{m}$ radius target particle with an effective stiffness of $\sim 10^{-7}\ \text{N m}^{-1}$ (standard deviation of $\sim 80\ \text{nm}$), is similar to a weak optical trap (the reader may refer back to Table 3.2 for results for different target sizes). But how does it compare to other micro-manipulation techniques?

Microfluidic hydrodynamic tweezers have been shown to suppress the Brownian motion of micro-particles down to standard deviations of $\sim 200\ \text{nm}$ (for a $0.5\ \mu\text{m}$ diameter particle) in a passive stagnation-point trap [17], and $\sim 80\ \text{nm}$ (for a $2.2\ \mu\text{m}$ diameter particle) in an active feedback Stokes trap [43] (we note that in the latter the particles were submerged in a fluid with viscosity ~ 10 times larger than that of water; assuming trapping tightness to be inversely proportional to diffusion, this result would scale up to a few hundreds of nanometers in our system). While microfluidic tweezers can isolate a single particle from a crowded environment, they unavoidably affect all the particles in the sample. Similarly, acoustic tweezers, while good for trapping live specimens (as they do not require high intensities), do not offer the ease of locally handling individual particles. Our reconfigurable hydrodynamic actuators, on the other hand, influence only the near-field and act equally well in all regions of the sample cell.

Magnetic tweezers work exclusively on magnetised particles, and do not have the material-independent versatility of hydrodynamic techniques. Nonetheless, they can be used to manipulate biomolecules by attaching them to magnetic beads [5, 138]. Furthermore, impressive rotation rates have been achieved with magnetic tweezers - beads of $4\ \mu\text{m}$ diameter can be spun at $110\ \text{Hz}$ [139], and therefore achieve very high fluid flow speeds. This feature was used to construct a single bead hydrodynamic ‘trap’, where a swimming bacterium would be sucked into the vortex flow and be confined

to stay at a certain distance away from the spinning actuator bead, whilst continuously revolving around it. The target bacterium could this way be transported across the sample. However, magnetic tweezers do not allow independent control of multiple particles, e.g. spinning several beads at different rates would not be possible, and are not suited for active trapping schemes.

Optoelectronic tweezers offer dielectrophoresis based control, and can reach trapping stiffness of $\sim 10^{-7} \text{ N m}^{-1}$ for particles a few microns in diameter [140] - same order of magnitude as hydrodynamic clamping we demonstrated. Optoelectronic tweezers are, however, capable of exerting control over significantly larger particles (of up to several hundred microns across [13]), which is not achievable with optical tweezers - and larger actuators would allow control of larger targets. It might therefore be tempting to transfer our hydrodynamic manipulation platform over to optoelectronic tweezers. We have, in fact, attempted this, but found that dielectrophoretic forces in optoelectronic tweezers are very long-ranged. This means that actuator traps inevitably affect the target particle, thus not allowing purely hydrodynamic interactions.

3.6.2 Summary

We began this chapter by setting out to design an optically actuated hydrodynamic control system. Having identified two actuator types, we developed a mathematical model which, capturing the hydrodynamic interactions in the system, provides a framework for predictably controlling the translational and rotational motion of a target particle. Backed up by information gained from computational simulations, we have implemented our technique into an optical tweezers platform in a fashion of closed-loop control, and have successfully demonstrated the viability of our method in multiple experiments.

By sculpting fluid flow patterns over tiny areas of tens of micro-metres, our approach offers near-range control which can be performed anywhere in the sample without significantly disrupting the surroundings. We have hydrodynamically clamped micro-sized objects by suppressing their Brownian motion to standard deviations of up to 79 nm, as well as exerted control over in-plane rotation of a yeast cell to within 7.5° . We have also transported particles over long-range distances and short-range intricate trajectories. Our scheme is independent of the material of the target particles and has a two-fold benefit in being used with biological systems - it minimises the light exposure by providing a non-invasive manipulation method and, depending on the system, can

circumvent the need of introducing foreign material into the sample by using ‘sacrificial’ cells as actuators. With holographic optical tweezers available commercially, our hydrodynamic manipulation technique is accessible to a wide scientific audience.

Our method is not without limitations. While its theoretical framework can be extended to an arbitrary number of controllable degrees-of-freedom, in practice this would face a number of challenges. From the need of a high speed tracking system capable of determining the target’s location in 3D as well as its orientation, to the complexity of fluid flow fields that need to be generated to allow simultaneous control of multiple particles. As far as hydrodynamic manipulation of sedimented particles is concerned, we have identified low optical trapping stiffness as one of the main limitations in the experiments that we have performed. This provides the motivation for Part II of this thesis - optimising the trapping light field with the aim of enhancing optical trapping stiffness.

PART II

Chapter 4

More background theory

‘Finally, from so little sleeping and so much reading, his brain dried up and he went completely out of his mind.’

Don Quixote — Miguel de Cervantes Saavedra

Motivated by one of the main limitations of our hydrodynamic manipulation method - the response time of the actuators to their optical traps - in Part II we aim to develop an approach that would enhance the trapping stiffness of optical tweezers. We are not the first ones to walk down this path - iterative optimisation approaches have been successfully employed to increase one-dimensional trapping stiffness [141, 142]. Some of these techniques, though, are limited by their long computation times (on the order of days), and they all are inherently unable to identify if the solution they converge to is the global, or just a local optimum. The optical eigenmodes (OEs) method surpassed the need for iterative optimisation by relating the far and near fields of the trapping light beam, to *globally* optimise optical force, size of a focused spot and transmission

through sub-wavelength apertures [53]. More relevant to us is that OEs have also been used to find fields with globally optimal one-dimensional stiffness [142]. The major drawback of the above accomplishments, though, is that they do not necessarily provide stable trapping in three-dimensions. In particular, there has been no attempt to ensure that the axial force exerted by the optimal fields does not push the trapped particle out of the transverse plane of the trap, and instead gravity is relied on to counter the repulsive optical force. The same OEs technique has been recently used to take the first steps towards indirectly enhancing stiffness in a 3D-stable optical trap by optimising the optical force in cylindrically symmetric light fields [143].

Of late, a new technique has found its way into the optical manipulation world from that of nuclear scattering [144–146] - the method of the generalised Wigner-Smith (GWS) operator. While equivalent to the OEs in terms of the results that it can produce [147], this method does not require the knowledge of the near field around the particle interacting with the light, which is often not available in disordered media. Instead, it makes use of the scattering information in the far field and its dependence on small variations in some property of the scatterer, e.g. its position, size or refractive index [147]. This is particularly convenient since we do have experimental implementations in mind, although not within the scope of this thesis. The GWS is one of the techniques we will be employing in simulations to generate light fields that optimally enhance optical trapping stiffness in 1D, while ensuring that the trap is still stable in 3D, or, in fact, fields which enhance optical trapping in all three dimensions; we describe this method in Section 4.1. To investigate the capabilities of the GWS operators in simulations we will need a much more rigorous framework for modelling light-matter interaction than the linear restoring force approximation we used in Part I. For this purpose we will be utilising the well-known generalised Lorenz-Mie theory, which we cover in Section 4.2.

4.1 The generalised Wigner-Smith operator

Originally introduced as a time-delay operator by Wigner [144] and Smith [145] and then generalised [146], the GWS operator is defined as follows:

$$Q_\alpha = -iS^{-1} \frac{dS}{d\alpha}, \quad (4.1)$$

where α is an arbitrary parameter characterising a system of interest (e.g. frequency of incident light, position or size of a scattering particle), and S is the scattering matrix

coupling the light fields incident on, $|\text{in}\rangle$, and exiting, $|\text{out}\rangle$, the system: $|\text{out}\rangle = S|\text{in}\rangle$. Here $|\text{in}\rangle$ and $|\text{out}\rangle$ are column vectors containing the complex electric field amplitude.

The exciting feature about the GWS operator is that its eigenstates describe light fields which, after interacting with the system, have outgoing fields invariant under small changes of α . The eigenvalues λ_α corresponding to these eigenstates are directly related to the conjugate quantity of α [146, 148]. For example, in the original Wigner-Smith operator with $\alpha = \omega$, where ω is the frequency of the light field, the eigenstates of Q_ω , which are also known as the principal modes, then have eigenvalues λ_ω which are related to the time delays that the principal modes experience between entering and exiting the system [149, 150]. Other examples of α and its conjugates include position and linear momentum, size and pressure, orientation and angular momentum, etc. α does not need to be a global parameter; in a system of multiple scattering particles it can describe some property of just one of them. Provided that S is a unitary matrix, the relationship between the real part of λ_α and the conjugate quantity of α is perfectly linear [147].

The principal modes of Q_ω have been demonstrated to exist in multimode fibers [151–153] and have been suggested to have applications in optical communication, imaging, and sensing, as well as quantum computing. What is more, the eigenstates of Q_α have been shown to be equivalent to the earlier mentioned optical eigenmodes [53, 147], and are therefore optimal states. The GWS method has been used to optimise the incident light fields to exert maximum torque or force, and create the most efficient focus within a light scattering system [147, 149].

We wish to use the GWS approach to create an optical trap for a single spherical particle with optimal trapping stiffness. We begin with a GWS operator that will provide us with information about the optical force. Since force is directly proportional to the rate of change of momentum, the α we need is the position of the particle, say x : $Q_x = -iS^{-1}\frac{dS}{dx}$. The eigenstates of Q_x are related to momentum, and therefore the force exerted on the particle. Now, stiffness is the gradient of the force, so an operator K_x can be defined as [147]:

$$K_x = \frac{dQ_x}{dx}, \quad (4.2)$$

which will have eigenvalues proportional to the optical trapping stiffness κ_x . We have explicitly presented the case for x , but the description is analogous for other directions.

Before discussing how the Q and K operators are to be estimated, it is beneficial to understand the construction of the scattering matrix. To this end we briefly discuss the idea of representing a light field by decomposing it into some basis. A visually clear

example of this is a multimode waveguide setup described in [146]. At one side of the waveguide are ten transmitting antennas; the field emitted from each constitutes one input basis mode. The fields generated by the antennas are linearly independent so that any field entering the waveguide can be represented as a weighted sum of the ten input modes: $|\text{in}\rangle = \sum_{i=1}^{10} a_i E_i^{\text{in}}$, where E_i^{in} is the complex amplitude of the i -th mode and a_i are the coefficients telling us how much of each mode is activated. Inside the waveguide there are particles which scatter the light, turning it from $|\text{in}\rangle$ to $|\text{out}\rangle$; this effect is then measured with another antenna at the other end of the waveguide. Ten measurements are done at separate locations, each of which represents an output basis mode. A 10×10 scattering matrix can thus be constructed by activating one input mode at a time and estimating the field at the ten output modes. Note that technically such a matrix contains only the transmission and not the full scattering information, since the output light is measured only at one end of the waveguide - but this example is here just to aid the understanding of the general idea of basis representation.

We, however, will not be working with waveguides but with optical tweezers, where the light field incident on the particle is emerging from a high NA lens. And while in our framework the input and output bases will be mathematically more complicated (which we will discuss in detail in this and the following chapters), the concept and principles remain the same. We now proceed to describe how the inverse and the differential of the scattering matrix S are calculated.

The inverse

Most preferably, the scattering matrix will be a unitary matrix, in which case its inverse is equal to its complex conjugate transpose (denoted with \dagger). A unitary S is not always available, though. In fact, the scattering matrix that we will be working with is not even a square matrix. Generally, our S will be of dimensions $N_{\text{out}} \times N_{\text{in}}$, where N_{in} is the number of modes in the input basis and N_{out} is the number of modes in the output basis. We will therefore be estimating an *effective* inverse, which we still denote as S^{-1} .

Computation of the effective inverse [146, 147] follows a protocol very similar to that of computing a pseudo-inverse, also known as the Moore–Penrose inverse [154, 155]. First, singular value decomposition (SVD) is performed on the matrix: $S = U\Sigma V^\dagger$, where Σ is a diagonal (not necessarily square) matrix with non-negative real entries, representing the singular values of S , and U and V are unitary matrices the columns of which contain the left-singular and right-singular vectors of S respectively. We perform the SVD with a MATLAB function ‘svd’ with an ‘econ’ option, which automatically removes zero rows

and columns from Σ and U respectively, making Σ a square $N_{\text{in}} \times N_{\text{in}}$ matrix. Next, to eliminate low-information channels, dimensionality of Σ is reduced by removing several of the smallest singular values, giving us $\tilde{\Sigma}$. Removing the corresponding columns from U and rows from V gives us \tilde{U} and \tilde{V} . The effective inverse is then defined as

$$S^{-1} = \tilde{V}\tilde{\Sigma}^{-1}\tilde{U}^\dagger. \quad (4.3)$$

It is not immediately clear how many singular values we should keep when reducing dimensionality, and it requires some trial and error to figure out what gives the best results. By evaluating the linearity of the force/stiffness vs. real eigenvalue trend for different numbers of removed singular values, we choose the one with the most linear trend (see Section 5.6.1 for more detail). Generally, we observed that if S is a full-rank matrix, it is enough to remove just a few (~ 10 out of 100-200) singular values, and this number grows as S becomes rank deficient.

The differential

To compute the operator Q_x we need to know the differential of the scattering matrix $\frac{dS}{dx}$. We evaluate this using the midpoint method. Two scattering matrices are estimated when the scattering particle is placed at different locations, $S(-\Delta x)$ and $S(\Delta x)$ (while S^{-1} is estimated at the midpoint of the two), then

$$\frac{dS}{dx} \approx \frac{S(\Delta x) - S(-\Delta x)}{2\Delta x}, \quad (4.4)$$

where Δx is a small distance by which the bead is displaced along the x -axis. Care must be taken to select the right size for Δx . Too big of a step will fail to capture the local gradient of the scattering matrix; if the step is too small the above differential might suffer from numerical errors. We find that a step size of $\Delta x = 0.01\lambda$, where λ is the wavelength of the incident light, is a good middle ground.

We further need to project this differential onto the subspace in which we estimated the effective inverse [146]. The projection is done using the operators obtained during the singular value decomposition: $\tilde{U}\tilde{U}^\dagger \frac{dS}{dx} \tilde{V}\tilde{V}^\dagger$; the Q_x operator can thus be estimated.

The stiffness operator

To obtain the stiffness operator K_x we need the differential of Q_x . Again, we use the same scheme as in Equation 4.4:

$$K_x \approx \frac{Q_x(\Delta x) - Q_x(-\Delta x)}{2\Delta x}. \quad (4.5)$$

For this we need a total of six scattering matrices - three for each Q_x : $S\left(-\frac{3}{2}\Delta x\right)$, $S(-\Delta x)$, $S\left(-\frac{1}{2}\Delta x\right)$, $S\left(\frac{1}{2}\Delta x\right)$, $S(\Delta x)$ and $S\left(\frac{3}{2}\Delta x\right)$.

The last thing remaining is the estimation of the scattering matrix itself. For this, we need to look at how the light scattering process can be modelled mathematically.

4.2 Generalised Lorenz-Mie theory

In Chapter 2 we have mentioned two frameworks within which light-matter interaction can be discussed. These are the dipole approximation for particles significantly smaller than the wavelength of light, and the ray optics for appreciably larger particles. Unfortunately, the two theories do not meet somewhere in the middle - they both fail when it comes to particles whose size is comparable to that of the wavelength of light. Reigning in this regime (although applicable to particles of any size) we have the generalised Lorenz-Mie theory (GLMT), first developed independently by Lorenz [156] and Mie [157] for plane waves, and later generalised to arbitrary light fields as presented by Stratton [158]. GLMT is a popular choice when it comes to modelling optical tweezing experiments.

Within the GLMT an arbitrary electromagnetic field is decomposed into some basis. Since this theory describes spherical particles, a particularly suitable basis to work in is the vector spherical wave-functions (VSWFs) \mathbf{N}_{mn} and \mathbf{M}_{mn} , as they can be split into radial and angular parts. Each VSWF is a solution of the Maxwell's equations and together they form a complete orthogonal basis. Full definition of VSWFs, which is overwhelmingly rich in the use of special functions, can be found in Section 4.2.5; for now we will only describe their properties conceptually.

4.2.1 Properties of VSWFs

First of all, we note that $n = 1, 2, 3, \dots$ denotes the degree of a VSWF and is related to the radial part, while $m = -n:n$ denotes the order and is related to the angular part.

We use two kinds of VSWFs: $\mathbf{N}_{mn}^{(1)}$, $\mathbf{M}_{mn}^{(1)}$ and $\mathbf{N}_{mn}^{(3)}$, $\mathbf{M}_{mn}^{(3)}$. Superscript (1) indicates that one of the building blocks of this VSWF is a spherical Bessel function of the first kind. This gives the VSWF the particular property of converging at the origin, as can be seen in Figure 4.1(a). For this reason $\mathbf{N}_{mn}^{(1)}$, $\mathbf{M}_{mn}^{(1)}$ are used to represent fields which are coming from infinity and have a finite value at the origin, for example the field of incident light. Superscript (3) means that the definition of a VSWF includes a spherical Hankel function of the first kind. This results in $\mathbf{N}_{mn}^{(3)}$, $\mathbf{M}_{mn}^{(3)}$ diverging at the origin (it is the imaginary part in particular that does this) as shown in Figure 4.1(b). This is a property of an *outgoing field*, i.e. a field, the source of which is located at the origin and which transports its energy towards infinity.

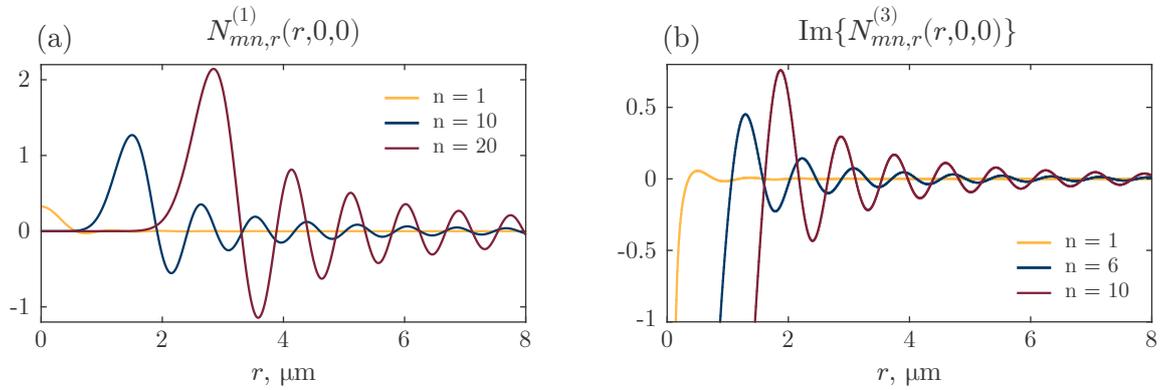


Figure 4.1: Properties of VSWFs. In (a) we see the variation of the radial part of $\mathbf{N}_{mn}^{(1)}$ with the distance from the origin r . Shown in (b) is the same variation but this time for the imaginary part of the radial component of $\mathbf{N}_{mn}^{(3)}$. The values are calculated at polar and azimuthal angles $\theta, \phi = 0$, for $m = 0$.

What is more, all of the VSWFs decay as they are approaching infinity, which is another important property of electromagnetic fields. Another characteristic we observe in Figure 4.1, and the significance of which will be made evident soon, is that the VSWFs with a higher n value peak further away from the origin.

4.2.2 Field decomposition into VSWFs

Since VSWFs of the same kind (as indexed by the superscript, e.g. (1)) form a complete orthogonal basis, any electromagnetic field can be written as a sum of VSWFs. To represent the external light field, which is the field of the incident light in the absence

of any particles, we use $\mathbf{N}_{mn}^{(1)}$ and $\mathbf{M}_{mn}^{(1)}$ as these functions are finite at the origin:

$$\begin{aligned}\mathbf{E}_{\text{ext}}(k\mathbf{r}) &= -i \sum_{n=1}^{\infty} \sum_{m=-n}^n \left(p_{mn} \mathbf{N}_{mn}^{(1)}(k\mathbf{r}) + q_{mn} \mathbf{M}_{mn}^{(1)}(k\mathbf{r}) \right), \\ \mathbf{H}_{\text{ext}}(k\mathbf{r}) &= -n_{\text{ext}} \sum_{n=1}^{\infty} \sum_{m=-n}^n \left(q_{mn} \mathbf{N}_{mn}^{(1)}(k\mathbf{r}) + p_{mn} \mathbf{M}_{mn}^{(1)}(k\mathbf{r}) \right),\end{aligned}\quad (4.6)$$

where \mathbf{E} and \mathbf{H} are the electric and magnetic fields respectively, n_{ext} is the refractive index of the external medium, k is the wavenumber of light in the external medium, \mathbf{r} is the position vector defined in a coordinate frame the origin of which is centred on the particle, and p_{mn} and q_{mn} are beam shape coefficients (BSCs), which are different for every light field.

The field which is scattered by the particle is represented in terms of $\mathbf{N}_{mn}^{(3)}$ and $\mathbf{M}_{mn}^{(3)}$ as these functions describe outgoing waves in the far field:

$$\begin{aligned}\mathbf{E}_{\text{scat}}(k\mathbf{r}) &= i \sum_{n=1}^{\infty} \sum_{m=-n}^n \left(a_{mn} \mathbf{N}_{mn}^{(3)}(k\mathbf{r}) + b_{mn} \mathbf{M}_{mn}^{(3)}(k\mathbf{r}) \right), \\ \mathbf{H}_{\text{scat}}(k\mathbf{r}) &= n_{\text{ext}} \sum_{n=1}^{\infty} \sum_{m=-n}^n \left(b_{mn} \mathbf{N}_{mn}^{(3)}(k\mathbf{r}) + a_{mn} \mathbf{M}_{mn}^{(3)}(k\mathbf{r}) \right).\end{aligned}\quad (4.7)$$

The total field outside of the particle is the sum of the external and scattered fields.

Finally, the total field inside the particle is described with the same VSWFs as for the external field:

$$\begin{aligned}\mathbf{E}_{\text{int}}(k\mathbf{r}) &= -i \sum_{n=1}^{\infty} \sum_{m=-n}^n \left(d_{mn} \mathbf{N}_{mn}^{(1)}(k'\mathbf{r}) + c_{mn} \mathbf{M}_{mn}^{(1)}(k'\mathbf{r}) \right), \\ \mathbf{H}_{\text{int}}(k\mathbf{r}) &= -n_{\text{ext}} \sum_{n=1}^{\infty} \sum_{m=-n}^n \left(c_{mn} \mathbf{N}_{mn}^{(1)}(k'\mathbf{r}) + d_{mn} \mathbf{M}_{mn}^{(1)}(k'\mathbf{r}) \right),\end{aligned}\quad (4.8)$$

where $k' = \frac{n_{\text{int}}}{n_{\text{ext}}} k$ is the wavenumber inside the particle. BSCs a_{mn} , b_{mn} , c_{mn} , d_{mn} are related to p_{mn} , q_{mn} and are dependent on the refractive index and size of the particle, as defined in Equation 4.30 in Section 4.2.5.

In practice we can only sum a finite number of VSWFs, so the infinite limit of the first sum is replaced with a cut-off value n_{max} . The choice of the cut-off value effectively determines the radius of the region over which the field is represented correctly, as illustrated in Figure 4.2. Empirically, the cut-off value is usually taken as [160]:

$$n_{\text{max}} = kr_{\text{roi}} + 4.05(kr_{\text{roi}})^{1/3} + 2, \quad (4.9)$$

where r_{roi} indicates the radius of the region of interest and is usually of a similar size to the particle radius. The reason why it is fine to neglect the larger n values is that VSWFs with high n peak further away from the origin and are effectively zero until they reach the base of that peak. They therefore do not contribute to the field in the proximity of the particle, as we saw earlier.

4.2.3 Beam shape coefficients and beam representation

In the most general approach the BSCs can be calculated for an arbitrary field \mathbf{E} by decomposing it into VSWFs [159]:

$$p_{mn} = i \frac{\int_S \mathbf{E} \cdot \mathbf{N}_{mn}^* dS}{\int_S |\mathbf{N}_{mn}|^2 dS}, \quad q_{mn} = i \frac{\int_S \mathbf{E} \cdot \mathbf{M}_{mn}^* dS}{\int_S |\mathbf{M}_{mn}|^2 dS}, \quad (4.10)$$

where S is a spherical surface of integration. To compute a sufficient number of coefficients p_{mn} , q_{mn} one would need to evaluate $\sim n_{\text{max}}^2$ surface integrals, which would be immensely time consuming. Luckily, for some particular light fields these integrals can be solved analytically, if not fully then at least partially. We discuss in more detail plane waves and Bessel beams.

Plane wave

Consider a plane wave with an electric field $\mathbf{E}(\mathbf{r}) = \mathbf{e}_0 e^{i\mathbf{k} \cdot \mathbf{r}}$, where $\mathbf{e}_0 = [e_r, e_\theta, e_\phi]$ is its polarisation vector, incident on a spherical particle located at $\mathbf{r}_0 = [r_0, \theta_0, \phi_0]$. Equation 4.10 has an exact analytical solution in this case [161, 162]:

$$\begin{Bmatrix} p_{mn} \\ q_{mn} \end{Bmatrix} = U_n \left(e_\theta \begin{Bmatrix} \tilde{\tau}(\cos \theta) \\ \tilde{\pi}(\cos \theta) \end{Bmatrix} - i e_\phi \begin{Bmatrix} \tilde{\pi}(\cos \theta) \\ \tilde{\tau}(\cos \theta) \end{Bmatrix} \right) e^{-im\phi} e^{i\mathbf{k} \cdot \mathbf{r}_0}, \quad (4.11)$$

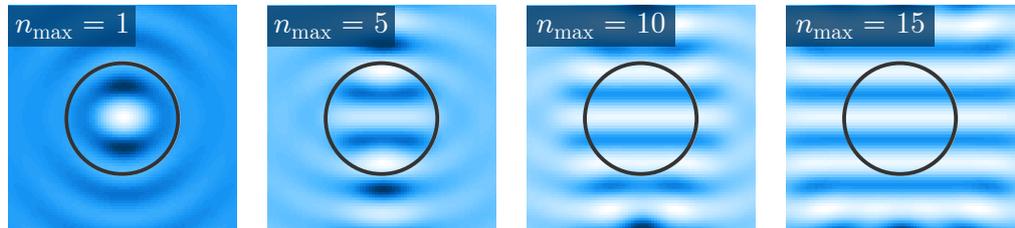


Figure 4.2: Effect of n_{max} . Including more terms in the VSWF expansion increases the area over which a light field, centred on the particle indicated with the black circle, is represented correctly. Shown here is the real amplitude of a plane wave. Figure adapted from [159].

where $U_n = \frac{4\pi i^n}{n(n+1)}$, θ and ϕ are the polar and azimuthal angles of the wavevector \mathbf{k} , and functions $\tilde{\tau}$, $\tilde{\pi}$ are defined in Section 4.2.5.

Bessel beam

Bessel beams will be of particular interest to us, so we present the results for their BSCs here. We build on the work in [163], extending it to a y -polarised beam and arbitrary values of orbital angular momentum (OAM) L .

In the *far field* an ideal Bessel beam is an infinitesimally thin ring of uniform intensity. The phase along this ring is uniform unless the beam has non-zero L value, in which case the phase changes along the ring from 0 to 2π a total of $|L|$ times, as dictated by $\exp(iL\phi)$, ϕ being the azimuthal angle, and shown in Figure 4.3(a,b). One can then think of a Bessel beam in the *focal plane* as being formed by plane waves originating from the ring in the far field and then being focused to travel along a cone, at an angle α^1 to the z -axis. In Figure 4.3(c) we show two of these waves interfering to create a Bessel beam. The two waves are π out of phase, so they interfere destructively along the beam axis, leaving a zero-intensity core. The dashed lines where the wavefronts cross correspond to the first and second intensity rings in the transverse intensity distribution (Figure 4.3(d)).

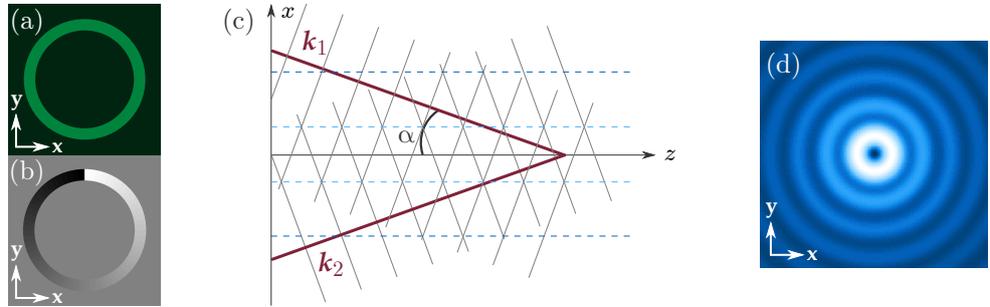


Figure 4.3: Formation of a Bessel beam. Intensity (a) and phase (b) in the far field of a Bessel beam with $L = 1$. In (c) plane waves travelling along a cone at an angle α to the z -axis, and originating from diametrically opposite points of the far field ring, interfere to create a Bessel beam in the focal plane (d). The light and dark blue dashed lines show the constructive interference regions which correspond to the first and second intensity rings respectively.

¹ not to be confused with the parameter α in the GWS context

To obtain the electric field at a point \mathbf{r} we sum over the plane waves which are travelling at a cone angle α to the z -axis:

$$\mathbf{E}_{\text{Bess}}(\mathbf{r}) = E_0 \int_0^{2\pi} \mathbf{e}(\alpha, \phi) e^{i\mathbf{k}(\alpha, \phi) \cdot \mathbf{r}} e^{iL\phi} d\phi, \quad (4.12)$$

where E_0 is the complex amplitude of the Bessel beam in the far field and we have also included the $e^{iL\phi}$ term to describe the OAM of the resulting beam.

It is not straightforward to derive the BSCs for arbitrary polarisation, so we do it for x and y -polarisations separately. The two can then be combined with appropriate amplitudes and a phase shift to describe a beam with any polarisation.

We begin by considering a single x -polarised plane wave travelling in the z -direction with wavevector $\mathbf{k}' = [0, 0, k]$ and polarisation $\mathbf{e}'^{(x)} = [1, 0, 0]$. Upon passing through point P at the base of the cone, as shown in Figure 4.4, the wave changes direction and is now described with wavevector \mathbf{k} and polarisation $\mathbf{e}^{(x)}$ (where the superscript (x) indicates that the initial polarisation of the plane wave was in the x -direction). We can determine \mathbf{k} and $\mathbf{e}^{(x)}$ by rotating the original direction and polarisation vectors by an angle α about a line which is *tangential* to the base of the cone at point P . This line can be described with a unit vector $\hat{\mathbf{u}} = [-\sin \phi, \cos \phi, 0]$ as shown in the inset of Figure 4.4. After the rotation we have:

$$\begin{aligned} e_x^{(x)} &= \cos \alpha + \sin^2 \phi (1 - \cos \alpha) & k_x &= k \sin \alpha \cos \phi \\ e_y^{(x)} &= -(1 - \cos \alpha) \sin \phi \cos \phi & k_y &= k \sin \alpha \sin \phi \\ e_z^{(x)} &= -\sin \alpha \cos \phi & k_z &= k \cos \alpha. \end{aligned} \quad (4.13)$$

And we can of course verify that $\mathbf{k} \cdot \mathbf{e}^{(x)} = 0$.

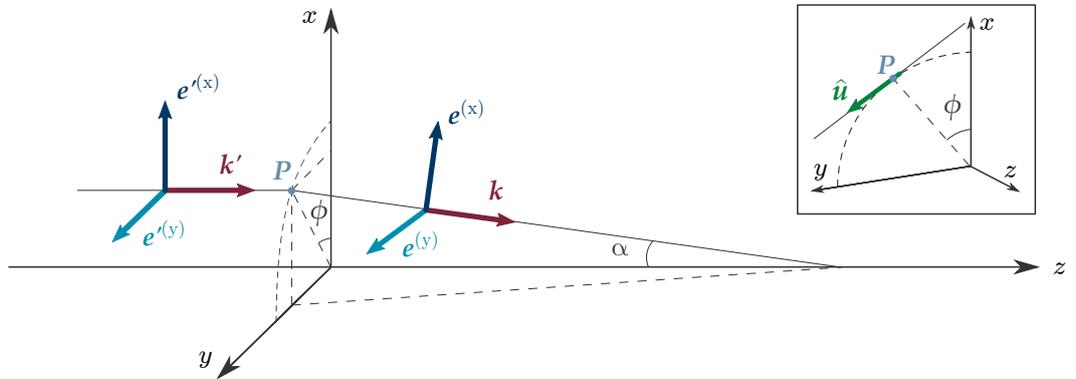


Figure 4.4: Description of a plane wave travelling at an angle to the z -axis. A plane wave with wavevector \mathbf{k}' in the far field is focused at an angle α . The new wavevector \mathbf{k} and polarisation are obtained via rotation about $\hat{\mathbf{u}}$, which is shown in the inset, and is tangential to the base of the cone at point P , from which the focused plane wave is emerging.

Polarisation vector $\mathbf{e}^{(x)}$ can then be expressed in the spherical basis of \mathbf{k} ($\mathbf{k} = [k, \alpha, \phi]$) using the following transform:

$$\begin{bmatrix} e_r^{(x)} \\ e_\theta^{(x)} \\ e_\phi^{(x)} \end{bmatrix} = \begin{bmatrix} \cos \phi \sin \alpha & \sin \phi \sin \alpha & \cos \alpha \\ \cos \phi \cos \alpha & \sin \phi \cos \alpha & -\sin \alpha \\ -\sin \phi & \cos \phi & 0 \end{bmatrix} \begin{bmatrix} e_x^{(x)} \\ e_y^{(x)} \\ e_z^{(x)} \end{bmatrix} = \begin{bmatrix} 0 \\ \cos \phi \\ -\sin \phi \end{bmatrix}. \quad (4.14)$$

We now take Equation 4.11 with BSCs for a plane wave and put it into Equation 4.12 for a Bessel beam to obtain:

$$\begin{Bmatrix} p_{mn} \\ q_{mn} \end{Bmatrix} = U_n E_0 \int_0^{2\pi} \left(\cos \phi \begin{Bmatrix} \tilde{\tau}(\cos \theta) \\ \tilde{\pi}(\cos \theta) \end{Bmatrix} + i \sin \phi \begin{Bmatrix} \tilde{\pi}(\cos \theta) \\ \tilde{\tau}(\cos \theta) \end{Bmatrix} \right) e^{-im\phi} e^{iL\phi} e^{i\mathbf{k}\cdot\mathbf{r}_0}. \quad (4.15)$$

We then expand the $e^{i\mathbf{k}\cdot\mathbf{r}_0}$ term, rewrite the sine and cosine in terms of exponentials and, defining $\rho = k\sqrt{x^2 + y^2} \sin \alpha$ and $\phi_0 = \arctan(-y/x) - \pi/2$, obtain:

$$\begin{Bmatrix} p_{mn} \\ q_{mn} \end{Bmatrix} = U_n E_0 e^{ikz \cos \alpha} \left(\begin{Bmatrix} \tilde{\tau}(\cos \theta) \\ \tilde{\pi}(\cos \theta) \end{Bmatrix} I_m^{(+)} + \begin{Bmatrix} \tilde{\pi}(\cos \theta) \\ \tilde{\tau}(\cos \theta) \end{Bmatrix} I_m^{(-)} \right), \quad (4.16)$$

where the integral $I_m^{(\pm)}$ is given by:

$$I_m^{(\pm)} = \frac{1}{2} \left[\int_0^{2\pi} e^{i(1-(m-L))\phi} e^{i\rho \cos(\phi+\phi_0+\frac{\pi}{2})} d\phi \pm \int_0^{2\pi} e^{i(-1-(m-L))\phi} e^{i\rho \cos(\phi+\phi_0+\frac{\pi}{2})} d\phi \right] \quad (4.17)$$

and it can be solved analytically in terms of Bessel functions of the first kind J :

$$I_m^{(\pm)} = \pi \left[e^{i(m-L-1)\phi_0} J_{1-(m-L)}(\rho) \pm e^{i(m-L+1)\phi_0} J_{-1-(m-L)}(\rho) \right]. \quad (4.18)$$

We can obtain $\mathbf{e}^{(y)}$ directly by rotating $\mathbf{e}^{(x)}$ by 90° about \mathbf{k} . Following the same procedure as above we arrive at the following result for the BSCs of a y -polarised Bessel beam travelling in the z -direction:

$$\begin{Bmatrix} p_{mn} \\ q_{mn} \end{Bmatrix} = -i U_n E_0 e^{ikz \cos \alpha} \left(\begin{Bmatrix} \tilde{\tau}(\cos \theta) \\ \tilde{\pi}(\cos \theta) \end{Bmatrix} I_m^{(-)} + \begin{Bmatrix} \tilde{\pi}(\cos \theta) \\ \tilde{\tau}(\cos \theta) \end{Bmatrix} I_m^{(+)} \right), \quad (4.19)$$

4.2.4 Optical force

In our earlier simulations of the hydrodynamics in a system of multiple optically trapped beads we used the linear approximation to model the optical restoring force in

a Gaussian beam trap. Now, however, we are aiming to find a new kind of optical trap that would have an enhanced stiffness. For this we need a general equation describing the force that a spherical particle experiences in an arbitrary light field. Barton [164] has derived exactly such an equation² in terms of the BSCs, which has been re-normalised to fit our definitions [159]:

$$\begin{aligned}
F_x + iF_y = \frac{\epsilon_r}{16\pi k^2} \sum_{n=1}^{n_{\max}} \sum_{m=-n}^n \left[-in(n+2) \sqrt{\frac{(n+m+2)(n+m+1)}{(2n+1)(2n+3)}} \right. \\
\times (a_{nm}p_{n+1,m+1}^* + b_{nm}q_{n+1,m+1}^* \\
+ p_{nm}a_{n+1,m+1}^* + q_{nm}b_{n+1,m+1}^* \\
- 2a_{nm}a_{n+1,m+1}^* - 2b_{nm}b_{n+1,m+1}^*) \\
- in(n+2) \sqrt{\frac{(n-m+1)(1-m+2)}{(2n+1)(2n+3)}} \\
\times (a_{n+1,m-1}p_{nm}^* + b_{n+1,m-1}q_{nm}^* \\
+ p_{n+1,m-1}a_{nm}^* + q_{n+1,m-1}b_{nm}^* \\
- 2a_{n+1,m-1}a_{nm}^* - 2b_{n+1,m-1}b_{nm}^*) \\
+ \sqrt{(n+m+1)(n-m)} \\
\times (a_{nm}q_{n,m+1}^* + b_{nm}p_{n,m+1}^* \\
+ q_{nm}a_{n,m+1}^* + p_{nm}b_{n,m+1}^* \\
- 2a_{nm}b_{n,m+1}^* - 2b_{nm}a_{n,m+1}^*) \left. \right], \quad (4.20)
\end{aligned}$$

$$\begin{aligned}
F_z = \frac{\epsilon_r}{8\pi k^2} \sum_{n=1}^{n_{\max}} \sum_{m=-n}^n \operatorname{Im} \left[n(n+2) \sqrt{\frac{(n-m+1)(n+m+1)}{(2n+1)(2n+3)}} \right. \\
\times (a_{n+1,m-1}p_{nm}^* + b_{n+1,m-1}q_{nm}^* \\
+ p_{n+1,m-1}a_{nm}^* + q_{n+1,m-1}b_{nm}^* \\
- 2a_{n+1,m-1}a_{nm}^* - 2b_{n+1,m-1}b_{nm}^*) \\
+ im(p_{nm}b_{nm}^* + q_{nm}a_{nm}^* \\
- a_{nm}b_{nm}^* - b_{nm}a_{nm}^*) \left. \right], \quad (4.21)
\end{aligned}$$

²This equation is obtained by integrating the Maxwell stress tensor, which represents the net flow of momentum across any given surface point, over the surface of the particle.

where ϵ_r is the relative electric permittivity of the external medium. With these equations we can obtain the information about the optical trapping stiffness as well, by numerically estimating the force gradient.

4.2.5 Full definition of VSWFs and BSCs

The remainder of this chapter is a lengthy list of special functions used to define the VSWFs³. We use the same definitions as in [159]. Presented in Table 4.1 are all the relevant functions, and below we list either the MATLAB functions that were used to calculate them or their definitions.

Function	Name	Defined in
$\mathbf{N}_{mn}^{(j)}, \mathbf{M}_{mn}^{(j)}$	VSWFs, $j = 1, 3$	4.22
P_n^m	Associated Legendre polynomials	4.24, 4.26
$\tilde{\pi}, \tilde{\tau}$	Unnamed angular functions	4.25
j_n	Spherical Bessel function of the first kind	4.27
h_n	Spherical Hankel function	4.28

Table 4.1: Functions used in defining the VSWFs. The third column lists the equations in this text which define the functions.

The vector spherical wave-functions are defined in a spherical polar coordinate system centred on the scattering particle and with unit vectors $\mathbf{i}_r, \mathbf{i}_\theta, \mathbf{i}_\phi$ as:

$$\begin{aligned}
 \mathbf{M}_{mn}^{(j)}(\mathbf{r}) &= [\mathbf{i}\tilde{\pi}_{mn}(\cos\theta)\mathbf{i}_\theta - \tilde{\tau}_{mn}(\cos\theta)\mathbf{i}_\phi] e^{im\phi} z_n^{(j)}(kr), \\
 \mathbf{N}_{mn}^{(j)}(\mathbf{r}) &= \frac{n(n+1)}{kr} \sqrt{\frac{2n+1}{4\pi} \frac{(n-m)!}{(n+m)!}} P_n^m(\cos\theta) e^{im\phi} z_n^{(j)}(kr) \mathbf{i}_r \\
 &\quad [\tilde{\tau}_{mn}(\cos\theta)\mathbf{i}_\theta + \mathbf{i}\tilde{\pi}_{mn}(\cos\theta)\mathbf{i}_\phi] e^{im\phi} \frac{1}{kr} \frac{d}{d(kr)} (kr z_n^{(j)}(kr)), \quad (4.22)
 \end{aligned}$$

where $\mathbf{r} = [r, \theta, \phi]$ is the position vector, r denotes the radial coordinate, θ is the polar angle measured from the positive z -axis towards the xy -plane, and ϕ is the azimuthal angle measured from the positive x -axis towards the positive y -axis. Superscript $j = 1, 3$ is not to be confused with j_n (see Table 4.1), and we used the following notation:

$$\begin{aligned}
 z_n^{(1)}(kr) &= j_n(kr), \\
 z_n^{(3)}(kr) &= h_n^{(1)}(kr). \quad (4.23)
 \end{aligned}$$

³ ‘Don’t panic’ — Douglas Adams, *The Hitchhiker’s Guide to the Galaxy*

Associated Legendre polynomials can be calculated for positive m values using MATLAB function ‘legendre’. For negative m we can use the relation:

$$P_n^{-m}(x) = (-1)^m \frac{(n-m)!}{(n+m)!} P_n^m(x). \quad (4.24)$$

Angular functions $\tilde{\pi}(\cos \theta)$, $\tilde{\tau}(\cos \theta)$ are defined as:

$$\begin{aligned} \tilde{\pi}_{mn}(\cos \theta) &= \sqrt{\frac{2n+1}{4\pi} \frac{(n-m)!}{(n+m)!}} \frac{m}{\sin \theta} P_n^m(\cos \theta), \\ \tilde{\tau}_{mn}(\cos \theta) &= \sqrt{\frac{2n+1}{4\pi} \frac{(n-m)!}{(n+m)!}} \frac{d}{d\theta} P_n^m(\cos \theta). \end{aligned} \quad (4.25)$$

We note that at $\theta = 0$ the $\tilde{\pi}(\cos \theta)$ function cannot be directly calculated because of the singularity in the $(1/\sin \theta)$ term. There are two distinct cases here: $m \neq 0$ and $m = 0$. After noting that $P_n^{m \neq 0}(1) = 0$ for all n , we can find the limit as $\theta \rightarrow 0$ using L’Hôpital’s rule, which states that for two functions $f(x), g(x)$, which are differentiable over an open interval, except maybe at some point o in that interval, the following is true: $\lim_{x \rightarrow o} \frac{f(x)}{g(x)} = \lim_{x \rightarrow o} \frac{f'(x)}{g'(x)}$. So we have:

$$\begin{aligned} m \neq 0: \quad & \lim_{\theta \rightarrow 0} \frac{P_n^m(\cos \theta)}{\sin \theta} = \lim_{\theta \rightarrow 0} \frac{\frac{d}{d\theta} P_n^m(\cos \theta)}{\cos \theta} = \frac{d}{d\theta} P_n^m(\cos \theta), \\ m = 0: \quad & \lim_{\theta \rightarrow 0} \frac{m}{\sin \theta} = \lim_{\theta \rightarrow 0} \frac{0}{\cos \theta} = 0. \end{aligned}$$

Derivatives of associated Legendre polynomials $\frac{d}{d\theta} P_n^m(\cos \theta)$ can be evaluated in a number of ways. We use the following relation:

$$\sqrt{1-x^2} \frac{d}{dx} P_n^m(x) = \frac{1}{2} [(n+m)(n-m+1)P_n^{m-1}(x) - P_n^{m+1}(x)].$$

We are after the derivative with respect to θ and not $\cos \theta$, so we let $x = \cos \theta$ and apply the chain rule to obtain:

$$\begin{aligned} \frac{d}{d\theta} P_n^m(\cos \theta) &= -\sin \theta \frac{d}{d(\cos \theta)} P_n^m(\cos \theta) \\ &= -\frac{1}{2} [(n+m)(n-m+1)P_n^{m-1}(\cos \theta) - P_n^{m+1}(\cos \theta)]. \end{aligned} \quad (4.26)$$

The above can be simplified for the cases $m = \pm n$, if we notice that for $m > n$ and $m < -n$, $P_n^m(\cos \theta) = 0$. Then:

$$\begin{aligned} m = n, & \quad \frac{d}{d\theta} P_n^n(\cos \theta) = -\frac{1}{2} \left[(2n)(2n+1) P_n^{n-1}(\cos \theta) \right] \\ m = -n, & \quad \frac{d}{d\theta} P_n^{-n}(\cos \theta) = \frac{1}{2} P_n^{-n+1}(\cos \theta). \end{aligned}$$

The spherical Bessel functions are defined as:

$$\begin{aligned} \text{of the first kind :} & \quad j_n(kr) = \sqrt{\frac{\pi}{2kr}} J_{n+1/2}(kr), \\ \text{of the second kind :} & \quad y_n(kr) = \sqrt{\frac{\pi}{2kr}} Y_{n+1/2}(kr), \end{aligned} \quad (4.27)$$

where J, Y are the Bessel functions of the first and second kind respectively and can be computed using MATLAB functions ‘besselj’ and ‘bessely’.

The spherical Hankel functions are in turn defined as:

$$\begin{aligned} \text{of the first kind :} & \quad h_n^{(1)}(kr) = j_n(kr) + iy_n(kr), \\ \text{of the second kind :} & \quad h_n^{(2)}(kr) = j_n(kr) - iy_n(kr). \end{aligned} \quad (4.28)$$

Derivatives of spherical Bessel and Hankel functions are defined as:

$$\frac{d}{d(kr)} z_n^{(j)}(kr) = -z_{n+1}^{(j)}(kr) + \frac{n}{kr} z_n^{(j)}(kr), \quad (4.29)$$

where z_n denotes any of the spherical functions as defined in Equations 4.23.

Definition of BSCs

The BSCs associated with the scattered and internal fields are defined in terms of p_{mn} and q_{mn} as:

$$\begin{aligned} a_{mn} &= \frac{\psi_n(x)\psi'_n(y) - n_{\text{rel}}\psi_n(y)\psi'_n(x)}{\psi'_n(y)\xi_n(x) - n_{\text{rel}}\psi_n(y)\xi'_n(x)} p_{mn}, \\ b_{mn} &= \frac{n_{\text{rel}}\psi'_n(y)\psi_n(x) - \psi_n(y)\psi'_n(x)}{n_{\text{rel}}\psi'_n(y)\xi_n(x) - \psi_n(y)\xi'_n(x)} q_{mn}, \\ c_{mn} &= \frac{-in_{\text{rel}}}{n_{\text{rel}}\psi'_n(y)\xi_n(x) - \psi_n(y)\xi'_n(x)} q_{mn}, \\ d_{mn} &= \frac{-in_{\text{rel}}}{\psi'_n(y)\xi_n(x) - n_{\text{rel}}\psi_n(y)\xi'_n(x)} p_{mn}, \end{aligned} \quad (4.30)$$

where $x = ka$, $y = n_{\text{rel}}x$, a is the bead radius, $n_{\text{rel}} = n_{\text{int}}/n_{\text{ext}}$ is the relative refractive index between the particle and the surroundings, and $'$ indicates a derivative with respect to the argument of the function (x or y). ψ_n and ξ_n are known as the Riccati-Bessel functions and are closely related to the spherical Bessel and Hankel functions:

$$\psi_n(x) = x j_n(x), \quad \xi_n(x) = x h_n^{(1)}(x). \quad (4.31)$$

Chapter 5

Optimising optical trapping stiffness

‘Well, I must endure the presence of a few caterpillars if I wish to become acquainted with the butterflies.’

The Little Prince — Antoine de Saint-Exupéry

To exert optical control on our hydrodynamic actuators, we used the conventional Gaussian traps. But we found that the trapping stiffness was not always sufficiently high to ensure that the actuator beads replicate the motion of their traps precisely enough. In particular, the actuators were too slow to react to high frequency changes in the trap motions, especially when the available laser power had to be shared amongst

multiple rotors. This, in turn, degraded the tightness with which a target particle could be hydrodynamically clamped. In this chapter, therefore, we will be searching for an alternative optical trap, one which would maximally enhance the optical trapping stiffness, given a fixed amount of laser power. Such a trap would by no means be limited in applicability to our hydrodynamic manipulation setup, but would be of general interest to the optical trapping community in applications such as trapping of high refractive index particles, force sensing, probing the properties of the surrounding environment, and micro-robotics. We will take a step beyond the research that has already been published in this field, by ensuring that the stiffness enhancement does not come at the expense of loss of stable three-dimensional trapping.

Alongside the GWS method described in the previous chapter, we will also approach the problem of enhancing optical trapping stiffness from another angle - iterative optimisation. With this additional outlook we hope to gain insight into how well our implementation of the GWS is performing, as well as augment the GWS in some problems that it is not well suited to tackle on its own. After we describe our methods in more detail, we will begin by considering stiffness enhancement along one dimension - in both the transverse and longitudinal planes of the laser beam. Following this, we will explore how the different dimensions can be combined to create a trap with desirable three-dimensional features.

5.1 Methods

We begin with the big picture of the setup of our problem. In a simplified optical tweezers system illustrated in Figure 5.1, the light field emerges from the far field (which we also refer to as the pupil plane), is incident on a high NA objective lens, and is then focused in the focal plane (also referred to as the object plane) where the bead is residing. In order to create a field with particular properties (i.e. high stiffness) in the focal plane, its features (i.e. intensity and phase distribution) first have to be ‘designed’ in the pupil plane.

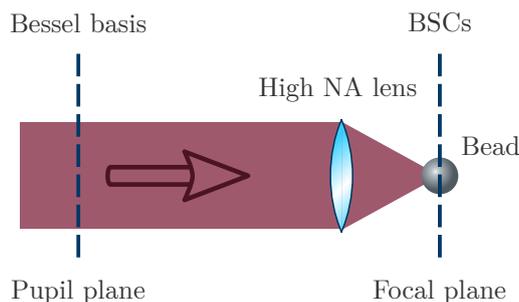


Figure 5.1: Relation between the pupil and focal planes via focusing by a high NA lens. Light fields in the pupil and focal planes are represented in terms of Bessel beams and BSCs respectively.

We now need a mathematical description of the light in this system. The VSWF representation is convenient when light is entering the system from all angles, but in our setup the light falling on the bead is limited to the angle subtended by the objective lens. Such directionality can be captured particularly well by Bessel beams. As we saw in Section 4.2.3 a Bessel beam in the far field is an annular ring described by the cone angle that it subtends upon being focused. The field in the pupil plane can thus be subdivided into concentric annular rings, each with a particular intensity and phase. Furthermore, Bessel beams are orthogonal to each other and form a complete basis - we therefore use this representation for the field in the pupil plane.

We already know how to describe Bessel beams in the focal region - we have derived the relevant BSCs in Section 4.2.3, which contain all the information about the light field and the force that it exerts on a spherical bead. In the following section we consider the Bessel basis in the pupil plane more rigorously.

5.1.1 The Bessel basis

In the far field, which we also refer to as the pupil plane, a Bessel beam is an annular ring with radius r and width dr , as shown in Figure 5.2(a). After being focused on the surface of the focal sphere \mathcal{F} , the rays forming the Bessel beam propagate along a cone and meet in the focal plane, Figure 5.2(b). The cone angle of each Bessel beam is related to its far field radius via $r = \sin \alpha$, assuming the radius of the focal sphere \mathcal{F} in Figure 5.2(c) to be 1.

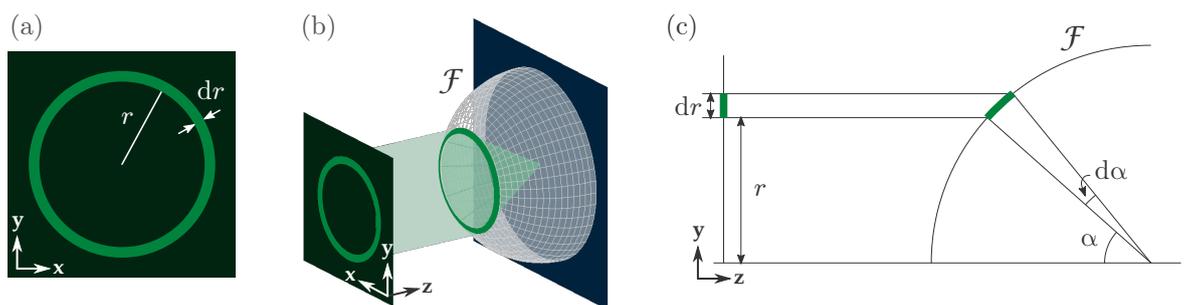


Figure 5.2: Geometry of Bessel beams. In the far field (a) the intensity of a Bessel beam is uniformly distributed along a ring of radius r and width dr . The rays originating from this ring change direction at the surface of the focal sphere (b) and propagate along a conical surface at angles in the range from α to $\alpha + d\alpha$ to the z -axis (c), to meet in the focal plane.

We allow the maximum cone angle of $\alpha_0 = 60^\circ$, which corresponds to NA of 1.3 of an oil immersion microscope (using $\text{NA} = \sin \alpha_0 n_{\text{oil}}$, where $n_{\text{oil}} = 1.5$ is the refractive

index of the immersion oil). We then divide up the pupil plane into N_α annuli which are evenly spaced in α , in the range from 0 to α_0 . For each cone angle we also allow N_L different OAM values, coming to a total of $N_b = N_\alpha \times N_L$ Bessel beams. This number further doubles if we allow polarisation to vary across the pupil, since we describe a Bessel beam of arbitrary polarisation by superposing two Bessel beams, one of x - and one of y -polarisation, with appropriate amplitude scaling and phase difference between the two, as explained in Section 4.2.3.

We describe each of the N_b Bessel beams in terms of the fraction f_i of the total power that they carry¹, and their phase Φ_i . The complex amplitude of the i -th beam in the pupil plane, propagating at cone angle α_i , is then given by:

$$E_{0i} = \sqrt{f_i P_{\text{tot}}} E_{N_i} e^{i\Phi_i}, \quad (5.1)$$

where E_{N_i} is the normalisation factor which we will derive shortly, and $\sum_{i=1}^{N_b} f_i = 1$.

Normalisation

Before we can use the above description in the optimiser, we must ensure that the power in the far field is conserved - otherwise the most straightforward route towards increasing stiffness is to ramp up the laser power. The total power in the pupil is given by integrating over r : $P_{\text{tot}} \propto \int_0^{r_0} |E_0(r)|^2 r dr$. However, since our beams are equally spaced in α , not r , it is more convenient to change the integration variable. Using $dr = \cos \alpha d\alpha$ and substituting for $E_0(\alpha)$ we have:

$$P_{\text{tot}} \propto \int_0^{\alpha_0} f(\alpha) P_{\text{tot}} E_N^2(\alpha) \sin \alpha \cos \alpha d\alpha. \quad (5.2)$$

We are, however, working with a discrete number of cone angles, so the above integral becomes a summation and $d\alpha \rightarrow \Delta\alpha$, which is the angular spacing between consecutive rings:

$$P_{\text{tot}} \propto \sum_i f_i P_{\text{tot}} E_{N_i}^2 \sin \alpha_i \cos \alpha_i \Delta\alpha, \quad (5.3)$$

which is a constant for any choice of f_i (provided that their sum is equal to one) if $E_{N_i} = 1/\sqrt{\sin \alpha_i \cos \alpha_i \Delta\alpha}$. Substituting this expression in Equation 5.1 we obtain

¹Note that an ideal Bessel beam carries an infinite amount of power in the *focal plane*, but here we are referring to the power available in the *pupil plane*.

the complex amplitude of the i -th Bessel beam:

$$E_{0i} = \sqrt{\frac{f_i P_{\text{tot}}}{\sin \alpha_i \cos \alpha_i \Delta \alpha}} e^{i\Phi_i}. \quad (5.4)$$

We note that we have not explicitly mentioned OAM in the above derivation, but the same normalisation applies to Bessel beams with non-zero L values too. It can be seen that if an extra factor of $e^{iL_i\phi}$ (where ϕ is the azimuthal coordinate) is included in Equation 5.1, it will not appear in the power normalisation equations because we are considering the square of E_{0i} , which eliminates the phase terms. This holds true even if we are adding several beams of different OAM values, propagating at the same α . Say we want to consider a *single ring* at some α , populated by several Bessel beams with different OAM values. The total complex amplitude on the ring is given by the sum of the complex amplitudes of all the beams, indexed with i :

$$E_{0,\text{ring}}(\phi) = \sum_i E_{0i}(\phi). \quad (5.5)$$

Assuming that each beam needs an OAM-dependent normalisation constant $E'_{Ni}(L)$ (which is different from the $E_{Ni}(\alpha)$ seen already), we write the complex amplitude of the i -th beam as $E_{0i}(\phi) = \sqrt{f_i P_{\text{tot}}} E'_{Ni} e^{i\Phi_i} e^{iL_i\phi}$, and then consider the total power in the ring:

$$P_{\text{tot},\text{ring}} \propto \int_0^{2\pi} \left| \sum_i \sqrt{f_i P_{\text{tot}}} E'_{Ni} e^{i\Phi_i} e^{iL_i\phi} \right|^2 d\phi. \quad (5.6)$$

Squaring the sum under the integral will produce cross-terms containing $e^{i(L_i-L_j)\phi}$, for $i \neq j$, which are zero when integrated over ϕ from 0 to 2π . The remaining terms will be of the form $f_i P_{\text{tot}} E'^2_{Ni}$, which do not have ϕ dependence, so they can be taken out of the integral and we have:

$$P_{\text{tot},\text{ring}} \propto \sum_i f_i P_{\text{tot}} E'^2_{Ni}. \quad (5.7)$$

We see that to ensure constant power in the ring there is no need for the OAM-dependent normalisation constant E'_{Ni} - it can be set to 1 for all beams and our condition requiring that $\sum_i f_i = 1$ is sufficient to guarantee constant power. Hence, our normalisation for Bessel beams is independent of their OAM values.

Next, we take care of the focal plane. Here the total field is obtained by adding the BSCs of the individual Bessel beams coming in from different cone angles - this takes the form of integration over the polar angle in the focal region: $p_{mn}^{(\text{tot})} = \int_0^{\alpha_0} p_{mn}^{(\alpha)} \sin \alpha d\alpha$.

Once again this becomes a sum, because we are dealing with discrete numbers of Bessel beams. Finally, we must ensure that the power is conserved when going from the pupil to the focal plane. As seen in Figure 5.2(b,c), the annular ring in the far field is projected onto a curved region of the focal sphere - the areas of the ring and its projection are not the same, yet they must be associated with the same power. This can be accounted for by adding an extra factor of $\sqrt{\cos \alpha}$ to the above equation [165]:

$$p_{mn}^{(\text{tot})} = \sum_i p_{mn}^{(i)} \sqrt{\cos \alpha_i} \sin \alpha_i \Delta \alpha, \quad (5.8)$$

and likewise for q_{mn} .

We check that power is indeed conserved by adding up different numbers of Bessel beams N_b in the range from 0 to α_0 to form an intensity spot in the focal plane, as seen in Figure 5.3(a). We then integrate over a square region with side length of $8 \mu\text{m}$ to obtain the total intensity. If power is conserved, the total intensity value should be independent of N_b . Indeed, we see in Figure 5.3(b) that in the limit of high N_b the total intensity approaches a constant value. And even if a small number of beams are used, e.g. 20, the intensity is still preserved to within 1.5 %.

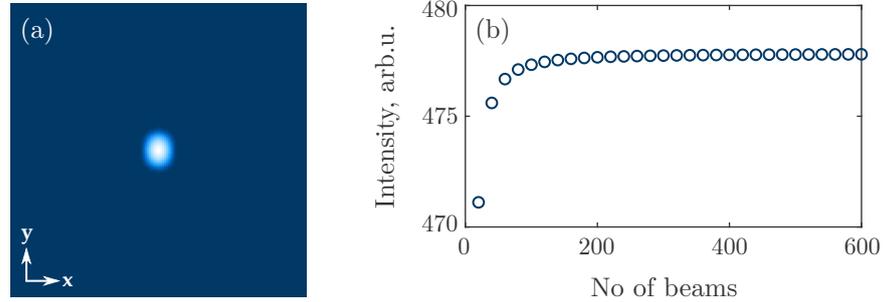


Figure 5.3: Conservation of power in Bessel beams. With appropriate normalisation, addition of multiple Bessel beams to create a bright intensity spot (a) conserves total intensity in the focal plane (b). The width and height of (a) are $8 \mu\text{m}$.

5.1.2 The optimiser

The task of the optimiser is to find the combination of Bessel beams in the far field that generates a field in the focal plane, which can most stiffly trap a spherical bead. The field that is being optimised, i.e. the one in the pupil plane, is described in terms of the fraction of power f_i and the phase Φ_i of each beam - these are the optimisation variables.

In every iteration the optimiser tests a different set of values for the optimisation parameters and estimates the optical trapping stiffness associated with that set. This is done by first computing the complex amplitude of each Bessel beam in the set using Equation 5.1, then plugging this into Equation 4.16 in order to obtain the BSCs, and directly estimating the force acting on the particle using Equations 4.20 and 4.21. We then compute the gradient of the force using a two-point finite difference scheme - the force is estimated once when the bead is residing at the origin, and once when the bead is displaced by a small distance from it, e.g. Δx if we are optimising for κ_x . Within the context of optimisation, the stiffness is referred to as the optimisation objective. The optimiser itself is built using MATLAB Optimisation Toolbox with the ‘fmincon’ function using the ‘interior-point’ algorithm, which minimises the optimisation objective subject to user specified constraints.

If we let the optimiser simply look for the field with the lowest (i.e. most negative) stiffness, it will stop as soon as it reaches any local minimum, as can be seen in Figure 5.4. We therefore add a constraint requiring that the objective function, i.e. the stiffness, reaches some specific value, which we set to be of a much larger magnitude than we think is realistic for the optimiser to actually achieve. This way the optimiser is able to ‘jump out’ of at least some of the local minima it encounters along the way. This works because within the interior-point method the initial problem of minimising a function subject to some constraints is rewritten in terms of a merit function - this function combines the initial objective function and the constraints [166, 167]. It is the merit function that the optimiser tries to minimise. So while our objective function is able to increase to escape local minima, behind the scenes the merit function is decreasing. This, however, does not ensure that the final solution will be a global minimum. We stop the optimiser if the stiffness has changed by less than 1% over the past 60 iterations. With an additional constraint we ensure that f_i of all beams add up to 1.

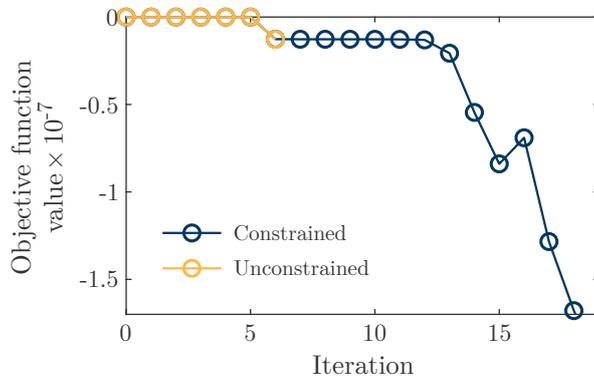


Figure 5.4: Constraining the objective function of the optimiser. An unconstrained optimiser will stop upon reaching the first local minimum (yellow). But if we include a constraint that requires the optimiser to reach a certain objective function value, it will be able to ‘jump over’ some of the local minima (blue).

5.1.3 GWS

For the GWS approach we first need to select the bases for the input and output fields, so that we can estimate the scattering matrix S . We use the same Bessel beam basis as for the optimiser, but now, instead of using the fractional power f_i and phase Φ_i , we define $|\text{in}\rangle$ directly in terms of the complex amplitude of each Bessel mode.

It would be most convenient to express $|\text{out}\rangle$ directly in terms of the BSCs that describe the light field around the particle after it has interacted with the light. However, as we saw in Section 4.2.2, the scattered and incoming fields are not expressed in the same VSWF bases, so we cannot just naively add the BSCs to describe the outgoing field. Furthermore, these bases are defined with respect to the location of the bead, so the scattering matrices S which are estimated with bead residing at different positions are effectively in different bases, and it is not straightforward to combine them². We therefore define the outgoing field as the vector light field on a spherical surface surrounding the bead, and centred at the origin, as shown in Figure 5.5(a) - this surface remains fixed in the same position when the particle is displaced for evaluation of the scattering matrices. In our simulation we are limited to how far away from the particle we can define this surface because of Equation 4.9. When n_{max} reaches the value of 86, the factorials involved in the definition of the VSWFs blow up numerically; this limits us to a spherical volume with an $8\mu\text{m}$ radius.

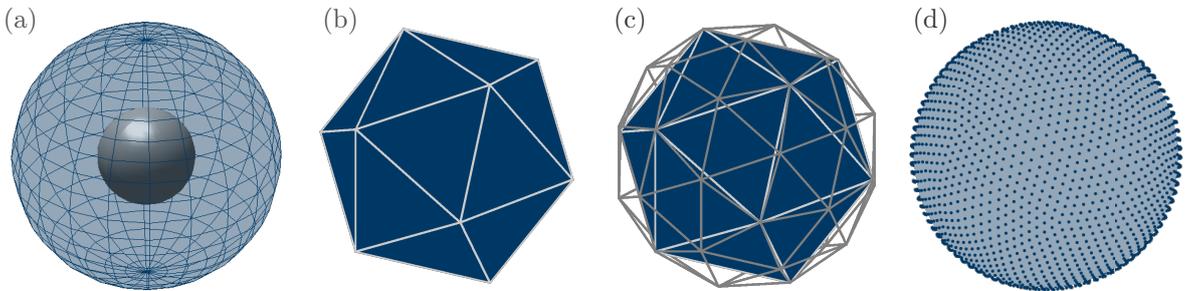


Figure 5.5: Output basis for the GWS method. The output state is defined as the complex light field amplitude on a spherical surface around the particle (a). This surface can be approximated as an icosahedron (b), which can further be divided up into a polyhedron (c) by splitting up each triangular face into four triangles and projecting the additional vertices onto a sphere. Each vertex then corresponds to a sampling point where the field is ‘measured’. The particular arrangement that we use has 2562 sampling points (d).

To obtain the $|\text{out}\rangle$ vector we sample the surface at different locations - this presents the issue of uniformly distributing points on a sphere. We choose a method which models

²It is not impossible, though. One can use translation matrices [159] to combine the different VSWF bases.

a sphere as a geodesic polyhedron. The particular code we used [168] starts with a regular icosahedron, the twelve vertices of which are located on our spherical surface, shown in Figure 5.5(b). Each of the twenty equilateral triangles is then subdivided into four and the new vertices are projected onto a sphere, as shown in Figure 5.5(c). The process can then be repeated until the number of vertices reaches the desired number of sampling points N_s . The main drawback of this technique is that N_s cannot be arbitrary - it can only take values of $N_s = 10 \times 4^n + 2$, where $n = 0, 1, 2, \dots$ is the number of triangle subdivisions performed. We choose to work with $n = 4$ which gives us 2562 sampling points, see Figure 5.5(d), to fully capture variations in the outgoing field. Effectively, N_s is further tripled because there are three electric field components associated with each sampling point.

The scattering matrix S connecting the incoming and outgoing fields will have dimensions of $N_s \times N_b$. The procedure for estimating S is fairly straightforward. We ‘activate’ the input basis modes one at the time, making $|\text{in}\rangle$ a column vector in which all but one entries are zero. For each active beam we then estimate the output field profile $|\text{out}\rangle$, which immediately gives us the columns of S corresponding to each $|\text{in}\rangle$ state. We illustrate this procedure with a simple example.

Suppose we have three input modes and four output modes. First, we set $|\text{in}\rangle = [a \ 0 \ 0]$, where a is some chosen number. Then, using the GLMT, we obtain the corresponding output state, $|\text{out}\rangle = [b_1 \ b_2 \ b_3 \ b_4]$. Explicitly writing out the elements in the scattering equation, we have:

$$\begin{bmatrix} b_1 \\ b_2 \\ b_3 \\ b_4 \end{bmatrix} = \begin{bmatrix} S_{11} & S_{12} & S_{13} \\ S_{21} & S_{22} & S_{23} \\ S_{31} & S_{32} & S_{33} \\ S_{41} & S_{42} & S_{43} \end{bmatrix} \begin{bmatrix} a \\ 0 \\ 0 \end{bmatrix}. \quad (5.9)$$

Immediately we see that $S_{11} = b_1/a$, and similarly we obtain the rest of the first column. We then set $|\text{in}\rangle = [0 \ a \ 0]$ and repeat the procedure, until the entire scattering matrix S is obtained. If we wish not to constrain the light input to a single polarisation state and instead allow the GWS to ‘choose’ it, we simply concatenate two scattering matrices - one estimated for x -polarised input, and one for y -polarised input.

5.1.4 Evaluating new traps, and additional notes

There are a few technicalities remaining that we have to cover before jumping straight to results. Our ‘test’ system consists of a silica bead of $3\mu\text{m}$ radius in water, illuminated with light of 1064nm wavelength. When estimating the differential of the

scattering matrix for the GWS we will be using a step size (as defined in Section 4.1) of $\Delta x = 10$ nm, unless specified otherwise. We will use the same step size when estimating the force gradient in the optimiser.

A note on semantics: here we define optical trapping stiffness as a negative quantity, so when we say ‘we want to minimise stiffness’ we mean that we want to find the most negative, not the smallest absolute, stiffness.

We will evaluate any new field that we obtain by comparing it to a standard Gaussian beam trap of equal power. For this purpose we define the enhancement factor \mathcal{R} as the ratio between the stiffness of any new trap we present and the stiffness of the Gaussian trap:

$$\mathcal{R}_\delta = \frac{\kappa_{\delta,\text{field}}}{\kappa_{\delta,\text{Gauss}}}, \quad (5.10)$$

where δ indicates the direction along which the stiffness is estimated, e.g. x .

There will be quite a few colour-map figures in this chapter. Most of these will represent light intensity (never amplitude) either in the pupil plane, or in the focal plane where the trapped bead resides. Figure 5.6 shows these colour-maps: a green colour-map for intensity distributions in the pupil plane, a blue one for intensity patterns in the bead plane for fields obtained with the GWS and a red colour-map for those obtained with the iterative optimiser. The colour in several intensity plots appearing in the same figure will not have the same scaling, unless specified otherwise. When discussing phase we will use a greyscale colour-map where the phase varies from π to $-\pi$.

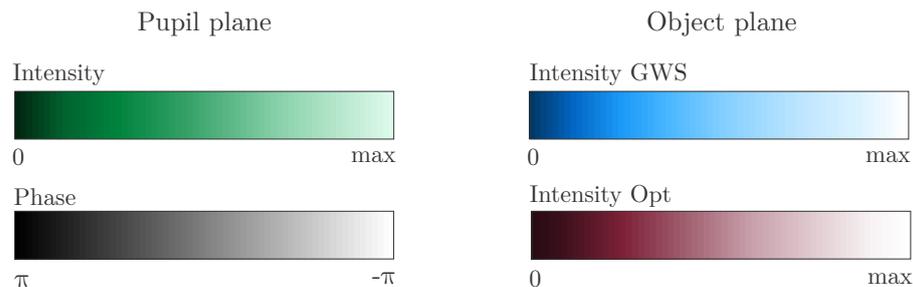


Figure 5.6: The main colour-maps used within Chapter 5.

To visualise the polarisation state of the input light we will use polarisation ellipses. A single polarisation ellipse is traced out in space by the vector describing the electric field magnitude during one cycle of the wave. For example, if the light is linearly polarised, the electric field oscillates along a single axis - its polarisation ellipse is then just a line parallel to this axis, as indicated in Figure 5.7. If the x and y components of the electric field are a $\frac{\pi}{2}$ out of phase and are of equal magnitude, the electric field vector will trace out a circle. For any other phase difference we will have elliptical

polarisation. The colour of each polarisation ellipse indicates the intensity of the light field. This completes our discussion of methods and we move on to implementing them.

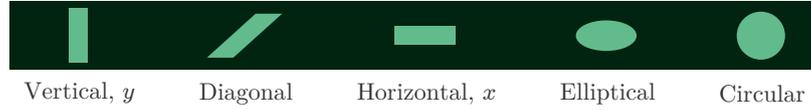


Figure 5.7: Polarisation ellipse representation.

5.2 Optimising 1D stiffness in the transverse plane

We begin by first considering stiffness enhancement along one dimension, in the plane transverse to the direction of the illuminating laser beam, namely the $z = 0$ plane. First, we will examine the performance of the GWS operators, exploring effects of polarisation, and making sure that optimisation behaves in the same manner along different directions in the transverse plane. We will then review the results achievable with the iterative optimisation scheme and compare the two approaches.

5.2.1 Using the GWS

Optimising the force

We first check our implementation of the GWS method by calculating the Q_x operator, as defined in Equation 4.1. The eigenvectors of Q_x , also called its principal modes, describe electric fields, the $|\text{out}\rangle$ state of which is invariant under small target displacements along x . The corresponding eigenvalues λ_x are directly related to the x -force that each principal mode exerts on the particle. We estimate this force F_x for each principal mode of Q_x , and plot the relationship between F_x and $\text{Re}(\lambda_x)$ in Figure 5.8(a). Note, that the force is given in arbitrary units (au), which are the same in all of the figures presented in this chapter, so that they can be compared directly. We observe a strong linear correlation as expected, but we also see manifestations of noise in the form of deviations from the line of best fit. However, we stated already in Section 4.1 that this trend is perfectly linear only if the scattering matrix is unitary, which is not the case for our S . Also, as expected, the minimum and maximum forces correspond to the minimum and maximum eigenvalues of Q_x .

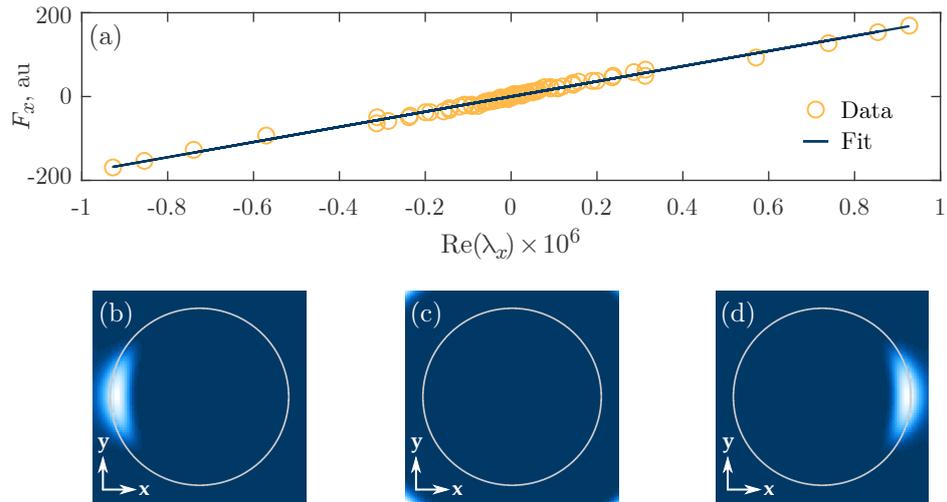


Figure 5.8: Optimising optical force with Q_x . (a) Shows the relationship between the x -force (given in arbitrary units, au) that each principal mode of Q_x exerts on the particle, and the corresponding real eigenvalues λ_x of Q_x , with a linear fit. Intensity distribution in the bead plane is shown for the principal modes exerting the most negative (b), smallest (c), and most positive (d) forces; the edge of the 3 μm bead is indicated with a grey ring.

The light fields that exert the maximum and minimum forces, shown in Figure 5.8(b,d), agree well with intuition we have from a Gaussian beam trap. The light imparts momentum only when it propagates in between media of different refractive indices - this happens at the bead-water interface. The particle will therefore experience the greatest force in the x -direction when the light is focused onto its right edge and the smallest force (which, to clarify, is the largest *negative* force) when the light is focused onto the left edge. For a full picture we also look at the field which would result in the smallest absolute value of the force. As we see in Figure 5.8(c), in this case all the light is simply bypassing the bead.

Optimising stiffness in 1D

We now estimate the operator K_x defined in Equation 4.2 and its principal modes with the aim of finding a light field with the best possible trapping stiffness. The field associated with the most negative eigenvalue is the one with the most negative stiffness. As shown in Figure 5.9(a) this field has a two lobe formation at the inner side of the edge of the bead. The force profile of this trap is shown in Figure 5.9(d), which presents the x -force that the light field exerts on the bead when the bead is displaced by a distance x from the origin. For completeness, we look at the field with the largest *positive* force gradient in Figure 5.9(c,e). We still see a two lobe formation, but now most of the intensity is thrown just outside the edge of the bead. In such a light field

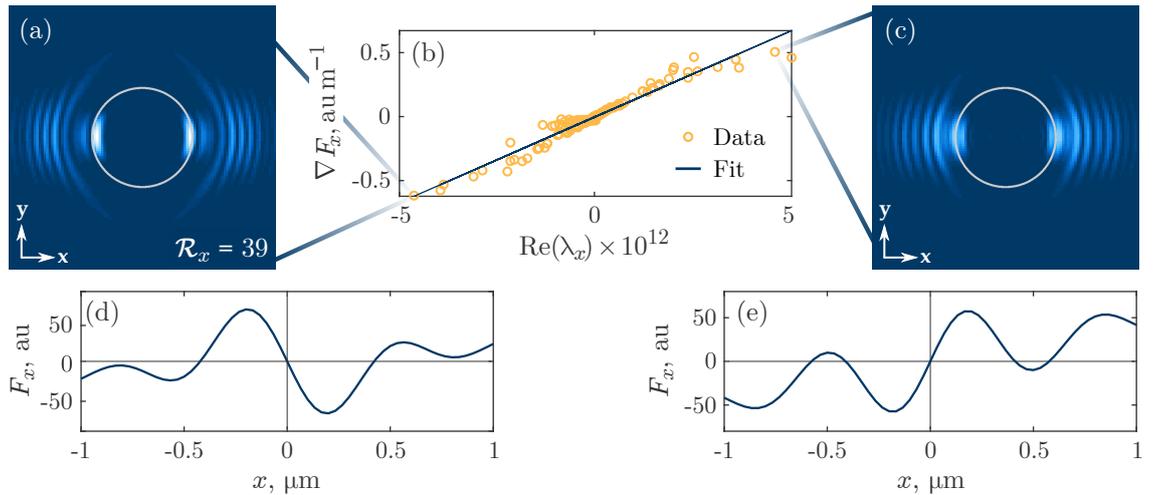


Figure 5.9: Optimising optical trapping stiffness with K_x . The principal modes of K_x with the best (a) and worst (c) trapping conditions and their corresponding force profiles (d,e). (b) Shows the relationship between the x -force gradient that each principal mode of K_x has, and the corresponding real eigenvalues of K_x .

a bead perturbed to move even slightly away from the origin will be drawn into one of the intensity lobes.

Unexpectedly, the field with the largest force gradient does not correspond to the largest eigenvalue of K_x . We also observe more noise in the relation between the force gradient and $\text{Re}(\lambda_x)$ than we did for Q_x . This is perhaps not entirely surprising, given that to obtain K_x we need to evaluate a second order differential and two effective inverses (while Q_x requires a first order differential and one effective inverse). We save a more detailed discussion of non-perfect linearity for Section 5.6.1.

In comparison to a standard Gaussian beam trap (which can be approximated by integrating over the BSCs of Bessel beams in the range from 0 to α_0^3 [163]) the two-lobe trap increases κ_x by 39 times. This substantial improvement comes at the expense of two other important trap features. First is the narrowing of the trapping range, which we define as the region around the trapping equilibrium (which is usually at $x = 0$) between the maximum and minimum forces, as indicated in Figure 5.10. The trapping range of the two-lobe trap is 16 times smaller than that of a Gaussian beam. Second is the reduction in the absolute maximum force that a light field can exert on the particle - this force dictates, for example, at what maximum speed a trap can drag a particle through the surrounding fluid. The maximum force in the two-lobe trap

³ Technically, having a hard cut-off at α_0 will produce an Airy pattern in the focal plane. The central disk of this pattern (where most of the intensity is concentrated) is well approximated by a Gaussian profile.

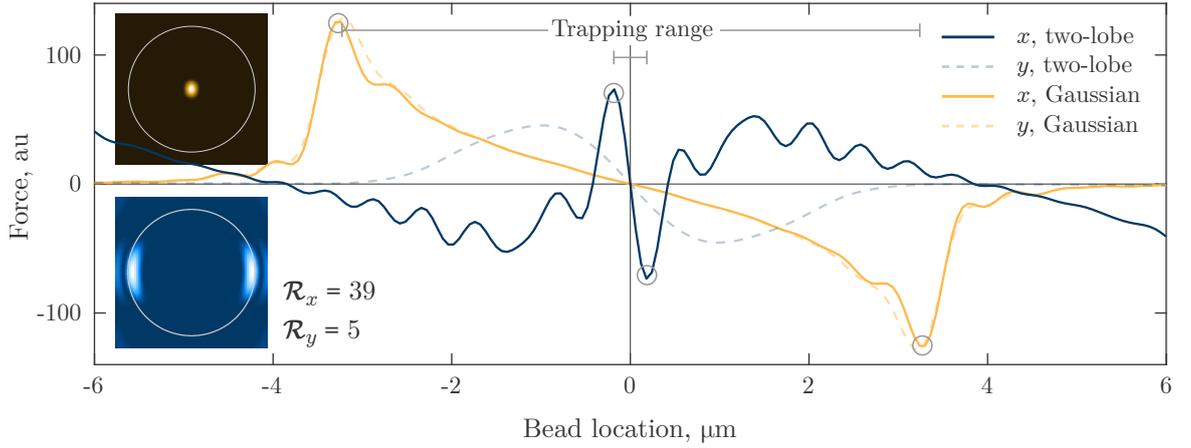


Figure 5.10: Force profiles of the two-lobe and Gaussian beam traps. The blue curves show the force that a bead experiences when it is moved along the x (solid line) and y (dashed line) directions with respect to the two-lobe trap, which is designed to enhance κ_x and is shown in the bottom inset. The yellow curves correspond to the Gaussian trap shown in the top inset. The horizontal grey lines indicate the trapping range for the Gaussian (upper line) and two-lobe traps (lower), and the grey circles point out the absolute maximum force that each trap can exert on the bead.

decreases by 1.7 times and is indicated in Figure 5.10. The two-lobe trap also shows secondary equilibria at $x = \pm 3.8 \mu\text{m}$ where the bead falls into one of the intensity lobes. Another, this time beneficial, byproduct of the two-lobe field designed to enhance κ_x is that κ_y increases as well - by nearly 5 times compared to a Gaussian beam. This trap is, however, not stable in the z -direction.

We now take a closer look at the input modes that make up the optimal field. In Figure 5.11 we see that the two-lobe trap is made up exclusively from modes with odd values of OAM. We observe symmetry in power distribution, as well as phase,

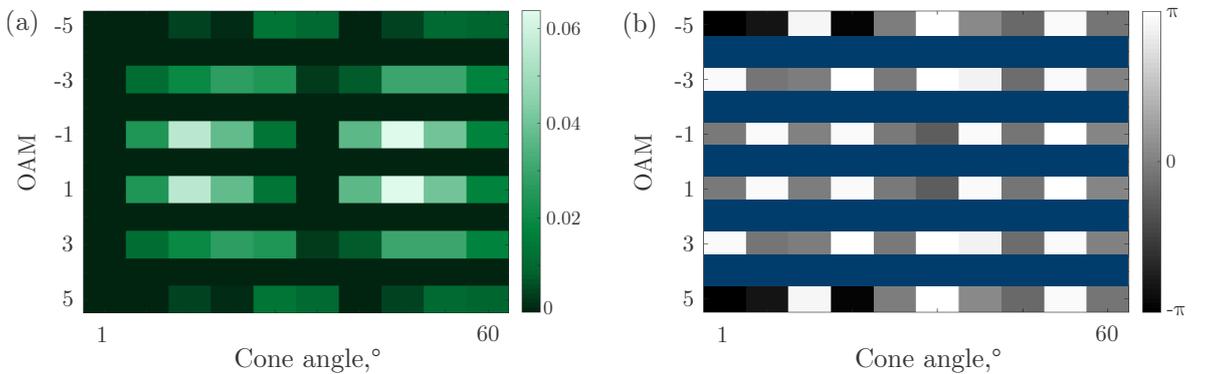


Figure 5.11: Make-up of the two-lobe trap. Heat maps showing the fractional power f_i (a) and phase Φ_i (b) of the input modes that constitute the two-lobe field. The phase of even-valued OAM has been deliberately hidden under the blue lines for ease of viewing, since the amplitude of these modes is zero and does not contribute to the final field.

between pairs of Bessel modes with positive and negative L - this ensures x, y mirror symmetry in the object plane. We explicitly state, however, that the input basis was not constrained in any way to enforce symmetry in the resulting light field - it emerges naturally in the best stiffness principal mode.

Polarisation

We take the opportunity at this point to consider the effects of polarisation. It is a well known fact that linearly polarised light produces the tightest focus along the direction orthogonal to its polarisation. We can see this effect quite clearly in our two-lobe trap. If we constrain our input basis exclusively to x -polarisation (Figure 5.12(a)) the resulting lobes are more ‘blurred-out’ in comparison to y -only polarisation (Figure 5.12(b)).

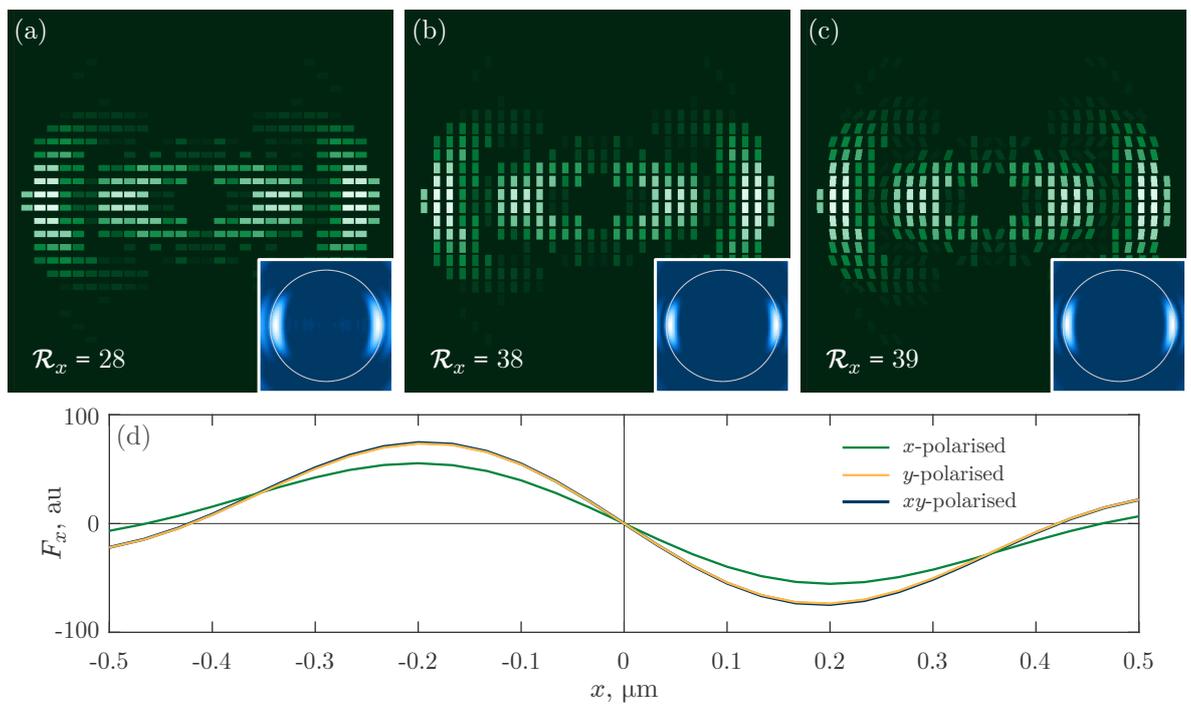


Figure 5.12: Effect of polarisation on κ_x . Shown here are the polarisation state maps in the pupil plane (green colour-map) with the corresponding focal plane intensity distributions (blue colour-map insets) for the highest stiffness principal modes of K_x when polarisation in the input basis is restricted to x (a), y (b), or if it is arbitrary (c). Colour is proportional to intensity. The corresponding force profiles are shown in (d).

Interestingly enough, if we allow the K_x operator to determine how polarisation varies across the pupil, it chooses y -polarised light along the x -axis, and polarisation becomes diagonal as we go further away from the x -axis, thus optimising the tightness of the lobes, as seen in Figure 5.12(c). We see only a very marginal improvement in κ_x

values, though, between y -only and mixed polarisation - they have nearly identical force profiles as shown in Figure 5.12(d). The improvement is 1%, to be exact, while x -only polarisation produces a 27% smaller stiffness than the other two.

Enhancing κ along an arbitrary direction

Naturally, we can use the K operator to optimise along any direction. Two example fields with the highest stiffness along y and $x=y$ directions are shown in Figure 5.13(a,b), which unsurprisingly end up being rotations of the intensity pattern giving the best κ_x . One might expect that the optimal fields along each direction would return identical stiffness, i.e. κ_x obtained with K_x should be the same as κ_y obtained with K_y (provided that polarisation does not have direction-dependent effect in the object plane, e.g. circular polarisation). But this is not the case - going from the x -axis through the $x=y$ axis the stiffness is gradually increasing, reaching the largest value at y -axis with κ_y being 0.8% larger than κ_x . This seems to point towards some systematic issue. We recall that the sampling points in the output basis are not entirely uniformly distributed. In fact, if we look specifically at the points lying in the xy -plane, we see (Figure 5.13(c)) that they are more densely clustered around the $y = 0$ axis - this might be the source of asymmetry in the results produced by K_x and K_y . We test this assumption by rotating our sampling surface by 90° about the z -axis and then repeating the optimisation procedure. The new scattering matrices give us stiffness operators K_x^{rot} and K_y^{rot} . We find that K_x^{rot} returns the same results as K_y did for the original sampling point distribution before it was rotated. The same is true for K_y^{rot} and K_x , which confirms our hypothesis. This is a clear illustration of how the choices made in selecting a sampling basis for the outgoing field affect the final optimisation results. We reiterate, however, that this is only a minor issue and optimisation can be done successfully along any direction in the transverse plane.

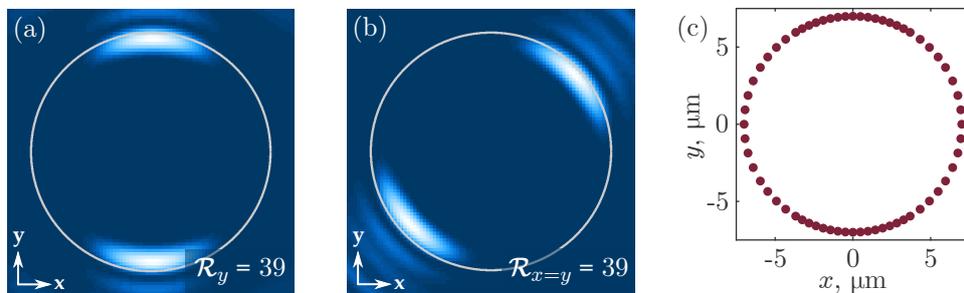


Figure 5.13: Optimising optical trapping stiffness with K . Principal modes with the best optical trapping stiffness along y (a) and $x=y$ (b) directions. Shown in (c) is the non-uniformity of the sampling points on the $z = 0$ plane.

5.2.2 Using the optimiser

Unlike the GWS method, the optimiser requires, and is strongly dependent on, a starting point - the initial light field distribution for the first iteration. Because of the complicated landscape that the optimiser has to explore (a total of 220 optimisation variables if we allow $N_\alpha = 10$ cone angles, OAM in the range from -5 to 5, and mixed polarisation), it is very difficult to predict what would be a good starting point. In Figure 5.14 we present the intensity of four different starting points and the solutions that the optimiser converged to in an attempt to find the stiffest trap along the x -direction. Initially, we took the approach of selecting starting points based on intuition. A Gaussian beam trap (Figure 5.14(a)) is a natural choice as it already has a high trapping stiffness. It can be created by activating only the $L = 0$ modes with a uniform phase (we use y -polarised Bessel beams here). From this initial configuration the optimiser arrived at a solution that spreads out the intensity along the x -axis, but leaves a significant portion of it in the middle. Next, we attempted to mimic the two-lobe structure of the GWS solution, by creating two intensity spots in the focal plane (Figure 5.14(b)), made up of $L = \pm 1$ modes in the pupil plane. This starting point resulted in a solution similar to the one obtained from a Gaussian beam. Both solutions have the same κ_x , which is roughly half of κ_x of the two-lobe trap obtained with the GWS. We have also tried a number of random starting points, one of which is shown in Figure 5.14(c), and which led the optimiser to a light field with an 11 % better stiffness than the previous two starting points. Finally, and somewhat surprisingly, setting the initial intensity to nearly zero (Figure 5.14(d)), resulted in the best solution - the already familiar two-lobe trap. What is particularly interesting, is that the solution

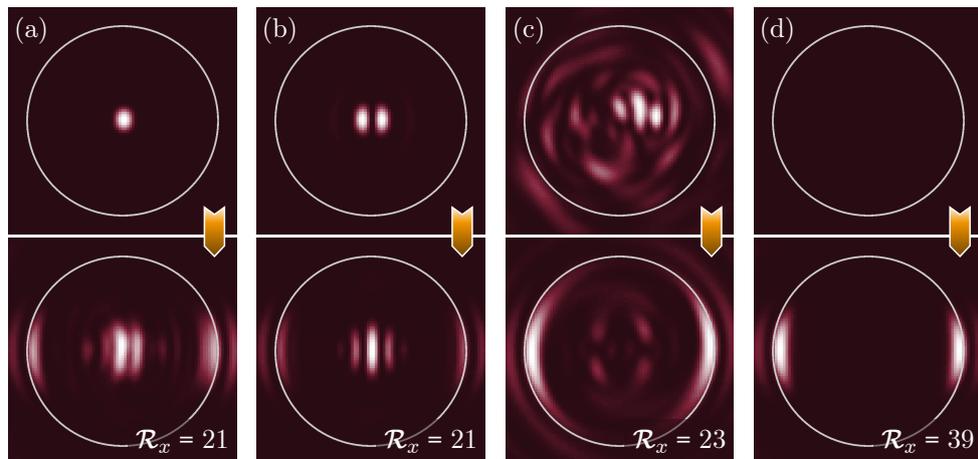


Figure 5.14: Dependency of the optimiser on the starting point: κ_x . The top row shows the starting conditions and the bottom row shows the solutions that the optimiser converged to, in the order of increasing stiffness in the x -direction.

that the optimiser converged to has a 3% larger κ_x than that obtained with the K_x operator. We saw already that our implementation of the GWS operators suffers from noise, which is especially evident in the non-perfect linearity of the force gradient vs. eigenvalue relationship. And now we have clear evidence that the best stiffness solution found with the GWS is not quite the optimum solution. We explicitly state that we do not believe this to be a fundamental issue of the method, but rather an issue of our implementation.

With the most favourable starting point it took the optimiser 38 min to reach the stiffness achieved by the GWS method, and another 114 min to improve this result by 3%. In comparison, finding the optimum field in the GWS framework (with the same number of input modes available) took a total of 11 min. Of this time, 10 min were used to calculate the six scattering matrices, 1 min was used to find the number of singular values which gives the best final result, and 7 s were used in obtaining the optimum field. This makes our implementation of the GWS method 3 times faster than the optimiser.

Unfortunately, the starting point which proved to be the best when optimising κ_x is not equally good for other directions. For example, the best strategy for x (starting with almost zero intensity) does not work nearly as well for y . See Figure 5.15(b), where κ_y of the final solution reaches just about half of that possible in a two-lobe trap. Meanwhile, starting with a bright intensity spot (now x -polarised, shown in Figure 5.15(a)) shows a 32% better stiffness in y than when optimising for x .

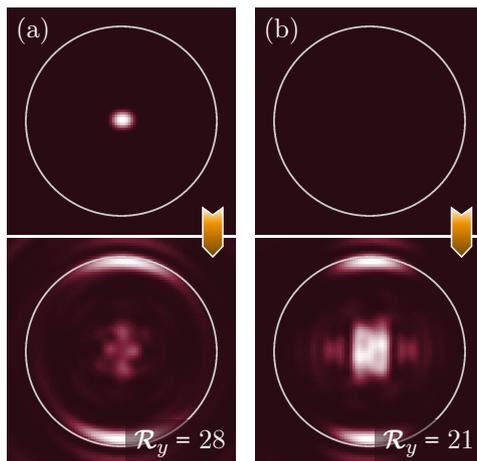


Figure 5.15: Dependency of the optimiser on the starting point: κ_y . The top row shows the starting conditions and the bottom row shows the solutions that the optimiser converged to.

It might seem that the optimiser should arrive at the same solution regardless of the direction along which we wish to optimise. However, while the link between a two-lobe trap aligned on the x -axis and one aligned on the y -axis appears clear to us - a rotation by 90° , from the point of view of the optimiser the two fields are very different. It is the

phase distribution in particular that determines the orientation of the final field, and it is very dissimilar for the two focal-plane intensity distributions we are considering. Therefore, the path that leads the optimiser to the best stiffness in x will be very different from the path that leads it to the best stiffness in y .

We can easily draw the conclusion that, while imperfect in our implementation, the GWS method for optimising stiffness in 1D in the transverse plane surpasses the iterative optimiser, both in terms of the solutions obtained and the computation time required.

5.3 Optimising 1D stiffness in the axial direction

We next move on to the z -direction, where we immediately face challenges. The method of the K operator does not explicitly find the best stiffness - it finds the steepest negative force gradient at $z = 0$. There are no guarantees that the force will cross from positive to negative in this region to create a stable trapping equilibrium. To illustrate this point, we have carried out the optimisation procedure with K_z , and found that the optimum field does indeed have a negative force gradient at $z = 0$, but the z -force itself is positive over the entire diameter of the bead, as can be seen in Figure 5.16(b). An equilibrium does exist, at about $z = -5 \mu\text{m}$ but the stiffness there is two orders of magnitude smaller than in a Gaussian beam trap.

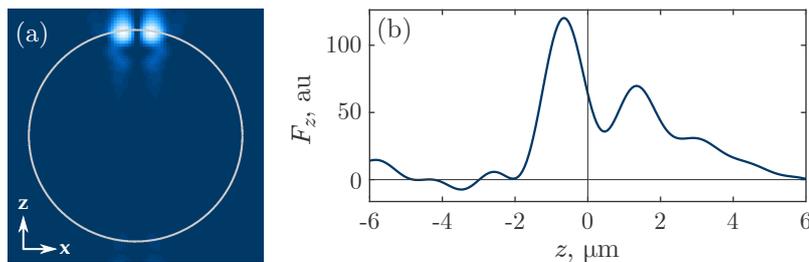


Figure 5.16: Enhancing κ_z with K_z . The principal mode (a) of K_z with the steepest negative force gradient, as seen in the z -force curve in (b).

Instead of optimising the field directly with K_z , we can look for a field with the most negative z -force using Q_z (for a bead located at $z = 0$). Such a field, while it does not guarantee enhanced stiffness, does guarantee a stable trapping location *somewhere* along the z -axis. Intuitively this can be understood by recognising that a bead can only be pulled against the direction of propagation of the beam if it is in a high intensity gradient neighbourhood, where the gradient force is larger than the scattering force. Such intensity gradients are generated near a beam focus. Further away from the focus,

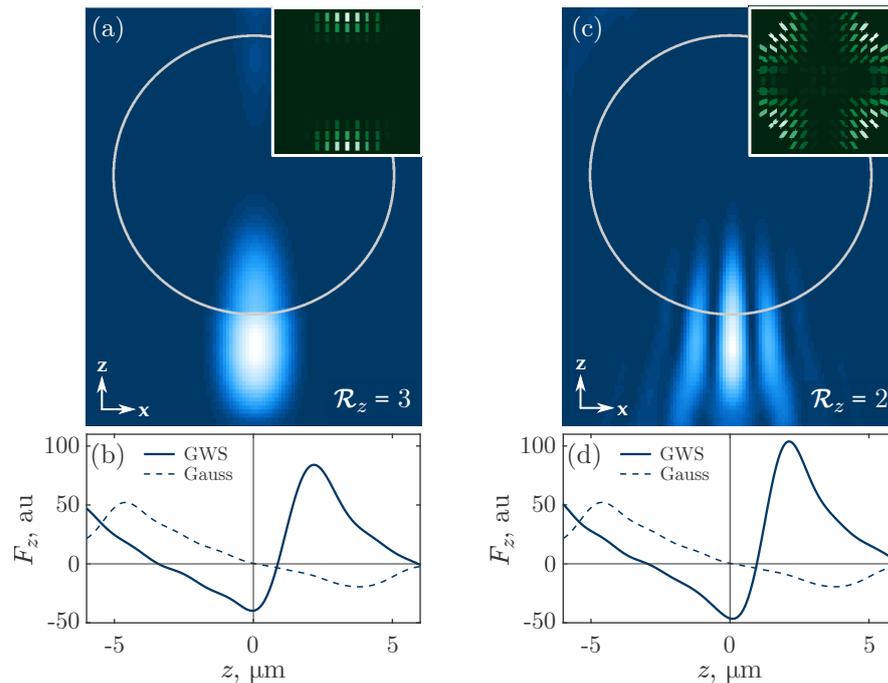


Figure 5.17: Enhancing κ_z indirectly with Q_z . By finding the principal mode of Q_z with the most negative z -force acting on the bead at $z = 0$, we identify a field which will have a stable trapping location. The intensity cross-section along the beam axis (a) for an input limited to y -polarisation (inset), and the corresponding force profile (b) are shown; (c,d) show results for mixed polarisation.

either in front or behind it, the scattering force dominates, pushing the bead along the beam. Since an intensity gradient cannot be maintained along the entire beam, a ‘pull’ region, where the force acting on the bead is negative, must be surrounded by two ‘push’ regions, where the force is positive, thus ensuring the presence of a stable equilibrium. The results with the most negative longitudinal forces are presented in Figure 5.17 for y -polarised (a,b) and mixed polarisation (c,d) inputs. While the mixed polarisation allowed a 17% more negative z -force, the linear polarisation resulted in a 32% higher trapping stiffness at the equilibrium.

Attempts to minimise κ_z directly with the iterative optimiser would run into the same issue as they did with the K_z operator - simply looking for a minimum gradient at $z = 0$ returns a solution which will only push the bead downstream the laser beam. Following the above approach of minimising the force instead, returns solutions which are similar to, but do not quite catch up with the ones obtained using Q_z . With the optimiser, however, we can impose additional constraints to try and enhance the stiffness, as well as have control over where the trapping equilibrium will be. We task the optimiser with minimising the z -force at some (positive) bead displacement Δz while constraining the force at $-\Delta z$ to be equal and opposite. Even if the constraint

condition cannot be met exactly, the resulting field will have an equilibrium in between $-\Delta z$ and Δz . Although, much like with the approach of minimising the force at $z = 0$ we cannot be certain of an increase in stiffness, as we are approaching the problem indirectly via the force.

The resulting field and its corresponding force profile for $\Delta z = 0.5 \mu\text{m}$ are shown in Figure 5.18(a). First, we observe that the absolute forces at $z = \pm 0.5 \mu\text{m}$ are not equal despite the constraint, but this is of little significance to us. The force follows a linear trend in nearly all of the range between $-\Delta z$ and Δz , with a steep gradient enhancing κ_z 16 times, compared to a Gaussian beam. In Figure 5.18(b) we see that using $\Delta z = 1 \mu\text{m}$ did not improve stiffness equally well. The force slope levels out just before crossing zero, leaving the stiffness improvement factor at the value of 3. Just like in the case of the transverse plane, enhanced stiffness comes at the price of a reduced trapping range. In addition to this, these traps have reduced stiffness in y , and are unstable in x . Both results were obtained in half an hour, using y -polarised light with a starting point of near-zero intensity.

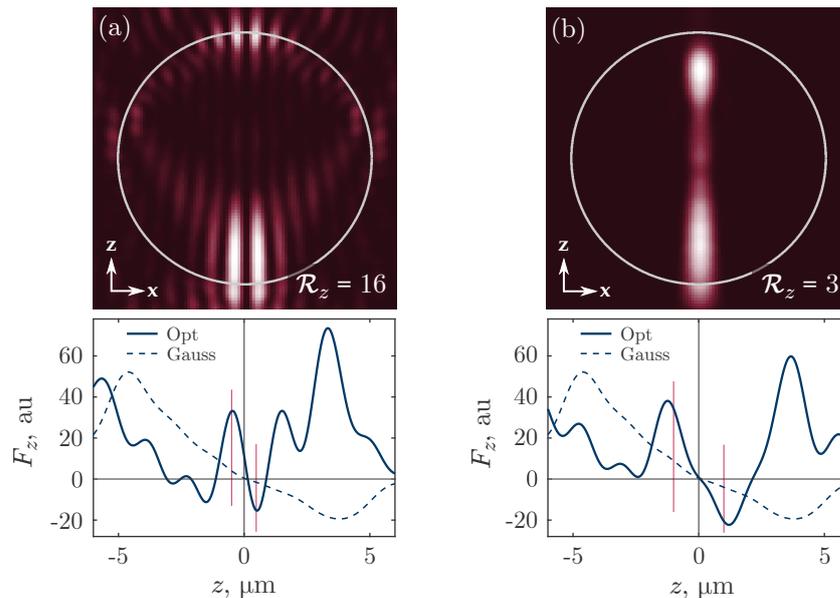


Figure 5.18: Enhancing κ_z with the optimiser. The light fields (top row) and corresponding z -force profiles (bottom row) obtained by the optimiser when minimising $F_z(\Delta z)$ with additional constraints. $\Delta z = 0.5 \mu\text{m}$ in (a) and $\Delta z = 1 \mu\text{m}$ in (b). The vertical pink lines indicate $\pm\Delta z$ and the blue dashed line corresponds to a Gaussian beam.

We can thus draw a conclusion opposite to that in the previous section - when it comes to enhancing stiffness along the axial direction, iterative optimisation is clearly capable of outperforming the GWS approach, by taking advantage of additional constraints.

In fact, the GWS operators are simply not suited for the task of simultaneously enhancing the longitudinal force gradient and ensuring stability at the point of the steepest gradient.

5.4 Multi-dimensional stiffness enhancement

One of the side effects of one directional stiffness optimisation is that trapping along other directions will be affected as well - be it for better or worse. A light field with good stiffness along one direction might not have an equally good stiffness, or even a stable equilibrium, in another. This is particularly true for κ_z when optimising κ in the transverse plane, and vice versa. Since stiffnesses in different dimensions are not independent from one another, the operators K_x, K_y, K_z do not commute and cannot be straightforwardly combined to simultaneously enhance stiffness in multiple dimensions. We therefore take the approach of combining fields which are optimal along each individual direction.

5.4.1 In the transverse plane

Once again, we begin by considering the transverse plane. We saw that, advantageously, the two-lobe trap which optimises κ_x also increases the stiffness in y . But what if we want to have control over the exact ratio between κ_x and κ_y - set it to 1, for example? The first option is to rely on symmetry. If we only allow input modes which have OAM of zero and carry out the familiar optimisation procedure with K_x , y -stiffness will automatically be the same as x -stiffness (assuming circular polarisation). We can see the resulting ring trap in Figure 5.19(a), which is circularly symmetric and has an enhancement factor of $\mathcal{R}_{x,y} = 17$. The inset to the right bottom of the intensity plot illustrates the magnitude of stiffness, which we can see to be the same in all directions.

The option of using only $L = 0$ modes, though, does not give us the freedom of changing the ratio between κ_x and κ_y . Alternatively, we can add the two-lobe solutions obtained with K_x and K_y , expecting minimal interference since they do not overlap a lot in the focal plane. Remembering to halve the power going into each solution (to keep the total power constant), we obtain a four-lobe trap seen in Figure 5.19(b). Here we are allowing the polarisation to vary across the pupil plane as dictated by the individual solutions. We recall that, as we saw in Section 5.2.1, our optimisation with K_x and K_y does not achieve identical stiffness along x and y . This is easily resolved by redistributing the amount of power that goes into each pair of lobes, creating a field that has the

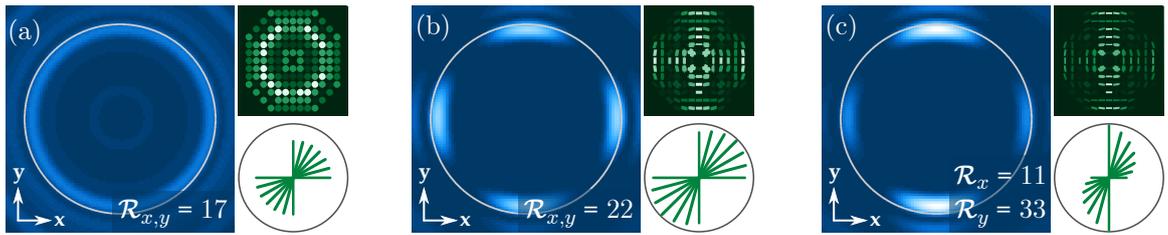


Figure 5.19: Enhancing 2D stiffness in the transverse plane with GWS. In (a) the input basis is restricted to modes with $L = 0$; along with circular polarisation this automatically ensures equal stiffness in all directions. In (b) the optimal fields for κ_x and κ_y are added together to create a four-lobe trap with $\kappa_x = \kappa_y$. And in (c) the power is redistributed so that $\kappa_y = 3\kappa_x$. The colour scaling is the same in the three intensity plots. The lower insets with green lines illustrate the variation of stiffness along different directions in the fields. The length and orientation of each line represent the magnitude of stiffness (not to scale in between (a-c)) and the direction along which it was estimated. The upper insets show polarisation ellipses in the pupil plane.

same stiffness in x and y . But we also observe that the stiffness is ever so slightly greater on the diagonals - by about 1% - making the trap nearly circularly symmetric in terms of stiffness. This also suggests that a more optimal field that simultaneously enhances κ_x and κ_y is this four-lobe field rotated by 45° . Should we need to change the ratio between the stiffnesses, we can simply redistribute the power. For example, in Figure 5.19(c) we see a four-lobe trap with $\kappa_y = 3\kappa_x$.

We expect that simply adding the two lobe pairs will not immediately give the best possible solution. To see if the four-lobe trap can be further improved we set it as a starting point in the optimiser set to minimise (i.e. find the most negative) $\kappa_x + \kappa_y$, while maintaining a ratio of $\frac{\kappa_x}{\kappa_y} = 1$. Initial attempts resulted in asymmetric fields and non-zero optical force at $x, y = 0$. We have therefore constrained the available input modes to odd values of OAM with equal power in opposite L values: $f_i(L) = f_i(-L)$.

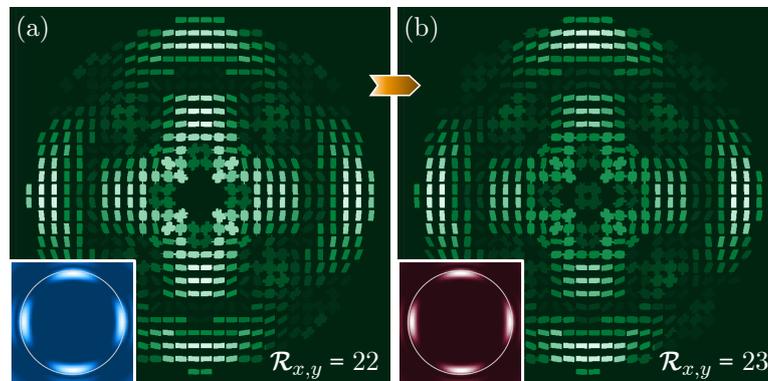


Figure 5.20: Optimising the four-lobe field. The polarisation ellipses of the starting point (a) and the solution (b) of the iterative optimiser. The insets show the intensity in the xy object plane.

This ensures rotational symmetry about the origin point in the xy -plane, and therefore zero transverse force acting on the particle. With symmetry ensured, the optimiser converged to a solution which marginally improved the stiffness sum by 4%, over a duration of 164 min. As seen in the insets of Figure 5.20 the change in intensity distribution in the xy -plane is barely sufficient to be visible. However, we observe in the polarisation maps, that the stiffness improvement was achieved by directing more intensity from the centre of the pupil plane towards the edges.

In terms of average stiffness (estimated along several different directions) the four-lobe trap outperforms the ring trap by 30%. Therefore it seems that it is more favourable to throw intensity into several strategic points, rather than spread it out evenly along the edge of the bead. However, it could also be the case that the four-lobe trap simply uses the available power more efficiently, while the ring trap throws some of it away unused. It is not immediately clear how the two trapping strategies would compare if all of the available intensity could be perfectly concentrated into the trap (provided, of course, that such a trap is compatible with Maxwell's equations). Furthermore, we have allowed varying polarisation in the four-lobe trap, while the ring trap is constrained to circular polarisation. We therefore also investigate a circularly polarised four-lobe trap. We see this field in Figure 5.21(a), which performs 14% worse than its mixed-polarisation counterpart, and 12% better than the ring trap. Further optimising this field returned a somewhat unexpected result, seen in Figure 5.21(b). Unlike in the mixed-polarisation case where the optimiser responded by pushing some of the available power from the middle to the edges of the pupil, we now see the four lobes being merged into two crescents positioned along one of the diagonals. This two-crescent trap improved the sum of κ_x and κ_y by 2.1% over 181 min.

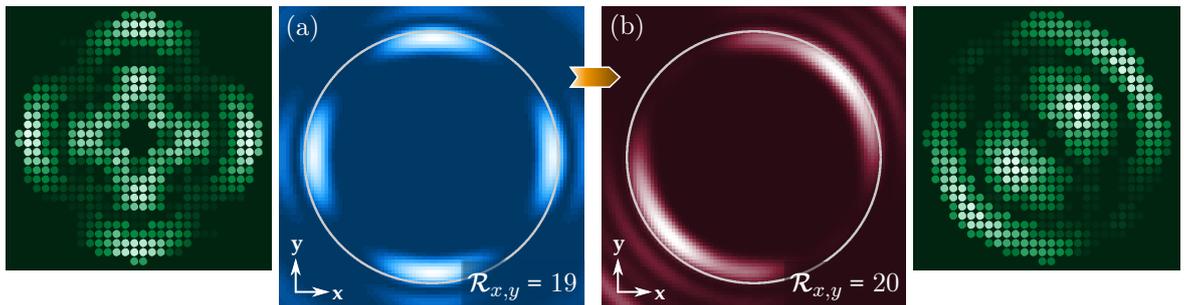


Figure 5.21: Optimising circularly polarised four-lobe field. The intensity in the xy -plane and the corresponding polarisation maps in the pupil plane of the starting point (a) and the solution (b) of the iterative optimiser aiming to improve the sum of κ_x and κ_y while ensuring they are equal.

We have so far used the optimiser to improve solutions that were already of very high stiffness. But what happens if we run the optimiser without the benefit of beginning its

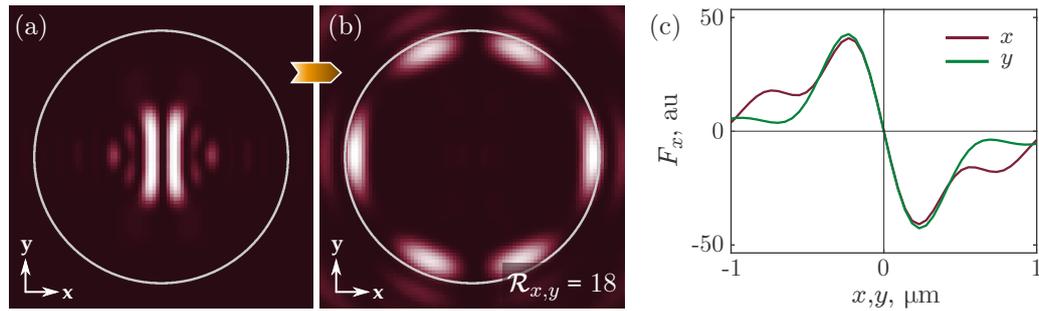


Figure 5.22: Enhancing $\kappa_x + \kappa_y$ with the optimiser. Starting the optimiser with equal power in each available mode (a) to minimise $\kappa_x + \kappa_y$ while requiring them to be equal, converged to a solution (b) with the force profile shown in (c).

journey from an already favourable point? In Figure 5.22 we present results achieved from a starting point where each available mode contains the same amount of power. Working with y -polarised enforced symmetry input, the optimiser was looking to minimise $\kappa_x + \kappa_y$ while keeping them equal. The optimisation took nearly three days to converge to the solution shown in Figure 5.22(b,c). Given that the input is limited to y polarisation, it is not surprising that the solution did not converge to a four-lobe trap, since it would not have equal stiffness along x and y . Instead we see a six lobe formation, with the familiar two lobes along the x -axis and four lobes along the y -axis. The solution enhances the stiffness along both directions 18 times, which is slightly worse than most of the fields we have seen in this section. We thus conclude that, as far as simultaneous 2D stiffness enhancement is concerned, the iterative optimiser does not quite reach the results attainable by combining the optimal x and y -stiffness fields obtained with K_x and K_y respectively, while taking significantly more time to produce a favourable result.

Summary

We summarise the results presented in this section in Table 5.1. To obtain a light field which traps equally well in x and y , we first constrained the symmetry of the field by restricting the input basis to circularly polarised Bessel beams with zero OAM. This ensured that using the GWS method to find the most negative x -stiffness would produce a circularly symmetric trap with isotropic stiffness in the transverse plane. This ring trap, shown in the first column of Table 5.1, offered stiffness enhancement of 17. Next, we approached the problem by superposing two two-lobe fields obtained with the GWS - one which offered the best trapping in x , and one with the best trapping in y . The result is a four-lobe structure, with nearly isotropic stiffness, and is show in

the second column of Table 5.1. The stiffness enhancement of this trap is 22 along x and y axes, with a further 1% increase along the diagonals.

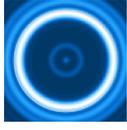
Trap					
$\mathcal{R}_{x,y}$	17	22	23	20	18
Polarisation	C	Mixed	Mixed	C	Y
Figure	5.19(a)	5.19(b)	5.20(b)	5.21(b)	5.22(b)
Duration	—	—	164 min	181 min	~ 3 days

Table 5.1: Summary of optimised 2D trapping in the transverse plane. Listed here are the enhancement factors $\mathcal{R}_{x,y}$ for the traps discussed in this section, along with their polarisation state in the pupil plane (C stands for circular) and the figure in which they appear in the text.

We aimed to improve this four-lobe trap by setting it as a starting point for the iterative optimiser and setting it to minimise $\kappa_x + \kappa_y$. This strategy achieved marginal improvement, bringing the enhancement factor up to 23 - the highest $\mathcal{R}_{x,y}$ we saw in this section. This field has a nearly identical four-lobe structure and is shown in the third column of Table 5.1. The two four-lobe traps take advantage of spatially shaped polarisation structures. To consider situations where polarisation shaping is not available, we used a circularly polarised four-lobe field obtained with the GWS as a starting point in the optimiser, and obtained a two-crescent trap, presented in the fourth column of Table 5.1. This field breaks the stiffness symmetry by concentrating the intensity along one of the diagonals, although it offers a 20-fold enhancement along x and y - a 12% increase compared to the, also circularly polarised, ring trap. Finally, we tested the capabilities of the optimiser without the convenience of having a starting point provided by the GWS. With the initial point where each available y -polarised Bessel mode has equal power, the optimiser converged to a six-lobe structured field, shown in the last column of Table 5.1, with $\mathcal{R}_{x,y} = 18$ - marginally better than the ring trap, but only after having taken nearly 3 days to converge. We conclude that the iterative optimiser on its own is not very well suited to deal with multidimensional stiffness optimisation, but it can accompany our strategy of combining GWS fields optimised for different directions.

5.4.2 Full 3D trap

The only step left now is to find a field which enhances transverse trapping stiffness while ensuring stable trapping along z . Perhaps the most obvious solution to this

problem is having a counter-propagating optical trap setup, where two beams travelling in opposite directions act to cancel out the scattering forces that destabilise the trapped particle in the longitudinal dimension [24, 169, 170]. Such setups are, however, fairly complicated and are not often encountered in labs, with the standard being a single Gaussian beam trap. We will therefore explore the options for achieving 3D stability with a single trapping beam.

Having witnessed the struggle of the optimiser to minimise stiffness along different dimensions simultaneously in the transverse plane, we again take the approach of combining minimally overlapping fields - one of which has a very high stiffness in the transverse plane and the other one which has a stable z -equilibrium. We first return to the field we saw in Figure 5.17(b) in Section 5.3, which minimises the z -force at $z = 0$ and is created using varying polarisation in the far field. We see in Figure 5.23(b) that this field has a four-lobe structure in the xy -plane which is ideally suited to be combined with our four-lobe trap. This particular tractor beam, however, is made up of modes with even-valued OAM, and adding it to our odd-OAM four-lobe field (Figure 5.19(b)) would break the rotational symmetry in the transverse plane of the beam. We therefore construct another four-lobe field with even- L Bessel modes, which can be seen in Figure 5.23(a). This particular trap has a 9% lower stiffness than its odd-valued counterpart.

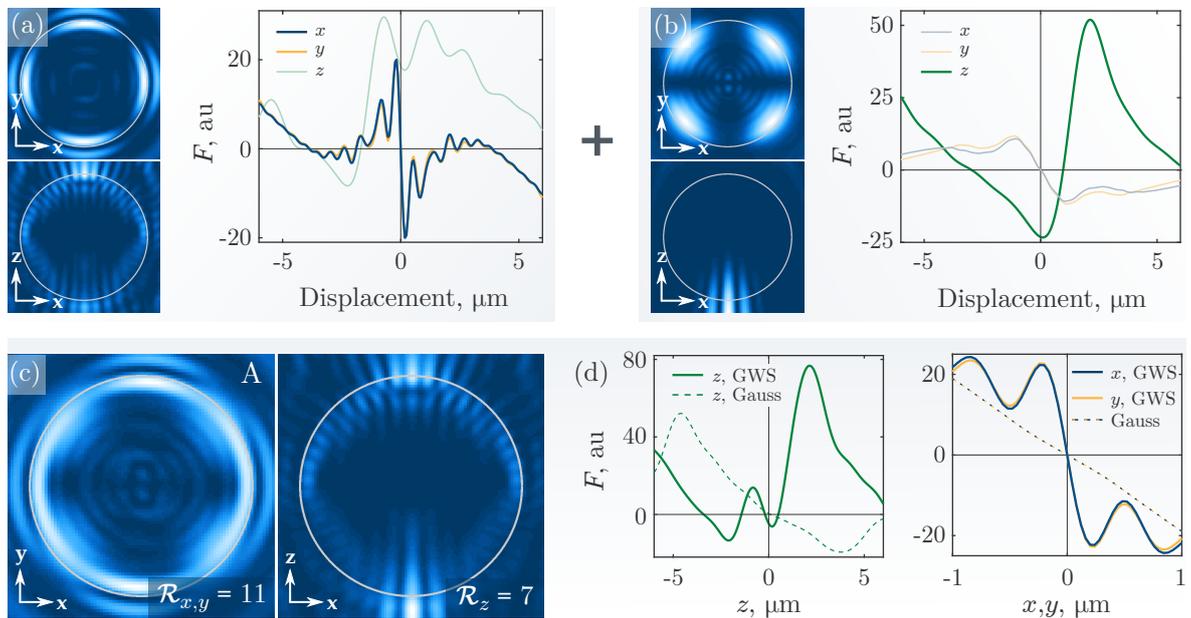


Figure 5.23: Combining fields for stable 3D trapping: A. The transverse and longitudinal intensity distributions with corresponding force curves for: a four-lobe field (a) with high transverse stiffness, a tractor beam (b) with a stable z -equilibrium, and the combination of the two (c,d). In (c,d) the fields and x, y -force profiles were estimated with the bead at the equilibrium.

Presented in Figure 5.23(c,d) is the result of adding the two above fields, with equal power in each - we refer to this new field as field A (or trap A). In the xy -plane we see the intensity being spread out along the edge of the bead, with two more prominent lobes along the y -axis. A more detailed visualisation of how the field changes upon propagation through the bead is presented in Figure 5.24. Compared to the four-lobe trap, $\kappa_{x,y}$ of trap A has been almost halved, as we would expect, given that the power is now shared between two beams. The stiffness in the transverse plane of this field is enhanced by 11 times. In the longitudinal cross-section of the beam we see the intensity being concentrated at the top and bottom of the bead, providing a stable equilibrium at $z = -0.2 \mu\text{m}$. Somewhat unexpectedly, κ_z has reached an enhancement factor of $\mathcal{R}_z = 7$ here, compared to $\mathcal{R}_z = 3$ of the original tractor beam. It seems that the effect of the negative z -force of the tractor beam at $z = 0$ was to ‘pull down’ the z -force curve features of the four-lobe trap (light-green line in Figure 5.23(a)), just enough to create a stable trapping point. We note however, that the z -force profile shows a rather shallow and narrow dip into the negative force region, which begs the question of how stable this trap really is. The distance between the equilibrium and the point of minimum force is about $0.35 \mu\text{m}$ - assuming a linear restoring force model, we can use this information in the equipartition theorem, $\kappa(\delta z)^2 = k_B T$, to determine that stiffness of at least $3.3 \times 10^{-8} \text{ N m}^{-1}$ is needed to ensure that a bead stays in its trap (assuming room temperature). This is well below what is achievable for a Gaussian trap with typical laser powers used in optical tweezers (we recall that stiffness is directly proportional to intensity), and with the enhancement factor of 7 will not pose any issues for stability of our trap.

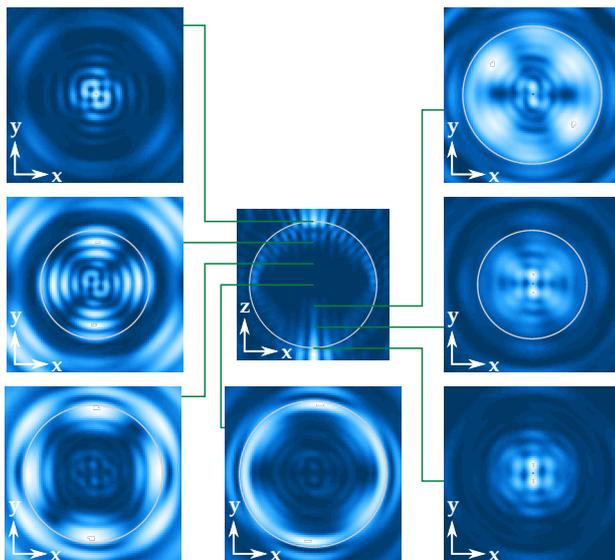


Figure 5.24: Intensity distribution in field A in different transverse planes along the beam axis. Note that intensity is not scaled in between the different planes so that the details of the light field can be easily resolved.

5.4.3 Optimal field for our hydrodynamic actuators

We now return to one of the motivations for Part II of this thesis - enhancing optical trapping stiffness in order to improve the performance of our hydrodynamic manipulation technique. The beads in the constellation rotors are restricted to move along a circular trajectory and therefore only need very high stiffness along the direction of their motion, not the direction orthogonal to it. It is therefore sufficient for us to be combining a two-lobe field with a compatible tractor beam. For this purpose we pick another field, from Figure 5.17(a), which is created with a y -polarised input and is shown again in Figure 5.25(b). Once again, the intensity structure of this tractor beam in the xy -plane is particularly well suited for us, as it has a two-lobe formation orthogonal to the already familiar two-lobe trap, seen in Figure 5.25(a). To our disappointment, adding the two fields, does not produce a z -stable trap. As can be seen in Figure 5.25(d), the z -force curve barely grazes the zero-force line (at the equilibrium closest to zero), which is also accompanied by minimal y -trapping (yellow curve).

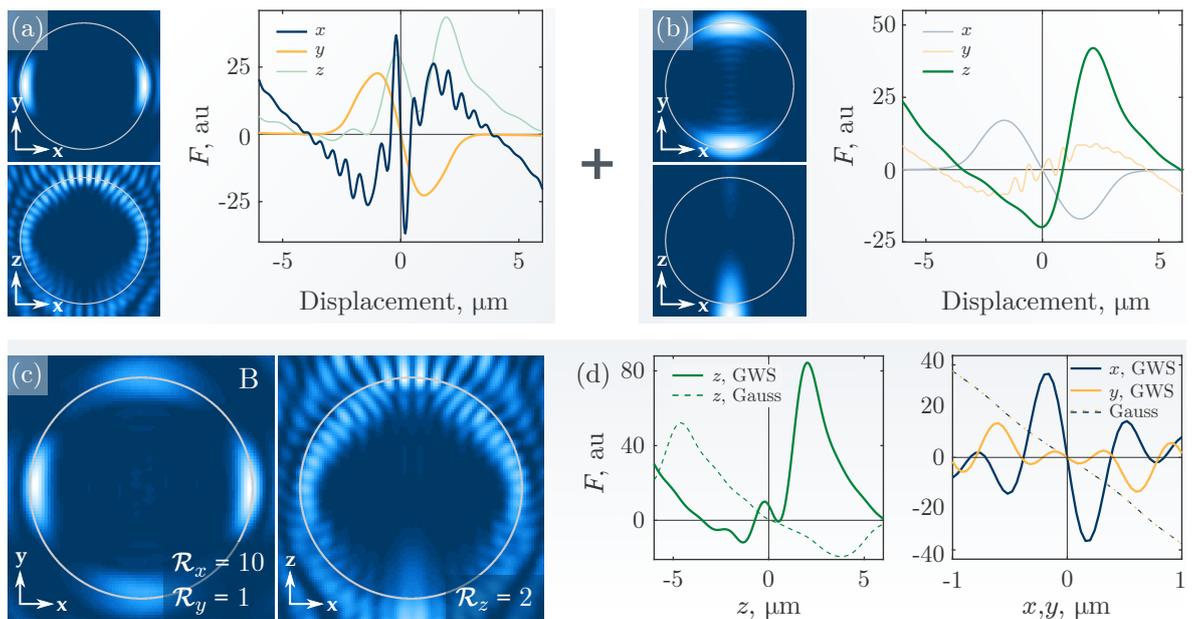


Figure 5.25: Combining fields for stable 3D trapping: B. The transverse and longitudinal intensity distributions with corresponding force curves for: a two-lobe field (a) with high x -stiffness, a tractor beam (b) with a stable z -equilibrium, and the combination of the two (c,d). In (c,d) the fields and x, y -force profiles were estimated with the bead at the equilibrium.

Hence, we turn to yet another z -stable beam, this time one obtained with the iterative optimiser and which first appeared in Figure 5.18(a). From Figure 5.26(b) we see that this field has a two-crescent structure just outside the bead, which overlaps with the two-lobe trap. It is not at all intuitive to predict how the two fields will interfere,

but we find that combining them results in a 3D-stable trap, which we call trap C, with the stiffness being 6 times larger along the main axis connecting the resulting two lobes, than along the line perpendicular to it. To be particular, the two axes exhibit enhancement factors of 12 (along the main axis) and 2 (along the normal to the main axis). The z -stiffness is also significantly enhanced, by 10 times, although we do point out that the peak of the positive z -force just before the equilibrium at $z = 0.1 \mu\text{m}$ is rather low and narrow. At the width of $0.3 \mu\text{m}$ this requires a minimum stiffness of $4.4 \times 10^{-8} \text{N m}^{-1}$ to be stable (using the same argument we did for field A). We also note that following the earlier intuition of using non-overlapping lobes, and rotating the z -stable field by 90° about the z -axis before adding it to the two-lobe field, does not result in a trap stable in 3D.

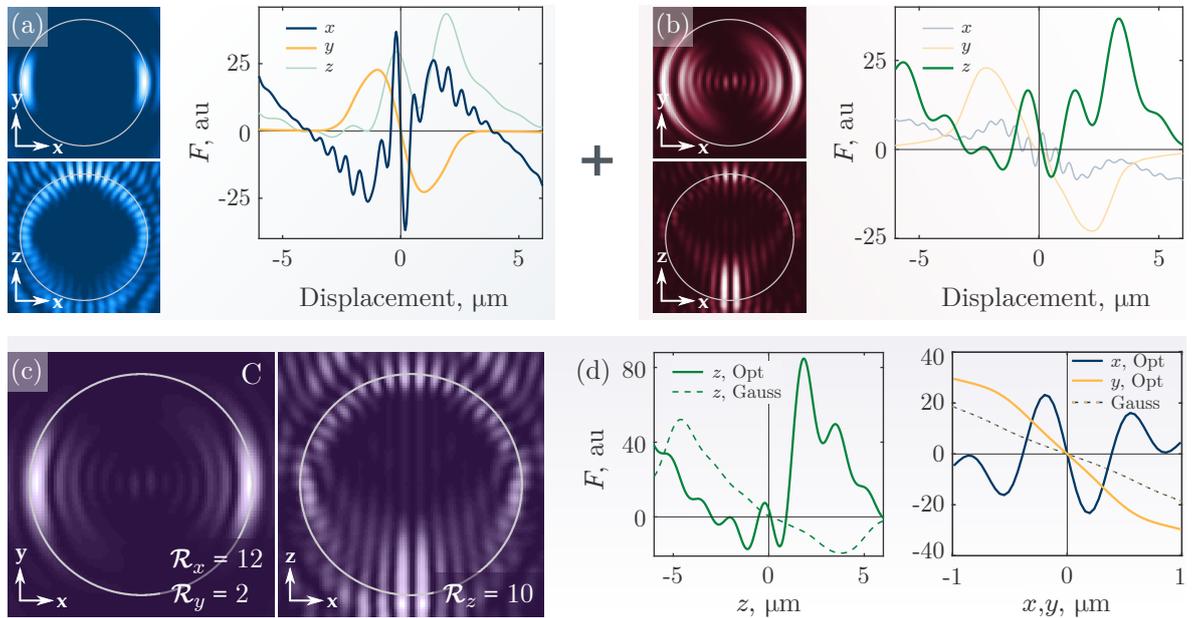


Figure 5.26: Combining fields for stable 3D trapping: C. The transverse and longitudinal intensity distributions with corresponding force curves for: a two-lobe field (a) with high x -stiffness, a beam with high z -stiffness (b) obtained with the optimiser, and the combination of the two (c,d). In (c,d) the fields and x, y -force profiles were estimated with the bead at the equilibrium.

In a hydrodynamic manipulation experiment the above trap could be dynamically rotated to adapt to the direction of motion of the bead in a constellation rotor. We must point out though, that because of polarisation effects this trap would not maintain constant stiffness under rotation. This is illustrated in Figure 5.27 where field C is rotated to coincide with the direction of motion, δ , of a bead in a constellation rotor (showing only one bead); the direction is defined by the angle it makes with the x -axis (i.e. 0° for x -axis, 90° for y -axis, etc.). We see that the stiffness is largest when the bead is moving along x , and reaches less than half of this value when it is moving

along y . This is accompanied by a narrowing of the trapping region - an appropriate amount of laser power would be needed to ensure that the force gradient is sufficient for the bead to stay in the trap. To equalise the performance of the trap in different orientations, it might be possible to distribute the power among the three beads in one rotor, depending on their direction of travel. κ_z , on the other hand, is increased by 40% when the bead is moving along y , compared to when it is moving along x , as seen in Figure 5.27(e). Advantage can be taken of this, by directing more power from the tractor beam part of field C into transverse trapping. Alternatively, the trap could be constructed with circularly polarised light to make it orientation-independent. Such a trap would replicate the enhancement factors of the y polarised trap oriented at 45° : $\mathcal{R}_\delta = 7$ (along the direction of motion of a bead in a constellation rotor) and $\mathcal{R}_z = 12$. Overall, it seems that trap C cannot quite compete with trap A in terms of the enhancement it offers. Field C, however, does not require polarisation control in the pupil to be created. However, if polarisation control is experimentally available, trap C can be adapted so that its polarisation is orthogonal to the direction of travel δ , thus eliminating orientation dependence and maintaining the maximum enhancement factor of 12. This makes trap C a better choice than trap A for our hydrodynamic actuators.

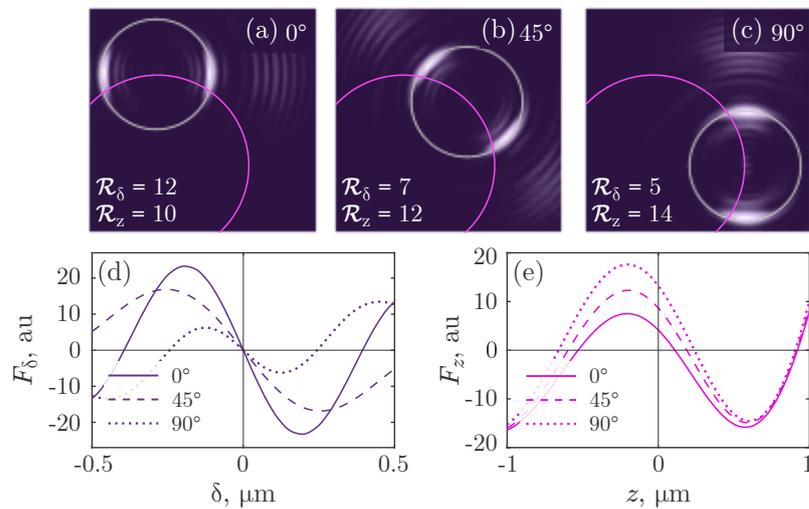


Figure 5.27: Field C in motion. Field C being applied to a bead in a constellation rotor to maximise the stiffness along the direction of motion δ of the bead. The direction is defined by the angle it makes with the x -axis; shown here are: 0° (a), 45° (b) and 90° (c). The pink circle indicates the trajectory of a bead in a constellation rotor. The force profiles along δ (d) and z (e) are shown for the three directions.

Implications for hydrodynamic clamping

But what exactly does such stiffness enhancement mean for hydrodynamic clamping? We take a look back to Section 3.3.3 where we have explored, within simulations, the influence of optical trapping stiffness on the standard deviation σ of a particle, hydrodynamically clamped using two constellation rotors. We reproduce a part of Figure 3.12 here and extend it to include larger absolute optical stiffness values. Assuming that with a conventional Gaussian beam we can trap the hydrodynamic actuators with stiffness of $1 \times 10^{-6} \text{ N m}^{-1}$ (which is indeed similar to the values in our experiments), a 12-fold stiffness enhancement (corresponding to the best case scenario of trap C) would reduce σ by 43%.

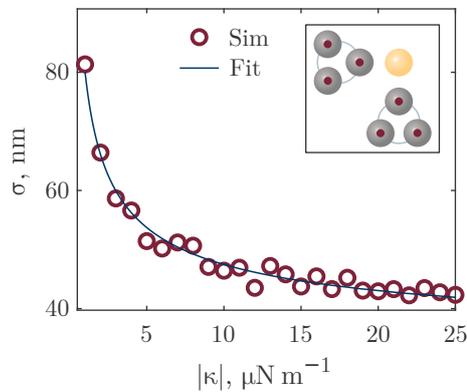


Figure 5.28: Dependence of hydrodynamic clamping on κ . Shown here is the standard deviation σ of a hydrodynamically clamped $5 \mu\text{m}$ radius silica bead, using constellation rotors formed from beads of the same size, obtained from simulated data. A fit of the form $\frac{c_1}{\sqrt{|\kappa|}} + c_2$ with $c_1 = 47.7 (\text{pNm})^{1/2}$ and $c_2 = 32.37 \text{ nm}$ is applied to the data.

However, these simulations assume a linear restoring force profile for trap-bead displacements of up to the *radius of the bead*. Our new traps do not have this feature - yes, we have substantially improved the stiffness, but at the same time we have significantly cut down the maximum force that can be exerted on the bead, by 5 times, as well as the trapping range, by 8 times (comparing the best case scenario of y -polarised trap C and a Gaussian beam of the same power). The maximum achievable optical force directly influences the highest speed at which a bead can be translated. For the simple case of a single optically trapped bead, we obtain the over-damped equation of motion by equating the optical and hydrodynamic drag forces: $F_{\text{opt}} = v/\mu$, where μ is the hydrodynamic mobility of the bead and v is the relative linear velocity between the bead and the fluid. From this we can obtain the maximum velocity at which a bead can travel when being subjected to an optical force F_{opt} . Since angular velocity is directly proportional to linear velocity, the maximum achievable rotation rate with trap C is five times smaller than that achievable with a Gaussian trap.

The above might sound like bad news, especially for applications which require the actuators to reach high rotation rates. For example, long-range speedy translation of a

single target across the holding sample cell or hydrodynamic manipulation of multiple targets, would very likely be worse off with our enhanced stiffness traps than with a Gaussian beam. On the other hand, the reduction of the maximum force that a trap can exert should still allow for rotation rates of $\sim 60^\circ \text{ s}^{-1}$ in a setup with two 3-bead constellation rotors. While it is not straightforward to quantify the combined effect of reduced rotation rate and increased trapping stiffness without actual experimental data, we believe that trap C should still perform well in hydrodynamic clamping and local translations of a single target bead.

We recall that the GWS method can be used to optimise the force, instead of stiffness, like we demonstrated in Section 5.2.1. It might seem that such force-optimised fields should be suitable for applications where the hydrodynamic actuators need to achieve high rotation rates. However, the fields we presented in Figure 5.8(b,d) only offer force enhancement of 1.34 compared to a Gaussian beam - this number would very likely drop even lower after making sure that the fields are stable in z . Therefore, there does not seem to be an advantage in using the force-optimised fields instead of the standard Gaussian traps.

5.5 Phase-only optimisation

The traps that we have demonstrated so far have non-trivial spatially varying intensity and phase structures in the far field. Creation of such fields would, therefore, require both amplitude and phase modulation in the pupil plane. While experimentally it is certainly possible to do both, conventional holographic optical tweezers setups favour phase-only modulation. In this section we will explore the options for stiffness enhancement with phase-only modulation. Here we will only be presenting preliminary results for circularly symmetric fields (created using only Bessel modes with $L = 0$), as this work is currently ongoing.

We will consider two approaches for enhancing stiffness by phase-modulating a uniform intensity beam. The most straightforward thing to do is to pick a target field (i.e. one of the optimum fields that we have demonstrated in this chapter) and imprint its phase onto a uniform amplitude beam. This should recreate the intensity distribution in the focal plane fairly well. Alternatively, we can devise an optimiser which directly optimises the phase of a uniform intensity beam to enhance optical stiffness - this strategy has been employed for 1D stiffness enhancement in [141] and [142].

First of all, we need to represent a uniform intensity beam in the Bessel basis. Since our input consists only of $L = 0$ modes, this task is fairly undemanding. Each mode already has uniform intensity and phase in the far field, we just need to make sure that the intensity is normalised with respect to the area that each mode takes up in the pupil (i.e. modes with a larger cone angle have greater area associated with them, so they need to carry more power). Once this is done, the phase of each mode can be changed as desired.

Our first target field of choice, T_{xy} , is the ring trap demonstrated in Section 5.4.1, which exhibits a 17-fold enhancement in the radial (transverse) stiffness, compared to a Gaussian beam. The far field intensity and phase structure of this field can be seen in Figure 5.29, along with the fields produced by phase modulation. Imprinting the phase of this field onto a uniform intensity beam (‘ T_{xy} phase’ in the figure) results in a 24% loss of stiffness, although this still represents a 13-fold enhancement factor. A phase-only optimiser (started with uniform phase across the pupil) converged to a solution which is 12% behind the original target field, making the phase-only optimiser more beneficial than simple target phase imprinting.

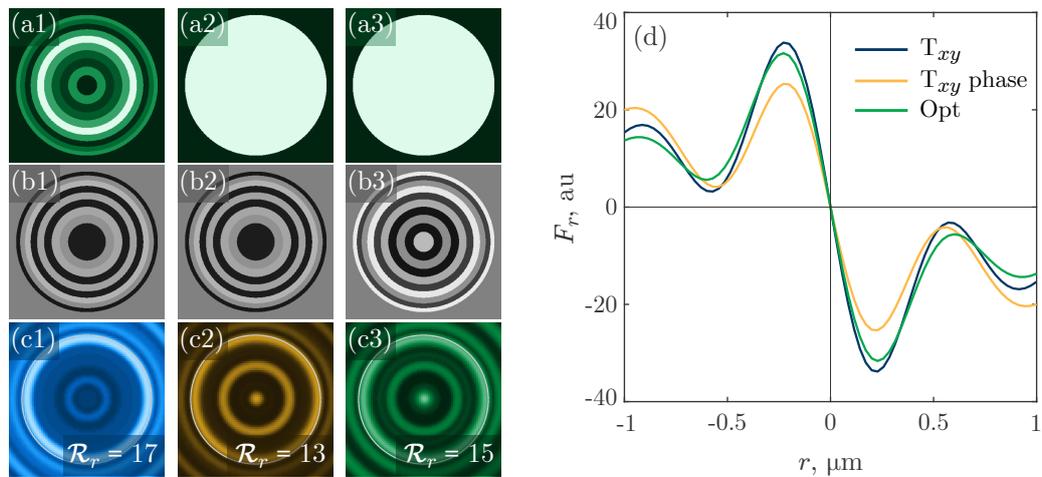


Figure 5.29: Phase-only stiffness enhancement in the xy -plane. The far field intensity (a) and phase (b) distributions followed by xy -plane intensity in the focal plane (c) of the target (1), target phase (2) and phase-only optimiser (3) fields, with the corresponding radial force profiles (d).

We next investigate how well phase-only stiffness enhancement works in the longitudinal direction. Our second target field, T_z seen in Figure 5.30(a1-c1), was obtained using the original optimiser (with y polarised input) with the technique described in Section 5.3 and has the enhancement factor of $\mathcal{R}_z = 13$. Taking the phase of this field and applying it to a uniform amplitude beam nearly halves the z -stiffness, as well as changes the equilibrium location. The phase-only optimiser shows a substantial

stiffness enhancement with $\mathcal{R}_z = 11$, leaving it at a 15% disadvantage compared to the target field. We also note that the corresponding transverse-stiffness enhancement factors are: $\mathcal{R}_x = 13$ and $\mathcal{R}_y = 5$ for T_z , $\mathcal{R}_x = 5$ and $\mathcal{R}_y = 7$ for T_z -phase, and $\mathcal{R}_x = 10$ and $\mathcal{R}_y = 8$ for the phase-optimised fields.

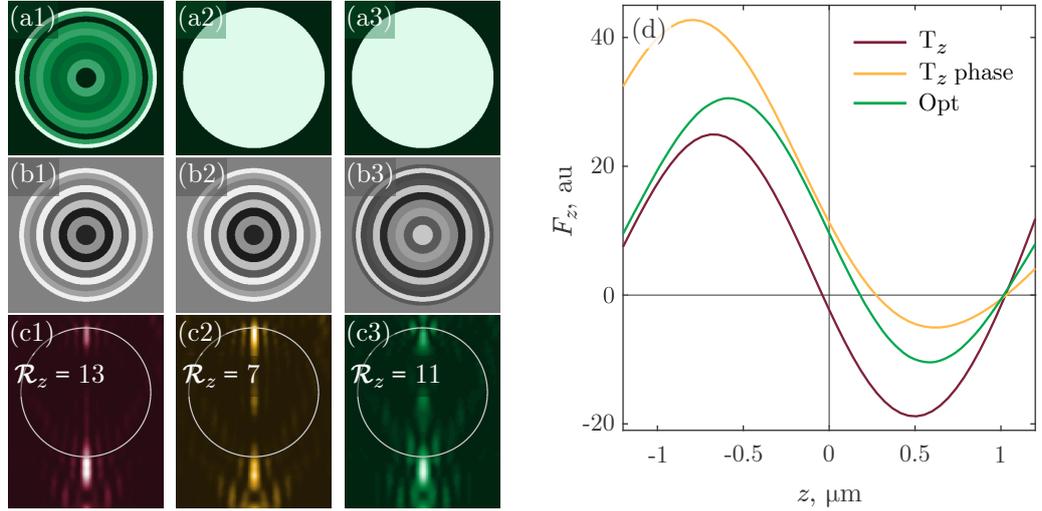


Figure 5.30: Phase-only stiffness enhancement along z . The far field intensity (a) and phase (b) distributions followed by xz -plane intensity in the focal region (c) of the target (1), target phase (2) and phase-only optimiser (3) fields, with the corresponding z -force profiles (d).

Given that in the above simulations the input modes were limited to a small selection because of the $L = 0$ condition, the phase-only optimiser took only a few minutes to converge to a solution. We expect that with a larger basis the convergence time would extend into several hours. While it is true that phase-only methods do not quite replicate what is achievable with full wavefront control, they still offer substantial enhancement factors. How well phase-only approaches will work with more complicated input bases remains to be seen.

5.6 Discussion

5.6.1 Irregularities in our GWS operators

Despite having achieved substantial stiffness enhancement with the GWS approach, we are troubled by the non-perfect linearity in the relationship between the eigenvalues λ of the K operator, and stiffness κ (the same is true for eigenvalues of Q and force, though to a lesser extent). As discussed in the previous chapter, if the scattering matrix S is

unitary, this trend should be perfectly linear. Our S is not unitary, firstly, because it is not a square matrix. However, our collaborators at the Vienna University of Technology have assured us that this should not, in principle, be a problem, and perfect linearity can be demonstrated with a rectangular S . And yet, instead of observing a perfect linear trend between λ and κ , or λ and F , we are seeing the relationship take the form shown in Figure 5.31, in the case of operators Q_x and K_x . We evaluate the linearity of the trends by using the coefficient of determination, R^2 , metric⁴. The closer to 1 the R^2 value, the more linear a trend is. Apart from small deviations from the line of best fit we can observe that more data points are accumulating around zero. This can be explained by the fact that many fields can be generated which will miss the particle altogether, exposing it to very small optical forces, while only a few fields capable of exerting large forces can be created.

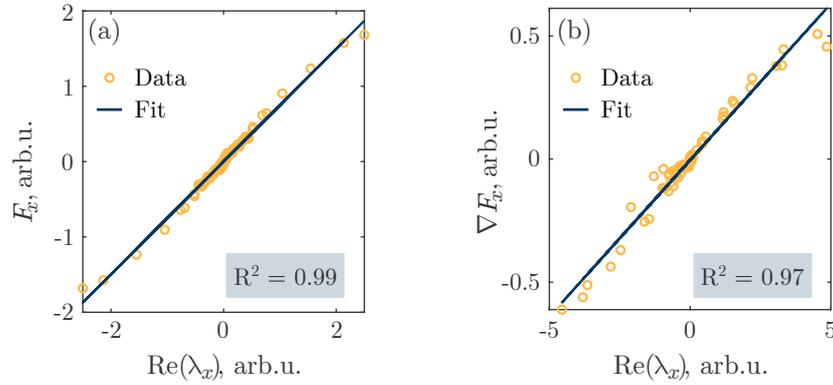


Figure 5.31: Non-perfect linearity: x . The relationship between eigenvalues of Q_x and F_x (a), and the eigenvalues of K_x and κ_x (b). Obtained with y -polarised input.

We have also observed that the κ_z vs. λ_z and F_z vs. λ_z trends are noticeably different from the ones in the transverse plane. First of all, we see in Figure 5.32(a) that the z -force is not symmetric about zero. This is not surprising, since the light is incident on the bead from one side only, and therefore the maximum pushing force will always be greater than the maximum pulling force. Second, we observe some clustering of the data points, which we did not see in the transverse plane (apart from the already discussed zero-force/stiffness). In particular, there is a big gap in F_z from 5 to 12. It may just be the case that none of the eigenvectors of Q_z represent fields which can exert z -forces in this range, given the particular selection of Bessel modes available in the pupil plane. The deviation from the linear fit is not significantly worse here than it

⁴ $R^2 = 1 - \frac{\sum_i (y_i - f_i)^2}{\sum_i (y_i - \bar{y})^2}$, where y_i is the value of the i -th data point, f_i is the fit value and \bar{y} is the mean of the data set y .

was for the x -force, and perhaps the clustering is caused by the one-directional nature of the incoming light.

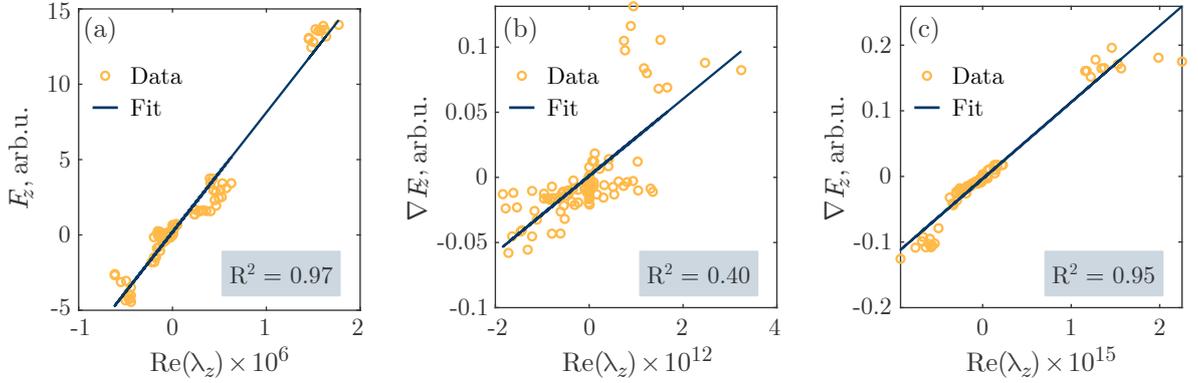


Figure 5.32: Non-perfect linearity: z . The relationship between eigenvalues of Q_z and F_z (a), and the eigenvalues of K_z and κ_z , using $\Delta z = 10$ nm (b) and $\Delta z = 500$ nm (c). Obtained with y -polarised input.

A surprising thing is that we get the best results, in terms of linearity and the largest negative and positive force gradients, with K_z when we use a very large step size for evaluation of the differential of S . Comparing the trends in Figure 5.32(b,c) we see that using a step size of $\Delta z = 10$ nm produces both less linearity and a decrease in the minimum and maximum values of the force gradient, than when $\Delta z = 500$ nm is used. It is not clear to us why this is the case, and we note that for F_z we do not observe an improvement with an increased step size. But why are these deviations from a straight line present at all, for both x and z directions?

We have ensured that the scattering matrix we are working with is full rank, by selecting an appropriate number N_α of the cone angles from which the Bessel modes are originating. If too close a spacing is allowed between neighbouring cone angles in the pupil plane, their scattering looks too similar on the sampling sphere (for identical values of OAM), which makes S rank deficient. Since we made sure that degeneracy is not an issue, we believe that the causes of non-perfect linearity between λ and κ are to be found in the output basis.

An example reported in literature [146] of how perfect linearity can be broken, is when not all parts of the scattering matrix are available, e.g. if only the transmitted or only the reflected light can be experimentally measured. Under such circumstances the eigenvalues of the K operator no longer correspond exactly to the stiffness, although the two quantities retain correlation. It therefore seems plausible that we are not sampling the field densely enough. However, even if the number of sampling points is increased four times up to 10 242, we do not observe improvement in λ vs. κ trend.

Another issue is, of course, that the distribution of the sampling points is not perfectly uniform, particularly in the $z = 0$ plane, as we have demonstrated in Section 5.2.1 in Figure 5.13(c). The areas on the sampling sphere where the points are closer together, are in a sense over-represented. Perhaps weighting each point with its associated area would fix our issue.

It might be the case that we would obtain better results if instead of discretely sampling the field we were to define the output directly in terms of the BSCs. This approach would require implementation of basis translation matrices; otherwise the scattering matrices $S(\Delta x)$ corresponding to different locations of the particle would be in different bases and could not be combined in a calculation.

It is also possible, that defining the output $|\text{out}\rangle$ in terms of the external and scattered electric fields, $\mathbf{E}_{\text{ext}} + \mathbf{E}_{\text{scat}}$, may not have been the best choice. Alternatively, within the GLMT, the electric field can be written in the convention of incoming \mathbf{E}_{in} and outgoing \mathbf{E}_{out} fields, where \mathbf{E}_{in} contains the part of the field which carries momentum towards the scattering particle and \mathbf{E}_{out} carries momentum away from it. It may well be the case that it is the \mathbf{E}_{out} part of the field that is needed for the GWS operator to work properly.

My supervisor Jonny Taylor has implemented the GWS operators in the incoming-outgoing field convention, but where both the input and output bases are the BSCs of the VSWFs (unlike the Bessel basis, this setup allows the light to enter the system from all directions, not limiting it to the pupil of an objective lens). In this description, he successfully recreated perfect linearity between the eigenvalues of the Q, K operators and the force and stiffness. We do not know for certain that combining this output basis with a Bessel input basis would still maintain the perfect linearity, but we do expect that to be the case. Further work should hopefully allow us to identify which of the reasons mentioned above contributed to the non-perfect λ vs. κ trend: non-uniform field sampling, failure to capture sufficient information about the $|\text{out}\rangle$ state, or the choice of sampling all of the field instead of the part which is carrying the outgoing momentum. Despite all the issues mentioned above, we have presented results which, without a doubt, demonstrate substantial stiffness enhancement - pointing towards graceful, rather than catastrophic failure of the GWS method under imperfect conditions.

5.6.2 Summary of results

Throughout this chapter we have demonstrated optical trapping stiffness enhancement in multiple dimensions. Employing the GWS approach we have first identified the optimal fields for trapping a spherical particle in the transverse plane. With their anisotropic two-lobe intensity profiles, shown in the first column of Table 5.2, these traps offer stiffness enhancement along the main trapping direction of 39 times, accompanied with a 5-fold increase along the perpendicular direction, compared to a conventional Gaussian beam trap. These traps, however, do not offer stable trapping in the longitudinal direction, meaning that in practice they would have to be used with very low laser power, so as to allow gravity to counteract the scattering force, as was done in [141, 142]. Given the proportionality between stiffness and laser power, a relative enhancement under these circumstances would not actually offer very high absolute stiffness.

The GWS does not allow for direct stiffness improvement along the direction of propagation of the trapping laser beam, as the principal mode with the steepest negative force gradient is not inherently constrained to contain a point of zero-force, where a particle could be stably trapped. We have presented results showing some success by approaching the stiffness problem indirectly, by searching for a field with the most negative z -force. Such a field ensures a stable z -equilibrium somewhere, but we do not have control over the exact location of it. With this method we were able to triple the longitudinal stiffness. What proved to be a more useful tool here, was an iterative optimiser. The functionality of additional constraints allowed us to shape the z -force profile so that it contains an equilibrium near a specified origin. Along with this came a 13-fold enhancement of longitudinal stiffness. But this particular trap, shown in the second column of Table 5.2, did not offer stability in x .

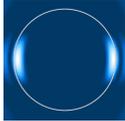
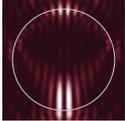
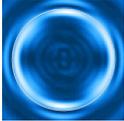
Trap				
\mathcal{R}_x	39	Unstable	11	12
\mathcal{R}_y	5	0.1	11	2
\mathcal{R}_z	Unstable	16	7	10
Polarisation	Mixed	Y	Mixed	Y
Figure	5.9(a)	5.18(a)	5.23(c)	5.26(c)

Table 5.2: Summary of the main traps in Chapter 5. Listed here are the enhancement factors \mathcal{R} for some of the traps presented in this chapter, along with their polarisation state in the pupil plane and the figure in which they appear in the text.

We approached the problem of creating a 3D-stable trap, by adding the fields which were individually optimised for each dimension. For this purpose we selected the fields which minimally overlap in the focal plane. However, we have also found, with the example of field C which merged two fields with overlapping intensity structures, that this is not a requirement for successfully combining different fields. We do not deny that some amount of luck plays a part in this approach, as it does not generally ensure that the resulting field will be 3D-stable and of high stiffness (consider field B, for example). Nevertheless, we were able to demonstrate traps, which are not only stable, but also enhance stiffness in all three dimensions. Field A, shown in the third column of Table 5.2, which took advantage of four-lobe structures and mixed polarisation input, achieved an enhancement of 11 in the transverse plane and 7 in the longitudinal direction. Field C, shown in the last column of Table 5.2, broke the symmetry between x and y by overlapping two-lobe fields to demonstrate a 12-fold enhancement in x , and a 10-fold enhancement in z directions. The biggest disadvantage of the fields we have presented is the reduction in the maximum force that they can exert on the trapped particle. These limitations make the fields best suited for applications which either require stationary trapping or fairly low travel speeds.

It is worth here to discuss the evaluation of the stiffness enhancement factor \mathcal{R} . When comparing our traps to a standard Gaussian beam, we have assumed that all the available laser power can be successfully transferred into a trap, which is not the case in reality. To achieve complex amplitude shaping with a single SLM it is necessary to spatially modulate the diffraction efficiency of the SLM [171] (which can be achieved by changing the depth of the hologram pattern and thus redirecting light from the desired regions of the first order into the zero order). The only way to then increase the efficiency of power transfer is by reducing the fidelity of the light shape being created. There is, of course, the option of using two coupled phase-only devices to achieve complex amplitude shaping [172], but such a setup is challenging to implement experimentally. Beam shaping required to create the non-trivial intensity distribution in one of our traps inevitably incurs loss of power - thus experimentally evaluated enhancement factors would be smaller than the ones presented here. However, even if substantial amounts of power are lost, say 50% when comparing Gaussian and our traps, our theoretically estimated values of \mathcal{R} suggest that we would still see a several-fold stiffness enhancement.

We have also demonstrated stiffness enhancement which does not require shaping of the complex amplitude, but relies only on phase modulation of the laser beam. Preliminary results suggest that high enhancement factors should still be attainable, even if not fully reaching (12-15% difference) those possible with full intensity and phase control.

5.6.3 Applications

We envision that multi-dimensional stiffness enhancement will be of great benefit in many different optical tweezers applications, a few of which we discuss here.

Trapping particles with high refractive index (such as titanium micro-spheres with $n = 1.73$ in water [47]) is troublesome, because they experience a high scattering force which is no longer balanced out by the gradient force, thus leaving the trapped particle without an equilibrium in z . Anti-reflection coatings have been employed to minimise the scattering force [47, 173, 174], but coating particles is a fairly involved procedure. The high z -stiffness traps that we demonstrated could prove useful for trapping high refractive index particles, as they increase the gradient force along the z axis which counteracts the scattering force. Care would have to be taken, of course, to ensure that these beams offer stable trapping in the transverse plane.

Optical tweezers are often used as force sensing instruments in biomolecular studies [175–179]. In such experiments one end of a molecule under investigation, e.g. a DNA spiral, is attached to an optically trapped bead, while the other end is held stationary. By moving the optical trap, force is applied to the molecule, which can be estimated by tracking the position of the trapped bead. The spatial resolution is inherently limited by thermal motion and can be improved by increasing the stiffness [180], which is where stiffness enhancement can be beneficial. Although this is accompanied by loss of sensitivity [180].

With appropriate modifications, the GLMT can be extended to model light interaction with non-spherical particles, opening up a whole new set of applications. Previously, stiffness of micro-rods, optically trapped with two Gaussian beams positioned at both ends of the particle, has been optimised [181]; with our technique this could be taken a step further, to find a whole new field that is particularly suited for trapping rod-like particles. Micro-rods would be well suited for force sensing applications as they exhibit smaller thermal motion than spheres of the same diameter [181]. Other micro-tools of more complex shapes are also frequently optically trapped for applications such as imaging surface topography [45], measuring fluid viscosity and detecting vibrations [182], assembling complex micro-systems [183], force sensing [184] and transporting materials e.g. for targeted drug delivery [185]. Such micro-tools would benefit from custom-tailored light fields that our technique can offer.

We reiterate that the enhanced stiffness fields we demonstrated can be used to trap and control the hydrodynamic actuators described in Chapter 3, provided that the rotors are used in situations where they are not required to achieve high rotation rates, e.g.

hydrodynamic clamping or short-range translation of a single target particle. Given the flexibility of SLMs, it would be possible to easily change the trap based on what task the trapped particle is required to carry out. We also discussed briefly in Chapter 3 that using spinning spheres as hydrodynamic actuators in place of constellation or micro-rotors would, in principle, be possible. However, fast enough spinning rates of optically trapped particles have not been demonstrated experimentally, for us to seriously consider this. The GWS approach could potentially be used to reach sufficiently high spinning rates by maximising the torque that a light field can apply on a spherical particle.

5.6.4 Comparison to the work of others

We are by no means alone in our efforts to enhance optical trapping stiffness, with important work reported by Michael Taylor in ‘Optimizing phase to enhance optical trap stiffness’ published in *Scientific Reports* in 2017 [142], and Michael Mazilu in his collection of conference presentations titled ‘Eigenmode beam optimisation for optical micro-manipulation’ and published on *arXiv* in 2019 [143].

Michael Taylor has tackled the 1D stiffness enhancement problem with the eigenmode method and an accompanying phase-only iterative optimisation. Qualitatively comparing the x -force profile of our two-lobe trap (in Figure 5.10) with his optimal eigenmode trap (orange curve in Figure 2(e) in [142]), we see that the two are very similar (as one would expect given the equivalence of the eigenmode and GWS methods demonstrated recently in [147]). For a more quantitative comparison we look at the stiffness enhancement factor under conditions equivalent to ours - a particle of $3\mu\text{m}$ radius, refractive index of 1.5, trapped in water with 1064 nm wavelength light. The number reported by Taylor lies somewhere in the range between 34 and 40^5 , compared to our 39, which is, again, in line with the similarity of the two methods used.

In the paper under discussion, Taylor did not attempt to stabilise the traps in the third dimension, instead relying on gravity to counteract the scattering force pushing the particle in the direction of the incident beam. In order to reduce the scattering force to the level of the gravitational force, one has to turn down the laser power. This, in turn, decreases the gradient force acting in the transverse plane, so that absolute stiffness is also reduced. The 39-fold enhancement factor is therefore limited to these conditions and might not actually represent a high stiffness value. We take a step further by

⁵ The exact value is not given and had to be inferred from the phase-only enhancement data in Figure 7(d), together with the quoted increase factor of 1.35 for the optimal eigenmode field.

ensuring that our traps are stable in all three dimensions, which does come at the expense of reduced enhancement in the transverse plane, but extends the versatility of our traps.

Michael Mazilu also uses the eigenmode approach and he is also employing the Bessel basis to represent light fields. Instead of enhancing the stiffness directly, however, he enhances the optical force (in the transverse plane) that a cylindrically symmetric light field (i.e. only modes with $L = 0$ are allowed) can exert on a particle - this coincides with the best achievable stiffness. This method should produce results equivalent to our ring trap. Mazilu reports a tenfold stiffness enhancement; however, we cannot compare this result directly with ours as the exact parameters used in [143] are not available.

Unlike Taylor, Mazilu takes an extra step in ensuring a 3D-stable trapping situation. Much like we did in Section 5.3, he finds the eigenmode with the most negative z -force and then combines it with the field that was optimal for transverse-plane trapping. Following this, he optimises the phase between the two fields to favour the transverse trapping. The resulting transverse stiffness, however, just barely exceeds that of a Gaussian trap, and the longitudinal stiffness does not appear to be enhanced (although we can only judge this by the visual information presented in Figure 5(left) in [143]). Meanwhile we were able to demonstrate (within the restrictions of using $L = 0$ modes only) simultaneous stiffness improvement with $\mathcal{R}_x = 13$, $\mathcal{R}_y = 5$, $\mathcal{R}_z = 13$ (with the T_z field presented in Section 5.5).

5.6.5 Improvements and future work

First and foremost, it is difficult to fully appreciate the power of the methods we have used, without having experimental demonstration of any of the fields we have presented in this chapter. It is our intention to create and test these fields experimentally in the near future at David Phillips' lab at the University of Exeter, and we aim to have these results published next year. To this end, we are developing the phase-only methods discussed in Section 5.5, so that we can fairly easily create the stiffness-enhanced traps without having to modulate the spatial structure of the intensity in the far field, as well as the phase.

We note that throughout this chapter we have only considered a single test particle of $3\ \mu\text{m}$ radius and a fixed refractive index of 1.5. The enhancement factor will, however, be dependent on the particle size and its refractive index. As the particle size changes, we expect the shape of the trapping beam to maintain the same structure as trap

A - intensity spread out along the inner edges of the bead - up until the diffraction limit, where the light can no longer be shaped with enough detail to contour small particles, and where Gaussian beam trapping cannot be improved on. Generally, the stiffness enhancement factor will be higher for larger particles, simply because they are more loosely trapped in a Gaussian beam. Similarly, as the refractive index of the particle changes, we would not expect to see major changes in the light structure in the transverse plane. However, with increasing refractive index the scattering force becomes more dominant, making it more difficult to stabilise the particle in the axial direction. We might therefore see some changes in the longitudinal shape of the field with a varying refractive index. Again, we expect the largest enhancement factors in the transverse plane for low refractive index beads, as they are less tightly trapped in a Gaussian beam. We also note, that we do not expect any substantial dependence of our light field shape on the wavelength of the trapping light, other than diffraction associated increase in intensity pattern features for longer wavelengths.

In addition to the above, it would be interesting to consider the trapping of arbitrarily shaped particles, or multiple particles at once. For example, instead of replicating an identical field on each bead in our constellation rotors, we could create a field that takes into account the light scattering between the three beads. My supervisor Jonny Taylor has estimated that such a field can increase the optical trapping stiffness of each bead by a further 10% (in the direction parallel to the motion of the bead, without taking into account z -trapping, and in the setup where incoming light is not limited to an aperture). The same field could be used for the wheel-shaped micro-rotors, as they are trapped by directing light onto their handle beads. An even stiffer trap could potentially be created for these micro-rotors, if light scattering information was available for the entire structure. It might even be possible to develop a scheme that optimises the light field and the rotor geometry simultaneously.

Computation speed was a factor that significantly limited iterative optimisation in multi-dimensional problems - convergence to a minimum could take up to several days. The bottleneck of our optimiser is the need to calculate the trapping force twice to obtain the stiffness along one dimension, for each evaluation of the objective function. We expect that a lot of speed could be gained by employing the GWS operators directly within the optimiser to estimate the objective function and its gradient, in a scheme similar to the one employed by Michael Taylor [142]. This leads us to the final discussion point.

As we saw in Section 5.6.1, our implementation of the GWS operators suffers from irregularities, which is evident from the non-perfect linearity between the eigenvalues

of the GWS operators and the corresponding force and stiffness values. This very likely means that we cannot rely on our Q, K operators to evaluate the force and stiffness correctly, and reformulation of our approach would be necessary. This would include determining which particular aspect of our implementation causes the issues: imperfect field sampling, defining the output state in terms of the total electric field rather than just the outgoing part, or the use of an unsuitable output basis altogether. Our current focus is on deducing which of the three (if either) or their combination is the root of the non-perfect linearity problem.

Provided that the above issue can be solved, and an optimiser incorporating the Q and K operators is developed, we could explore multi-dimensional stiffness enhancement with a more general approach, rather than adding fields optimal for each individual dimension. We remind the reader that doing this with our current optimiser requires days to converge to a solution. With a speedy optimisation algorithm available, we could explore some additional problems. For example, how stiffness can be enhanced while making sure that the resulting trap is stable for a range of bead sizes. This could prove useful in an experiment, since commercially available optically trappable beads usually have some size variation and the optimal field obtained within simulations might not be equally good, or even stable, for small deviations in the bead size. Another example problem could be how to achieve maximum stiffness, while minimising light intensity incident on an object. This could prove useful for trapping live specimens which are photosensitive.

Perhaps the biggest drawback of stiffness enhancement is that it comes at a cost of the maximum achievable force and the trapping range. This limits the possible applications of the traps we have demonstrated. It might be possible to extend the trapping range using the super-principal mode technique [186]. This method was developed in relation to multimode waveguides to expand the frequency range of the signal that the waveguide can carry through undistorted. It makes use of the original time delay operator Q_ω and could potentially be adapted for K to extend the trapping range.

Chapter 6

Conclusion

‘That’s how every good story must end, Max. When a person stops understanding something, he’s on the right track.’

The Stranger — Max Frei

Non-invasive manipulation methods which rely on contactless interactions have provided deep insight into the micro-world. In this thesis we have considered and integrated two such techniques - optical and hydrodynamic - within the context of enhancing the capabilities of optical tweezers.

In Part I we set out to merge the benefits of optical and hydrodynamic tweezers in the form of optically actuated hydrodynamic manipulation. Our technique extends the range of influence of an optical tweezers platform to materials which do not yield to optical trapping, or which can be damaged by direct light exposure, while simultaneously overcoming the requirement of very dilute samples for microfluidic hydrodynamic tweezing. By optically trapping spherical beads and specially designed

micro-rotors, we used them to actuate the fluid, thus exerting control over nearby particles freely diffusing within it. With our technique we have demonstrated two-dimensional suppression of the Brownian motion of a single sedimented $5\ \mu\text{m}$ radius silica bead down to a standard deviation of $79\ \text{nm}$ - which is equivalent to holding it in a weak optical trap of $6.4 \times 10^{-7}\ \text{N m}^{-1}$ stiffness. Unlike optical tweezing, hydrodynamic clamping is less effective for smaller particles, since they diffuse faster and have reduced tracking accuracy. We have clamped the motion of $2.5\ \mu\text{m}$ and $1\ \mu\text{m}$ radii silica beads down to $186\ \text{nm}$ and $588\ \text{nm}$ respectively. Much like with optical trapping, though, hydrodynamic clamping gets worse when control is exerted over several particles simultaneously. The reason for this is twofold. First, simultaneous fluidic control of multiple particles requires complex fluid flow patterns, where the local currents are diverging from one another. Since this is only possible with the presence of stagnation points - points of zero fluid velocity - the overall achievable velocities are reduced, resulting in slower target response. Second, more targets means that more actuators are required; to accommodate them, the target-actuator separation has to be increased, thus reducing the strength of hydrodynamic interaction between the targets and the actuators. Increased number of rotors also means reduced actuator optical trapping stiffness, since the available laser power has to be shared amongst more traps - this in turn means reduced rotation rates, and hence, a reduced effect of the rotors on the targets. Furthermore, low optical trapping stiffness has the effect of low-pass filtering the actuator motion in comparison to their trap motion, thus reducing the precision of hydrodynamic control as well. In fact, we have identified insufficient optical trapping stiffness as one of the main limitations of our method.

Hydrodynamic manipulation is independent of the material of the target particle, as we have demonstrated by clamping a piece of chromium - a material which cannot be trapped optically - and suppressing its thermal motion down to a standard deviation of $482\ \text{nm}$. As we do not expose the target to direct laser illumination, our technique is well suited for living cell manipulation. Furthermore, if these cells are also optically trappable, a few of them can be ‘sacrificed’ to be used as the actuators, thus avoiding introduction of foreign material into the cell sample. We have demonstrated the latter with yeast cells, which we have also used to establish the viability of in-plane orientation control, achieving standard deviation of 7.5° , simultaneously with position clamping down to $385\ \text{nm}$. With our technique we can perform local small-scale manipulations as well as long-range translations of the target through the holding sample cell. Both can be done while only minimally disturbing the rest of the surrounding particles. By making use of the reconfigurable nature of optical tweezers our hydrodynamic manipulation approach is a step closer towards non-invasive microrobotics.

In Part II we explored optimisation of the trapping light field to enhance three-dimensional trapping stiffness. While this work was originally inspired by one of the main limitations of our hydrodynamic manipulation method - insufficient optical trapping stiffness - its potential reaches far beyond that. We believe that stiffness enhancement will find home in many optical tweezers applications, amongst which are optical trapping of high refractive index particles, force sensing in biomolecules, and trapping of micro-tools for applications such as surface imaging, quantification of the properties of the surrounding environment, assembly and dismantling of specially designed micro-structures, and material transport.

Using the relatively new approach of the generalised Wigner-Smith operators we have demonstrated, within simulations, optical stiffness enhancement of 39 times (compared to a Gaussian beam trap, for a $3\ \mu\text{m}$ radius silica bead, with refractive index of 1.5, in water, illuminated with light of 1064 nm wavelength) in 1D in the plane transverse to the beam axis; this is in line with the results recently published by others [142]. With the help of an iterative optimiser we have obtained a field that gives a 16-fold stiffness enhancement in the direction longitudinal to the beam - we believe that this is a unique result, both in terms of the enhancement factor achieved and the technique used, which has never been demonstrated before.

Fields which are optimised to increase stiffness in one direction only, however, often do not offer trapping stability in one of the other dimensions. In previous works by others, z -instability has been dealt with by relying on gravity to counterbalance the pushing optical force [141, 142]. This method requires reduction of the trapping laser power which, by reducing the optical force and its gradient, does not showcase the full potential of enhanced stiffness trapping. To stabilise the trapping in 3D we took the approach of combining fields, separately optimised for each dimension. This allowed us to achieve simultaneous stiffness enhancement factors of 11 and 7 in the transverse and longitudinal directions respectively, when requiring equal trapping strength along x and y . If enhancement along one of the transverse directions was prioritised over the other, the maximum enhancement factor in the transverse plane was 12, accompanied by 10-fold stiffness increase in the longitudinal direction.

However, such order-of-magnitude stiffness enhancement comes at the cost of a reduced trapping range and a decreased maximum optical force that the trapped particle can be subjected to. This is of little consequence to stationary trapping (provided that the trapping laser power is sufficient to overcome the thermal motion of the trapped particle), but is potentially problematic if the trapped object is to be translated, like in the case of our hydrodynamic rotors. This limits the applicability of the 3D stiffness

enhanced fields that we demonstrated to situations where fast translation speeds are not required.

The story of enhanced stiffness optical trapping is not yet complete. There are still irregularities in our implementation of the generalised Wigner-Smith operators, causes of which we do not understand quite fully. While they did not prevent us from achieving very good results, they might be in the way of achieving great ones. We did not have the time (is this not so often the case with PhDs?), or the means (due to world-wide shut down caused by Covid-19), to test our simulated enhanced stiffness traps experimentally. Without a doubt, this will be done in the near future and I look forward to seeing the results.

Contributions

I, the author of this work, developed the mathematical model for hydrodynamic control feedback as described in Section 3.1.2, wrote the codes and carried out the simulations presented in Sections 3.2 and 3.3, updated the existing Red Tweezers software to include the feedback loop for hydrodynamic manipulation experiments, performed all of the hydrodynamic manipulation experiments presented in Chapter 3 as well as analysed the data, extended the work of Jonathan M. Taylor in [163] to derive the BSCs for a y -polarised Bessel beam, as well as Bessel beams with an arbitrary value of OAM, wrote the code which simulates light-matter interaction based on the GLMT, devised the optimiser which minimises the stiffness of an optical trap as explained in Section 5.1.2, implemented the GWS operator within, and carried out, the simulations presented in Chapter 5. I did all of this under invaluable supervision and guidance from Jonathan M. Taylor and David B. Phillips.

Other people have contributed to the work presented in this thesis as follows.

Dr. [David B. Phillips](#) conceived the idea of optically actuated hydrodynamic manipulation and did the preliminary work on the subject in [122], designed and contributed to the manufacturing of the laser-printed micro-rotors and wrote the MATLAB code which renders the micro-rotor structures (e.g. as seen in Figure 3.2(b)).

Dr. [Jonathan M. Taylor](#) developed the idea for Part II of this thesis together with David B. Phillips, laid out the theoretical basis of the GLMT in [159] that I used to

write my own GLMT code, and aided in validating my implementation of GLMT by providing data obtained with his codes.

Dr. [Graham M. Gibson](#) maintained the optical tweezers setup used in the experiments of Part I, in the lab of Prof. [Miles Padgett](#) at the University of Glasgow. He designed and constructed the beam block mask and its holder for the hydrodynamic manipulation experiments and aided in expanding the optical tweezers setup to include a second camera. He also made the optical tweezers setup figure, which Figure 3.16 was adapted from.

Prof. [John Rarity](#) has generously allowed us to use the Nanoscribe printer at the University of Bristol, and Dr. [Mike Taverne](#) and Dr. [Ying-Lung D. Ho](#) assisted in laser-printing of the micro-rotor structures.

Prof. [Stefan Rotter](#) and [Michael Horodyski](#) at the Vienna University of Technology developed the GWS method and guided us through the details of its theoretical framework.

Bibliography

- [1] Robert Hooke. *Micrographia: or some physiological descriptions of minute bodies made by magnifying glasses, with observations and inquiries thereupon*. J. Martyn and J. Allestry, 1665.
- [2] Jürgen Fröhlich and Helmut König. ‘New techniques for isolation of single prokaryotic cells’. In: *FEMS microbiology reviews* 24.5 (2000), pp. 567–572.
- [3] LI Sly and Vullapa Arunpairojana. ‘Isolation of manganese-oxidizing Pedococcus cultures from water by micromanipulation’. In: *Journal of microbiological methods* 6.3 (1987), pp. 177–182.
- [4] Jürgen Fröhlich and Helmut König. ‘Rapid isolation of single microbial cells from mixed natural and laboratory populations with the aid of a micromanipulator’. In: *Systematic and applied microbiology* 22.2 (1999), pp. 249–257.
- [5] Iwijn De Vlaminck and Cees Dekker. ‘Recent advances in magnetic tweezers’. In: *Annual review of biophysics* 41 (2012), pp. 453–472.
- [6] Andreas R Bausch, Winfried Möller and Erich Sackmann. ‘Measurement of local viscoelasticity and forces in living cells by magnetic tweezers’. In: *Biophysical journal* 76.1 (1999), pp. 573–579.
- [7] Charlie Gosse and Vincent Croquette. ‘Magnetic tweezers: micromanipulation and force measurement at the molecular level’. In: *Biophysical journal* 82.6 (2002), pp. 3314–3329.

- [8] Elizabeth Mirowski et al. ‘Manipulation and sorting of magnetic particles by a magnetic force microscope on a microfluidic magnetic trap platform’. In: *Applied Physics Letters* 86.24 (2005), p. 243901.
- [9] Kerstin Ramser and Dag Hanstorp. ‘Optical manipulation for single-cell studies’. In: *Journal of biophotonics* 3.4 (2010), pp. 187–206.
- [10] Stephen C Chapin, Vincent Germain and Eric R Dufresne. ‘Automated trapping, assembly, and sorting with holographic optical tweezers’. In: *Optics Express* 14.26 (2006), pp. 13095–13100.
- [11] K Dholakia and T Čižmár. ‘Shaping the future of manipulation’. In: *Nature Photonics* 5.6 (2011), p. 335.
- [12] David G Grier. ‘A revolution in optical manipulation’. In: *Nature* 424.6950 (2003), pp. 810–816.
- [13] Ming C Wu. ‘Optoelectronic tweezers’. In: *Nature Photonics* 5.6 (2011), pp. 322–324.
- [14] Lifeng Zheng, James P Brody and Peter J Burke. ‘Electronic manipulation of DNA, proteins, and nanoparticles for potential circuit assembly’. In: *Biosensors and Bioelectronics* 20.3 (2004), pp. 606–619.
- [15] Thomas Schnelle et al. ‘Three-dimensional electric field traps for manipulation of cells – calculation and experimental verification’. In: *Biochimica et Biophysica Acta (BBA)-General Subjects* 1157.3 (1993), pp. 127–140.
- [16] Adam E Cohen and WE Moerner. ‘Suppressing Brownian motion of individual biomolecules in solution’. In: *Proceedings of the National Academy of Sciences* 103.12 (2006), pp. 4362–4365.
- [17] Melikhan Tanyeri and Charles M Schroeder. ‘Manipulation and confinement of single particles using fluid flow’. In: *Nano letters* 13.6 (2013), pp. 2357–2364.
- [18] Tristan Petit et al. ‘Selective trapping and manipulation of microscale objects using mobile microvortices’. In: *Nano letters* 12.1 (2011), pp. 156–160.
- [19] Narayan Sundararajan et al. ‘Three-dimensional hydrodynamic focusing in polydimethylsiloxane (PDMS) microchannels’. In: *Journal of Microelectromechanical systems* 13.4 (2004), pp. 559–567.
- [20] Xiaoyun Ding et al. ‘On-chip manipulation of single microparticles, cells, and organisms using surface acoustic waves’. In: *Proceedings of the National Academy of Sciences* 109.28 (2012), pp. 11105–11109.
- [21] Jinjie Shi et al. ‘Acoustic tweezers: patterning cells and microparticles using standing surface acoustic waves (SSAW)’. In: *Lab on a Chip* 9.20 (2009), pp. 2890–2895.
- [22] HM Hertz. ‘Standing-wave acoustic trap for nonintrusive positioning of microparticles’. In: *Journal of applied physics* 78.8 (1995), pp. 4845–4849.

- [23] Mikael Evander et al. ‘Noninvasive acoustic cell trapping in a microfluidic perfusion system for online bioassays’. In: *Analytical chemistry* 79.7 (2007), pp. 2984–2991.
- [24] Arthur Ashkin. ‘Acceleration and trapping of particles by radiation pressure’. In: *Physical Review Letters* 24.4 (1970), p. 156.
- [25] Arthur Ashkin and JM Dziedzic. ‘Optical levitation by radiation pressure’. In: *Applied Physics Letters* 19.8 (1971), pp. 283–285.
- [26] Asimina Arvanitaki and Andrew A Geraci. ‘Detecting high-frequency gravitational waves with optically levitated sensors’. In: *Physical Review Letters* 110.7 (2013), p. 071105.
- [27] Levi P Neukirch et al. ‘Observation of nitrogen vacancy photoluminescence from an optically levitated nanodiamond’. In: *Optics Letters* 38.16 (2013), pp. 2976–2979.
- [28] Mor Habaza et al. ‘Tomographic phase microscopy with 180 rotation of live cells in suspension by holographic optical tweezers’. In: *Optics Letters* 40.8 (2015), pp. 1881–1884.
- [29] Min-Cheng Zhong et al. ‘Trapping red blood cells in living animals using optical tweezers’. In: *Nature Communications* 4.1 (2013), pp. 1–7.
- [30] Arthur Ashkin and James M Dziedzic. ‘Optical trapping and manipulation of viruses and bacteria’. In: *Science* 235.4795 (1987), pp. 1517–1520.
- [31] Tudor N Buican et al. ‘Automated single-cell manipulation and sorting by light trapping’. In: *Applied Optics* 26.24 (1987), pp. 5311–5316.
- [32] Steven M Block, David F Blair and Howard C Berg. ‘Compliance of bacterial flagella measured with optical tweezers’. In: *Nature* 338.6215 (1989), pp. 514–518.
- [33] DB Phillips et al. ‘Surface imaging using holographic optical tweezers’. In: *Nanotechnology* 22.28 (2011), p. 285503.
- [34] Peter John Rodrigo et al. ‘Actuation of microfabricated tools using multiple GPC-based counterpropagating-beam traps’. In: *Optics Express* 13.18 (2005), pp. 6899–6904.
- [35] L Ikin et al. ‘Assembly and force measurement with SPM-like probes in holographic optical tweezers’. In: *New Journal of Physics* 11.2 (2009), p. 023012.
- [36] David M Carberry et al. ‘Calibration of optically trapped nanotools’. In: *Nanotechnology* 21.17 (2010), p. 175501.
- [37] Zdeněk Pilát et al. ‘Effects of Infrared Optical Trapping on *Saccharomyces cerevisiae* in a Microfluidic System’. In: *Sensors* 17.11 (2017), p. 2640.
- [38] Keir C Neuman et al. ‘Characterization of photodamage to *Escherichia coli* in optical traps’. In: *Biophysical journal* 77.5 (1999), pp. 2856–2863.

- [39] Satoru Ayano et al. ‘Quantitative measurement of damage caused by 1064-nm wavelength optical trapping of *Escherichia coli* cells using on-chip single cell cultivation system’. In: *Biochemical and biophysical research communications* 350.3 (2006), pp. 678–684.
- [40] Utkur Mirsaidov et al. ‘Optimal optical trap for bacterial viability’. In: *Physical Review E* 78.2 (2008), p. 021910.
- [41] Michael D Armani et al. ‘Using feedback control of microflows to independently steer multiple particles’. In: *Journal of Microelectromechanical systems* 15.4 (2006), pp. 945–956.
- [42] Roland Probst et al. ‘Flow control of small objects on chip: Manipulating live cells, quantum dots, and nanowires’. In: *IEEE Control Systems* 32.2 (2012), pp. 26–53.
- [43] Anish Shenoy, Christopher V Rao and Charles M Schroeder. ‘Stokes trap for multiplexed particle manipulation and assembly using fluidics’. In: *Proceedings of the National Academy of Sciences* 113.15 (2016), pp. 3976–3981.
- [44] Vigneswaran Narayanamurthy et al. ‘Microfluidic hydrodynamic trapping for single cell analysis: mechanisms, methods and applications’. In: *Analytical Methods* 9.25 (2017), pp. 3751–3772.
- [45] DB Phillips et al. ‘Shape-induced force fields in optical trapping’. In: *Nature Photonics* 8.5 (2014), p. 400.
- [46] J Trojek, V Karasek and P Zemanek. ‘Extreme axial optical force in a standing wave achieved by optimized object shape’. In: *Optics Express* 17.13 (2009), pp. 10472–10488.
- [47] Anita Jannasch et al. ‘Nanonewton optical force trap employing anti-reflection coated, high-refractive-index titania microspheres’. In: *Nature Photonics* 6.7 (2012), p. 469.
- [48] Péter Galajda and Pál Ormos. ‘Rotors produced and driven in laser tweezers with reversed direction of rotation’. In: *Applied Physics Letters* 80.24 (2002), pp. 4653–4655.
- [49] Grover A Swartzlander Jr et al. ‘Stable optical lift’. In: *Nature Photonics* 5.1 (2011), p. 48.
- [50] Tomáš Čižmár, Michael Mazilu and Kishan Dholakia. ‘In situ wavefront correction and its application to micromanipulation’. In: *Nature Photonics* 4.6 (2010), pp. 388–394.
- [51] Kyoohyun Kim and YongKeun Park. ‘Tomographic active optical trapping of arbitrarily shaped objects by exploiting 3D refractive index maps’. In: *Nature Communications* 8.1 (2017), pp. 1–8.

- [52] Yoonkyung E Lee et al. ‘Computational inverse design of non-intuitive illumination patterns to maximize optical force or torque’. In: *Optics Express* 25.6 (2017), pp. 6757–6766.
- [53] Michael Mazilu et al. ‘Optical eigenmodes; exploiting the quadratic nature of the energy flux and of scattering interactions’. In: *Optics Express* 19.2 (2011), pp. 933–945.
- [54] Robert Brown. ‘XXVII. A brief account of microscopical observations made in the months of June, July and August 1827, on the particles contained in the pollen of plants; and on the general existence of active molecules in organic and inorganic bodies’. In: *The Philosophical Magazine* 4.21 (1828), pp. 161–173.
- [55] Philip H Jones, Onofrio M Maragò and Giovanni Volpe. *Optical tweezers: Principles and applications*. Cambridge University Press, 2015.
- [56] Bernt Øksendal. ‘Stochastic differential equations’. In: *Stochastic differential equations*. Springer, 2003, pp. 65–84.
- [57] Rep Kubo. ‘The fluctuation-dissipation theorem’. In: *Reports on progress in physics* 29.1 (1966), p. 255.
- [58] David B Thomas and Wayne Luk. ‘Multivariate Gaussian random number generation targeting reconfigurable hardware’. In: *ACM Transactions on Reconfigurable Technology and Systems (TRETS)* 1.2 (2008), p. 12.
- [59] RB Jones. ‘Hydrodynamic interaction of two permeable spheres I: The method of reflections’. In: *Physica A: Statistical Mechanics and its Applications* 92.3-4 (1978), pp. 545–556.
- [60] GJ Kynch. ‘The slow motion of two or more spheres through a viscous fluid’. In: *Journal of Fluid Mechanics* 5.2 (1959), pp. 193–208.
- [61] BU Felderhof. ‘Hydrodynamic interaction between two spheres’. In: *Physica A: Statistical Mechanics and its Applications* 89.2 (1977), pp. 373–384.
- [62] Peter Mazur and W Van Saarloos. ‘Many-sphere hydrodynamic interactions and mobilities in a suspension’. In: *Physica A: Statistical Mechanics and its Applications* 115.1-2 (1982), pp. 21–57.
- [63] CWJ Beenakker, W Van Saarloos and P Mazur. ‘Many-sphere hydrodynamic interactions: III. The influence of a plane wall’. In: *Physica A: Statistical Mechanics and its Applications* 127.3 (1984), pp. 451–472.
- [64] P Mazur and G Van der Zwan. ‘Brownian motion in a fluid close to its’ critical point’. In: *Physica A: Statistical Mechanics and its Applications* 92.3-4 (1978), pp. 483–500.
- [65] P Mazur and D Bedeaux. ‘A generalization of Faxén’s theorem to nonsteady motion of a sphere through an incompressible fluid in arbitrary flow’. In: *Physica* 76.2 (1974), pp. 235–246.

- [66] MGR Urbanczyk. ‘Solar Sails-A Realistic Propulsion for Spacecraft’. In: (1967).
- [67] Arthur Ashkin et al. ‘Observation of a single-beam gradient force optical trap for dielectric particles’. In: *Optics Letters* 11.5 (1986), pp. 288–290.
- [68] Arthur Ashkin. ‘Forces of a single-beam gradient laser trap on a dielectric sphere in the ray optics regime’. In: *Biophysical journal* 61.2 (1992), pp. 569–582.
- [69] Edward M Purcell. ‘Life at low Reynolds number’. In: *American journal of physics* 45.1 (1977), pp. 3–11.
- [70] Luc P Faucheux and Albert J Libchaber. ‘Confined brownian motion’. In: *Physical Review E* 49.6 (1994), p. 5158.
- [71] MIM Feitosa and ON Mesquita. ‘Wall-drag effect on diffusion of colloidal particles near surfaces: a photon correlation study’. In: *Physical Review A* 44.10 (1991), p. 6677.
- [72] Andy WC Lau and Tom C Lubensky. ‘State-dependent diffusion: Thermodynamic consistency and its path integral formulation’. In: *Physical Review E* 76.1 (2007), p. 011123.
- [73] P Lançon et al. ‘Drift without flux: Brownian walker with a space-dependent diffusion coefficient’. In: *EPL (Europhysics Letters)* 54.1 (2001), p. 28.
- [74] Donald L Ermak and JA McCammon. ‘Brownian dynamics with hydrodynamic interactions’. In: *The Journal of chemical physics* 69.4 (1978), pp. 1352–1360.
- [75] M De Corato et al. ‘Hydrodynamics and Brownian motions of a spheroid near a rigid wall’. In: *The Journal of chemical physics* 142.19 (2015), p. 194901.
- [76] Steven Delong et al. ‘Brownian dynamics without Green’s functions’. In: *The Journal of chemical physics* 140.13 (2014), p. 134110.
- [77] R Perez-Carrasco and JM Sancho. ‘Stochastic algorithms for discontinuous multiplicative white noise’. In: *Physical Review E* 81.3 (2010), p. 032104.
- [78] Giovanni Volpe et al. ‘Influence of noise on force measurements’. In: *Physical Review Letters* 104.17 (2010), p. 170602.
- [79] Thomas Brettschneider et al. ‘Force measurement in the presence of Brownian noise: Equilibrium-distribution method versus drift method’. In: *Physical Review E* 83.4 (2011), p. 041113.
- [80] Eugene Helfand. ‘Brownian dynamics study of transitions in a polymer chain of bistable oscillators’. In: *The Journal of Chemical Physics* 69.3 (1978), pp. 1010–1018.
- [81] Rebecca L Honeycutt. ‘Stochastic runge-kutta algorithms. i. white noise’. In: *Physical Review A* 45.2 (1992), p. 600.
- [82] A Iniesta and J Garcia de la Torre. ‘A second-order algorithm for the simulation of the Brownian dynamics of macromolecular models’. In: *The Journal of Chemical Physics* 92.3 (1990), pp. 2015–2018.

- [83] AC Brańka and DM Heyes. ‘Algorithms for Brownian dynamics simulation’. In: *Physical Review E* 58.2 (1998), p. 2611.
- [84] Giorgio Volpe and Giovanni Volpe. ‘Simulation of a Brownian particle in an optical trap’. In: *American Journal of Physics* 81.3 (2013), pp. 224–230.
- [85] RF Hay et al. ‘Lissajous-like trajectories in optical tweezers’. In: *Optics Express* 23.25 (2015), pp. 31716–31727.
- [86] Wesley P Wong and Ken Halvorsen. ‘The effect of integration time on fluctuation measurements: calibrating an optical trap in the presence of motion blur’. In: *Optics Express* 14.25 (2006), pp. 12517–12531.
- [87] Christopher TH Baker and Evelyn Buckwar. ‘Numerical analysis of explicit one-step methods for stochastic delay differential equations’. In: *LMS Journal of Computation and Mathematics* 3 (2000), pp. 315–335.
- [88] Donald R Barr and Norman L Slezak. ‘A comparison of multivariate normal generators’. In: *Communications of the ACM* 15.12 (1972), pp. 1048–1049.
- [89] Joseph A Curcio and Charles C Petty. ‘The near infrared absorption spectrum of liquid water’. In: *JOSA* 41.5 (1951), pp. 302–304.
- [90] George M Hale and Marvin R Querry. ‘Optical constants of water in the 200-nm to 200- μm wavelength region’. In: *Applied Optics* 12.3 (1973), pp. 555–563.
- [91] Erik Fällman and Ove Axner. ‘Design for fully steerable dual-trap optical tweezers’. In: *Applied Optics* 36.10 (1997), pp. 2107–2113.
- [92] Maarten C Noom et al. ‘Visualizing single DNA-bound proteins using DNA as a scanning probe’. In: *Nature Methods* 4.12 (2007), pp. 1031–1036.
- [93] Gary J Brouhard, Henry T Schek and Alan J Hunt. ‘Advanced optical tweezers for the study of cellular and molecular biomechanics’. In: *IEEE Transactions on Biomedical Engineering* 50.1 (2003), pp. 121–125.
- [94] JJJM Durnin, JJ Miceli Jr and JH Eberly. ‘Diffraction-free beams’. In: *Physical Review Letters* 58.15 (1987), p. 1499.
- [95] J Arlt and MJ Padgett. ‘Generation of a beam with a dark focus surrounded by regions of higher intensity: the optical bottle beam’. In: *Optics Letters* 25.4 (2000), pp. 191–193.
- [96] H He et al. ‘Direct observation of transfer of angular momentum to absorptive particles from a laser beam with a phase singularity’. In: *Physical Review Letters* 75.5 (1995), p. 826.
- [97] Marcus Reicherter et al. ‘Optical particle trapping with computer-generated holograms written on a liquid-crystal display’. In: *Optics Letters* 24.9 (1999), pp. 608–610.
- [98] Yoshio Hayasaki et al. ‘Nonmechanical optical manipulation of microparticle using spatial light modulator’. In: *Optical review* 6.1 (1999), pp. 24–27.

- [99] Don A Gregory, James C Kirsch and Eddy C Tam. ‘Full complex modulation using liquid-crystal televisions’. In: *Applied Optics* 31.2 (1992), pp. 163–165.
- [100] Jan Liesener et al. ‘Multi-functional optical tweezers using computer-generated holograms’. In: *Optics Communications* 185.1-3 (2000), pp. 77–82.
- [101] Ralph W Gerchberg. ‘A practical algorithm for the determination of phase from image and diffraction plane pictures’. In: *Optik* 35 (1972), pp. 237–246.
- [102] Eric R Dufresne et al. ‘Computer-generated holographic optical tweezer arrays’. In: *Review of Scientific Instruments* 72.3 (2001), pp. 1810–1816.
- [103] Graham M Gibson et al. ‘Measuring the accuracy of particle position and force in optical tweezers using high-speed video microscopy’. In: *Optics Express* 16.19 (2008), pp. 14561–14570.
- [104] Oliver Otto et al. ‘Optical tweezers with 2.5 kHz bandwidth video detection for single-colloid electrophoresis’. In: *Review of Scientific Instruments* 79.2 (2008), p. 023710.
- [105] John C Crocker and David G Grier. ‘Methods of digital video microscopy for colloidal studies’. In: *Journal of colloid and interface science* 179.1 (1996), pp. 298–310.
- [106] Richard Bowman et al. ‘Stereoscopic particle tracking for 3D touch, vision and closed-loop control in optical tweezers’. In: *Journal of optics* 13.4 (2011), p. 044003.
- [107] Simon F Tolić-Nørrelykke et al. ‘Calibration of optical tweezers with positional detection in the back focal plane’. In: *Review of scientific instruments* 77.10 (2006), p. 103101.
- [108] Robert M Simmons et al. ‘Quantitative measurements of force and displacement using an optical trap’. In: *Biophysical journal* 70.4 (1996), pp. 1813–1822.
- [109] Keir C Neuman and Steven M Block. ‘Optical trapping’. In: *Review of scientific instruments* 75.9 (2004), pp. 2787–2809.
- [110] Alexis I Bishop et al. ‘Optical microrheology using rotating laser-trapped particles’. In: *Physical Review Letters* 92.19 (2004), p. 198104.
- [111] Daryl Preece et al. ‘Optical tweezers: wideband microrheology’. In: *Journal of optics* 13.4 (2011), p. 044022.
- [112] Manlio Tassieri et al. ‘Microrheology with optical tweezers: data analysis’. In: *New Journal of Physics* 14.11 (2012), p. 115032.
- [113] Iva Marija Tolić-Nørrelykke et al. ‘Anomalous diffusion in living yeast cells’. In: *Physical Review Letters* 93.7 (2004), p. 078102.
- [114] R Di Leonardo et al. ‘Multipoint holographic optical velocimetry in microfluidic systems’. In: *Physical Review Letters* 96.13 (2006), p. 134502.

- [115] Boaz A Nemet and Mark Cronin-Golomb. ‘Microscopic flow measurements with optically trapped microprobes’. In: *Optics Letters* 27.15 (2002), pp. 1357–1359.
- [116] Gregor Knöner et al. ‘Characterization of optically driven fluid stress fields with optical tweezers’. In: *Physical Review E* 72.3 (2005), p. 031507.
- [117] Martin Andersson et al. ‘Using optical tweezers for measuring the interaction forces between human bone cells and implant surfaces: System design and force calibration’. In: *Review of Scientific Instruments* 78.7 (2007), p. 074302.
- [118] Scot C Kuo and Michael P Sheetz. ‘Force of single kinesin molecules measured with optical tweezers’. In: *Science* 260.5105 (1993), pp. 232–234.
- [119] Giovanni Volpe, Gregory Kozyreff and Dmitri Petrov. ‘Backscattering position detection for photonic force microscopy’. In: *Journal of Applied Physics* 102.8 (2007), p. 084701.
- [120] Kirstine Berg-Sørensen and Henrik Flyvbjerg. ‘Power spectrum analysis for optical tweezers’. In: *Review of Scientific Instruments* 75.3 (2004), pp. 594–612.
- [121] Richard W Bowman et al. ‘Red Tweezers: Fast, customisable hologram generation for optical tweezers’. In: *Computer Physics Communications* 185.1 (2014), pp. 268–273.
- [122] David B Phillips et al. ‘Optically controlled hydrodynamic micro-manipulation’. In: *Optical Trapping and Optical Micromanipulation XII*. Vol. 9548. International Society for Optics and Photonics. 2015, 95481A.
- [123] MEJ Friese et al. ‘Optical torque controlled by elliptical polarization’. In: *Optics Letters* 23.1 (1998), pp. 1–3.
- [124] MEJ Friese et al. ‘Optical angular-momentum transfer to trapped absorbing particles’. In: *Physical Review A* 54.2 (1996), p. 1593.
- [125] E Higurashi, Renshi Sawada and T Ito. ‘Optically induced angular alignment of trapped birefringent micro-objects by linearly polarized light’. In: *Physical Review E* 59.3 (1999), p. 3676.
- [126] Peter Banzer et al. ‘The photonic wheel: demonstration of a state of light with purely transverse angular momentum’. In: *arXiv preprint arXiv:1210.1772* (2012).
- [127] Kevin J Mitchell et al. ‘High-speed spatial control of the intensity, phase and polarisation of vector beams using a digital micro-mirror device’. In: *Optics Express* 24.25 (2016), pp. 29269–29282.
- [128] Brian H Cumpston et al. ‘Two-photon polymerization initiators for three dimensional optical data storage and microfabrication’. In: *Nature* 398.6722 (1999), pp. 51–54.

- [129] B. Carrasco and J. García de la Torre. ‘Hydrodynamic properties of rigid particles: comparison of different modeling and computational procedures’. In: *Biophysical Journal* 76.6 (1999), pp. 3044–3057.
- [130] A Ortega and J Garcia de la Torre. ‘Hydrodynamic properties of rodlike and disklike particles in dilute solution’. In: *The Journal of chemical physics* 119.18 (2003), pp. 9914–9919.
- [131] Jens-Christian Meiners and Stephen R Quake. ‘Direct measurement of hydrodynamic cross correlations between two particles in an external potential’. In: *Physical Review Letters* 82.10 (1999), p. 2211.
- [132] Eric R Dufresne et al. ‘Hydrodynamic coupling of two Brownian spheres to a planar surface’. In: *Physical Review Letters* 85.15 (2000), p. 3317.
- [133] Hilding Faxén. ‘Der Widerstand gegen die Bewegung einer starren Kugel in einer zähen Flüssigkeit, die zwischen zwei parallelen ebenen Wänden eingeschlossen ist’. In: *Annalen der Physik* 373.10 (1922), pp. 89–119.
- [134] Unè G Bütaitè et al. ‘Indirect optical trapping using light driven micro-rotors for reconfigurable hydrodynamic manipulation’. In: *Nature Communications* 10.1 (2019), pp. 1–10.
- [135] Jonathan Leach et al. ‘Interactive approach to optical tweezers control’. In: *Applied Optics* 45.5 (2006), pp. 897–903.
- [136] Vincent Ricardo Daria et al. ‘Arbitrary multisite two-photon excitation in four dimensions’. In: *Applied Physics Letters* 95.9 (2009), p. 093701.
- [137] R Bowman et al. ‘Optimisation of a low cost SLM for diffraction efficiency and ghost order suppression’. In: *The European Physical Journal Special Topics* 199.1 (2011), pp. 149–158.
- [138] Rupa Sarkar and Valentin V Rybenkov. ‘A guide to magnetic tweezers and their applications’. In: *Frontiers in Physics* 4 (2016), p. 48.
- [139] Zhou Ye and Metin Sitti. ‘Dynamic trapping and two-dimensional transport of swimming microorganisms using a rotating magnetic microrobot’. In: *Lab on a Chip* 14.13 (2014), pp. 2177–2182.
- [140] Steven L Neale et al. ‘Optoelectronic Tweezers (OET) trap stiffness with HeLa cells’. In: *Optical Trapping and Optical Micromanipulation V*. Vol. 7038. International Society for Optics and Photonics. 2008, 70381K.
- [141] Michael A Taylor et al. ‘Enhanced optical trapping via structured scattering’. In: *Nature Photonics* 9.10 (2015), p. 669.
- [142] Michael A Taylor. ‘Optimizing phase to enhance optical trap stiffness’. In: *Scientific Reports* 7.1 (2017), p. 555.
- [143] Michael Mazilu. ‘Eigenmode beam optimisation for optical micro-manipulation’. In: *arXiv preprint arXiv:1908.01179* (2019).

- [144] Eugene P Wigner. ‘Lower limit for the energy derivative of the scattering phase shift’. In: *Physical Review* 98.1 (1955), p. 145.
- [145] Felix T Smith. ‘Lifetime matrix in collision theory’. In: *Physical Review* 118.1 (1960), p. 349.
- [146] Philipp Ambichl et al. ‘Focusing inside disordered media with the generalized Wigner-Smith operator’. In: *Physical Review Letters* 119.3 (2017), p. 033903.
- [147] Michael Horodyski et al. ‘Optimal wave fields for micromanipulation in complex scattering environments’. In: *Nature Photonics* (2019), pp. 1–5.
- [148] Stefan Rotter, Philipp Ambichl and Florian Libisch. ‘Generating particlelike scattering states in wave transport’. In: *Physical Review Letters* 106.12 (2011), p. 120602.
- [149] Stefan Rotter and Sylvain Gigan. ‘Light fields in complex media: Mesoscopic scattering meets wave control’. In: *Reviews of Modern Physics* 89.1 (2017), p. 015005.
- [150] Shanhui Fan and Joseph M Kahn. ‘Principal modes in multimode waveguides’. In: *Optics Letters* 30.2 (2005), pp. 135–137.
- [151] Joel Carpenter, Benjamin J Eggleton and Jochen Schröder. ‘Observation of Eisenbud–Wigner–Smith states as principal modes in multimode fibre’. In: *Nature Photonics* 9.11 (2015), p. 751.
- [152] Wen Xiong et al. ‘Spatiotemporal control of light transmission through a multimode fiber with strong mode coupling’. In: *Physical Review Letters* 117.5 (2016), p. 053901.
- [153] Benoît Gérardin et al. ‘Particlelike wave packets in complex scattering systems’. In: *Physical Review B* 94.1 (2016), p. 014209.
- [154] Eliakim H Moore. ‘On the reciprocal of the general algebraic matrix’. In: *Bull. Am. Math. Soc.* 26 (1920), pp. 394–395.
- [155] Roger Penrose. ‘A generalized inverse for matrices’. In: *Mathematical proceedings of the Cambridge philosophical society*. Vol. 51. 3. Cambridge University Press, 1955, pp. 406–413.
- [156] Louis Lorenz. *Lysbevægelsen i og uden for en af plane Lysbølger belyst Kugle*. Vol. 6. 6. Det Kongelige Danske Videnskabernes Selskabs Skrifter, 1890, 1D62.
- [157] Gustav Mie. ‘Contribution to the optical properties of turbid media, in particular of colloidal suspensions of metals’. In: *Ann. Phys.(Leipzig)* 25 (1908), pp. 377–452.
- [158] Julius Adams Stratton. *Electromagnetic theory*. Vol. 33. John Wiley & Sons, 2007.
- [159] Jonathan M Taylor. *Optical Binding Phenomena: Observations and Mechanisms*. Springer Science & Business Media, 2011.

- [160] Warren J Wiscombe. ‘Improved Mie scattering algorithms’. In: *Applied Optics* 19.9 (1980), pp. 1505–1509.
- [161] Craig F Bohren and Donald R Huffman. *Absorption and scattering of light by small particles*. John Wiley & Sons, 2008.
- [162] Daniel W Mackowski. ‘Calculation of total cross sections of multiple-sphere clusters’. In: *JOSA A* 11.11 (1994), pp. 2851–2861.
- [163] JM Taylor and GD Love. ‘Multipole expansion of Bessel and Gaussian beams for Mie scattering calculations’. In: *JOSA A* 26.2 (2009), pp. 278–282.
- [164] JP Barton, DR Alexander and SA Schaub. ‘Theoretical determination of net radiation force and torque for a spherical particle illuminated by a focused laser beam’. In: *Journal of Applied Physics* 66.10 (1989), pp. 4594–4602.
- [165] B Richards and E Wolf. ‘Electromagnetic diffraction in optical systems, II. Structure of the image field in an aplanatic system’. In: *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* 253.1274 (1959), pp. 358–379.
- [166] Richard A Waltz et al. ‘An interior algorithm for nonlinear optimization that combines line search and trust region steps’. In: *Mathematical Programming* 107.3 (2006), pp. 391–408.
- [167] Richard H Byrd, Jean Charles Gilbert and Jorge Nocedal. ‘A trust region method based on interior point techniques for nonlinear programming’. In: *Mathematical Programming* 89.1 (2000), pp. 149–185.
- [168] Wil O.C. Ward. *icosphere*. <https://www.mathworks.com/matlabcentral/fileexchange/50105-icosphere>. Accessed: 2020-04-01. 2015.
- [169] Richard Bowman et al. ‘Position clamping in a holographic counterpropagating optical trap’. In: *Optics Express* 19.10 (2011), pp. 9908–9914.
- [170] Mike Woerdemann, Konrad Berghoff and Cornelia Denz. ‘Dynamic multiple-beam counter-propagating optical traps using optical phase-conjugation’. In: *Optics Express* 18.21 (2010), pp. 22348–22357.
- [171] Jeffrey A Davis et al. ‘Encoding amplitude information onto phase-only filters’. In: *Applied Optics* 38.23 (1999), pp. 5004–5013.
- [172] Luiz Gonçalves Neto, Danny Roberge and Yunlong Sheng. ‘Full-range, continuous, complex modulation by the use of two coupled-mode liquid-crystal televisions’. In: *Applied Optics* 35.23 (1996), pp. 4567–4576.
- [173] Volker Bormuth et al. ‘Optical trapping of coated microspheres’. In: *Optics Express* 16.18 (2008), pp. 13831–13844.
- [174] Ying Hu et al. ‘Antireflection coating for improved optical trapping’. In: *Journal of Applied Physics* 103.9 (2008), p. 093119.

- [175] Karel Svoboda and Steven M Block. ‘Optical trapping of metallic Rayleigh particles’. In: *Optics Letters* 19.13 (1994), pp. 930–932.
- [176] Steven M Block, Lawrence SB Goldstein and Bruce J Schnapp. ‘Bead movement by single kinesin molecules studied with optical tweezers’. In: *Nature* 348.6299 (1990), pp. 348–352.
- [177] Karel Svoboda et al. ‘Direct observation of kinesin stepping by optical trapping interferometry’. In: *Nature* 365.6448 (1993), pp. 721–727.
- [178] Mathai Mammen et al. ‘Optically controlled collisions of biological objects to evaluate potent polyvalent inhibitors of virus-cell adhesion’. In: *Chemistry & Biology* 3.9 (1996), pp. 757–763.
- [179] Michelle D Wang et al. ‘Force and velocity measured for single molecules of RNA polymerase’. In: *Science* 282.5390 (1998), pp. 902–907.
- [180] Keir C Neuman and Attila Nagy. ‘Single-molecule force spectroscopy: optical tweezers, magnetic tweezers and atomic force microscopy’. In: *Nature Methods* 5.6 (2008), p. 491.
- [181] David B Phillips et al. ‘Optimizing the optical trapping stiffness of holographically trapped microrods using high-speed video tracking’. In: *Journal of Optics* 13.4 (2011), p. 044023.
- [182] Weronika Lamperska et al. ‘Micro-Dumbbells – A Versatile Tool for Optical Tweezers’. In: *Micromachines* 9.6 (2018), p. 277.
- [183] Jannis Köhler et al. ‘Optical screw-wrench for microassembly’. In: *Microsystems & Nanoengineering* 3.1 (2017), pp. 1–7.
- [184] David B Phillips et al. ‘Force sensing with a shaped dielectric micro-tool’. In: *EPL (Europhysics Letters)* 99.5 (2012), p. 58004.
- [185] Mark Jayson Villangca et al. ‘Light-driven micro-tool equipped with a syringe function’. In: *Light: Science & Applications* 5.9 (2016), e16148–e16148.
- [186] Philipp Ambichl et al. ‘Super- and anti-principal-modes in multimode waveguides’. In: *Physical Review X* 7.4 (2017), p. 041053.